UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Using Brain-Machine Interfaces to Study Motor Cortical Population Activity

Permalink
https://escholarship.org/uc/item/380070zf

Author
Khanna, Preeya

Publication Date
2017

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Using Brain-Machine Interfaces to Study Motor Cortical Population Activity

by

Preeya Khanna

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Joint Doctor of Philosophy
with University of California, San Francisco

in

Bioengineering

in the

Graduate Division
of the
University of California, Berkeley

Committee in charge:

Professor Jose M. Carmena, Chair
Professor Richard Ivry
Professor Robert Knight
Professor Joshua Berke

Fall 2017

Using Brain-Machine Interfaces to Study Motor Cortical Population Activity

Copyright 2017
by
Preeya Khanna

Abstract

Using Brain-Machine Interfaces to Study Motor Cortical Population Activity

by

Preeya Khanna
Joint Doctor of Philosophy
with University of California, San Francisco in Bioengineering

University of California, Berkeley

Professor Jose M. Carmena, Chair

Motor actions constitute the way in which we interact with the world, and are driven by millions of neurons in our distributed motor system. Studying how patterns of activity in motor cortical populations of neurons give rise to withholding movement, generating fast and accurate movements, and generating sequences of movements is the topic of this thesis. One challenge in studying how patterns of population activity support features of movement is that population patterns are difficult to manipulate in experiments with current neuroscience techniques. Current methods allow for anatomically-specific and cell-specific activation or inhibition, but do not, for example, allow for the careful manipulation of correlated versus uncorrelated activity pattern. We turn to closed-loop brain-machine interfaces as tools to perturb population activity patterns and to study the consequences of these perturbations on motor behavior.

The first part of this thesis focuses testing how tightly linked a specific feature of motor cortical local field potential signals are with withholding of movement. We use a non-human primate model system where subjects learn to control this neural feature, termed beta band oscillations, through a closed-loop brain-machine interface. Subjects perform tasks where they volitionally bring their internal beta band oscillatory state to a specified level, and immediately afterwards perform a motor task. The sequential task design allows for testing how tightly linked beta band oscillations are to movement onset, more so than can be claimed by correlational studies. We use a similar approach to investigate the relationship between beta band oscillations and movement in parkinsonian subjects.

Sequential task designs shed light on how a specific neural signal contributes to a feature of natural movement. Studying the complete link between cortical neural signals and natural movements, however, is challenging given i) experimenters' limited access to neural signals driving movement ii) the number of non-linearities in the neural to movement map, iii) the challenge in fully capturing a complete picture of natural movements. One approach to simplifying the problem of studying sensorimotor control is to study control of a fully characterized virtual plant, such as a 2D velocity-controlled cursor, that is controlled neurally through an experimenter-defined transform and with observed neu-

ral activity patterns. Such systems have the advantage of allowing the experimenter to define mathematically which types of population activity influence the movement of the plant. We take advantage of this feature to investigate how different decompositions of population activity support fast versus accurate movements. Finally, we use this system to study principles of how neural population activity is generated for different orderings of cursor movements, or action sequences. We find that for action sequences that are constituted from the same commands but in a different ordering, subjects have different ways of generating the same command. We test how large these differences are by decomposing the population activity that updates the movement of the plant and assess how cursor movements are influenced. With this approach, we describe a model of how neural activity is generated that captures a majority of the neural variance observed across different action sequences.

As the motor systems neuroscience field increasingly collects simultaneously acquired population neural activity, hypotheses about how features of the population support movement will continue to emerge. Testing these hypotheses will require manipulation of possibly abstract population decompositions, a challenging feat to do precisely with current stimulation methods, but possible with closed-loop brain-machine interfaces.

Dedicated to:
Sebastian, Jeeves, Cartman, Keller, Gromit, Obi-Wan, Yoda
BRPD03, BRPD04, BRPD05, BRPD06, and
HUD1

Contents

Publications and Presentations Resulting from Thesis

Chapter 2+3

- **Khanna P.**, and Carmena J.M. (2017) "Beta Band Oscillations Drive Population Signals that Inhibit Movement in the Motor System." *ELife* doi: 10.7554/eLife.24573.
- **Khanna P.** and Carmena J.M. (2015). "Neural oscillations: beta band activity across motor networks." Current Opinion in Neurobiology, 32: 60-67.
- **Khanna P.** and Carmena J.M. (2015). "Changes in Reaching Reaction Times Due to Volitional Modulation of Beta Oscillations." IEEE Neural Engineering Conference, Montpelier (France).
- **Khanna P.**, and Carmena J.M. (2015). 'Effects of Volitional Modulation of Beta Oscillations on Reaching Tasks.' Society for Neuroscience annual meeting, Chicago IL (poster)

Chapter 4

- **Khanna P.**, Swann N. C., Hemptinne C., Miller, A., Starr P. A., and Carmena J.M. (2016). "Volitional Control of Beta Band Power Using the Medtronic Activa PC + S and Nexus-D Streaming." *IEEE Transactions on Neural Systems and Rehabilitation Engineering.* doi: 10.1109/TNSRE.2016.2597243
- **Khanna P.**, Swann N.C., Hemptinne C.d., Miocinovic, S., Miller, A., Starr P.A., and Carmena J.M. (2016). "Volitional Control of Beta Band Power Using the Medtronic Activa PC + S and Nexus-D Streaming." International BCI meeting, Pacific Grove, CA (poster).
- **Khanna P.**, Stanslaski S., Xiao Y., Ahrens T., Bourget D., Swann N., Starr P., Carmena JM., Denison T. (2015). "Enabling Closed-Loop Neurostimulation with Downloadable Firmware Upgrades." IEEE Biomedical Circuits And Systems Conference, Atlanta

- **Khanna P.**, Swann N.C., Starr P.A., and Carmena J.M. (2017). "Effects of Neurofeedback Control of Beta Band Oscillations in Motor Cortex on Finger Tapping in PD Patients." Society for Neuroscience annual meeting, Washington DC (poster).

Chapter 5

- **Khanna P.**, Athalye V.R., Carmena J.C. (2016) "Leveraging Emergent Coordinated Neural Dynamics in Improving Brain-Machine Interface Control" *IEEE Engineering in Medicine and Biology Conference*, Orlando (talk)

Chapter 6

- **Khanna P.**, Athalye V. R., Costa R. M., Carmena J. M. (2017). "Distinct neural encoding schemes emerge for actions generated by the same effector." Computational and Systems Neuroscience, Salt Lake City UT (poster).

- **Khanna P.**, Athalye V. R., Gowda S., Costa R. M., Carmena J. M. (2016). "Distinct subspaces emerge in neuroprosthetic control during different tasks." Society for Neuroscience annual meeting, San Diego CA (poster).

Other

- You, A., Athalye, V. R., Gowda, S. **Khanna P.**, Moorman, H. G., and Carmena J.M. (2017). 'Co-modulation of joint angles in control of a kinematically redundant brain-machine interface.' Society for Neuroscience annual meeting, Washington DC, (poster)

- Summerson S.R., **Khanna P.**, Rich, E.L, Wallis, J.D., and Carmena J.M. (2015). 'Stimulation in primate caudate nucleus modulates action selection in probabilistic reward task.' Society for Neuroscience annual meeting, Chicago IL (poster)

- **Khanna P.**, So K., and Carmena J.M. (2012). "Volitional phase control of neural oscillations using a brain-machine interface". IEEE Neural Engineering Conference, San Diego (poster)

Acknowledgments

When I began graduate school, I didn't know much but I knew I was excited to study function in healthy brains and use engineering tools to improve the condition of unhealthy ones. Luckily for me I had decided to study at UC Berkeley-UCSF where I quickly found a community of engineers, neuroscientists, and clinicians studying health and neurological disease at the scale of individual neurons, brain regions, full brains, and complex human behaviors – plenty to fuel my broad interests. Throughout my time in this community, my path has led me to interact and collaborate with colleagues who have fed my enthusiasm, humored my ideas, kindly mentored me, and shared in the inevitable struggles of performing new scientific research. Without these people helping me refine my broad interests into tractable investigations, and inspiring and encouraging me to persist, little of the work in this dissertation would be here and my continued excitement and desire to continue studying neural engineering may have vanished.

Firstly, I want to thank Jose M. Carmena, my mentor throughout this journey. Jose has an infectious enthusiasm for science and technology that all who interact with him know well. I always left meetings from his office feeling calmer and more excited about my project than going in. From the beginning, Jose gave me the freedom to pursue ideas I found most interesting, and then followed up this intellectual freedom that with fruitful discussions where he divulged experience and larger context to my results (or often non-results). Thank you also for recommending opportunities outside of UC Berkeley-UCSF to enrich my education whether that be a summer at Medtronic, a consultant job at Cortera, or a semester in Spain.

I also would like to thank my qualifying exam and thesis committee consisting of Richard Ivry, Robert Knight, Michel Maharbiz, Philip Sabes, and Joshua Berke for their your guidance throughout these last years. I especially thank Rich and Bob who have served on both committees. Rich, I thank you for your incredible knack for finding the temporal confounds in my analyses and suggesting controls I need to do. Bob, thank you for welcoming me into your lab – it has been my go-to resource for neural oscillations over the years.

Most of the work presented within this thesis was done in close collaboration with colleagues. Work from Chapter 4 was done in collaboration with Nicki Swann and Phil Starr at UCSF. Thank you to you both for taking interest in and supporting my ideas over these last three years. Nicki, thank you for patiently training me to work with patients, helping me navigate the Nexus-D, and for insightful discussions about the thalamo-cortico-basal ganglia loop in parkinsonian patients. You've been a role model to me over the years as a down-to-earth person who magically managed to balance your work in lab, applications for your future career, family, and still were a compassionate researcher who developed real connections with patients and their families.

Work from Chapter 5 and 6 were done in close collaboration with Vivek Athalye. Vivek, working with you has been one of the most enjoyable times in my graduate school journey. Whenever we're discussing new analyses, figures, or future research directions,

your creativity is unbounded, your communication is crystal clear, and your knowledge and understanding of relevant literature blows me away. Thank you for working with me and kindly guiding me into the world of studying sensorimotor control with brain-machine interfaces. Aside from our joint work together, thank you for being a close friend through the years. I will dearly miss our dancing to Matt Steffanina videos, Sandwich Zone lunches, and push-up breaks with Albert on the Barker lawn.

Lastly, unpublished here but constituting a majority of my last semesters is work on the ISMORE project in collaboration with University of Tuebingen researchers. Thank you to Andrea Sarasola and Nerea Irastorza for tolerating and patiently correcting my Spanish, for literally welcoming me into your respective houses and families in Donostia and Zarautz, and for making awesome teammates during the ISMORE project. Thank you also for being amazing friends who watch silly movies with me and make scrumptuous potlucks to enjoy together. Thank you to Leire Santisteban, Cris Amador, and Ana Bengoetxea for teaching me so much about the peripheral nervous system and being warm colleagues. Finally, thank you to Ander Ramos for leading the work on the insane amounts of paperwork to make this project a reality, and for continuing to deliver enthusiasm and support throughout the long four (plus!) months of the trial. Thank you especially to Ander for the reminder that at the end of the day, if we seek to make the lives of patients better, scientific papers are are the beginning of the process, not the end. Your leadership and dedication to developing the skills necessary to translate experimental therapies to reality is inspiring.

In addition to the above collaborators, I've benefited immensely from the guidance, wisdom, and camaraderie of the graduate students and post-docs in the Carmena lab and others throughout the years. To the wiser generation of Amy Orsborn, Maryam Shanechi, Erin Rich, Simon Overduin, and Aaron Koralek, thank you for welcoming me into the lab and continuing to be mentors, role models, and friends. Amy, thank you for your guidance on monkey-training, baking tips and equipment, and career advice. Aaron, thank you for hiking and camping with me, teaching me the true utility of a frisbee (fanning the fire), and hilarious subject-less emails. Erin, thank you for introducing me to primate research, tolerating my early experimental stubbornness, and continuing to be a friend and sounding board throughout the years. Thank you to Kelvin So and Sid Dangi for directly training me in primate handling, and for humoring my initial research ideas. Along with Vivek, I'll always remember our initial days in our Sutardja Dai office and my utter failure at office basketball. Thank you to Helene Moorman and Suraj Gowda for making our days in the basement of Li Ka Shing entertaining. Thank you for drinking beer when experiments didn't work, and also when they did! Thank you to Ryan Neely and Samantha Summerson for being my longest colleagues through my years in Berkeley. Ryan, as safety officer, thank you for always having my back when I committed lab safety violations – even ones caught on camera. Sam, thank you for training me on new procedures and being patient as I fumbled my way around. I literally can never believe how you manage to have so many projects going on at once, yet you still manage to be a thoughtful friend (always remembering birthdays!) and kind mentor

ready to edit abstract and posters or discuss experiments. Finally, thank you to the new generation of Tanner Dixon, Albert You, David Piech, and Nuria Vendrell Llopis! Thank you to Albert for being my monkey-training partner. You love and respect of animals is enormous and an attribute all animal researchers can learn from. Thank you for teaching me how to play badminton and paint my bike!

Tool-building and grunt work doesn't always get the credit it deserves, but is especially present in primate research. I especially want to especially thank James Gao, Helene, Suraj, and Sid who spent large portions of their PhDs writing an open source software platform that today is used not just for my experiments, but for countless in the lab. Thank you for writing the code with the intention of developing a general tool. Thank you to Suraj for documenting it beautifully, and for being available to answer questions or confusions about it past your due time. Thank you also to Sid and Suraj for their extension of the codebase developed specifically for the ISMORE project. This code set up me, Andrea, and Nerea to build the experiments we are using in the ISMORE clinical trial today. In addition to software tools, this thesis contains a number of primate experiments. Primate work comes with the unglamorous initial process of training a primate to sit in a chair and to engage in behavioral tasks. This process can take from months to years of daily work. Thank you to Amy, Helene, Suraj, Alejandra Dominguez, Kelvin, and Sid for all of your work in training the primates used in the work in this thesis. I was extremely lucky and saved enormous time in my PhD by being able to inherit well-trained animals and set-up experimental rigs from your hard work. Thank you.

One benefit of living in the Bay Area, besides access to the amazing neural engineering community, is the always sunny weather that is perfect for outdoor activities! Thank you to my friends in the BioEngineering Association of STudent (the BEAST!) and beyond graduate school for playing outside with me! Whether intramural soccer, rock climbing, bike rides, or hikes, I'm thankful for being surrounded by such a lively, athletic, well-balanced group of friends. Thank you to my many frisbee teammates in the Bay Area Disc Association and on FAB ultimate for continuing to foster my love of ultimate frisbee. Thank you to all the roommates at 2530 Grant Street for being wonderful friends and making our apartment a warm home. And thank you for all my colleagues at the Prison University Project for encouraging and supporting the math classes I taught and tutored in. Without you all I would be less joyful and enthusiastic.

Finally, the last thanks goes to my parents, Bakul and Chander Khanna. Your commitment to me and my education has not wavered from day one. It began with the sacrifices you made when I was young to drive me to soccer practice, continued with your persistence in getting me to wake up on time in high school (7:24...7:25....), and is still present in our long phone conversations. Your commitment to a better life in America, and the struggle associated with making it here, has given me opportunities and freedoms that you didn't have the luxury of experiencing yourself. Still, you've never pressured me into doing or being anything but myself, and have not given me anything but support, love, and positivity regardless of whether I'm soaring high or am grumpy. Thank your

for supporting me despite the distance we've lived over the last ten years. Thank you for always being there when I need to call (literally, Dad). And thank you for being the kind, silly, hard-working, and honest role models that I have always looked up to.

I would not be here, finishing at UC Berkeley - UCSF, if it were not for you all.

Chapter 1

Introduction

1.1 The Distributed Motor System

Whether performing simple, necessary activities of daily living, expressing ourselves through language, or performing exquisite feats of strength and precision in professional athletics, the ability to actuate our limbs and vocal cords is the basis of most of our lives. While the motor systems neuroscience field has a detailed anatomical description of the distributed motor system including the motor areas of cortex, subcortical structures, brain stem, cerebellum, spinal cord, and muscles, the concepts of how and where actions are learned, stored, and initiated from is not understood.

1.1.1 Anatomy

We begin an older view of how the motor system generates movement, and although parts have been debated and refined over the last several decades, it serves a useful starting point. Imagine attempting a habitual action in response to a visual stimulus, say reaching for a cup of coffee. First, the retina receives visual information about the object of interest. This information travels down the optic nerve, through the thalamus to visual cortex where the features of the visual scene are parsed. The information then is thought to travel according to the dorsal-visual-stream, or 'vision-for-action' stream [1, 2]. It reaches higher order visual areas and parietal cortex where the goal is recognized and evidence of a decision to initiate an action can be decoded [1]. This decision or plan is then communicated to premotor cortices (see 1.1 for the macaque motor areas), where high-level motor planning and preparation is thought to take place. These commands then are sent to motor cortex where high-level command are translated to muscle commands and then muscle commands travel down the spinal cord to the muscles via the corticospinal pathway. Thus, motor cortex has been labeled the "final common pathway" for cortically-controlled movement [3]. If we instead perform an internally-generated action, activation of supplementary motor areas (SMA) and pre-SMA will drive motor cortical activity [1].

There are many lines of evidence contradicting the above generalization. In the above description, the spinal cord is described only as a path for transmitting motor cortical commands to the musculature, ignoring complex spinal-level computations including reflexes [5], central pattern generators for rhythmic movements [6], and even modulation of spinal cord segments during preparation of an upcoming movement [7, 8]. Further, in addition to the corticospinal tract, there are contributions to movement from spinal cord tracts that originate in the midbrain such as the rubrospinal tract (originating in the red nucleus), the vestibulospinal tract (originating in the vestibular nuclei), and reticulospinal tract (originating in the reticular nucleus) [9]. These tracts also contribute to movement related modulations. For example, red nucleus neurons projecting down the rubrospinal tract have been shown to encode inter and intra limb coordination [10] and the reticulospinal tract has been shown to contribute to distal hand function [11, 12].

Another modification of the described model is the identification of many separate areas of premotor cortices, shown in 1.1, and the finding that many of them with have direct

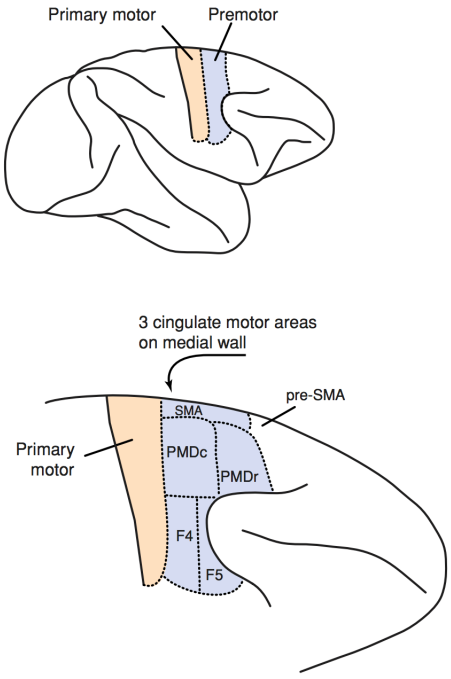

Figure 1.1: Pre-motor and motor cortical areas in a macaque monkey, from [4]

corticospinal projections [3], challenging the notion that primary motor cortex is the "final common pathway". When performing retrograde mapping from spinal cord segments that control arm motions, approximately 50% of neurons originated in primary motor cortex, but about 12-20% originate in SMA, 15-20% originate in the cingulate motor areas (CMAs). It is unclear exactly which areas contribute to which types of movements through their corticospinal connections, but lesions and stimulation studies provide some insight. Standard microelectrode stimulation in M1 evokes grasping movements, supporting its contribution to distal control [13]. Long trains of electrical stimulation of M1 give rise to reaching movements or movements resembling manipulation of objects. On the other hand, stimulation in dorsal and ventral premotor cortices (PMd, PMv) give rise to more complex hand-to-mouth movements or defensive movements [14]. Stimulation in supplementary motor area (SMA) gives rise to primarily proximal movements and combinations of proximal and distal movements, but requires more electrical current to evoke movement than primary motor cortex [15]. Further, when primary motor cortex is removed, SMA stimulation only triggers proximal movements [3]. Stimulation in pre-SMA rarely evokes movements, but lesions affecting pre-SMA, but not SMA, give rise to a deficit in inhibiting competing motor plans [16]. Finally, the cingulate motor areas (CMAs) exhibit segmentation by area as well, breaking into a rostral, dorsal, and ventral subsection. Stimulation of the rostral section does not elicit movement, whereas stimulation of the dorsal and ventral sections do evoke movement [17, 18].

In addition to evoking different movements when electrically stimulated, premotor and primary motor cortex are also innervated by different parietal areas [2], and have different connectivity with the basal ganglia and thalamus [19]. Different premotor areas also have different degrees of projection to primary motor cortex [3]. This structural organization suggests that instead of viewing premotor cortices as a level above primary motor cortex in the hierarchy of motor processing, it may be more appropriate to view primary motor cortex on an equal level as the premotor cortices [3], organized in parallel loops each of which has a different contribution to movement [14]. Further evidence for parallel motor loops controlling different aspects of movement can be found by studying neural activations from each area. Primary motor cortex has a high representation of hand-tuned neurons. PMd is strongly connected to parietal area VIP, which typically associated with visual information in the dorsal stream. PMd activity is high in stimulus-response actions. A class of PMv neurons called 'mirror neurons' exhibit activity during observation of movements. SMA and pre-SMA neurons tend to exhibit activity prior to movement [17]. These findings point to involvements of premotor cortices and motor cortices in the mental rehearsal, planning, and direct control of movement.

Overall, the motor system's structure and hierarchy is far more complex than the simple description given at the beginning, just based on structure and connectivity alone. To probe the functional connectivity, coordination, and generation of movement, measurement methods varying in spatial and temporal scale are used.

1.1.2 Electrophysiology Measurements

Many types of measurements can be made from cortical, subcortical, spinal cord, and musculature circuitry that yield a insight into the computations performed in each area by single neurons, ensemble of neurons, or brain regions. Here we limit our discussion to electrophysiological approaches, though note that many new measurement methods are emerging that have high temporal and spatial resolution. One of these methods is magnetoencelography (MEG), a non-invasive method that captures changes in magnetic field due to the electrical activity of neurons, and has a high temporal resolution of ~1ms and spatial resolution of 2-3 mm. Another method used in animal models uses voltage-sensitive dyes (VSDs) which are proteins that bind to the membrane of cells and change optical properties with voltage. These changes in voltage can be measured with high-resolution fast-speed digital cameras positioned over a window where the skin, skull, and dura have been removed [20]. Both MEG and VSDs are promising measurement methods that will facilitate our understanding of motor function.

Currently, the four primary electrophysiological recording modalities that are used in awake and behaving subjects are electroencelography (EEG), electrocorticography (ECoG), local field potentials (LFPs), and extracellular single-neuron action potential recordings (single and multi-units). These four methods vary in invasiveness of method and spatial scale of recording, with smaller spatial scale requiring the most invasive methodology as shown in 1.2. The EEG method involves placing electrodes outside of

the skull and measuring signals that are summed contributions from over ~10 cm swaths of cortex and have been low-pass filtered due to passing through the hard tissue of the skull. Benefits of the EEG method are the non-invasiveness. Drawbacks include the extensive spatial averaging due to the electrodes being so far from the signal source and the day-to-day changes in electrode positioning and impedance. The ECoG method involves placing electrodes on the surface of the cortex, below the skull, and thus records activity from within ~0.5 cm or less, with the exact spatial spread depending on the size of the electrode contact. Depending on the preparation, ECoG grids can be in place chronically or temporarily. ECoG contacts are much closer to the signal source than EEG contacts, yielding a more spatially specific, and less attenuated signal. While EEG and ECoG are biased towards sampling cortical activity in superficial layers [21], recording LFP signals requires penetrating cortex with microelectrodes and can yield signals from deeper cortical layers (such as layer V, the layer containing corticospinal projecting cells in motor areas). The small microelectrode recording tip yields signals from a smaller population of ~ 1mm radius, though the actual spatial extent relies on the neuron morphology, the distribution of synapses along the neuron, and the correlations present in synaptic activity of neurons within the recorded population [22]. If synapse activity is uncorrelated, the LFP represents cells within a radius of a few hundred micrometers. High frequency activity from LFP signals ($>$ 80 Hz, often termed 'high-gamma activity') has been shown to be correlated with spiking activity [23, 24], whereas lower frequency activity (0 Hz - 100 Hz) can be attributed to a variety factors including recording electrode position, frequency content of synaptic inputs, and neuron morphology [25]. Finally, extracellular recordings can be acquired with the same method used to collect LFP signals. Instead of examining the low-pass filtered signal, neural activity is high-pass filtered between 250 Hz - 10 kHz and if lucky, one or more characteristic action potential waveforms can be identified as standing out from the baseline high-pass filtered activity [26]. Even if a high SNR action potential cannot be picked out, recently groups have characterized "threshold-crossing" activity by setting a static threshold on the high-pass filtered signal, and using the crossings of this threshold as analogs for population spiking activity regardless of whether this activity results from more than one, easily isolatable neuron [27].

In this thesis, we study LFP signals, ECoG signals, and simultaneously recorded populations of single and multi units to garner how neural ensembles coordinate to accomplish motor tasks.

1.2 Coding in the Motor System

1.2.1 Representation of Movement

Examining how neural activity relates to movements throughout the motor system may yield insight into principles of how the distributed motor system generates actions. Early work in studying neural coding was based on the idea that motor intentions are represented

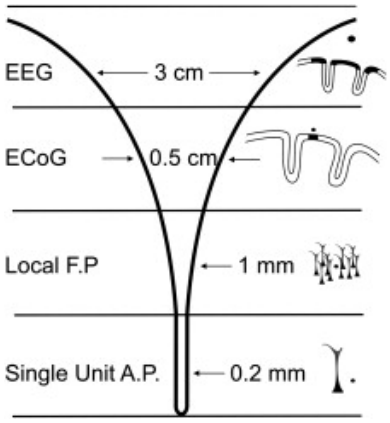

Figure 1.2: Figure from [28] comparing the spatial scale of recording by EEG, ECoG, LFPs, and single-unit action potentials.

Figure 1.3: Figure from [29] describing a series of sensorimotor transforms that could characterize the computations in visual, parietal, premotor, and motor cortices to generate an action in response to a presented target.

throughout the distributed motor system, but in different coordinate frames. For example, posterior parietal cortex represents goal information in retina-centric coordinates, whereas premotor cortices represents information in terms of hand kinematics, and primary motor cortex represents joint torques [29]. An example of a cascade of sensorimotor transforms is shown in 1.3.

A challenge to this approach is that the issue of which coordinate frame M1 represents action within has remained largely unresolved [30, 31]. There have been proposals that M1 encodes hand direction, hand velocity, joint angle velocity, joint torques, muscle activity (reviewed in [30]). Many of these studies have employed electrophysiological recordings of single units to make these inferences, but even simultaneously recording from a full population does not assist in resolving the question. One issue in discriminating which variables are most reliably encoded is that during standard reaching tasks, many of the proposed variables are highly correlated.

1.2.2 Generation of Movement

A different approach to studying the generation of action is leave the question of which movement parameters the motor cortex represents and instead focus on identifying principles of how movement is generated. Early work from Fetz demonstrated that an artificial neural network with input units, hidden units, and output units could be trained to produce specific EMG activity in response to an input signal [32]. Fetz noted that there were multiple coding schemes that allowed combinations of hidden units to generate the correct EMG output signal, including schemes that would result in a poor correlation between the hidden unit and trained output EMG patterns explicitly. Since Fetz, other groups have embarked on studying properties of populations of single and multi-unit activity and developed proposals of how these properties generate movement.

One group has modeled simultaneously recorded populations of single and multi-unit activity recorded from motor cortical areas during very well-practiced, fast arm reaches, as observations from a low-dimensional linear dynamical system. Here, the activity of tens to hundreds of single and multi-units are well-summarized by only a few dimensions (low dimensional population state), where the low-dimensional observations abide by lawful temporal rules [33, 34]. When subjects initiate their reaches, the low dimensional population state starts at a reach-specific initial position and temporal rules evolve the population state to produce a pattern. This temporal pattern has the correct frequency content to produce muscle electromyography (EMG) activity [33]. This model of how neural activity evolves captures substantial neural variance, and explains specific characteristics of population activity that other typical representation models do not capture. Whether a consistent dynamical system can also explain complex behaviors such as feedback corrected reaches, or delayed reaches is beginning to be explored [35, 36, 37].

1.2.3 Studying Principles of Movement Encoding

In this thesis, we aim to study questions about how population activity within motor cortex codes for withholding movement, generating fast movement, and generating action sequences. One approach to these questions is to design behaviors that capture our questions of interest (e.g. moving vs. non-moving, fast vs. slow movements, sequences vs. single actions), record neural data during these ongoing behaviors, and look for patterns in the recorded neural data that may support the behavior. This approach has been fruitful for countless studies, and yielded insights that form the basis of our theories today. One challenge with the approach though, is if one desires to causally show how a specific pattern of neural activity relates to the behavior of interest, this will be very difficult unless the identified neural pattern of interest is gross activation or gross inhibition of a large swath of brain tissue. Below, we review Brain Machine Interfaces (BMIs), a tool we will use throughout this thesis to try to use to make claims stronger than correlations.

1.3 Closed loop Brain Machine Interface (BMI) Systems

Closed loop brain machine interfaces (BMIs) are systems that map a subject's neural activity to the behavior of an effector. The subject then senses the effect their own neural activity has on the behavior of the effector, and can adjust their neural activity in response, closing the loop as in 1.4. In all closed-loop BMI systems there are four vital components, illustrated in 1.4. First, real-time recording or imaging of neural activity is used to estimate the state that a subject is in with respect to the BMI system. Many approaches exist for capturing neural activity from the brain, including electrophysiology at various scales as shown in 1.4, or various neuroimaging approaches. Activity from the peripheral nervous system can also be recorded as a BMI control signal via surface or intramuscular electromyography (EMG). Peripheral nerves can be recorded with cuff electrodes or more recently, miniature wireless motes [38]. Second, a mapping between neural activity and the BMI effector extracts relevant features from the neural activity, and dictates how the effector changes its behavior. In neuroprosthetic BMI systems, where the goal is to estimate a subject's motor intent, this neural activity to BMI effector mapping is called a decoder since the mapping "decodes" brain activity and translates it into the desired action. On the other hand, in neurofeedback training systems where the goal is to provide the subject feedback about their current brain state, the mapping is only designed to make the feedback interpretable and does not require estimations of subject intent. Third, the virtual or real effector changes behavior according to the output from the neural-to-effector mapping. Effectors can include virtual cursors and robots, real robots, muscle stimulation settings, brain stimulation settings, wheelchairs or other transporters, spellers for communication, and more. Finally, the subject receives feedback of the effector either visually, through somatosensory feedback, or through changes in stimulation. The feedback update rate can be rapid (200 Hz), or very slow (1 Hz) depending on the system and the goal.

1.3.1 BMIs for Clinical Applications

Closed loop BMIs appear in many clinical applications. Closed loop BMIs can be neuroprosthetic systems where users control virtual effectors such as cursors [39, 40, 41, 42, 43] or modern day spellers [44]. Combined with a modern predictive text completion, accessibility to the host of applications, games, and resources available through the Web, and paralyzed users controlling a cursor could communicate, read, and even play music [45]. Control of robotic arms has also been demonstrated with closed loop BMIs, allowing paralyzed subjects to feed themselves and perform other reach-to-grasp motions [46, 47]. Closed loop BMIs have also been used for control of wheelchair movements [43]. There have been very few studies where wearable exoskeletons have been controlled in multiple degrees of freedom as an assistive device, owing to the complexity of the robot-human

interaction. There have, however, been many studies investigating wearable exoskeletons in closed loop BMIs as rehabilitation tools [48, 49], and studies using cortically driven electrical stimulation of the natural arm [50, 51, 52]. Further, there are exoskeleton devices that can by controlled with muscle signals [53, 54, 55].

In addition to using BMIs to control virtual cursors and robots, closed loop BMIs can also be closed-loop brain stimulation systems [56, 57, 58] where brain signals or other biometric signals such as accelerometers are used to titrate the therapeutic stimulation level to avoid delivering too much or too little therapy. Closed-loop peripheral nervous system stimulation and closed-loop drug delivery systems also exist for conditions such as urinary incontinence, chronic pain, modulation of spinal cord circuits, and modulation of cortical activity [59, 60, 61, 62].

Finally, closed loop BMIs can be used for brain training applications. Online monitoring of brain activity during learning or gaming can allow for a closed loop training environment where neural signals contribute to progression through the material to be learned or the game [63, 64]. Alternatively, if a particular neural signal is known to be related to a desired behavior or state, subjects can perform neurofeedback tasks where they learn to volitionally increase or decrease that particular signal [65, 66, 67, 68, 69, 70].

Overall, closed loop BMIs thus have the promise of augmenting natural human function, assisting injured populations, contributing to rehabilitation, and enhancing existing open-loop therapies.

1.3.2 BMIs as Scientific Tools

In addition to closed loop BMIs offering promising clinical therapies, they also serve as valuable tools for scientific investigation. In this thesis, we use closed loop BMIs first as a tool to probe the relationship between specific neural activity patterns and natural behaviors, and second to study sensorimotor control principles in a simplified setting. Notable other scientific uses of closed loop BMIs are to study neural plasticity [71, 72, 73] and motor learning [74].

1.3.2.1 BMIs for Probing the Brain-Behavior Relationship

One challenge in classical neurophysiological experiments is that experimenters only observe a subset of the neural activity relevant for the generation of behaviors studied. Even full brain imaging methods suffer in temporal resolution and/or spatial resolution. Thus, the best an experimenter can do to explain how a particular behavior may arise from observed neural activity is to correlate the two observations – the brain activity and the behavior. While there are manipulations to silence, remove, or stimulate brain areas hypothesized to be related to ongoing behavior, it is difficult to do so precisely. In this thesis we consider relationships between specific decompositions of population neural activity and behavior, and few tools exist to perturb only specific decompositions of neural activity. For example, we consider oscillations in a particular frequency band

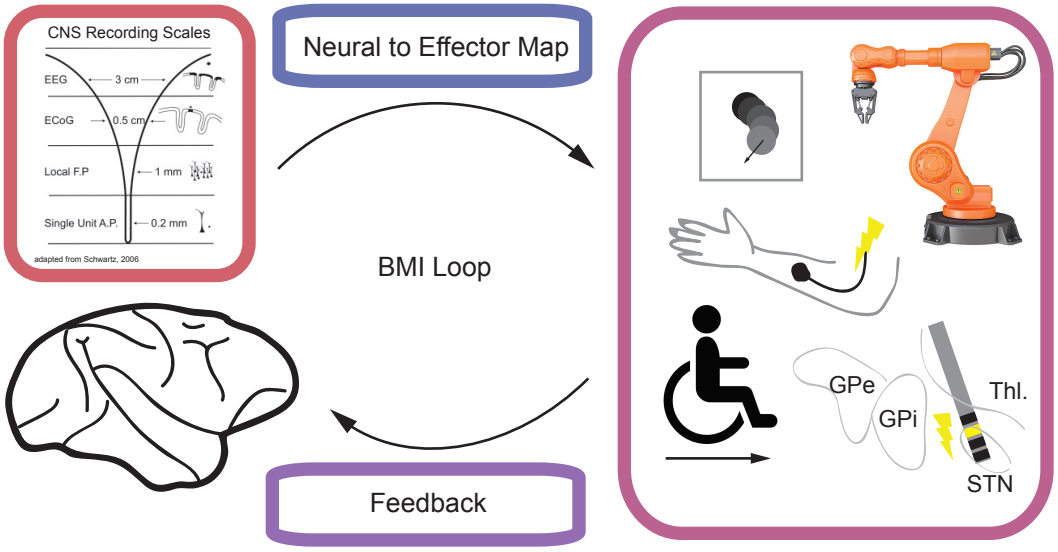

Figure 1.4: A schematic demonstrating the four central components of a closed loop BMI. In red is the neural recording modality pictured as varying scales of electrophysiological measurements, but which also can include non-invasive brain imaging methods, optical sensing of voltage sensitive dyes, and recording from the peripheral nervous system. In blue is the "decoder", or neural to effector map. In magenta is the effector, which can be a real assistive device, virtual effector, or therapeutic stimulation. Finally, in purple is the feedback which can include visual feedback, somatosensory feedback, or feedback via changes in therapy.

whose activation has long been thought to be related to motor inhibition. Thus, to manipulate frequency specific population activity patterns, we instead use a closed-loop BMI.

One method to study how tightly coupled a neural signal is to a behavior is to have a subject volitionally manipulate the neural signal of interest through a closed-loop BMI, and then to perform a behavioral task to probe the effect of the BMI manipulation. Although this is not a strictly causal manipulation since volitionally modulating a neural signal may involve unobserved co-modulations that end up being truly responsible for any resultant changes in behavior, this becomes less likely the more behavioral probes that are done. It also become less likely if bidirectional modulation (increasing the presence of the signal and decreasing the presence of the signal) have opposite effects on behavior.

Presenting feedback of a neural signal of interest for the purpose of learning to volitionally modulate the signal is referred to as "neurofeedback". Neurofeedback has been used as a tool to study the relationship between neural signals and behavior in a variety of brain areas. For example, in the frontal eye fields (FEF) of the rhesus macaque, there are neurons that are thought to play a role both in preparation of saccadic eye movements and visual spatial attention. When their activity is volitionally decreased, subjects exhibit worse performance in a visual attention task, but not in oculomotor preparation [75]. This points to activity of FEF neurons being tightly linked to visual attention, and is suggestive of neurofeedback as a potential therapy for disorders of attention. Besides visual attention, neurofeedback has been used to influence visual perception [76], attention deficit disorder [66], motor improvement in parkinsonian patients [65], emotional regulation networks [77], pain [67], and more.

When neurofeedback manipulations cause changes in behavior, it does not definitively prove that the conditioned neural signal drives the behavior, but it is evidence for a tighter relationship between the signal and behavior than correlations alone.

1.3.2.2 BMIs for Studying Sensorimotor Control Principles

Another scientific use of closed-loop BMI systems is to study sensorimotor control principles. BMIs where a virtual effector is controlled by populations of single and multi-unit activity simplify many aspects of studying the endogenous sensorimotor system, yet still incorporate fundamental principles relevant for studying control [78] and learning [79]. In studying natural control, the mapping between observed neural signals and the movement of the limb is undefined and likely non-linear, whereas in a BMI it is fully defined by the experimenter and linear. In natural motor control the experimenter only observes a subset of the relevant neurons for driving the musculature, whereas in a BMI system the experimenter can define a subset of observed neurons as the "output neurons" thereby allowing full observation all relevant neurons. Further, in our bodies, the biomechanical properties of limbs are complex and non-linear, whereas the dynamics of a virtual effector in a BMI can be made simple and linear. Finally, in natural control, feedback about the state of the limb arrives through both visual and sensory feedback, complicating the study

of how feedback may influence control. In a BMI system, feedback is usually restricted to visual feedback, thus allowing for a simpler study of feedback manipulations.

Using populations of single and multi unit activity from primary and pre-motor cortices mapped to the movement of a 2D or 3D cursor has been a popular approach for studying sensorimotor control and learning due to the ease in which non-human primates can learn such control (days or weeks), and the compelling evidence that controlling such a BMI may resemble performing habitual motor actions. Evidence has been found for internal models of how neural commands influence the cursor during BMI control [80], which are also thought to exist for natural control of limbs [81]. Neural-to-cursor BMI mappings can be learned over days, and once learned are stable across time, can be readily recalled, and are resistant to interference [71], resembling a putative motor memory. Further, BMI cursor movement speeds and accuracy are approaching natural arm speeds [82], so the statistics of kinematics are also becoming comparable.

BMIs using other neural control signals (LFPs, ECoG signals, EEG signals) or other types of effectors (real robots) may also be used for study of sensorimotor control, but have not been so extensively studied or analogized to natural control as the typical BMI using populations of multi-units mapped to the activity of a cursor.

1.4 Open Questions and Chapter Previews

The following thesis uses BMIs that are driven by population-level (LFP signals and single and multi-unit populations) motor cortical activity in an effort to investigate principles of population-level motor cortical computation.

One open question relates to how motor areas transition from withholding movement to initiating movement. Individual neural responses in motor areas have been shown to exhibit preparatory activity and movement activity, but how the preparatory activity does not cause actual movement is not well understood. Some hypotheses posit the presence of a downstream gate, perhaps in the spinal cord, that blocks or allows movement using an independent signal. Another posits that motor cortical population patterns may modulate during preparation, but in a dimension that results in a non-moving readout [83]. A similar hypothesis proposes that temporally synchronized patterns at a specific frequency, evidenced in LFP recordings, may reflect population patterns that encode a non-moving state [84].

Chapter 2 sets out to test the hypothesis that a specific LFP feature of motor cortical population activity, beta band oscillations, is tightly linked to withholding movement. For the last two decades, the behavioral correlates of beta band oscillations have been debated and proposed to be related to behaviors like attention, co-contraction, or idling. Without technology to reliably perturb population signals, the only evidence that can be gathered are correlations between beta band oscillation strength and motor behavior. In chapter 2, we approach this problem with a sequential BMI-motor task. Subjects use a closed loop BMI to learn to control the state of their motor cortical beta band

oscillations. On each trial, they make their beta oscillatory state match the cued state (either high, medium-high, medium-low, or low). Once their beta band state matches the cued state, the closed loop BMI turns off and subjects perform a typical motor reaching task. Analyzing how motor reaching changes depending on which beta band state was achieved prior to the reach is able to shed light on how tightly linked beta band oscillations are to movement inhibition. One benefit of the BMI paradigm is that different subjects may differ in the manner in which they generate or dissipate beta band activity during the closed loop BMI. If subjects perform the task differently, characterized by differences in other neural features, yet the effect of the pre-reaching beta state on motor behavior is consistent across subjects, evidence for the link between beta band oscillations and movement inhibition is stronger. One downside of this paradigm is that it is not a causal manipulation. Subjects do need to perform an unnatural closed loop BMI task prior to executing their motor reaching task, so it is possible that the manner in which they perform the closed loop BMI may differ from the way beta band oscillations occur during natural reaching behavior. Thus, in chapter 3 we address the question of how similar neural populations of multi and single-unit activity are during natural motor reaching and the BMI-reaching task.

While beta band oscillations represent one statistic of frequency-specific population neural activity on the scale of millimeters, single-unit action potentials and multi-units offer another perspective into the state of the population on the scale of hundreds of microns, capturing high frequency spiking activity. In chapter 3, unit activity during ongoing beta oscillations is compared in a natural reaching setting versus the closed-loop BMI task setting. If populations of units show consistent relationships to ongoing beta oscillations during both contexts, the probability that the beta oscillations generated during the closed-loop BMI task are qualitatively the same as the ones during natural motor control is higher, strengthening the argument that the behavioral results found in chapter 2 reflect a relationship that is present in natural motor control contexts. Overall, analyses using different statistics of population activity are investigated to assess how similar high beta neural states are in a natural and closed loop BMI context. Similarity of these contexts supports the finding from chapter 2 that beta oscillations are tightly linked to movement inhibition states.

While closed loop BMIs are useful scientific tools to perturb neural states, they also may have clinical applications. In Parkinson's patients, beta oscillations in the basal ganglia are known to be exacerbated when patients are not receiving pharmacological or stimulation therapy. While motor cortex does not exhibit the same increase in beta oscillatory patterns, it does exhibit greater coupling of high-frequency, putatively spiking activity, to ongoing beta oscillations [85, 86, 87]. It is also known that cortex serves as the greatest synaptic input into the basal ganglia, making it likely that correlated cortical activity has a role in the exaggerated basal ganglia oscillations. Thus, in chapter 4 we investigate first whether parkinsonian patients can perform closed loop BMI tasks given their compromised basal ganglia, and second whether there is any behavioral benefit of learning to reduce motor cortical beta power on motor behavior in a similar BMI-motor

task. Evidence of motor cortical beta state being linked to changes in motor behavior would not only establish a basis for closed loop BMI training as a possible therapy, but also may yield insight into how motor cortical beta oscillations are related to the basal ganglia beta oscillations, behavior, and parkinsonian symptoms. This chapter concludes our study of beta band oscillations and their relationship to movement inhibition.

We next turn our attention to questions focusing on how population activity of tens single and multi-units coordinate in a closed-loop BMI cursor control tasks as a starting point for trying to understand how the cortical motor system may generate different types of movement. Representational theories of the motor system posit the premotor and motor cortex contain thousands of neurons, all with tuning to a specific goal-related, kinematic, or kinetic variables. Downstream systems perform transformations on these variables, until they are transformed into muscles commands. Given that the musculature of the arm contains tens of muscles to control precisely, one might expect that high dimensionality to be reflected in the computations in motor cortices. However, recent findings highlight that tens to hundreds of simultaneously single and multi-units exhibit low dimensional firing patterns during arm reaching tasks [33, 88, 34], and even that the arm itself only uses a fraction of the dimensionality available to it [89]. In an analogous BMI system, when expert subjects perform a 2D BMI task driven by tens to hundreds of neurons, [90, 91] uncover low dimensional structure as well. How might low dimensional population activity emerge from inputs in a system that receives inputs from thousands of neurons each with their own respective independent neural noise? And why might low dimensional encodings be useful for control? In chapter 5, we develop a generative model of BMI population activity by modeling inputs into a population of simulated BMI neurons. Two different mechanisms are identified for how low dimensional, coordinated activity can emerge in population through inputs that are not explicitly designed to be correlated. Further, given the well defined neuron-to-cursor mapping in BMI experiments, it is possible to study how low dimensional coordinated activity contributes to online control with respect to speed and accuracy. Discovering which components of motor cortical population activity drive fast and accurate movements suggests how low-dimensional structure contributes to control during BMI tasks, and may suggest how low-dimensional natural motor cortical activity contributes to control of limbs.

Finally, natural actions are not comprised only of fast, accurate reaching movements to peripheral targets, but rather of individual actions strung together in sequences. How similar are neural population patterns when generating sequences with different orderings of the same subactions? Does the temporal arrangement of a series of actions influence the manner in which neural activity is produced? In chapter 6 we use two cursor BMI tasks, requiring different cursor trajectories, to investigate the coding, low-dimensional structure, and temporal constraints in neural population activity. Essential to these questions is having a well-defined neural-to-behavioral mapping to allow for interpretation of how specific components of neural commands do or do not directly contribute to movement.

BMIs are used in this thesis to study motor population activity and their relation to behavior. Specifically, they allowed for the manipulation of abstract types of neural

activity that are currently impossible to manipulate with optogenetics, electrical stimulation, or pharmacology. BMIs enabled the study of different frequency decompositions of local field potentials, as well as different decompositions of populations of single and multi-units into correlated and uncorrelated patterns. As the field of systems neuroscience shifts away from studying one neuron at a time to studying the population, BMIs serve as tools to perturb populations in specific ways, and gain a stronger understanding of how different pieces of their activity sum to drive the full movements we rely on daily.

Chapter 2

Neurofeedback Control of Beta Band Oscillations Affects Movement Onset

2.1 Introduction

2.1.1 Beta Band Oscillations in the Motor System

In motor output areas and the basal ganglia, beta (15 – 40 Hz) oscillations have emerged as a prominent signal in LFP and ECoG recordings. However, the behavioral correlates of beta activation across different motor tasks are still poorly understood. A better characterization of motor oscillations has potential to directly improve our ability to design neural prosthetic devices such as BMIs driven by motor field potential activity [92], BMIs driven by single unit activity but that take into account field potential activity during decoding [93, 94], and create closed-loop deep brain stimulators for movement disorders triggered by field potential signals [56, 57].

Before diving into the review of the existing hypotheses, we clarify what is meant when stating that 'x is a correlate of beta oscillations', where x may be 'attention', 'motor preparation', or 'cue processing', for example. LFP activity is generated by firing of neurons, and specific features in the field potential emerge due to the network connectivity, spatial arrangement of cells, and statistics of synaptic inputs to a region [21, 22, 95]. In this sense, the oscillatory activity measured in the field potential is an epiphenomenon of local neuronal activity, and the field potential is a population summary statistic. When asking what behaviors correlate with beta oscillations, we really are asking what behaviors correlate with these specific neural patterns that can be measured through the presence of an oscillation in the field potential.

The neural patterns that co-occur with cortical beta oscillatory events have been somewhat categorized. There is evidence that firing rate of recorded motor layer V pyramidal neurons does not significantly change during beta oscillations [95]but increased regularity of firing is observed [96]. Neuronal spike timing locked to a particular phase of the ongoing oscillation is observed, and in particular, spike timing in response to a visual stimulus becomes more predictable when taking into account a neuron's preferred phase of firing during ongoing beta activity [97]. Across different tasks though, neurons may exhibit different preferred beta phases [98]. Also, increases in synchronous 'unitary events', or simultaneous firing of multiple neurons, has been observed [99]. A subtlety is that since the LFP represents an electrical potential measured in the extracellular space, it will in turn affect the probability of firing of local neurons. Such ephaptic effects have been demonstrated to affect neuronal firing in vitro [100]. Thus, it is not clear whether the previously described neural firing patterns are indeed causal in creating the oscillation, or are a product of an ongoing oscillation that is generated elsewhere. Different generators of oscillatory activity (remote, local, local network) are schematized below. Which one of these best categorized beta activity is an issue that remains unresolved. Below, we will consider hypotheses that fall into each of these possibilities.

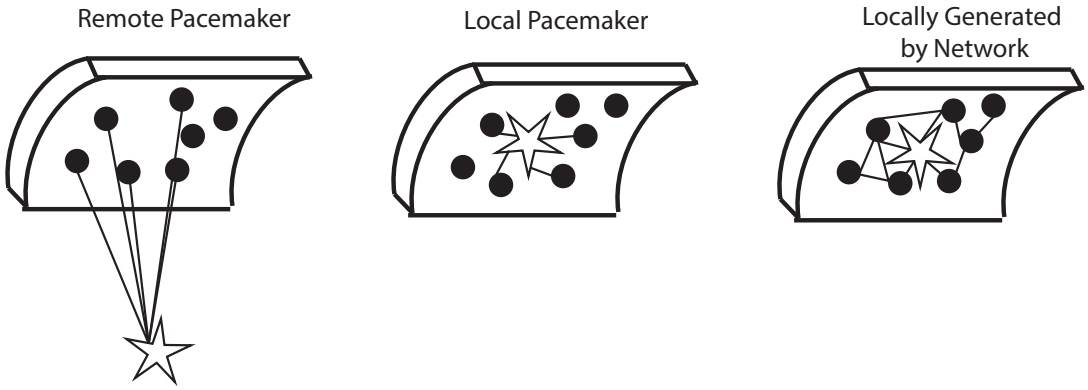

Figure 2.1: Possible drivers of oscillatory activity. Remote (a) or local (b) 'pacemaker' neurons could be responsible, or the oscillation could be an emergent property of the tissue (c). Modified from Whittington et al., 2000 with permission.

2.1.2 Beta Band Oscillations *In Vitro*

In cortex, locally generated oscillations in LFP signals are attributed to cycles of synchronous excitation followed by synchronous inhibition [101, 102]. There are wide ranges of network structures and cell types that give rise to different frequencies and amplitudes of oscillatory activity [102]. Through stimulation studies, beta oscillations in motor cortex have been found to be generated in part by corticospinal projecting pyramidal neurons [103]. Modeling and in vitro slice work has identified two specific mechanisms as responsible for sustained beta activity in pyramidal tract neurons in the presence of excitatory drive: gap junctions and the M-current [104, 105]. Roopun and colleagues took slices of rat somatosensory cortex and pharmacologically activated excitatory glutamate receptors. They found an emergence of beta oscillations in layer V, and that the oscillations were robust to applications of various neurotransmitter blockers. Notably, three manipulations changed the oscillatory activity: GABAA receptor blockers, gap junction blockers, and an M-current blocker. GABAA receptor blockers at high enough dosage and gap junction blockers resulted in the disappearance of the oscillation, and the M-current blocker affected the peak frequency of the oscillation in a dose-dependent manner. This finding suggests that due to the structure of sensorimotor cortical regions, in the presence of a constant excitatory input, oscillations in beta frequencies will emerge making them an emergent property of the tissue, as schematized on the right in 2.1. While this study did identify mechanisms that explain why beta frequencies are dominant, it did rely on a constant excitatory drive, a condition that may not be physiologically reproduced *in vivo*.

2.1.3 Behavioral Correlates of Beta Oscillations

2.1.3.1 Motor Cortical Beta Oscillations are generated by a local network in motor cortex

Early studies found beta oscillations to emerge reliably during precision reach tasks [106], pressing of force plates with digits [107], and isometric contraction [108]. Clear onset and offset of muscle stiffening measured by electromyography (EMG) predicts the emergence and cession of oscillatory beta activity [106]. In contrast, during natural reaching and grasping where limbs are in motion, beta oscillations are not consistently correlated with EMG [95]. It is interesting that during pressing, pushing, and contraction tasks, when the upper arm is not moving but the muscles are stiffened that oscillations are well correlated with EMG. One possible reason for this is that during more dynamic movements muscle synergies are activated in a temporal sequence whereas when stiffening, all muscles in a given area receive simultaneous excitatory drive [109]. Given evidence that a strong excitatory drive in slice is required to give rise to robust beta oscillatory activity [104], it is possible that in stiffening tasks these conditions are consistently met *in vivo*. The "stiffening hypothesis" corresponds best with the architecture described in the right of 2.1since it suggests that beta oscillation is an emergent property of subjects having to co-activate their muscles. Another possibility is that beta oscillatory activity could be transmitted from elsewhere (left of 2.1, [110]) or that beta activity could be driven by specific local pacemaker cells instead of driven by the entire network (middle of 2.1).

2.1.3.2 Motor Cortical Beta Oscillations are generated by few pacemakers in the motor system

Another class of hypotheses has emerged after the findings that beta oscillations are not just limited to motor cortex and EMG activity, but are concomitantly present in somatosensory and parietal cortices [111, 112]. Beta oscillations emerge in somatosensory cortices during somatosensory stimulation [113]. Even the earliest studies of beta oscillations in NHPs have noted consistent beta oscillations in M1 in response to 'cutaneous stimulation of the hand or wrist' [114]. Further, a number of studies have noted the emergence of beta oscillations in S1 concurrently with M1 oscillations, and even noted an anterior-posterior directional flow of activity [111, 115], implying an interaction between somatosensory and motor cortical beta oscillations. One theory as to how M1 and S1 could be interacting postulates that beta activity is a correlate of active peripheral sensory sampling. Proponents theorize that beta oscillations synchronize and serve as a 'test pulse' descending the spinal cord and emerging in motor and somatosensory cortex just prior to movement in order for motor commands to be updated based on the current state of the periphery. At times when motor cortex has just sent a signal, beta oscillations emerge to ensure that the feedback arrives at a time when motor cells are 'most receptive' to receive it so that they can quickly adapt to feedback [113], p. 20.

If the case, it would be expected to find coherence in the beta range between the initial

motor cortical oscillations and the ascending spinal cord as well as motor cortex and somatosensory cortex, which are both found [112, 116]. Further, if motor and sensory cells participate in this network wide oscillation, cells will likely have phases where they better perceive sensory feedback. In human EEG studies, a relationship between ongoing beta phase and a reliance of the magnitude of sensory evoked responses is found [117]. Additionally, numerous studies on 'sensory gating' have demonstrated that somatasensation is reduced during movement [118, 119, 120], as are beta oscillations, though these studies use cutaneous somatosensation, and spinal cord recordings suggest a greater involvement of motoneurons than cutaneous neurons in muscular coherence [116]. The peripheral sensory sampling hypothesis relies on an internal driver pacemaker of beta activity to entrain cells such that their spiking probability is higher at times when they receive sensory feedback, an architecture resembling the middle of 2.1. As further evidence, potential pacemaker cells that exhibit after-spike hyperpolarization properties useful for entraining 25-35Hz activity have been identified in primate motor cortex for 25-35 Hz activity [121]. Thus far, we have discussed the biophysical generation of beta oscillations and noted that periods of constant excitatory drive ought to reliably evoke beta activity. Below, we consider a final hypothesis that posits that distal generation of beta may drive the observed oscillatory activity.

2.1.3.3 Distal Areas Generate Motor Cortical Beta Oscillations

Substantial work has found beta frequencies in other parts of cortex involved in attention and long –distance synchronization[110]. Examples include beta synchronization of dorsal prefrontal cortex and posterior parietal cortex during top-down search for a visual stimulus [122, 123, 124]. Further, modeling work has suggested that beta oscillations are best well suited for communication between distal areas, compared to gamma oscillations which are better suited for local processing [125, 110]. In addition, it has been suggested that beta rhythms could represent subject attention to motor tasks [95, 126, 127], and could be generated in a manner suggested by the left of 2.1. The link between beta in motor cortex and attention is best illustrated by an elegant human study.

Saleh and colleagues developed a task that isolates attention from motor preparation, postural holding, and muscular contraction [128]. Here experimenters instruct human subjects to observe five sequential cues, each of which instructs a reach to a different target. After observation of all five cues, subjects then execute the fourth cue, which they have to have remembered from the prior presentation. During this experiment, transient beta oscillatory activity emerged just prior to the first through fourth cues 2.2. Since the cues are presented rhythmically, they are predictable. The authors interpret the transient oscillatory activity as a proxy for subjects paying attention to the upcoming cue. For the first through third cues, subjects must attend to them in order to count their occurrence, and for the fourth cue, subjects must remember the instruction. Since the action that subjects are performing occurs much later in time, it is unlikely that the beta activity in this case represents maintenance of a postural state or any kind of muscular stiffening.

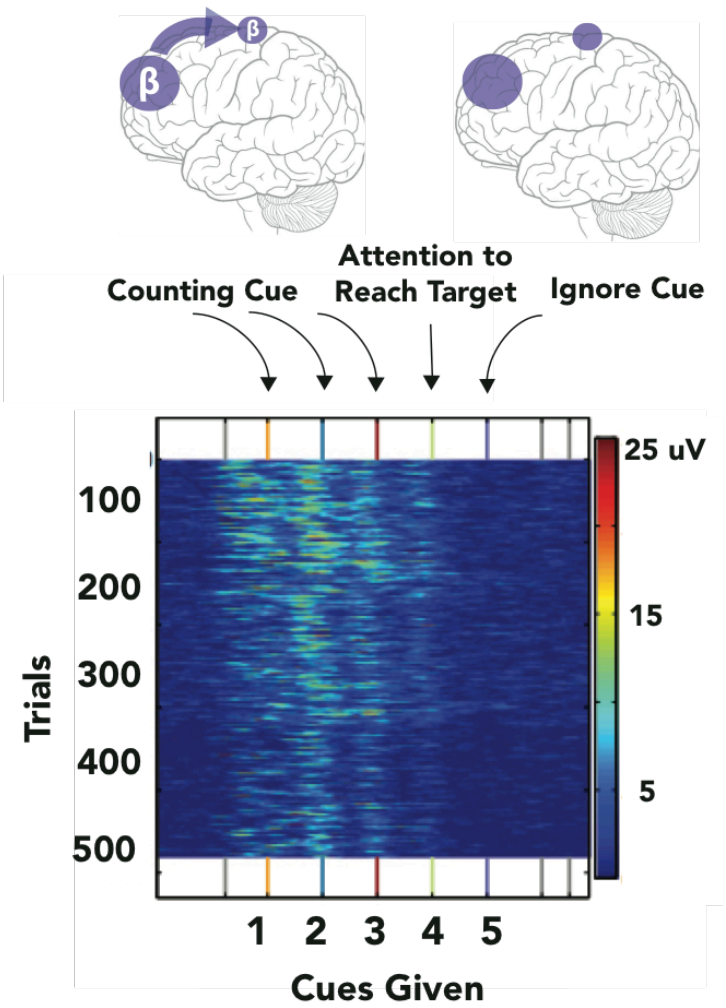

Figure 2.2: Single trials where beta power is represented by color, and the task cues are shown on the horizontal axis. In this task, a human subject is instructed to watch five cues without responding and afterwards must execute the action instructed by the fourth cue [128]. At the top, schematics of neural activity demonstrate an attentionally driven beta oscillation from prefrontal areas during times in the trial where the subject must pay attention. Attention is needed either to count the cue (cues 1-3), or to remember the instruction given by the cue (cue 4). The lack of a driving oscillation is shown when the subject no longer needs to attend to the cue (as in cue 5). Modified from Saleh et. al., 2010 with permission.

Another possible way that the distal beta source can influence computation is to selectively coordinate local cell assemblies that are task-relevant. Probabilistic models of neuronal spiking have shown a significant dependence on distal LFP phases and even distal LFP-LFP phase coupling [129]. Further, cells that form a functional cell assembly may be simultaneously co-activated by exhibiting strong coherence to distal field potentials. This phenomenon has been observed in motor cortical cells that are responsible for controlling a BMI. They exhibit enhanced coherence to striatal field potentials compared to cells in the same region that were not responsible for controlling the BMI [130]. The BMI-controlling cells also show increase locking to slow-wave activity during sleep [131]. While the above evidence is not specific to beta activity, it is possible that beta is being transmitted from a distal region to bind cells together much in the same way oscillatory activity has been proposed to solve the 'binding problem' in perception [132].

2.1.4 Perturb Beta Band Oscillations to Infer Relationship to Behaviors

Ideally, in order to test these hypotheses, experiments could be designed that perturb beta oscillatory activity during muscular activation tasks, movement tasks, and tasks that require motor-related attention. In addition to measuring behavioral variables such as reaction time, reaching error, and perception errors for example, ideally experimenters would also be able to recording neural activity during perturbations. Electrical stimulation is one approach to experimentally inducing oscillations. Slowly oscillating cortical macrostimulation (up to 1.7 Hz) has been shown to entrain single unit neural activity in anesthetized rodents, though it was overpowered by endogenous rhythms in an awake preparation [133, 134]. Non-invasive stimulation using transcranial alternating current stimulation (tACS) and repeated transcranial magnetic stimulation (rTMS) has been used to induce changes in cortical oscillations [135], but the frequency of the induced oscillations is not solely dependent on the stimulation frequency [136, 137]. Further, the change in motor behavior from stimulation at beta frequencies has been non-congruent [138, 139, 140].

A non-stimulation based approach is to use neurofeedback to have subjects learn to manipulate their own beta oscillations in their motor cortex. While neurofeedback was pioneered by rewarding changes in firing rates of single motor cortical cells [70], learning to control cortical local field potential (LFP) features has been proposed [126, 141] and more recently has been demonstrated [142, 143]. Thus, if subjects perform a sequential neurofeedback-behavior task where they are trained to modulate their beta oscillations to an instructed level, and immediately afterwards perform a behavioral task, the relationship between beta oscillations and the probed behavior can investigated. Sequential neurofeedback-behavior task designs have been used before with a variety of neural recording modalities, neural signal features, and behaviors [66, 144, 75, 65]. For example, one study examines the effects of operantly increasing or decreasing the firing rate of a single

neuron in frontal eye field region on perception of targets in the receptive field of that neuron [75]. Sequential neurofeedback-behavioral paradigms have the benefit of first, perturbing neural signals endogenously instead of driving the tissue with an artificially large amount of current, second the ability to record neural activity simultaneously without stimulation artifacts, and finally the ability for experimenters to specify the time length and cortical location of the neural activity they wish to study.

2.2 Methods

To explore how beta oscillations reflect changes in arm movements, we trained three macaque monkeys to perform a typical center-out arm-reaching task (CO tasks) and a novel sequential beta neurofeedback arm-reaching task (NR task) while recording from bilateral intracortical microelectrode arrays.

2.2.1 Surgery, electrophysiology, and experimental setup

Three male rhesus macaques (Macaca mulatta, RRID: NCBITaxon:9544) were chronically implanted with arrays of 128 Teflon-coated tungsten microwire electrodes (35 μm in diameter, 500 μm separation between microwires, 16 x 8 configuration, 6.5 mm length Innovative Neurophysiology, Durham, NC) in left upper arm area of primary motor cortex (M1) and posterior dorsal premotor cortex (PMd). Localization of target areas was performed using stereotactic coordinates from a neuroanatomical atlas of the rhesus brain (Paxinos et al., 2013). LFP activity was recorded at 1 kHz using either the 128-channel Multichannel Acquisition Processor (Plexon, Inc., Dallas, TX) (Monkeys S, G) or the 256-channel Omniplex D Neural Acquisition System (Plexon, Inc.) (Monkey C). Single unit and multi-unit activity from Monkey G was manually sorted offline using Offline Sorter (Plexon, Inc). Channel-level activity [27] from Monkey C was defined using OmniPlex's auto-threshold procedure to set each channel threshold to 5.5-standard deviations from the mean signal amplitude. Thresholds were set at the beginning of each session based on 1-2min of neural activity recorded as the animal sat quietly (i.e. not performing a behavioral task). Monkeys S and G were trained to perform a center-out delayed reaching task using a KINARM exoskeleton (BKIN Technologies, Kingston, ON, Canada) fitted to their right arm. Monkey C was trained using a custom right-arm sleeve with a red LED marker on the hand that was tracked in real-time with an Impulse X2 motion capture system (PhaseSpace, San Leandro, CA). For all monkeys and tasks in this study, visual feedback of hand position was shown by a circular cursor on the task screen. Monkey S and G's right arm movements were restricted to the horizontal plane by the KINARM. Monkey C could rest and move his right arm on a horizontal plane like Monkeys G and S, but could also move his arm above the plane. Prior to this study, Monkey S was trained at reaching tasks and spike-based brain-machine interface (BMI) cursor tasks for 4 years, Monkey G was trained at joystick tasks and spike-based BMI cursor tasks for 1 year, and Monkey C

was trained at reaching and spike-based BMI cursor and virtual exoskeleton tasks for 3 years. All procedures were conducted in compliance with the NIH Guide for the Care and Use of Laboratory Animals and were approved by the University of California, Berkeley Institutional Animal Care and Use Committee.

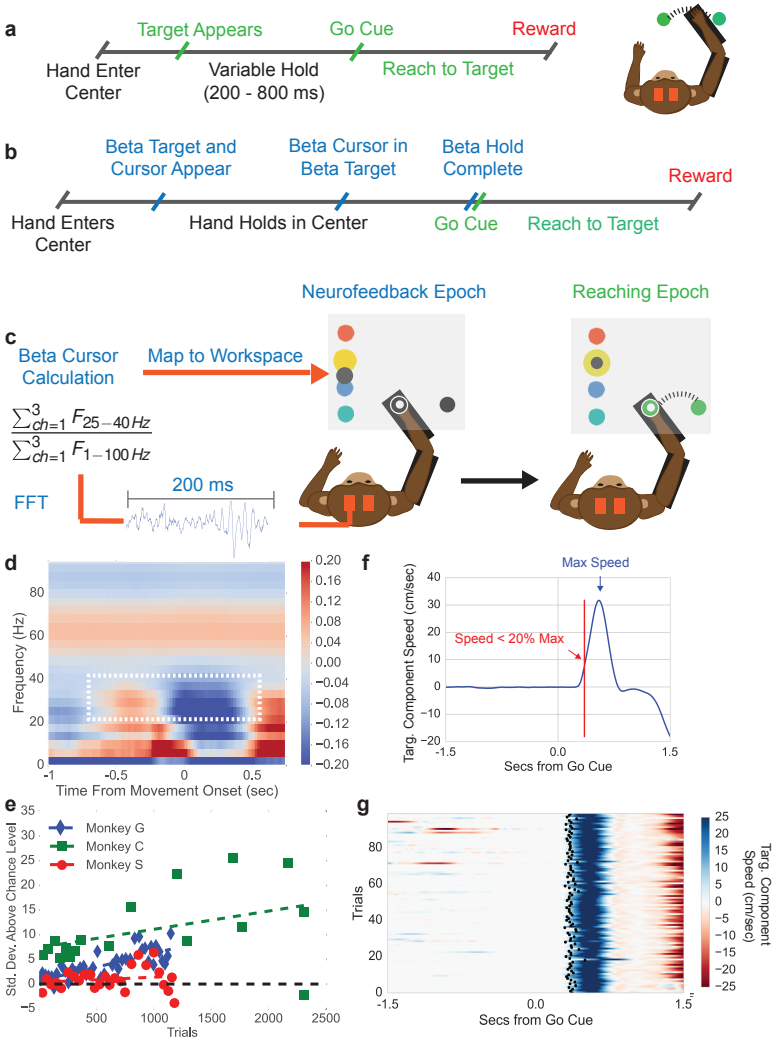

Figure 2.2.1 (a) Timeline of center-out reaching task (CO task) with variable hold times (200-800 ms) (b) Timeline of neurofeedback-reaching task (NR task) where blue text indicates the neurofeedback epoch and green text indicates the reaching epoch. (c) The NR task feedback loop. Subject keep their right hand held in a central target throughout the task. They are then shown a single beta target (shown in yellow here) and beta cursor (shown in gray here) on the screen which is updated every 100 ms. Once the beta cursor is held in the beta target for 450 ms, the beta cursor and beta target disappear and the subject reaches to a peripheral target 6.5 cm away. (d) Trial averaged spectrogram of movement onset aligned motor cortical LFP signals for Monkey C, with a mean 1/f trend

estimated with first-order linear regression and subtracted away. White box highlights the beta desynchronization in the 25-40 hz range. (e) All three subjects perform the neurofeedback epoch part of the task above chance. The x axis corresponds to all trials from all sessions concatenated, each point corresponds to a session, and the x axis position of each point corresponds to the first trial that falls within that session. Position on the Y axis indicates standard deviations above mean chance level (shown with black dotted line) (f) Illustration of the metric termed movement onset time (MOT) throughout the text. Trial-averaged hand speed in the direction of the target is shown in blue with an arrow pointing out the time of maximum hand speed. To find the MOT, step backward in time along the hand speed trace until the hand speed falls below 20% of the maximum speed value. (g) 100 trials (rows) of hand speed are shown where time prior to 0.0 sec is the neurofeedback epoch and time after 0.0 sec is the reaching epoch in (b). Black dots indicate the calculated MOT. Increasing blue corresponds to increasing hand speed.

2.2.2 Behavioral Tasks

2.2.2.1 Center-Out Task (CO task)

Subjects performed a CO reaching task consisted of right hand movements from a center target to a peripheral target distributed over a 13 cm diameter circle, panel a of 2.2.1. The workspace was created to minimize any requirement for postural changes during task performance. Target radius was typically 1.2 cm in the workspace. Trials were initiated by entering the center target and holding for a variable time (uniformly distributed within 200–800 ms). The go cue after the hold period was indicated by the center target changing color and the peripheral target illuminating, cuing a reach to that target. A liquid reward was provided after a successful reach to each target and a peripheral hold period of 200 ms.

2.2.2.2 Neurofeedback-Reaching Task (NR Task)

Prior to training subjects to perform the NR task, beta frequency band limits used in the neurofeedback portion of the NR task were computed from the CO task. A movement onset-aligned trial-averaged spectrogram from the intracortical recordings in contralateral motor and premotor cortex (e.g. Monkey C in panel d, 2.2.1) showed that the clearest movement related desynchronization was in the 25 - 40 Hz band for all monkeys, consistent with early reports of beta oscillations in macaque motor cortex [106, 107]. Thus, the beta band limits for the neurofeedback epoch of the NR task were set to 25-40 Hz. Subjects were then trained to perform the NR task. Trials were initiated by moving the right arm (co-located with a cursor on a screen) such that the cursor fell within a central target. Holding in the center target initiated the neurofeedback epoch where a beta neurofeedback cursor and one of four possible beta neurofeedback targets appeared on the screen (all blue text in panel b, 2.2.1 falls in the neurofeedback epoch). Subjects modulated endogenous motor cortical local field potential signals to move the vertical

position of the beta cursor. Specifically, the cursor was controlled by a spectral estimate of beta band power normalized by a spectral estimate of broadband (1-100 Hz) power:

$$ß_{est} = \frac{\frac{1}{3}\sum_{ch=1}^{3}\sum_{f=25Hz}^{40Hz} PSD_f^{ch}}{\frac{1}{3}\sum_{ch=1}^{3}\sum_{f=1Hz}^{100Hz} PSD_f^{ch}}$$

where PSD_f^{ch} is the power spectral density estimate of a particular channel (ch) at a particular frequency (f) assessed using the multi-taper method (5 tapers) in windows of 200 ms [145, 143]. Once was computed, a subject-specific linear transform was used to map $ß_{est}$ to a vertical screen position. A two-timestep (200 ms) moving average (boxcar) filter was then used to smooth out the displayed beta cursor position. Subject-specific values were finalized after ~1 week of training the subjects on the NR task. Initial training began with more lenient beta neurofeedback requirements (values to achieve low beta and high beta target were closer to the mean beta cursor value). As subjects improved in performance, the beta targets moved further apart until the top and bottom targets had a mean time to target of 5-10 seconds.

After successful beta target completion, both the beta cursor and beta target disappeared, cueing that the reaching epoch had begun (all green text in panel b of 2.2.1 falls in the reaching epoch). Subjects then executed a right-arm reach from the central target to a peripheral target to receive a liquid reward. During NR task performance from these days, subjects exhibited above chance performance as shown in panel e of 2.2.1. Chance performance was computed by shuffling the beta target order from a single block, designed to assess whether the subjects' performance during the beta neurofeedback epoch of the NR task was merely due to chance fluctuations in beta power or was due to volitional changes in neural activity that were specific for the target on the screen. In the simulation, after the beta cursor entered the beta target and held for the 450 ms beta target hold time, an average arm-reaching time, the constant reward time, and the constant inter-trial interval time transpired to simulate the natural pacing of the task. At the end of the simulation, a metric of chance performance was the mean number of successful beta targets acquired over the length of the session. For example, if one target-shuffled performance yields 10 successful trials in 10 minutes, the chance rate for that simulation would be 1 rewarded target / minute. One hundred simulations were run per session (each session ~ 10-40 min) yielding a distribution of rewarded targets/minute. The mean and standard deviation of the distribution was calculated, and used to z-score the actual number of rewarded trials. The resultant z-scores for each session are plotted in panel e of 2.2.1where each point corresponds to a session (session i), and each point's position on the x axis (xi) corresponds to the first trial that falls within session i amongst the concatenated trials over all sessions.

Overall, monkeys achieved average success rates of 60% and performed on average ~4 successful trials per minute in the NR task. Errors almost entirely came from accidentally moving the right hand outside the center target during the neurofeedback epoch.

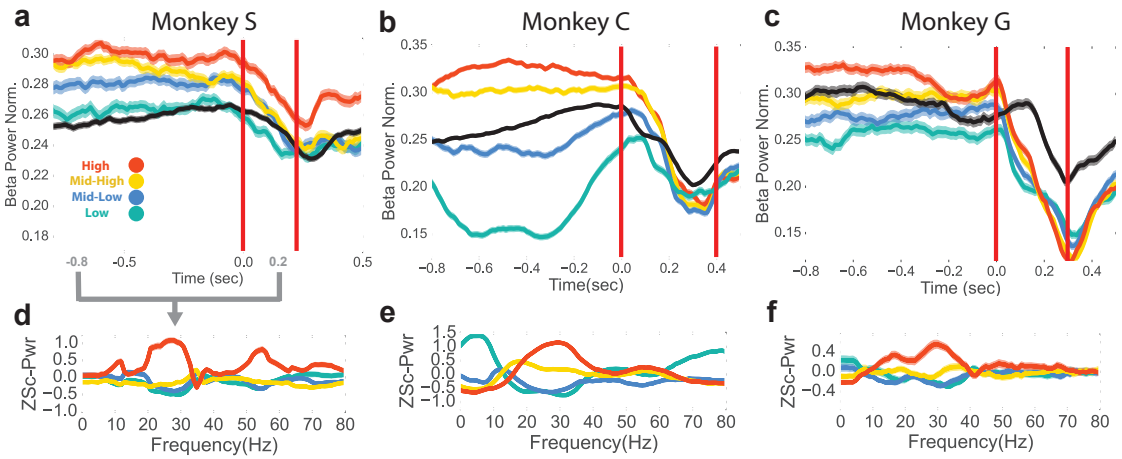

Figure 2.3: (a-c) Mean (s.e.m) of normalized beta power for Monkey S, C, and G aligned to the end of neurofeedback epoch. High, mid-high, mid-low, and low beta targets are in red, yellow, blue, and green, and go-cue aligned CO trials are in black for reference (d-f) Z-scored PSDs estimated from a time slice 0.8 seconds prior to the end of the neurofeedback epoch (labeled as 0.0 sec in a-c) and 0.2 seconds after the end of the neurofeedback epoch. This time slice is displayed in gray arrows below the time axis in (a).

2.3 Results

2.3.1 Neural Dynamics During the NR task

The neurofeedback epoch of the NR task accomplished the goal of bringing beta power to different levels shown by plotting mean normalized beta power for the last 1 second of the neurofeedback epoch and the first 0.5 seconds of the reach epoch for rewarded trials to each of the four beta targets in 2.3, panels a-c, averages (s.e.m.) over all trials (Monkey S: total n = 1184, Monkey C: total n = 2328, Monkey G: n = 1042).

The first vertical red line indicates the end of the neurofeedback epoch, or go cue for the reaching epoch. The second vertical red line indicates the mean movement onset time of the reach. The mean normalized beta power of CO trials is shown in black for reference. At the time of the go cue there is a significant group difference between normalized beta power for the four different neurofeedback target conditions (two-tailed Kruskal-Wallis test, Monkey S: n = 1184, H = 47.44, p = 2.803e-10, Monkey C: n = 2328, H = 250.1, p<5e-20, Monkey G: n = 1042, H = 48.11, p = 2.023e-10). To assess how subjects co-modulate other frequency bands in addition to beta band, and to ensure that the beta cursor changes were not a product of the normalization in panels a-c in 2.3, non-normalized, z-scored power spectral densities (PSDs) were computed over the last 0.8

seconds of the neurofeedback epoch to the first 0.2 seconds of the reaching epoch (total window is 1.0 second) is shown in panels d-f in 2.3. Mean traces show that high and low beta targets were achieved by increasing and decreasing beta power. Since calculation of the beta cursor position involved an estimate of broadband power, changes in non-beta frequencies also affected beta cursor position. In some subjects (Monkeys C, G), increases and decreases in beta power were accompanied with reliable decreases and increases in low frequencies (1-10 Hz). A final time-domain metric was computed to confirm that the occurrence of beta oscillations was changing across the different beta targets in the neurofeedback epoch. Instead of using PSD estimators over a window (as in panels a-f, 2.3), a time-domain method was used to extract time segments with bursts of beta oscillations. Briefly, instantaneous beta amplitude was measured by bandpass filtering the raw LFP with a 5th order Butterworth filter to isolate 25-40 Hz components, and taking the Hilbert transform. If the amplitude exceeded the 60th percentile of beta amplitude (computed each day) for a period of at least 125 ms (3-5 cycles of beta oscillations), the time points within that period were labeled as "on-beta". The percent of time points that were labeled as on-beta was computed in the same time window as Figure 2d-f for all trials. All subjects exhibit increasing percentages of on-beta time points for the low to high beta targets. The mean (s.e.m) of percent of time points labeled as on-beta for low, mid-low, mid-high, high beta targets respectively in Monkey S is 9.78 (0.783), 11.2 (0.778), 17.6 (0.913), 39.4 (1.35), Monkey C is 8.65 (0.523), 17.6 (0.708), 35.9 (0.840), 37.8 (0.868), and Monkey G is 18.3(1.04), 16.6 (0.994), 22.5 (1.22), 33.4 (1.43)). These three metrics (normalized beta power in a-c, 2.3, non-normalized PSDs in d-f 2.3, and percentage of time labeled as on-beta) demonstrate that the neurofeedback epoch served to increase and decrease beta oscillatory power prior to the arm-reaching epoch.

2.3.2 Beta Band Oscillations Delay Movement Onset

In the NR task, movement onset times, movement onset speed, peak reach speed, and movement onset acceleration were calculated for the reaching epoch of each successfully completed trial from days with a fixed beta-to-cursor transform. The two-tailed nonparametric Wilcoxon-like Cuzick's test was used to test for increasing or decreasing ordering of trial across the four beta target groups. Cuzick's test [146] is a non-parametric test for significant ordering of groups in an increasing or decreasing manner (two-tailed) and was used to assess significance of ordering of behavioral metrics according to the four beta targets. A test statistic (Z) is calculated for the hypothesis that groups follow a designated ordering (Cuzick, 1985). Z is calculated using the ranks of individual points and the group assignment (assignments used here: low beta target: 1, mid-low beta target: 2, mid-high beta target: 3, high beta target: 4 (except where noted) to determine if there is a significantly increasing or decreasing metric following the group ordering. Z follows a standard normal distribution (confirmed for data here by shuffling group labels 10,000 times and comparing the resultant Z distribution to a standard normal distribution with the KS test), so a p-value can be calculated using the cumulative standard normal

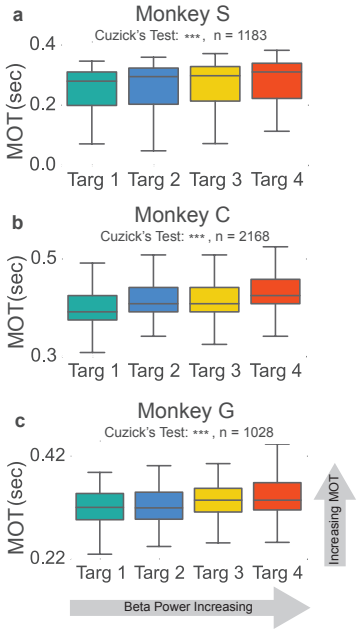

Figure 2.4: (a-c) Boxplot of reaching movement onset times grouped by proceeding beta target. Subjects exhibit an increase in movement onset time (MOT) when modulating normalized beta power to higher targets. Gray line at center of boxplot is median. *** p $<$ 5e-09, Cuzick's two-tailed test.

distribution.

We tested whether rewarded trials preceded by low, mid-low, mid-high, and high beta targets exhibit increasing or decreasing behavioral metrics. In all three animals, trials with high beta power targets had subsequent reaches with slower movement onset times (panels a-c 2.4Two-tailed Cuzick's test, Monkey S: z = 5.763, p =8.267e-09, n = 1183, Monkey C: z = 11.987, p $<$ 5e-20, n = 2168, Monkey G: z = 5.856, p = 4.729e-09, n = 1028). Note that trials with movement onset times greater than 0.7 seconds or less than 0.0 seconds were removed from this and all subsequent analyses (Monkey S: 1 trial, Monkey C: 160 trials, Monkey G: 14 trials). Other groups have found correlations between increased beta power and reduced onset speed, peak speed, and onset acceleration [147, 140] which we do not find consistently across subjects when comparing metrics grouped based on our proxy for beta power, the preceding beta target (see table 2.1).

Table 2.1: Z-statistic (and p-values) for Cuzick's two-tailed test for ordered grouping performed on onset speed, peak speed, and onset acceleration. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, Monkey S: n = 1183, Monkey C: n = 2168, Monkey G: n = 1028, Combined: n = 4379

Metric	Monkey S	Monkey C	Monkey G	Combine
Onset Speed	1.378, p=0.168	-3.593 (***)	0.0734, p=0.941	-1.2443, p=0.213
Peak Speed	4.476 (***)	4.142 (***)	-1.330, p=0.187	1.8339, p=0.067
Onset Accel.	1.901, p=0.057	2.573 (*)	0.159, p=0.873	0.8178, p=0.414

2.3.3 Sequential Neurofeedback -Reaching Task Controls

2.3.3.1 Beta target acquisition difficulty does not correlate with movement onset time

To ensure the cognitive effort required to increase and decrease beta power during the neurofeedback epoch did not result in increasing movement onset time observed in 2.4, we compared the amount of time it took to acquire each beta target as an approximate measure of each target's difficulty. For Monkeys S and G the time to acquire beta target did not significantly predict MOT in a linear regression but it did for Monkey C and when data was combined across monkeys (panels a-c in two-sided Student's t-test for non-zero slope in linear regression, Monkey S: t = 0.7119, p = 0.476, n = 1183, Monkey C: t = 2.352, p = 0.0188, n = 2168, Monkey G: t = 1.651, p = 0.0991, n = 1028, Combined across monkeys: t = 2.0832, p = 0.0373, n = 4379). When linear regression was used to predict MOT (MOTpred) from time to beta target, and was subtracted from the actual MOT (MOTres = MOT - MOTpred), increasing MOTres with increasing beta power target remained (two-tailed Cuzick's test on MOTres, Monkey C: z = 13.191, p < 5 e-20, n = 2168, Combined data across monkeys: z = 13.615, p < 5e-20, n = 4379). Thus, the slight predictive power of time to beta target on MOT does not account for increasing MOT with increasing beta target from Figure 3.

We controlled for whether looking at the top part of the screen (where the high power beta target is displayed) was effortful for subjects and resulted in slower movement onset times. In Monkey C the relationship between beta target and screen location was reversed by mapping increased beta power to the bottom of the Y-axis for a set of trials analyzed separately. For these trials, the high beta power target became the green Target 1 instead of the red Target 4. 2.5e shows that increasing beta-target versus increasing movement onset time reverses relationship (z = -5.971, p = 2.354e-09, n = 2113) demonstrating that increasing beta power, not the beta target position on the screen, consistently correlates with the observed increasing movement onset time. To test that the beta target versus movement onset time relationship generalizes to more than a single arm reaching target, we show that the same task with a different arm reaching target location produces the same effect (2.5d, Monkey S: z= 3.972, p =7.117e-05, n=735).

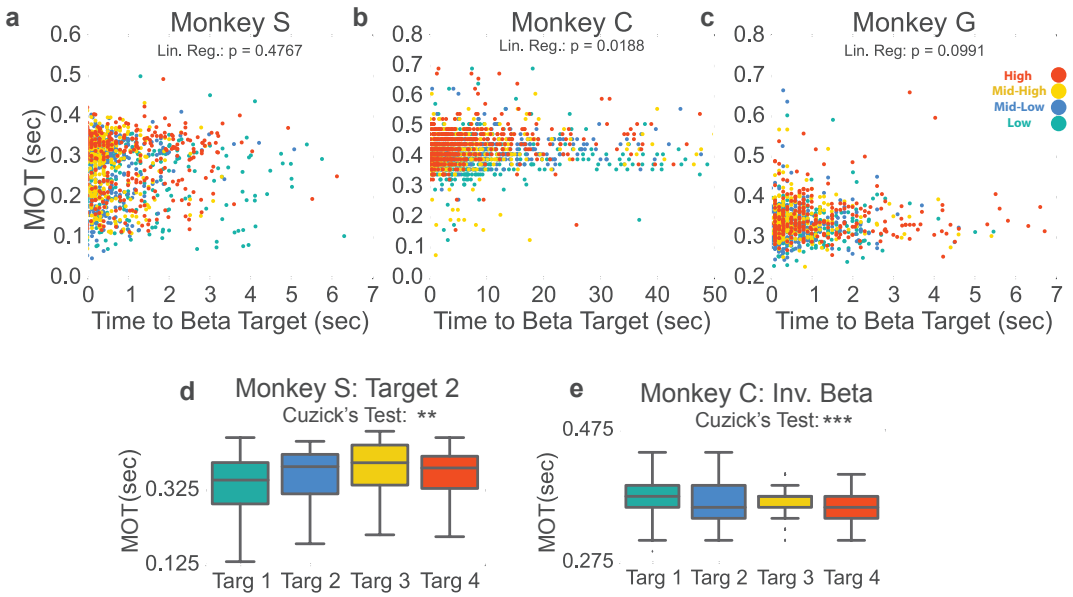

Figure 2.5: (a-c) Movement onset time (MOT) from 2.4 is plotted against time to beta target for Monkey S, C, G. Colors correspond to the beta target for that trial following the same colormap as 2.3 and 2.4. Linear regression is performed to assess if time to beta target is predictive of MOT. Non-significant p-values for Monkey S, G show time to beta target (interpreted as beta target difficulty) does not significantly predict MOT. Monkey C exhibits a significant relationship, but when MOTres is computed by subtracting predicted MOTs from time to beta target from actual MOT, MOTres exhibits the same increase with increasing beta target as seen in 2.4 (d) Changing location of the manual control reaching location from 6.5 cm to the right of the central target to 6.5 cm above the central target does not change the observed relationship between increasing movement onset times and increasing beta power target, Monkey S. (e) Changing the vertical ordering of beta targets on the screen (green is the high-beta target, blue is mid-high beta target, yellow is mid-low beta target, and red is low beta target) also shows the same increase as in 2.4B. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, Cuzick's two-tailed test for significant increases and decreases by grouping except where noted.

2.3.3.2 Movement onset time increase is specific to beta band frequencies

Using other methods to compute beta power shows the same movement onset relationship To confirm that the correlation between lower beta power targets and faster movement onset times was not due to the beta cursor calculation method that normalizes beta power by total broadband power, offline we sought to account for the increase in movement onset time with a beta power calculation method that was non-normalized. Non-normalized beta power was computed in the window spanned by the last 0.8 seconds of the neurofeedback epoch. Trials were then re-labeled by which quartile their un-normalized beta power fell (e.g., if the un-normalized beta power falls in the 0-25th percentile of all trials, the trial would be assigned to group 1). Trials with movement onset times less than 0 seconds or greater than 0.7 seconds were removed, as before. The mean movement onset time for each of these new groups is plotted by monkey (2.6A-C, darkest and lightest bars correspond to lowest and highest non-normalized power respectively). All three monkeys exhibit significantly increasing movement onset times with increasing non-normalized beta power (Two-tailed Cuzick's Test, Monkey S, z = 7.162, p = 7.945e-13, n = 1183, Monkey C, z = 7.767, p = 7.994e-15, n = 2168, Monkey G, z = 7.709 p =1.266e-14, n = 1028, Combined Across Monkeys: z = 6.168, p = 6.924e-10, n = 4379). This same procedure was performed except instead of relabeling by non-normalized beta power, trials were re-labeled by the percentage of on-beta time points using the previously explained time-domain method. Indeed, the same increase in movement onset time follows where trials with a larger percentage of on-beta time points exhibit slower movement onset times (Two-tailed Cuzick's Test, Monkey S, z = 7.575, p = 3.597e-14, n = 1183, Monkey C, z = 5.488, p = 4.068e-08, n = 2168, Monkey G, z = 7.890, p =3.108e-15, n = 1028, Combined Across Monkeys z = 5.1301, p = 2.895e-07, n = 4379). Thus, the normalization of the cursor does not account for the increase in movement onset time.

Non-beta frequencies are co-modulated during the beta neurofeedback epoch of task Since the neurofeedback epoch required control of normalized beta power, it is possible for subjects to have neurofeedback strategies that involve modulation of non-beta frequency bands to move the cursor. Using the same trial re-labeling procedure as described above, individual trial labels were re-assigned depending on normalized power in non-beta frequency bands (1-10 Hz, 10-25 Hz, 40-65 Hz, and 65-100 Hz) for the same time window as above. The resulting movement onset times are plotted by re-labeled group, frequency band, and monkey in 2.6D-F. For the 10-25 Hz and 40-65 Hz frequency bands there is ordering correlated with the 25-40 Hz band, and no consistent ordering across monkeys respectively. (Two-tailed Cuzick's Test, Monkey S 10-25 Hz: z = -0.6660, p = 0.5054, n = 1183, 40 – 65 Hz: z = -1.599, p = 0.1097, n = 1183, Monkey C 10-25 Hz: z = 5.717, p = 1.082, n = 2168, 40-65 Hz: z = 0.1569, p = 0.8753, n = 2168, Monkey G 10-25 Hz: z = 5.477, p = 4.33e-08, n = 1028, 40-65 Hz: z = -0.9347, p = 0.3500, n = 1028, Combined across Monkeys 10-25 Hz: z = 2.9728, p = 0.002951, n = 4379, 40-65 Hz: z = -0.4443, p = 0.6568, n = 4379). The consistent ordering in the 10-25

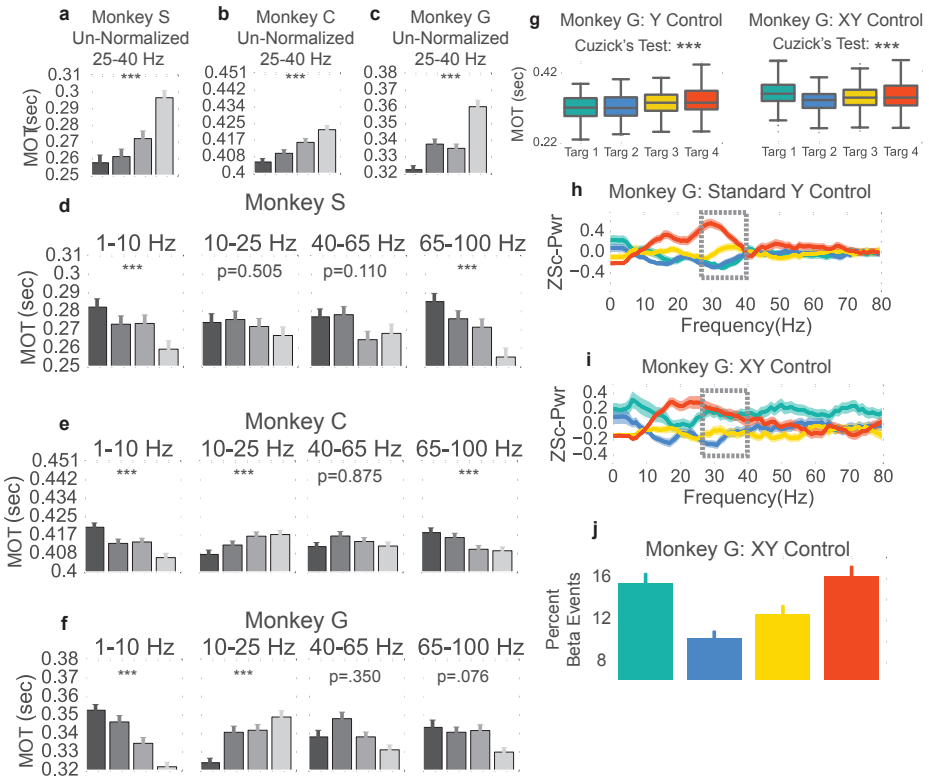

Figure 2.6: (a-c) Trials from Monkey S, C, G were re-labeled as low, mid-low, mid-high, and high according to the non-normalized beta power during time slice -0.8 to 0.0 sec with respect to the end of the neurofeedback epoch. The movement onset times (MOTs) of the resorted trials were compared and the mean (s.e.m) are plotted in each subplot. Below titles, p-values are shown for Cuzick's test (d-f) Same as a-c except using normalized non-beta frequencies indicated at top of plot (g) Right shows MOTs for version of task where subjects control X axis with normalized 1-10 Hz power in addition to Y axis with normalized 25-40 Hz. Note that Cuzick's test for the right plot assumes ordering is mid-low, mid-high, high, low targets (instead of low, mid-low, mid-high, high). Left, same MOT plot for Monkey G from Figure 3c for comparison (h, i) Z scored PSD plots (same method as Figure 2d-f) for different beta targets in the standard beta neurofeedback task (h) and the neurofeedback task that incorporates delta power modulation on the X axis (i). Dotted lines point out the ordering of targets in the beta range, following the movement onset time ordering in (g). (j) Percent of time points within the last 0.8 sec of the neurofeedback epoch that are part of beta oscillatory episodes during beta and 1-10 Hz XY control. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, Cuzick's two-tailed test for ordered grouping.

Hz band (increased power correlated with increased MOT) is likely due to the natural beta band for each animal extending into frequencies below 25 Hz. 2.3D-F shows that Monkey C and Monkey G exhibit increases in low beta frequencies that are match those in the 25-40 Hz range. The 65-100 Hz (gamma) band does exhibit consistently decreasing power for higher beta targets across monkeys (Two-tailed Cuzick's test, Monkey S: z = -5.0279, p = 4.96e-07, n = 1183, Monkey C: z= -4.227, p = 2.368e-05, n = 2168, Monkey G: z = -1.775, p = 0.0759, n = 1028, Combined across Monkeys z = -3.1079, p = 0.001884, n = 4379). Indeed, beta power and gamma power have been shown to be anti-correlated in motor-related regions during tasks involving movement [148, 149] in prefrontal cortex during working memory tasks [150], and in parkinsonian subjects at rest [151]. Increased gamma power may then be a physiological pattern that emerges with reduced beta power. It is unlikely that subjects are relying on changes in gamma power, which would change the denominator term in the beta cursor computation, to drive their neurofeedback strategy since gamma power constitutes less than 3% of the total broadband estimate, as shown in table 2.3. The correlation between increased gamma power and reduced MOT was further investigated with a model selection analysis. MOTs were either linearly estimated using normalized gamma power as a predictor (Model 2, table 2.2), normalized beta power as a predictor (Model 1, table 2.2), or both normalized gamma power and normalized beta power as predictors (Model 4, table 2.2) from the last 0.8 seconds of the neurofeedback epoch. The normalized beta power model explained more MOT variance than the normalized gamma power model (table 2.2), and the F-test demonstrated that adding normalized beta power as a predictor in a model with normalized gamma power resulted in significant improvement (Model 2 vs. Model 4, table 2.2). Thus, while gamma power is negatively correlated with MOT, beta power explains more MOT variance than gamma power, and addition of beta power to a model predicting MOT with gamma power significantly improves prediction. The 1-10 Hz band also shows a consistent across-monkey decrease in movement onset time with increased power, discussed below (Two-tailed Cuzick's Test, Monkey S z = -4.290, p = 1.785e-05, n = 1183, Monkey C: z = -6.8774, p = 6.097 e-12, n = 2168, Monkey G: z = -8.4548, p = 2.795e-17, n = 1028, Combined Across Monkeys z= -5.3049, p = 1.127e-07, n = 4379).

Modified beta neurofeedback task shows 1-10 Hz band power does not account for movement onset time increase The 1-10 Hz power subplot of 2.6D-F shows reduced movement onset time with increasing 1-10 Hz power for all three subjects. To investigate whether the movement onset time increase observed was truly due to changes in beta power and not changes in the 1-10 Hz band power, we performed an experimental manipulation as well as a regression analysis, as above with gamma power. In the experimental manipulation, Monkey G performed a NR task variant where beta power continued to move the beta cursor up and down the Y axis, but now instead of having a fixed X axis position, 1-10 Hz power controlled the cursor on the X axis. The targets were in the same positions in as in the standard NR task, but now Monkey G had to ensure that his 1-10 Hz power was nether too low nor too high else the horizontal position

Table 2.2: Predicting MOT with Beta and non-Beta Frequency Bands

	Model 1: MOT ~ Beta (R2)	**Model 2: MOT ~ Gamma (R2)**	**Model 3: MOT ~ LF (R2)**	**Model 4: MOT~ Beta+ Gamma (R2)**	**Model 5: MOT ~ Beta + LF (R2)**
Monkey S	0.0403	0.0148	0.006417	0.05806	0.04321
Monkey C	0.0138	0.004384	0.01253	0.01733	0.01456
Monkey G	0.0767	0.005467	0.05866	0.10561	0.07963
Combined	0.04363	0.01383	0.01820	0.05677	0.04522

	F statistic Model 2 vs. Model 4	**p (Model 4 >> Model 2)**	**F statistic Model 3 vs. Model 5**	**p (Model 5 >> Model 3)**
Monkey S	F (1181, 1180) = 54.162	p < 1e-16	F (1181, 1180) = 45.373	p < 1e-16
Monkey C	F(2166, 2165) = 28.531	p < 1e-16	F(2166, 2165) = 4.477	p < 1e-16
Monkey G	F(1026, 1025) = 114.77	p < 1e-16	F(1026, 1025) = 23.356	p < 1e-16
Combine Monkeys	F(4377, 4376) = 199.2	p < 1e-16	F(4377, 4376) = 123.83	p < 1e-16

Table 2.3: Percentage of Broadband Power Estimate Comprised by 65-100 Hz

	Mean (std) Percentage
Monkey S	2.43 +/- 1.13 %
Monkey C	1.50 +/- .518 %
Monkey G	2.14 +/- .825 %

of his cursor would not fall within the width of the beta target. Monkey G learned this task and after 3-4 days of practice achieved similar performance to the standard beta-only task of 5-10 sec to each beta target. In the beta task variant, Monkey G adopted a new strategy for getting to the lowest target. 2.6H and 2.6I show PSDs from the last 0.8 seconds of the neurofeedback epoch to the first 0.2 seconds of the reach epoch. For the lowest (green) target in the beta task variant, Monkey G managed to increase the power of his beta frequencies to similar levels as the highest beta target (red) but since he concomitantly increased the power of other frequencies, the denominator term in the normalized beta metric increased more, making the cursor move downwards (2.6I). To ensure that the PSD plot reflected the presence of beta oscillations, we also calculated the percent of on-beta time points using the previously explained time-domain method for the beta task variant (2.6J). This metric reflects the same ordering as in the PSD that the lowest beta target had a comparable percentage of on-beta time points to the highest beta target. This task variant effectively decoupled beta and 1-10 Hz power. In the original task, beta and 1-10 Hz power were inversely correlated (low beta power occurred with high 1-10 Hz power and vice versa) but in this modified task, high 1-10 Hz power and high beta power co-occurred during the low, green target and low 1-10 Hz power and high beta power co-occurred during the high, red target. We analyzed whether movement onset times followed the beta power or the 1-10 Hz power ordering. The movement onset times (2.6G, right) for the green target rose to match the movement onset times of the red target, indicating that the movement onset times followed beta power ordering, not 1-10 Hz power ordering (One-tailed Cuzick's test for significant ordering of beta targets in 2.6G, right assesses increasing movement onset times per the group order 2, 3, 4, 1 instead of group order of 1, 2, 3, 4 used in all other Cuzick's tests. Monkey G: $z = 7.1359$, $p = 4.807\text{e-}13$, $n = 1164$). If the 1-10 Hz power target ordering is used then MOTs show no significant trend (One-tailed Cuzick's test for MOTs increase with decreasing 1-10 Hz power, group order of 1, 2, 3, 4, Monkey G: $z = -2.5713$, $p = 0.9949$, $n = 1164$). Although 2.6D-F show a strong co-modulation of 1-10 Hz frequencies with beta frequencies, the 1-10 Hz band does not explain the observed ordering of movement onset times.

In addition to the above experimental manipulation, we also assessed the contribution of the 1-10 Hz band on explaining MOT variance. MOTs were either linearly estimated using the normalized 1-10 Hz power as predictor (Model 4, table 2.2), normalized beta power as a predictor (Model 2, table 2.2), or both normalized 1-10 Hz power and nor-

malized beta power as predictors (Model 5, table 2.2). The normalized beta power model explained more MOT variance than the normalized 1-10 Hz power model (table 2.2), and the F-test demonstrated that adding normalized beta power as a predictor in a model with normalized 1-10 Hz power resulted in significant improvement (Model 4 vs. Model 5, table 2.2). Thus, while 1-10 Hz power is negatively correlated with MOT, beta power explains more MOT variance than 1-10 Hz power, and addition of beta power to a model predicting MOT with 1-10 Hz power significantly improves prediction. Finally, we investigated which sub-frequency bands within the 1-10 Hz band were most closely correlated with MOT. We divided the low frequencies into the delta band (1-3 Hz), theta band (4-7 Hz), and alpha band (8-12 Hz). By performing the same analysis as in Figure 5d-f with the narrower bands, we find that delta and theta bands, but not alpha band, strongly correlate with MOT in all three animals (Monkey S: 1-3 Hz: z = -3.276, p = 1.052e-03, n = 1183, 4-7 Hz: z = -3.245, p = 0.001174, n = 1183, 8-12 Hz: z = -1.3085, p = 0.1907, n = 1183, Monkey C: 1-3 Hz: z = -6.9864, p = 2.821e-12, n = 2168, 4-7 Hz: z = -6.1334, p = 8.602e-10, n = 2168, 8-12 Hz: z = -1.779, p = 0.0753, n = 2168, Monkey G, 1 – 3 Hz: z = -8.2012, p = 2.379e-16, n = 1028, 4 – 7 Hz: z = -6.1655, p = 7.0239e-10, n = 1028, 8-12 Hz: z = 0.0669, p = 0.9467, n = 1028, Combined over Monkeys: 1-3 Hz: -5.0064, p = 5.55e-07, n = 4379, 4 – 7 Hz: z = -4.355, p = 1.33e-05, n = 4379, 8-12 Hz: z = -0.9782, p = 0.328, n = 4379).

2.4 Conclusion

2.4.1 Behavioral Correlates of Beta Oscillations

Here we have shown evidence that volitionally increasing and decreasing beta power in the motor system with neurofeedback achieves neural states that precede slower and faster movement onset times respectively in three monkeys. How do these results jive with previous hypotheses?

We first consider how beta oscillations may be generated in the motor system. Our result that volitionally increasing and decreasing motor cortical beta power with neurofeedback precedes slower and faster movement onset times supports the hypothesis that beta oscillations are linked to neural patterns that slow onset of new movements [152]. Our results add to previously reported findings of elevated beta power prior to and after well-trained movements [153, 107, 127]. While we cannot identify the mechanism that drives the beta oscillations observed, modeling and in vitro slice work shed light. Recent modeling of striatal neural populations show increased medium spiny neuron (MSN) excitation can result in beta oscillations within the striatum, which can propagate through output structures of the basal ganglia [154]. MSN excitability is affected by many neuromodulators such as acetylcholine [155] which notably drives increased MSN excitability primarily in D2 MSNs, or the MSNs responsible for the indirect pathway activation [156]. The findings from our study show that when cells generate beta oscillations they encode

a slower movement state, which could reflect indirect pathway activation. Recent work has also shown that shifts in attention due to salient stimuli are thought to involve the intralaminar nuclei of the thalamus [157], which projects to the striatum [158], potentially resulting in transient increases in beta power in D2 MSNs as well [154]. This common striatal beta-generating mechanism would explain how increases in attention have been reported to evoke beta oscillations [126, 159, 128], and could be used to pause current motor programs in response to salient stimuli [156]. This mechanism is also a plausible explanation for evidence of beta oscillations occurring during untrained, free reaching movements [114, 153]. These oscillations could be driven by salient stimuli that subjects encounter as they execute and update their internally generated motor plan. Our results and proposed mechanism of beta generation do not predict oscillatory events occurring during isometric contraction [106, 108], however it is becoming increasingly common to find different mechanisms for generating similar frequency oscillations [155]. In vitro slice work has identified that with sufficient excitatory drive to slices of sensorimotor or motor cortex, beta frequency oscillations emerge in deep cortical layers [104, 160]. It is possible that the strong excitatory drive needed to stiffen muscles during an isometric contraction task, in contrast to reaching movements that require temporal coordination of antagonist muscle groups [109], is sufficient to generate beta oscillations by the same means described by Roopun et al., 2006 and Yamawaki et al., 2008. Further evidence for this proposed mechanism comes from computational models of driving motoneuron recruitment with pyramidal tract neurons. When pyramidal tract neurons fire at beta band or higher frequencies, motoneurons increased recruitment and hence muscular force production [161], as would be required in an isometric contraction task. Potentially beta oscillations are observed during isometric contraction tasks because of large muscular force requirements in the task, not because of the same striatal beta-generation pathway previously described. Finally, our results and the proposed striatal beta-generating mechanism do not predict the correlation of beta power with other behavioral metrics such as movement onset speed, peak speed, and onset acceleration that were correlated with beta power in other studies [147, 140]. Note though, that the findings of Joundi et al., 2012 and Pogosyan et al., 2009 are from experiments using transcranial alternating current stimulation (tACS) at beta frequencies applied to motor cortical areas. The mechanisms of tACS are still unclear [162], so it is possible that the reported behavioral effects are due to evoked neuronal activity patterns that are specific to tACS stimulation and do not occur endogenously. While identifying mechanisms that generate beta oscillations can shed light on how certain types of behavior such as attention or isometric contraction may be correlated with onset of beta oscillations, they do not inform how the dynamics of the underlying neural population generate movement change.

2.4.2 Benefits of the Sequential Neurofeedback-Reaching Task

Using a neurofeedback paradigm to investigate behavioral and population neural correlates of oscillations has several advantages. First, since the neurofeedback epoch only

requires subjects to modulate beta power and to be seated at rest, they choose their own subject-specific strategy for generating or quenching beta activity. These strategies may include co-modulating other frequency bands, imagining movement, or performing other internal behaviors that generate beta activity. For example, while Monkeys C and G inversely modulated low frequency (1-10 Hz) power with beta frequency power, Monkey S did not modulate 1-10 Hz power as drastically but did exhibit increased 50-60 Hz power with increased beta power (2.3D-F). Despite the different approaches that were taken to increase and decrease beta power across animals there is still a consistent effect of high versus low beta power on movement onset times, increasing confidence that the oscillation is a reliable marker of the observed behavior. In contrast, many motor tasks engage motor preparation, increased attention, cue expectation, and possibly muscular stiffening all the same time within the task. These overlapping behaviors make it challenging to deliver a parsimonious explanation for the behavioral correlates of beta oscillations using correlational studies. Another advantage of using neurofeedback over other approaches to perturb neural oscillations such as non-invasive transcranial alternating current stimulation [147, 140] or invasive transcranial electrical stimulation [133] is that the recorded neural signal is not tarnished with a stimulation artifact. In this study, simultaneously recorded units were analyzed and shown to exhibit different activity during neurofeedback-induced beta oscillations compared to natural beta oscillations occurring during typical reaching tasks. Despite differences at an individual unit level, population level analyses show that beta oscillations promote a consistent movement-slowing state during both the CO and NR tasks, matching what is observed behaviorally. This analysis may be compromised if a stimulation artifact prevented recording of local field potentials or single units.

Another report, [144], used a similar but non-invasive beta band neurofeedback method prior to a movement task and reports comparable behavioral results. The authors find that 3 of the 8 subjects exhibit significant reductions in movement onset time following reduced beta power, consistent with our findings. Possibly the movement onset increase was not in all subjects because there was substantially more temporal smoothing in their neurofeedback task setup (neural signals averaged in window of 1 seconds, compared with our window of 200 ms), and no hold requirement for their neurofeedback cursor (compared to our hold requirement of 450 ms). Thus, their subjects could be in a greater range of neural states prior to beginning the movement task, making the movement onset versus beta target relationship less robust. Finally, the authors do report a group-level significant increase in movement accuracy following beta reduction, a metric that did not change in our experiment likely due to the subjects' overtraining of arm reaches in our study.

Finally, neurofeedback is a tool that if effective at introducing a change in subsequent behaviors could possibly be a directly translatable therapy for patients. For example, if excessive synchronization of motor neurons is pathological in PD, learning to reduce neural coupling with neurofeedback of beta oscillations may improve bradykinesia symptoms. Evidence suggests that PD patients do indeed exhibit a stronger movement-related beta power reduction prior to movement onset than non-PD patients [163], implying that some patients may already reduce beta power to initiate movement more easily. Thus,

neurofeedback could be a tool to teach patients to cognitively modulate their beta power for symptom improvement, as discussed in chapter 4.

Chapter 3

Individual and Population Neural Activity during Neurofeedback Control and Naturally Occurring Beta Band Oscillations

3.1 Introduction

3.1.1 Neural Signals are Characterized During Beta Band Oscillations and During Motor Behavior, but Not During Both

In chapter 2, we discussed the behavioral correlates of beta oscillations, and an experiment designed to test the correlations between motor cortical beta oscillations and motor behavior. We now turn to the question, of how the presence of beta oscillations reflects underlying population signals in motor cortex. In the study described in chapter 2, a neurofeedback-reaching (NR) task was used in concert with simultaneous, multi-scale, high-count neural recordings to first study how the presence of beta oscillations influence arm-reaching behavior, and second, how underlying neuronal population patterns shift when beta oscillations are generated. It has been suggested that the generation of the beta oscillation itself, either locally or distally, could influence neuronal computation [164, 97] since slowly oscillating ephaptic fields have been shown to entrain spiking behavior in vitro [100]. However, little evidence exists showing that beta oscillations in the local field potential influence spiking activity through ephaptic mechanisms. Thus, in the subsequent analyses and discussions, we interpret beta oscillations as a statistic of synchronization of the underlying neural signals, not as a signal that can independently and causally influence neural spiking through ephaptic effects. We aim to investigate how the underlying neural signals change their encoding during epochs when they generate beta oscillations and do not make claims about the causality of beta oscillations on spiking.

Many proposed behavioral correlates of beta oscillations exist, as described in chapter 2, but to link oscillations to a behavior rigorously it is necessary to understand how oscillations reflect the underlying neural activity that ultimately drives the behavior. Prior studies investigating neural activity changes during beta oscillations have used acute, single-electrode recording preparations and shown that single cells are synchronized to ongoing oscillations but that the strength of this synchronization is unrelated to the involvement of the neuron during the motor task [95]. Further, individual cells do not change their mean spike firing rate but do exhibit a reduction in spiking variability during oscillations compared to before the oscillation [95]. How might these changes in individual units relate to attention, motor preparation, or idling? Modeling groups have aimed to bridge this gap by showing how beta oscillations could be a signal generated by cells conveying top-down information [164, 125], could reflect a pattern of firing used to activate specific cell assemblies [165, 98], or could reflect a specific spatiotemporal recruitment of cells [166, 115]. However, experimental evidence supporting that beta-generating spiking patterns are used to accomplish the proposed functions is lacking. In contrast, if one were to omit the role that beta oscillations may play in motor behavior, there is substantial work linking spiking patterns to specific aspects of motor behavior such as movement onset [83], reaction time [167, 168], movement angle [169, 170], and movement speed [171, 172], to list just a few. We have already reported changes in motor behavior following

performance of neurofeedback during a sequential neurofeedback-reaching task, and now we turn to analyzing a neural population shift that mirrors the change in motor behavior observed. Notably, this shift in neural population was also seen during naturally occurring beta oscillations during reaching tasks suggesting that beta oscillations reflect a common underlying subset of spiking patterns even in different task contexts. Taken together, the behavioral results plus the population analysis ties together existing works on behavioral correlates of beta oscillations with hypotheses of how motor cortex encodes movement onset through the lens of population level neural activity.

3.2 Methods

In the experiments described in chapter 2 with Monkeys G and C, single and multi-unit activity was recorded. All population unit analysis was only conducted with these two animals.. Monkey S had arrays that had been implanted for > 3 years and no longer were usable for recording single-units. All neural data from sessions from Monkey G were offline sorted using the Plexon Offline Sorter. Isolated single units and multi units were included in analysis. For Monkey C, channel activity (Chestek et al., 2011) was used (see Surgery, electrophysiology, and experimental setup). Analyses were performed within day to prevent day-to-day recording instability from influencing analysis.

3.3 Results

3.3.1 Relationship of Individual Units to Beta Band Oscillations During NR Task and CO Task

First we consider that if the neurofeedback-induced beta oscillations are qualitatively the same as naturally occurring beta oscillations during reaching tasks, it might be unsurprising that increasing beta power biases subjects toward slower movement onset based on previous studies. We investigate exactly how similar the beta oscillations in the different tasks are through the lens of unit neural activity. On most days, subjects performed 5-10 minutes of the CO task prior to beginning the NR task. Only days when the CO task was performed were used for subsequent analysis (Monkey G: 6 days, Monkey C: 4 days). Simultaneous single-unit and multi-unit activity were recorded throughout Monkey G's task sessions, and multi-unit and channel level activity [27] were recorded throughout Monkey C's task sessions. Both single-unit and multi-unit activity were manually sorted, whereas channel level activity used the auto-thresholding function in Omniplex-D software. In subsequent analyses, time bins (100ms or 25ms depending on analysis) will be labeled as on-beta or off-beta referring to whether they fall within or outside a beta oscillation. A beta oscillation is defined as periods in which beta amplitude is above 60th percentile for at least 125 ms (same definition used in chapter 2 when computing percentage of time point on-beta, e.g. in 2.6J). Bins will also be referred to as slow or fast referring to

whether the mean hand speed within the bin is below or above 3.5 cm / sec. The 'slow' versus 'fast' bin distinction was made to separate bins that were before movement onset from ones after movement onset (approximate movement onset time occurred when hand velocity crossed 3.5 cm / sec). The point in the trial corresponding the cue for movement onset is referred to as the go cue in both tasks (corresponding to the end of neurofeedback epoch in the NR task).

We first assess whether units fire at similar rates during CO task beta oscillations and NR task beta oscillations. Go cue aligned trials were aggregated for each task with each trial lasting 2.5 sec (1.5 sec before go cue through 1.0 sec after go cue). Unit activity was binned into 100 ms bins yielding 25 bins per trial. For every trial, bins that were labeled as slow and on-beta were selected. The distribution of spike counts for these slow, on-beta bins from the NR task was compared to the slow, on-beta events from the CO task. Counts from fast bins were not used in the analysis because there were very few fast bins that were also on-beta. Example mean firing rates of four consecutive time bins in a row aligned to onset of on-beta are shown for two example single unit recordings (unit 101a and unit 1a) from Monkey G in 3.1A where red is the mean firing rate for NR slow, on-beta bins and blue is the mean firing rate for CO slow, on-beta bins. 3.1B shows the fraction of units exhibiting significantly different mean firing rates between the slow, on-beta bins from the two tasks on each day (Mann-Whitney test, $p < 0.05$, number of units recorded per day displayed above bar). Each day, 40-60% of units from Monkey G (15-20% of units from Monkey C, inset) exhibited significantly different firing rates for NR versus CO slow, on-beta bins. While many individual cells showed changes in mean firing rate during slow, on-beta bins across the two different tasks, it is possible that units could still exhibit a consistent spike rate change in response to beta amplitude changes. We used methods adapted from [98] to fit a continuous beta amplitude-to-spike rate mapping for the NR and CO task to determine if the units' beta amplitude-to-spike rate correlations are consistent across tasks. Briefly, the logarithm of instantaneous beta amplitude was calculated for the entire 2.5 second epoch and was then correlated against the firing rate of each cell. Three example units are shown in 3.1C-E where the red and blue traces are the relationship between cell firing and beta amplitude for the NR task and the CO tasks respectively. Some units exhibit similar mean firing rates but different beta-to-spike rate slopes (3.1C), some exhibit different mean firing rates but similar beta-to-spike rate slopes (3.1D), and some exhibit different mean firing rates and different beta-to-spike rate slopes (3.1E). To assess whether units exhibit similar beta amplitude-to-spike rate slope across the tasks, we compared within-task and across-task slope estimates. First, two non-overlapping subsets of the CO (CO1, CO2) and NR (NR1, NR2) tasks were used to estimate separate beta amplitude-to-firing rate slopes per unit (CO1, slope, CO2, slope and NR1, slope, NR2, slope). Note that slightly overlapping CO sets were used for Monkey C due to limited CO data. Then, the two slope estimates for each task are correlated to assess within-task slope estimate stability (3.1F: CO1, slope vs. CO2, slope and 3.1G: NR1, slope, vs. NR2, slope). In 3.1F-G, each plotted marker corresponds to a unit and its color corresponds to day on which it was recorded (Monkey G: main

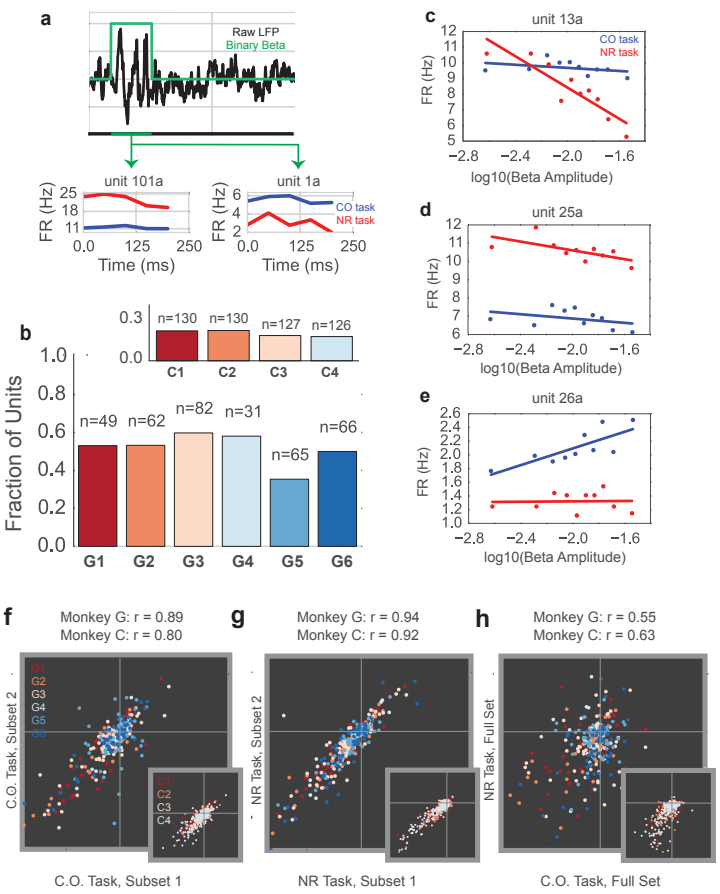

Figure 3.1: (a) Schematic of an LFP trace (in black) with a time period corresponding to an on-beta time period (in green). During oscillatory events, example mean firing rates are shown for two example single units from Monkey G (unit 101a, unit 1a) where the red trace is for on-beta bins during the neurofeedback epoch and the blue trace is for slow, on-beta bins during the CO task. Graphs are aligned to starting bin of beta event (b) For each day (main plot Monkey G: days G1 – G6, subplot Monkey C: days C1 – C4), a bar plot indicates the fraction of units that exhibit significantly different firing patterns during slow, on-beta time points in the CO and NR task assessed by the Mann-Whitney U test ($p < 0.05$). Number of units recorded per day are printed above each bar (c-e) Example beta amplitude-to-spike rate mappings for single-units from a day. Mappings in red are from the CO task. Mappings in blue are from the same unit on the same day during the NR task. (f, g) Stability of beta-to-rate slope estimates from subset #1 versus subset #2 of CO (f) and NR (g) tasks (Main plot Monkey G, subplot Monkey C). R values indicate mean correlation coefficient between slopes computed from subset #1 and from subset #2 across days (h) Comparison of slopes from subset #2 CO versus subset #1 NR tasks. R values indicate mean correlation between slopes from subset #2 CO task versus slopes computed from subset #1 NR task across days.

plot, Monkey C: inset, note same colormap as 3.1B). The printed correlation coefficients are the mean correlation coefficient across days and describe how well a linear regression captures the correlation between slope estimates from subset 1 vs. slope estimates from subset 2. For both tasks, correlation coefficients exceed 0.8. These high correlation coefficients suggest a stable within-task beta-to-rate mapping.Across task slope estimates are visually compared in Figure 6h (COall data vs NRall data). In contrast to the stable within-task slope estimates, across task slope estimates are less correlated across units.. To assess whether the within-task and across-tasks slope differences are significant, a paired Student's t-test is performed to assess the differences between within-task and across-tasks between CO1 vs. CO2 slopes, NF1 vs. NF2 slopes, and CO2-NF1 slopes where units from each day are treated as independent observations (Monkey G: CO1 vs. CO2, t = 1.275, p = 0.2032, n = 355 units, NF1 vs. NF2, t = 0.0169, p = 0.9866, n = 355 units, CO2-NF1, t = 2.3403, p = 0.0198, n = 355, Monkey C: CO1 vs. CO2, t = -1.5075, p = 0.1323, n = 513 units, NF1 vs. NF2, t = -1.309, p = 0.1910, n = 513 units, CO2-NF1, t = 9.880, p =3.473e-21, n = 513, Combined Across Monkeys: CO1 vs. CO2, t = 0.6258, p = 0.5316, n = 868 units, NF1 vs. NF2, t = -0.5547, p = 0.5793, n = 868 units, CO2-NF1„ t = 5.1782, p = 2.786e-07, n = 868 units). The subset comparison of CO2-NF1 was randomly chosen to report -- CO2-NF2, CO1-NF1, and CO1-NF2 also show the same difference in CO vs. NF slopes. Thus, both the mean firing rates of units during slow, on-beta time points as well as the continuous beta amplitude-to-spike rate mappings across tasks are different for many units. Given that the behavioral effect from neurofeedback induced beta oscillations matches well with hypotheses claiming a movement-slowing role of beta oscillations during natural movements [152], it was surprising that individual unit responses during beta activity were so different across tasks.

3.3.2 Relationship of Population Signals to Beta Band Oscillations During NR Task and CO Task

While many individual units exhibit different spiking patterns during beta oscillations in the CO and NR tasks, population-level activity across the two tasks could still exhibit consistent patterns. Specifically, since increased beta oscillations during the NR task exhibit slower MOT times, beta oscillations may reflect shifting neural population patterns in a way that affects movement onset. To assess the relationship between CO and NR spike activity patterns with respect to movement onset, we train a classifier on within-day CO spiking activity to discriminate bins occurring pre and post movement onset (PreMO, PostMO). We then use the preMO and postMO neural population activity from the same day's NR task and assess first whether the same CO-trained classifier successfully distinguishes PreMO and PostMO in the NR task, and second how the presence of beta oscillations influences the separation of the two labeled classes. The approach used to discriminate PreMO and PostMO neural population activity was to train a logistic regression classifier on the first two-thirds of the CO spike counts. CO spiking activity was

binned in 25 ms, each unit was z-scored according to its mean and standard deviation during the CO task, and each 25 ms bin was labeled as PreMO or PostMO. A logistic regression classifier was trained on the binned spike counts with an additional 2 bins of history (number of spike features per observation equal to 3 x number of neurons, and each 3-bin set is referred to as a 'chunk'). The trained classifier yields a probability of each chunk being PreMO and PostMO:

$$p(y_i = 1) = \frac{1}{1 + e^{(\beta 0 + \beta 1 \times x_i)}}$$
$$p(y_i = 0) = 1 - p(y_i = 1)$$

where β0 and β1 are the intercept and neural weights found by the logistic regression classifier respectively. By setting a threshold on these probabilities we can assign predicted PreMO or PostMO labels. For example, if the threshold is 0.5 and the probability of an observation being PostMO is greater than 0.5, the chunk would be assigned as PostMO else it would be assigned as PreMO. Typically in logistic regression, a threshold of 0.5 is used to classify the two classes, where observations with greater than 0.5 would be assigned the label of '1' and less than 0.5 would be assigned the label of '0'. Training with unbalanced groups can result in other threshold values being optimal which are typically discovered with an ROC curve analysis (Bradley, 1997). We find optimal thresholds for maximizing percent correct classification are 0.5 and 0.315 for Monkey G and Monkey C. These values are the MO thresholds in Figure 7.

We found that actual PreMO and PostMO chunks exhibited significantly different distances to MO threshold for the held-out 1/3 of data from the CO task (3.2A, Blue: CO task, paired two-tailed Student's t-test on mean within-day probabilities for CO PreMO and PostMO, Monkey G: n = 6, T = -20.899, p = 4.65e-6, Monkey C (inset): n = 4, T-4.2711, p = 0.0236, Combined Across Monkeys: n = 10, T = -4.1097, p = 0.002638). Thus, population spike count chunks reliably encode before and after movement onset in the CO task. Note that Monkey G does exhibit about an order of magnitude more reliable separation between PreMO and PostMO than Monkey C (y axis in 3.2A), and this is likely due to the lower neural signal quality in Monkey C (implanted ~3 years prior to study without resolvable single or multi-units) than Monkey G (implanted only ~1 year prior to study, with resolvable single and multi-units). To assess whether the same spiking patterns were present during PreMO and PostMO in the NR task, the CO-trained classifier was used to predict the PreMO and PostMO labels of z-scored spiking activity chunks from NR trials. Note that in the NR task, chunks are labeled as 'PreMO' during the neurofeedback epoch of the NR task and before movement onset during the reaching epoch of the NR task, and labeled as 'PostMO' after movement onset during the reaching epoch. Given the differences in individual unit firing patterns across the CO and NR task (3.1), it is possible that the neural population activity also varies drastically across tasks and that the CO-trained classifier may not perform well when given NR population activity. Instead, we confirm that the same CO-trained classifier does yield significantly different distances to MO threshold for the NR task (3.2A, Red: NR task,

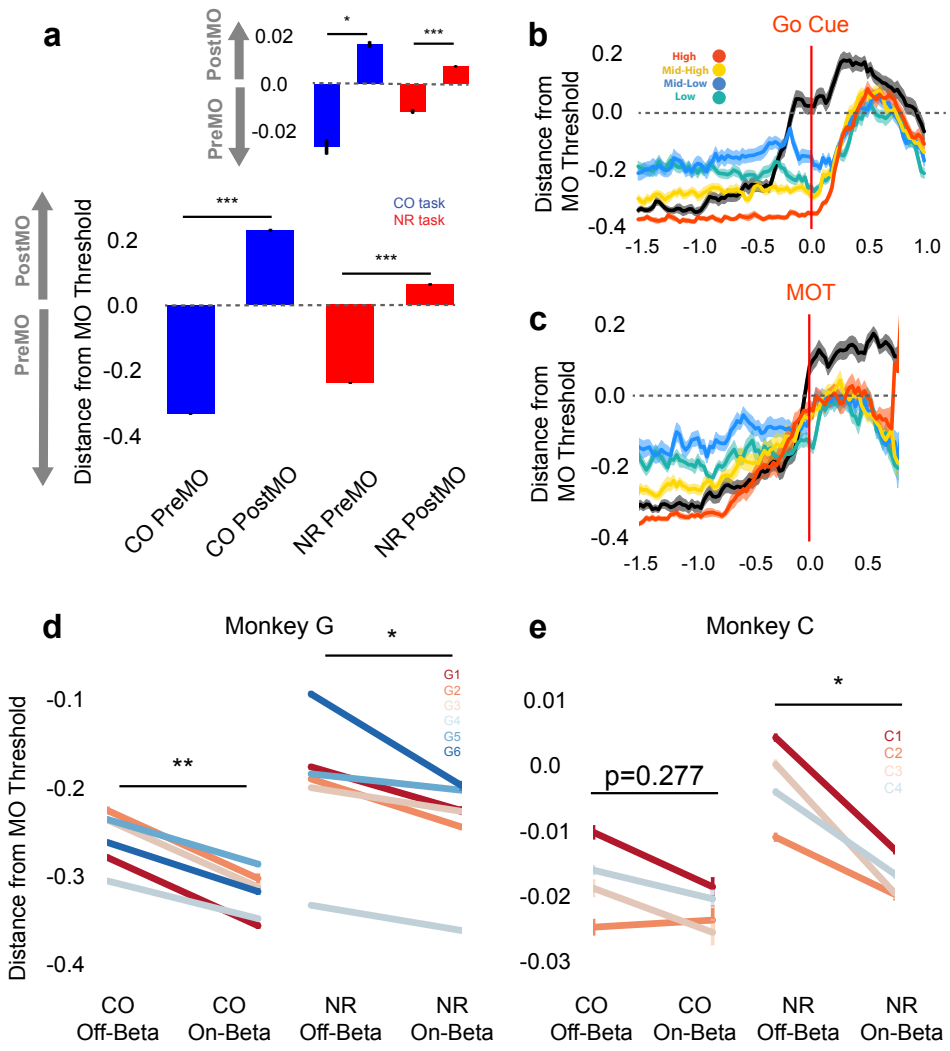

Figure 3.2: (a) Distance to MO threshold for the CO (blue) and NR (red) tasks for PreMO and PostMO time points from Monkey G and Monkey C (inset). Bars less than and greater than zero indicate mean scores predicting PreMO and PostMO respectively. (b) Example mean (s.e.m) of distance from MO threshold as a function of time to go cue for CO trials (black) high beta target NR trials (red), mid-high beta target NR trials (yellow), mid-low beta target NR trials (blue), and low beta target NR trials (green). At the go cue, distances to MO threshold are greatest for high beta target, and lowest for low beta targets. (c) Similar to (b) but aligned to MOT instead of go cue. Trials converge to MO threshold at MOT. (d-e) The mean distance from MO threshold for slow, preMO, off-beta and slow, preMO, on-beta time points during the CO (left) and NR (right) tasks for Monkey G and Monkey C respectively. Individual lines connect mean off-beta and on-beta distances (s.e.m) for individual sessions.

paired two-tailed Student's t-test on mean within-day probabilities for CO PreMO and PostMO chunks, Monkey G: n = 6, T = -7.4593, p = 6.83e-4, Monkey C (inset): n = 4,T = -13.591, p = 8.62e-4, Combined Across Monkeys: T = -3.6193, p = 5.578e-3, n = 10). This finding validates that the population reliably encodes gross kinematics similarly across tasks despite the individual unit activity changes observed in the previous analysis.

Given that the CO-trained movement onset classifier maintains its predictability in the NR task, we can now ask how the presence of beta oscillations influences the distance to the MO threshold. It is possible in both the CO task and NR task that the presence of beta oscillations does not affect the distance to MO threshold since it is known that cells more involved in movement are no more likely to be entrained to beta oscillations than signals not involved in movement [95]. The classifier thus may have captured reliable movement signals from units that are not entrained by beta oscillations, making the distance to MO threshold unaffected by the presence of the oscillations. A second possibility is that the presence of beta oscillations keeps neural activity further away from the MO threshold only during CO trials, but not NR trials. Since the classifier weights were trained on the CO task the classifier weights may reflect beta-related structure in the population that is useful for predicting CO bins to be PreMO. If the neurofeedback beta-related structure in the population is different than it is in the CO case, as is suggested on an individual unit basis by 3.1, then the population spiking patterns occurring during beta oscillations during the NR task may not exhibit similarly greater distances from the MO threshold. A final possibility is that despite the differences in individual unit firing patterns during the CO and NR beta oscillations, production of beta oscillations requires a consistent shift in population activity that is also related to movement onset. By comparing the signed distance to the MO threshold of each of the beta targets during NR trials and CO trials during on-beta and off-beta periods we can begin to discriminate amongst the three possibilities. 3.2B-C show the mean signed distance to MO threshold for beta targets and CO trials from one representative day (Monkey G, session 5). In 3.2B, traces are aligned to the go cue, showing that during the neurofeedback epoch prior to the go cue, the high beta target (red) maintains a much greater distance from the MO threshold than the low and mid-low beta targets (green and blue). At the time of the go cue the high beta target is further from the MO threshold than the lower beta targets, suggesting that subjects must traverse a greater neural distance to arrive at the MO threshold, which may take longer resulting in a longer MOT. When aligning the same trials to the MOT, all traces converge around the MO threshold showing that all NR trials must arrive at the same MO threshold to initiate movement. These examples suggest that the presence of beta oscillations is an indicator of subjects' neural population being far from movement initiation. It is possible, however, that the decoder has identified discriminative firing patterns that are unrelated to the presence of beta oscillations. Since subjects are performing a neurofeedback task with different strategies for different beta targets, the differences observed in 3.2B-C could be related to their distinct beta-target strategy instead of the actual presence of beta oscillations. In 3.2B-C, the green line corresponding to the lowest beta target is further from the blue line corresponding to the

Table 3.1: Unit Properties of Chosen vs. Unchosen and Chosen+ vs. Chosen- : All tests are Kruskal Wallis test for differences in median

	Chosen	**Unchosen**	**Classifier Wt.**	**Beta-to-FR Slope**	**Mean Modulation (Hz)**	**Mean Firing Rate (Hz)**	**Beta Rhythmicity**
Monkey G	N = 104	N = 251	H = 32.42, p = 1.24e-8	H = 24.78, p=6.41e-07	H = 0.97, p=0.324	H = 14.40, p=1.48e-04	H = 0.229, p = 0.632
Monkey C	N = 171	N = 342	H = 244.2, p=4.72e-55	H = 145.0, p = 2.1e-33	H=1.195, p = 0.274	H = 139.6, p = 3.3e-32	H = 0.364, p = 0.546
Combine	N = 275	N = 593	H = 213.4, p = 2.5e-48	H = 132.3, p=1.30e-30	H = 0.004, p = 0.947	H = 99.16, p = 2.3e-23	H = 0.016, p = 0.897
	Chosen +	**Chosen -**	**Classifier Wt.**	**Beta-to-FR Slope**	**Mean Modulation (Hz)**	**Mean Firing Rate (Hz)**	**Beta Rhythmicity**
Monkey G	N = 85	N = 19	n/a	H = 35.72, p = 2.27e-9	H = 0.009, p = 0.92	H = 7.13, p=0.00756	H = 1.56, p = 0.2116
Monkey C	N = 170	N = 1	n/a	H = 2.50, p = 0.114	H = 0.673, p = 0.412	H = 2.965, p = 0.0851	H = 0.059, p = 0.8079
Combine	N = 255	N = 20	n/a	H = 39.30, p = 3.6e-10	H = 1.164, p = 0.281	H = 3.702, p = 0.054	H = 1.765, p = 0.184

mid-low beta target, showing that the presence of more beta oscillations in the mid-low beta target is not the only factor in determining distance from the MO threshold. To directly test if the presence of beta oscillations affects distance to the MO threshold, we collapse all NR data across beta targets, isolate slow chunks that occur prior to actual movement onset, and compare distance to MO threshold for on-beta and off-beta bins in both the CO and NR task. 3.2D-E show mean on-beta and off-beta distances to MO threshold for individual days in the CO (left) and NR (right) tasks. In all cases but one, off-beta slow chunks exhibit significantly closer distances to MO threshold than on-beta slow chunks (paired Students' t-test of within-day means 3.2D: Monkey G: CO off-beta vs. on-beta, n = 6, T = 6.7423, p = 0.001089, , NR off-beta vs. on-beta, n = 6, T = 2.6073, p = 0.0478, 3.2E: Monkey C: CO off-beta vs. on-beta, n = 4, T = 1.3245, p = 0.2772 NR off-beta vs. on-beta n = 4, T = 4.7632, p = 0.0176, Combined Across Monkeys: CO off-beta vs. on-beta, n = 10, T = 4.2458, p = .00216, NF off-beta vs. on-beta, n = 10, T = 3.1851, p = 0.0111). The CO and NR tasks exhibit common population level activity changes reflected by the onset of beta oscillations, and specifically, these population changes encode a shift further away from the MO threshold.

To determine which types of units contributed most to the success of the logistic classifier, we selected units that fulfilled two criteria. First, a logistic classifier trained

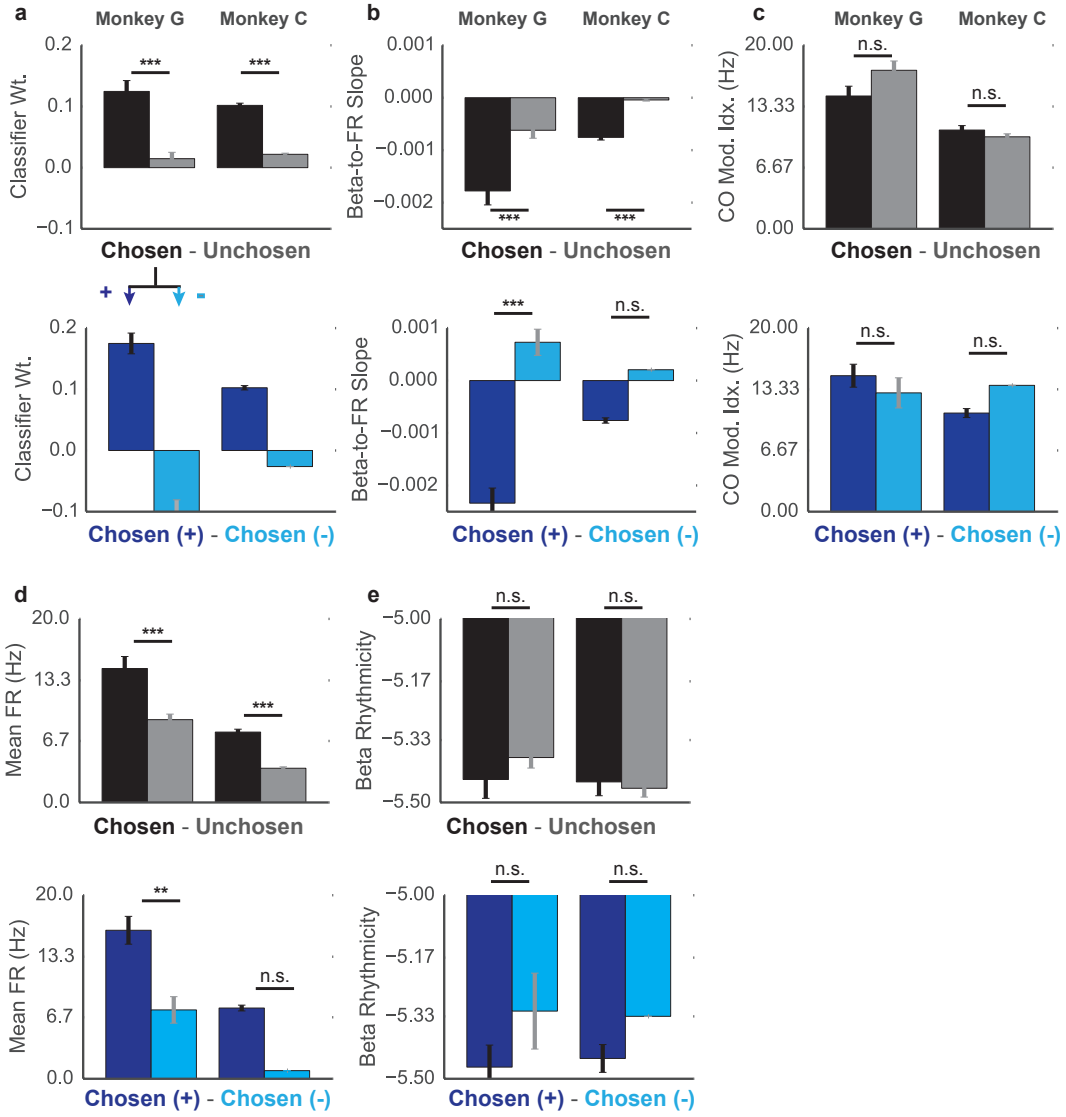

Figure 3.3: Top shows differences between classifier weight β1 for chosen and unchosen units (top) and chosen+ and chosen- units (bottom) for Monkey G (left) and Monkey C (right). Same layout for (b) beta-to-FR slope (as in 3.1f-h), (c) CO task modulation, (d) mean firing rate, and (e) beta rhythmicity. Differences are assessed with Kruskal Wallis test and reported along with sample sizes in 3.1. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

only on chunks from the individual unit in 2/3 of the CO task had to predict significantly lower scores for PreMO than PostMO in held-out data in the CO and NF tasks. The second criterion was that units exhibit lower predicted scores for 'on-beta', slow chunks than 'off-beta', slow chunks in both tasks. The units that fulfilled both criteria were referred to as 'chosen' units, and were compared to all other 'unchosen' units (units collapsed across days -- Monkey G: 104 chosen units, 251 unchosen units, Monkey C: 171 chosen units, 342 unchosen units). We analyzed the differences between chosen and unchosen weights using five different metrics—classifier weight, beta-to-firing rate slope during CO task, mean modulation during the CO task, mean firing rate during the CO task, and beta rhythmicity during the CO task (3.3, table 3.1). We found the weights of chosen units exhibited significantly higher and more positive weights than unchosen units, indicating that most chosen units increased their firing rate during PostMO compared to PreMO. Chosen units also tended to have higher firing rates. Further, chosen units show significantly lower beta-amplitude-to-firing-rate slopes during the CO task than unchosen units, consistent with the finding that firing rates increase during movement concomitantly with beta amplitude decreases. Chosen units and unchosen units showed no difference in their mean task modulation, and no difference in their beta rhythmicity (3.3, table 3.1). Lastly, to assess any difference between the chosen units with a positive $\beta 1$ classifier weight and the chosen units with negative $\beta 1$ classifier weight, the same analyses were performed between these groups (termed 'Chosen +' units and 'Chosen –' units, Monkey G: 85 Chosen+ units, 19 Chosen – units, Monkey C: 170 Chosen + units, 1 Chosen – unit). The Chosen+ units and Chosen- units respectively increase and decrease firing rate during movement. Thus, we also find that the Chosen+ group exhibits significantly lower beta-to firing rate slopes. The groups do not show any difference in mean task modulation. The positive units exhibit significantly higher firing rate. Finally, the negatively modulated units have a non-significantly higher 'beta rhythmicity' than the positive units. Overall, units that contribute most to the classification of PreMO versus PostMO tend to be high firing units that are positively modulated with movement onset, yet are no more task modulated. Within this group are a few units that instead fire less with movement that are generally lower firing rate and possibly more entrained to ongoing beta oscillations in the local field potential. These Chosen- units are possibly a distinct subpopulation of single and multi-units that may act as pacemaker cells for the population [121], though more data and a more careful characterization of firing properties, ISIs, and waveform would be needed to make this claim.

3.4 Conclusion

In chapter 2 we presented evidence that volitionally increasing and decreasing beta power in the motor system with neurofeedback achieves neural states that precede slower and faster movement onset times respectively in three monkeys. These results support the hypothesis that beta oscillations in the motor system reflect neural patterns that are far

from a movement onset neural state. Importantly, we use simultaneously recorded single and multi-unit activity during the NR and CO tasks to characterize how the presence of beta oscillations reflects changes in underlying neural population activity that ultimately drives the behavioral changes observed. During the neurofeedback epoch of the NR task, population neural activity exhibits greater distances from the computed MO threshold in the presence of beta oscillations, and shorter distances when there are no oscillations. We emphasize that this result is not merely driven by the observation that there are often more beta oscillations at rest than during movement. Rather, when subjects are performing neurofeedback and they are at rest, their underlying neural population is shifting further away from the MO threshold when beta oscillations are observed. We discuss mechanisms that generate beta oscillation in the conclusion of chapter 2 which can shed light on how certain types of behavior such as attention or isometric contraction may be correlated with onset of beta oscillations. However, they do not inform how the dynamics of the underlying neural population generate the movement change.

We find here that populations of neurons generate beta oscillations, their patterns are further from movement onset threshold than when they do not generate oscillations. In both the CO and NR tasks, on-beta PreMO bins exhibit further distances from the MO threshold than off-beta PreMO bins 3.2 emphasizing the common population shift that occurs in the CO and NR tasks during beta oscillatory periods. We propose that neural populations must stop generating beta oscillations before they can instantiate specific preparatory and movement generating patterns that may occur closer to the MO threshold. Thus, the generation of beta oscillations favors a low-risk state where neural populations will not accidentally create patterns that cause movement, and in exchange compromise their readiness for upcoming movements. Generation of these oscillations may possibly be implemented by a distinct subpopulation of 'pacemaker' cells, as discussed in 3.3, though data with more clearly isolatable single units is needed to describe how distinct subpopulations may each contribute to the pacing of beta oscillations versus the encoding of kinematic information. The hypothesis that encoding of specific movements is compromised with beta oscillations is further supported by experiments showing that in a delayed reaching task similar to the CO task here, movement cues associated with more uncertainty are correlated with higher beta power during preparatory periods than movement cues associated with certainty [173]. Uncertain stimuli bias the subject against preparing movements, corroborating that periods of beta oscillations are associated with compromised preparation. Further, during spike driven cursor brain-machine interface tasks when subjects are not moving, periods of beta oscillatory activity correspond to inferior neural decoding [93] and slower cursor movements [94], suggesting that the population contains less specific directional information that can be used to move the prosthetic cursor. How does this hypothesis of reduced movement encoding mesh with results showing linearly separable neural activity patterns for different movements occurring during preparatory periods when beta oscillations are prominent [174]? Since beta oscillations occur in transient bursts and only show elevated power for trial averages [175], it is possible that movement encoding during bursts is compromised, but outside of bursts

is intact. Further, not all cells engage in beta oscillatory events [95] making it possible that some cells are still encoding movement information during ongoing oscillations [97].

Chapter 4

Neurofeedback Control in Parkinsonian Patients

4.1 Introduction

4.1.1 Beta Band Oscillations in Parkinsonian Patients

In addition to neurofeedback being a useful scientific tool for perturbing neural signals of interest as demonstrated in chapter 2 and 3, neurofeedback has also been used as a therapy for various neurological disorders including attention deficit disorder [66], chronic pain [67], and even epileptic seizure frequency [176]. Below we first consider whether neurofeedback control can be performed in Parkinsonian patients, and second whether repeated practice of a neurofeedback-behavioral task can reveal how cortical beta oscillations are related to motor behaviors in Parkinsonian patients.

Parkinson's disease is characterized by a denervation of dopaminergic inputs to the basal ganglia, resulting in a variety of motor and cognitive deficits including tremor, bradykinesia, rigidity, and disorders of executive function, impulse control, and mood. Two common therapies for Parkinson's disease are pharmacological (dopamine replacement therapy) and surgical (deep brain stimulation). Pharmacological treatment is initially an effective therapy, but eventually the required dose for efficacy becomes too high. Deep brain stimulation is an effective therapy for patients responsive to dopamine replacement therapy, though its mechanism of action is not clear. One consistent finding in Parkinsonian patients is that they exhibit increased beta oscillatory power in the basal ganglia when they are not receiving deep brain stimulation therapy or dopamine replacement therapy, compared to when they are on these therapies [177]. The mechanisms through which dopamine administration or deep brain stimulation ameliorates patient symptoms and reduces beta oscillations remains to be determined.

Primary motor cortex, a primary source of input to the motor basal ganglia does not exhibit the same increase in beta power when patients are off therapy [86], nor does it exhibit a difference between PD and non-PD patients [85, 163]. However, individual motor cortical neuron recordings from parkinsonian non-human primates and electrocorticography recordings in parkinsonian humans support the hypothesis that excessive beta oscillations in the LFP of the basal ganglia may be reflected in hyper synchronized spiking activity or high-frequency LFP activity of motor cortical signals. This hyper synchronization may prevent successful encoding of kinematic information during movement. Specifically, simultaneously recorded activity from many motor cortical cells was found to be more correlated in motor disabled non-human primates after systemic treatment with 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine (MPTP, a drug causing Parkinson's disease) than before treatment [178, 179]. Cells also fire less specifically to passive limb movements after MPTP treatment [178] suggesting a link between parkinsonian symptoms and highly synchronous, uninformative spiking activity. It has also been found that coupling between low frequency beta phase and high frequency broadband gamma power is elevated in PD patients who had their deep brain stimulation therapy turned off [86] or were in an off state following an extended period of no levodopa medication [87]. Both stimulation and levodopa administration therapies improve symptoms of bradykinesia and rigidity for PD

patients, suggesting that synchronization of high gamma LFP activity, a proxy for spiking activity, to beta frequencies may be related to akinesia symptoms. This evidence supports the hypothesis that beta oscillations reflects reduced capability of kinematic encoding in motor cortical population activity, and that in Parkinsonian patients, reducing excessive beta activity in the basal ganglia, or reducing synchronization of spiking activity to beta frequencies in motor cortex may be therapeutic. One possible strategy for overcoming excessive beta synchronization in the basal ganglia and motor system may be to learn to reduce motor cortical beta power, which may serve to reduce the beta-synchronized spiking activity. Since cortex constitutes the majority of input to the basal ganglia, this manipulation may also reduce excessive beta activity in the basal ganglia. Recent work has found that PD patients do exhibit stronger sensorimotor cortical beta desynchronizations in early movement preparation compared to Essential Tremor (ET) patients [163], suggesting some patients may already use this cortical beta desynchronization strategy. Finally, the technical reason to test the cortical beta desynchronization approach instead of directly reducing basal ganglia beta power is that cortical neural signals are much greater in amplitude, and less likely to be affected by stimulation artifacts than basal ganglia signals [180].

4.2 Methods

4.2.1 Activa PC + S Use

Neurofeedback studies to date either utilize invasive or non- invasive neural recordings. Studies making use of invasive recording systems often require subjects to be in a hospital setting (e.g. [181]) and have not been developed for patients to use while going about their normal daily activities. On the other hand, acquiring non-invasive signals is more convenient but the signal itself suffers from reduced spatial resolution and attenuated high frequency activity. Further, non-invasive systems are sensitive to ambient electronic noise as well as biological artifacts from muscle and eye movements. An ideal neurofeedback system would combine the high signal quality of invasive studies with the convenience of a non-invasive setup. In addition, it would support real-time data streaming and not exert excessive power demands on the device.

The Medtronic Activa PC + S neurostimulator coupled with the Medtronic Nexus-D communication link system enables such capabilities [182, 183, 184]. This device is an investigational pulse generator capable of delivering continuous constant frequency electrical stimulation, similar to commercially available pulse generators used in DBS for movement disorders such as PD, dystonia, and essential tremor. In addition to delivering therapeutic subcortical stimulation, bipolar ECoG neural signals can also be recorded. Recently, recording invasive cortical ECoG signals chronically has been demonstrated in a nonhuman primate by implanting a quadripolar lead over sensorimotor cortical regions and routing leads (with lead extensions) into the neurostimulator [185]. Coupled with

Table 4.1: Neurofeedback task parameters for three study subjects

Patient #	Months Post-Surgery	Home or Clinic	Stim On or Off	Streamed Data	Beta Band (Hz)	Power Est. Method	Cursor Predict Alg.
1	13.5	Clinic	Off	Time Domain	10-20	Multi-Taper	Linear Reg.
2	11	Home	On	Power Estimate	12.5-17.5	On Chip	Linear Reg.
3	18	Clinic	On	Time Domain	20-30	Welch	Kalman Filter

the Nexus-D communication link, the system allows for real- time, wireless transfer of ECoG signals. Here, we leverage the fully embedded, wireless streaming capabilities of the device to investigate whether PD patients can learn to perform a neurofeedback task driven by invasively recorded cortical signals.

4.2.2 Activa PC + S Configuration

We first sought out to determine if 3 PD patients could perform neurofeedback control with the Activa PC + S device. This protocol was approved by the UCSF institutional review board (protocol # 13-10878) under a physician sponsored investigational device exemption (IDE #G120283 to Dr. Phil Starr). The study was registered at Clinical Trials.gov (NCT01934296). Informed consent was obtained under the Declaration of the Principles of Helsinki.

Study subjects were evaluated by a movement disorders neurologist and met criteria for a diagnosis of PD (i.e. presence of bradykinesia and at least one other parkinsonian cardinal symptom and responsiveness to levodopa). Baseline motor function in the on medication and off medication states were characterized using the Unified Parkinson's Disease Rating Scale, motor subscale (UPDRS III). Patients were evaluated by a neuropsychologist to exclude significant cognitive impairment or untreated mood disorder. In all three patients, the 4-contact cortical ECoG lead (Medtronic model 3587a) was placed in the subdural space through the same burr hole used for the therapeutic subthalamic nucleus (STN) lead. The STN lead was placed with standard methods [186]. At least one contact covered the posterior precentral gyrus (presumed primary motor cortex). Localization of the ECoG strip was confirmed using intraoperative CT merged to the patients preoperative MRI, as previously described [187]. Then, the free ends of the cortical and subthalamic leads were each connected to 40 cm lead extender (Medtronic model 37087) and tunneled down the neck to a Medtronic Activa PC+S bidirectional neural interface placed in a pocket over the pectoralis muscle. Sessions occurred at 13.5, 11, and 18 months post-surgery for patients 1-3 respectively (see table 4.1).

The Activa PC + S accommodates different sampling rates, channel streaming con-

figurations, and filtering options. The sampling rate utilized in this study was 422 Hz (maximum sampling rate supported when device is streaming data). At this sampling rate, the device can stream one channel of time-domain data per lead (packets of 169 points sent every 400 ms). It can also stream pre-calculated power in a frequency band (2 points sampled at 5 Hz, sent every 400 ms). The benefit of the on-chip power calculation option is that it substantially reduces the power consumption of the neurostimulator device. In contrast to streaming time-domain data which consumes 2.5 mA, streaming power channels consumes only 90 uA (both estimates do not include additional drain of 220 uA required for stimulation) [188]. In order to validate that our experimental paradigm would function in a setting where we used the lower-power on-chip spectral power estimate, we chose to use the on-chip power estimate of beta power for Patient 2's online control instead of using beta power calculated from the streamed time-domain channel. Prior work has shown that with stimulation on in the Activa PC+S device, artifacts can be introduced into time domain data and power estimate data at multiple frequencies, corrupting the underlying neural signal [188]. In order to avoid Patient 2's stimulation artifact from compromising beta band recordings, we selected a bandwidth of 15 +/- 2.5 Hz (recommended for a stimulation frequency of 130 Hz) instead of optimizing the beta band frequency limits with the baseline calibration task (as was done for Patients 1, 3). The Activa PC + S is also outfitted with a number of filter and gain options. Our real-time streamed data was filtered through a low-pass anti-aliasing filter at 260 Hz and a 0.5 Hz high pass filter. We used the maximum gain (2000) that the device allowed. The Activa PC+S also requires a data compression feature to be used when streaming data in real- time with Nexus-D.

4.2.3 Neurofeedback Task Calibration

All neural recordings were bipolar. Contact selection occurred prior to the training session using both recorded neural data from previous neural recordings sessions and anatomical location. Contacts that were over sensorimotor areas and exhibited a strong beta desynchronization during overt movement tasks were selected. For Patients 1 and 3, to estimate the beta band limits to be used for online control, patients performed a 1-2 minute movement task involving elbow flexion and extension in response to quasi-random auditory cues. These tasks elicit strong beta synchronizations in anticipation of the movement cue and after movement, and desynchronizations at movement onset [189]. Using the neural recording during this movement task, the power spectral density was estimated (see below) in windows of 400 ms at steps of 400 ms (no overlap). The variance of the spectrum at each frequency was plotted. The frequency band in the beta region (10-40 Hz) that yielded the local maxima on the frequency versus variance plot was selected.

For patient 1, the multi-taper method (using 5 tapers) was used to estimate the baseline spectrum as well as beta power online [143]. In patient 2, however, STN stimulation was on which resulted in a noisier signal. Due to artifacts from stimulation and the

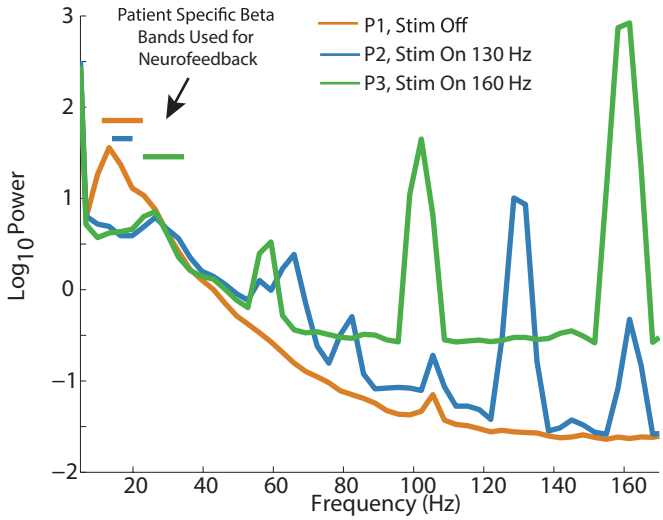

Figure 4.1: Patient power spectral densities during online neurofeedback control. Despite stimulation on in Patients 2 and 3 (130 Hz, 160 Hz respectively), beta peaks are still resolvable. Colored horizontal lines denoted by black arrow show the beta frequency range used for online control for each patient (Patient 1: 10-20 Hz, Patient 2: 12.5-17.5 Hz, Patient 3: 20-30 Hz). Note that the beta band used for neurofeedback control in Patient 2 was specially configured for streaming power estimates (instead of time domain data), and does not match the actual beta peak.

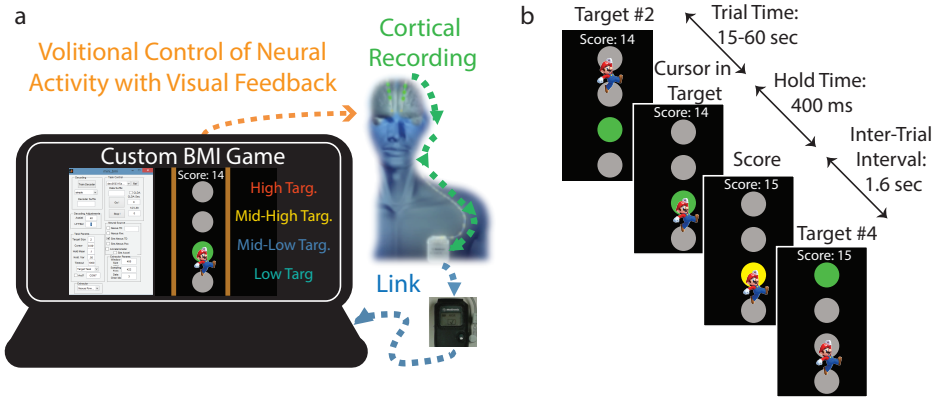

Figure 4.2: (A) Patients implanted with the Activa PC + S and cortical leads have the implantable pulse generator (IPG) located over the pectoralis muscle. A telemetry module has an antenna that sits on the skin surface in close proximity to the IPG and wirelessly acquires neural data and transmits the data to a Windows 7 machine via serial port. The Medtronic Nexus-D application program interface provides functions called from Matlab 2014b to acquire data from the serial port. Neural data is then translated into cursor position. (B) Task timeline begins with a target appearing. The patient then must make the cursor enter the target and hold (in all sessions reported, hold <400 ms making the hold time effectively 0 ms) after which the target turns yellow and the score count is incremented. An inter-trial interval of 1.6 seconds follows before the next trial begins.

processor clock (at multiple frequencies, see [188]), the smoothing from the multi-taper method spread artifacts to other frequency band estimates, yielding a noisy estimate of the spectrum. To mitigate this issue in future patients, estimates of beta power using the time domain signal used a method without smoothing (Welch's method with a Hamming window, pwelch in Matlab).

To fit the mapping between beta power estimates and cursor position, either simple linear regression (Patients 1, 2) or a Kalman filter (Patient 3) was used. To fit the simple linear regression model, a distribution of baseline beta power estimates from a 1-2 minute movement task was first acquired. The slope and offset of the linear regression were fit using two points ((x1 , y1) and (x2 , y2)) where the 25th and 50th percentiles of the beta power distribution from the movement task were x1 and x2 and the cursor positions at the bottom of the screen and middle of the screen were y1 and y2. For Patient 3, the Kalman filter utilized had constant A, W, C, and Q matrices, and the state-space model and observation model were formulated typically:

$$y_{t+1} = Ay_t + w_t, w_t \sim N(0, W)$$

$$\beta_t = Cy_t + q_t, q_t \sim N(0, Q)$$

Here, y_t represents cursor position, and β_t represents beta power estimate. C was fit by estimating the target position from the baseline movement task where it was assumed that at all times when beta power was in the 0-25th percentile of the overall baseline beta power distribution, the 'intended' cursor position was the position of the low beta target. Similarly, beta power in the 25th-50th percentile, 50th-75th percentile, and 75th-100th percentile were assumed to be aiming at the mid-low, mid-high, and high target respectively. Then beta power (β_t) was regressed against these inferred target positions (y_t). Q was calculated as the covariance of ($\beta_t - Cy_t$). A was calculated by adding minimal Gaussian noise to the inferred target position and calculating the correlation between time t and $t-1$. W was calculated as the covariance of ($y_{t+1} - Ay_t$). Matrices remained constant throughout the task, and the standard time-update and measurement-update steps were utilized to estimate cursor position from neural input and previous cursor state [190].

Once a mapping was set, patients began to play the neurofeedback game (4.2b). The task includes 4 targets each centered at -6, -2, 2, and 6 on a y axis that extends from -10 to 10 (arbitrary units). Targets had a radius of 2 (1.75 in early training for Patient 1 only) and the cursor was a point that was represented by Mario, a popular video game character. The cued target for that particular trial was indicated by the target turning green. Once subjects got their cursor (Mario) in the target, an auditory cue sounded, and the target turned yellow to indicate success. The score counter on the task interface incremented. Finally, an inter- trial interval time of 1.6 seconds elapsed before the next target was cued (turned green). The GUI and custom code for interfacing with the Activa PC + S using Nexus D and the Medtronic provided Matlab API functions is available on http://github.com/pkhanna104/nexusbmi

In order to help patients reach targets in early training an assistive feature was used to bring the cursor closer to the final target position that was currently trying to be acquired. The assist was an additive offset to the decoded cursor position:

$$cursorpos = (1 - \alpha) cursorpos_{decoder} + \alpha finalpos$$

where α is a value between 0 (no assist) and 1 (full assist) that corresponds to how much the assistive feature determines the position of cursor. α parameters used for each patient). Patients provided verbal input during breaks in between blocks indicating what value of α they wanted for the next block.

4.3 Results: Evidence of Neurofeedback Control

Three patients completed the cortical beta power driven neurofeedback task (4.2). Two patients visited the UCSF Movement Disorders and Neuromodulation Center, and one patient completed the task at home. All patients had one Activa PC + S device that supports two quadripolar leads. The lead for DBS therapy was placed in the subthalamic nucleus, while the lead for recording covered sensorimotor cortex. Patients with bilateral therapy also had a separate Activa SC unit (for clinical therapy only). Depending on

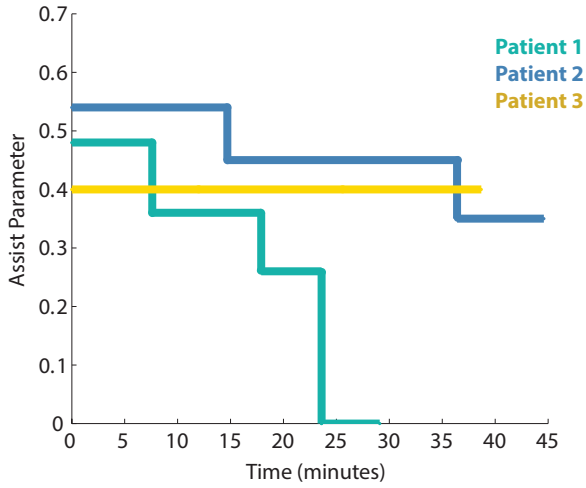

Figure 4.3: Assist parameter (α) used over the course of training for Patients 1-3. Training blocks are concatenated together for visualization, even if time elapsed between blocks. All patients either reduce their reliance on the assist, or maintain a constant assist level throughout the course of training.

severity of patient condition, patients were tested either on or off stimulation for the entirety of the neurofeedback training session (table 4.1). Because ECoG contact locations and the limits of the beta frequency band vary slightly by subject, contact selection and band limits were defined through an initial calibration procedure at the beginning of each patients training session. After contact selection and beta band frequency limits were selected, a linear mapping between beta band power and the one-dimensional height of a cursor was fit. The mapping was fit such that mean beta power during the baseline task positioned the cursor in the middle of the 1D axis of cursor movement, and beta power above and below baseline levels moved the cursor higher and lower respectively. This position-based mapping is in contrast to many velocity-driven neurofeedback systems. Patients then proceeded to play a neurofeedback game where they controlled a cursor shaped like Mario, a video game character from the Nintendo Mario franchise. The Mario cursor could move in one dimension and was controlled by patients changing their endogenous sensorimotor beta power. On each trial one of the four targets that were uniformly spaced along the 1D workspace was illuminated. Patients goal was to move the cursor to a position within the illuminated target before the trial timed out (4.2B). Patients completed 5-10 blocks of 5-15 minutes each where they practiced moving the cursor into targets. In early blocks, an assistive feature that moved their cursor closer to the desired target was utilized to make the task easier for patients. As training progressed, the assist was either maintained or reduced (4.3). Patients completed 1-2 hours of training each. During online control Patients 1 and 3 exhibited beta power peaks that were similar to their baseline recording peak (4.1) illustrated by the match between the beta peak in the

Table 4.2: Mean time to target (s.d.) in late training (secs)

Patient #	Low	Mid-Low	Mid-High	High
1	**5.6 (4.9)**	0.5 (0.3)	4.0 (4.6)	**29.3 (32.5)**
2	**20.4 (29.6)**	1.7 (0.9)	1.6 (0.6)	**15.5 (15.3)**
3	n/a	**6.2 (12.6)**	1.3 (0.3)	1.5 (0.2)

power spectral density plot calculated from online control sessions and the horizontal line representing the beta band used for online control (calculated from baseline recordings). Patient 2's baseline beta peak was not used to fit the beta band limits used for online control, which is why the peak in the PSD from online control appears mismatched to the horizontal line illustrating the beta band used for neurofeedback.

When examining the average time to target, there was a natural division where trials to one or two of the four targets were acquired on average in <2 sec and trials to the other of the four targets were acquired on average in >4 seconds (see table 4.2). It is possible that the trials to targets that were acquired on average in <2 sec were just acquired due to the natural variation of the cursor, in contrast to the patient trying to get the cursor in the target. To determine whether these two groups of times-to-target were better described by a single Gaussian or a mixture of two Gaussians we fit a Gaussian mixture model to the time-to-target data, and compared the one-Gaussian model to the two-Gaussian model. Specifically, first, all time-to-target data across 3 subjects from their late learning sessions was pooled. The mean and standard deviation was estimated for the pooled data, and the resultant Gaussian probability distribution function was used to calculate the log-likelihood of the data. Then, a second model was fit that used one Gaussian for the quickly acquired targets, and a second Gaussian for the more slowly acquired targets (see table 4.2for target designation for each patient) and log-likelihood of the data was estimated for the second model. Finally, Akaike information criterion (AIC) and Bayesian information criterion (BIC) were calculated to compare the models account for the greater number of parameters in the second model. We found that times-to-target were better explained by two Gaussians (AIC, BIC of one Gaussian: 947.13 and 948.01, and of two Gaussians: 547.06, 548.82). In addition, patients reported that they did not need to control the cursor in order to acquire these easier targets. Because of our analysis showing that these times-to-target are better described by two Gaussians and because patients reported a lack of voluntary control to certain targets, we proceeded to only include the targets with mean acquisition times of >4 sec in the subsequent chance analyses (table 4.2 bolded targets are used for each patient). In order to assess whether task performance was above chance level, the cursor trajectory from late training blocks with constant assist level was replayed through a target- shuffled version of the task 1000 times. We re-simulated the patient cursor trajectory from late sessions through many task simulations with shuffled target orders. Included in the task simulation is the target and cursor size, hold time, timeout times, and inter-trial intervals. In 4.4, we only report counts for those of the four targets that took longer than an average of 2 seconds for patients to acquire. For

example, if the mid-low and mid-high were deemed to be spuriously acquired by the previous analysis, and a patient has a final target count of 30 in a particular block, but 15/30 of the simulated targets were trials to the spuriously acquired targets, we would only count the other 15 targets in their final count for that block. Then, if a simulation of the block was performed with target shuffling and the simulation scored 40 targets but 30/40 of the targets are trials to the mid- low and mid-high targets, we would only count the 10 trials to other targets. Thus, the final count for the actual performance would be 15 and the simulated performance would be 10, in contrast to the original 30 to 40 comparison.

The actual number of successes by each patient was compared to the distribution of successes in simulated performance for targets deemed to be acquired by voluntary control (4.4A-C).

Subject performance significantly exceeded the distribution of chance performance (4.4A-C, Patient 1: $p < 0.05$, Patient 2: $p < 0.01$, Patient 3: $p < 0.001$, one-tailed test, Bootstrapped distribution).

Through the course of training, subjects had access to an assistive feature that they could use at a level varying between 0 (no assist) to 1 (full assist). Subjects' beta cursor was determined by a linear summation of the cursor position estimated by their neural activity weighted by 1- α, and the desired target position weighted by α. The assist features thus served to nudge the cursor closer to the cued target. Patients provided verbal input in between blocks indicating what value of α they wanted for the next block. All patients exhibited a trend of reducing or maintaining a constant assist level over blocks indicating that patients felt they were either improving or doing the task proficiently over the course of the session (4.3). In addition to the reduced reliance on the assist feature, we also analyzed improvement in time to target for late training blocks where patients had a constant assist level. Patient 3 improved significantly in time to target for the mid-low target which was the lowest target achieved by this patient (4.4F). Other patients exhibited non-significant reduction in time to target over the course of late training evidenced by linear regression slopes less than zero (4.4D-E). These improvements are promising but not significant, likely because subjects did too few trials with assist at a constant level.

Interestingly, Patient 1 and 3 exhibited few overt movements during their training sessions, but Patient 2 explored the movement space and converged on a movement strategy (raising his right hand and holding it up) to make their beta power increase. It is possible that he was activating the beta rebound generated post-movement, or persistent beta activity produced during posture maintenance [108]. Discovery of this strategy may have contributed to Patient 2's reduction in time to the highest target whereas Patient 1 and 3 exhibited improvements with mental strategies.

As subjects relied less on the assist, they learned to generate distinct target-specific neural patterns. We analyzed the beta band input signal preceding successful target acquisition for each patient from early in the session and late in the session (early: 4.5A-C, late: 4.5D-F). Early in the session patients relied more heavily on the assist feature to move the cursor close to the target, and thus only needed to change their endogenous beta

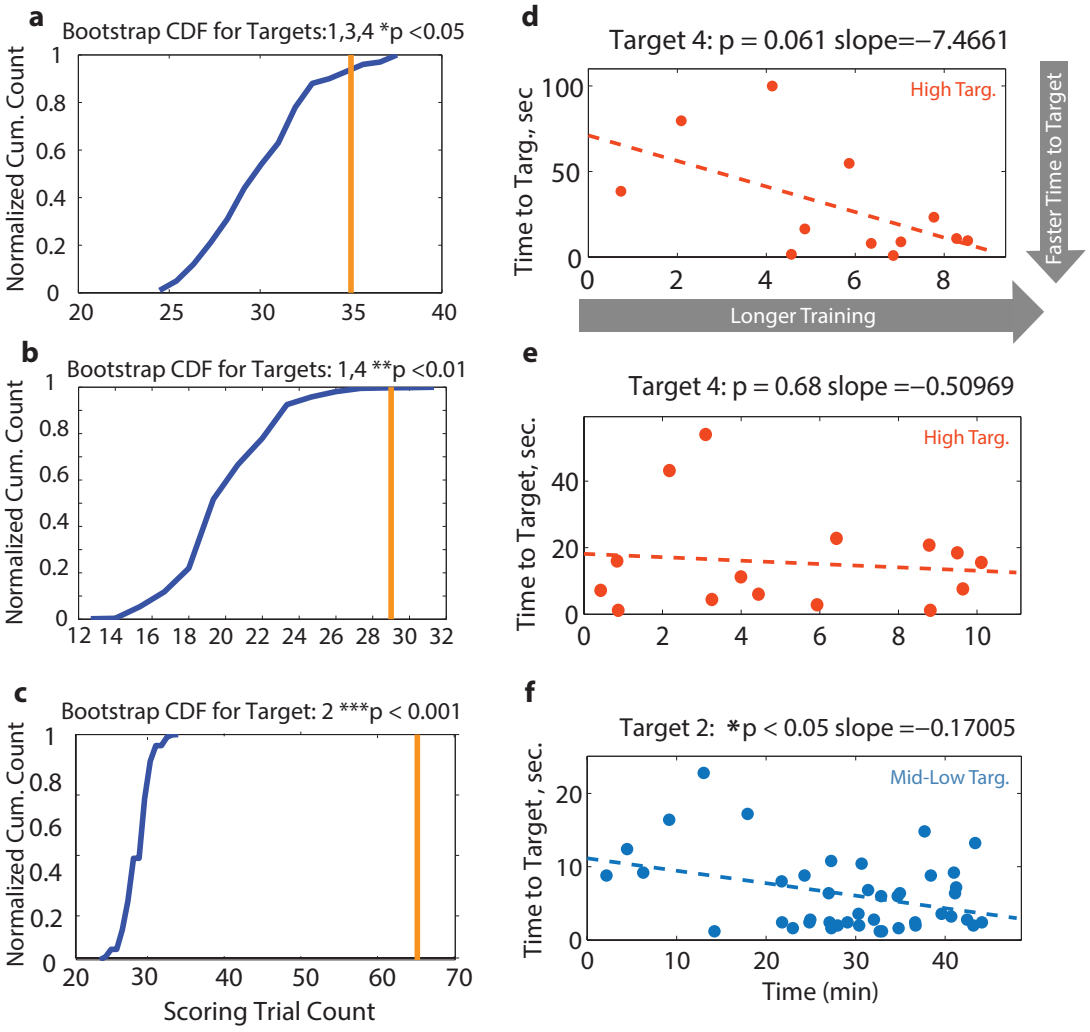

Figure 4.4: Patients perform the neurofeedback task above chance levels. (A- C) Patient chance level is illustrated by the blue cumulative distribution (see Experimental Procedures for calculation method). The x-axis is total rewards from simulated performance and the y-axis is a cumulative normalized count of how many simulations yield that number of rewards. Actual patient performance (total rewards) is shown with the vertical orange line. P-values are printed and are the percent of chance level simulations greater than actual performance (one-tailed test). Only data from late learning (constant assist) and from targets with mean time to target greater than four seconds are included in the chance level performance calculation (targets included are denoted in title, see Experimental Procedures). (D-F) For each patient, the time to target is plotted versus session time for the target with most improvement (restricted to late session data). Note that for patient 3, the negative slope indicative of reduced time to reach target is significantly different than zero (Students t-test, p <0.05).

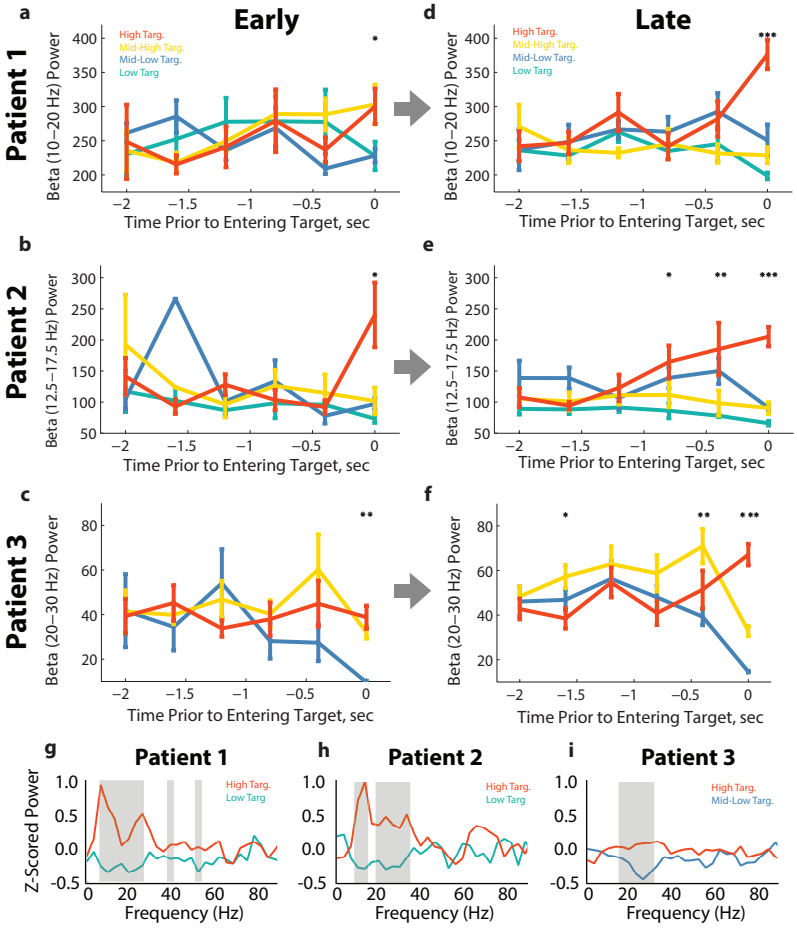

Figure 4.5: Neural changes emerge with training. (A-F) Trial averaged beta power estimates used to drive the cursor are plotted for each patient two seconds before target acquisition to time of targets acquisition (0 sec on x axis). Different colored traces are for the different targets as indicated by the target color key in (A) and (D). (A-C) show Patient 1-3 neural activity for early in the session (high assist levels). (D-F) show Patient 1-3 neural activity for late in the session (lower constant assist level, same data as Fig 3.). Asterisks indicate significant group differences (Kruskal-Wallis test, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$). (G-I) Modulation of full spectrum during neurofeedback task. Traces show trial-averaged z-scored power spectral densities (z-scored by subtracting mean and dividing by standard deviation of aggregated data from late training session) calculated in the 800 ms before target acquisition. Red traces are for the high beta target, teal traces are for the lowest beta target (G, H), and blue trace is for the mid-low beta target (I). Shaded gray indicates significant difference between the top and bottom target plotted for each subject (Kruskal-Wallis two-tailed test, $p < 0.05$).

power by a small amount to acquire the target. Thus, the distribution of the beta power produced for each target was not very different. By late in the session though (4.5D-F) subjects learned to generate beta band power at distinct levels: they produced their highest beta power for the highest target, and lowest beta power for the lowest target. Further, late in the session, this target-based separation of beta band power occurs for longer periods of time prior to patients entering the target. Patients 2 and 3 exhibited a significant difference in beta band power at 0.4 seconds (and 0.8 seconds for Patient 2) before their cursor entered the target, compared to early learning where there was no difference (Patient 2: $p < 0.05$ at 0.8 seconds before target acquisition, $p < 0.01$ at 0.4 seconds before target acquisition, Patient 3: $p < 0.01$ at 0.4 seconds before target acquisition, in contrast to no significant differences at these time points in early session data). Overall, patients learned to generate sensorimotor cortical signals containing different amounts of beta power to move the cursor to each of the four targets.

Finally, we investigated how other frequencies co-modulated during the beta power task for the top most and bottom most targets (4.5G-I). We calculated the full power spectrum for trials to the highest target (in red) and trials to the lowest target achieved (lowest beta target in teal for Patients 1, 2 and mid- low target in blue for Patient 3) for late session performance. The spectrum was calculated from the 800 ms prior to the time of target acquisition. Grey rectangles represent significant differences between the power spectrum of the highest target achieved and lowest target achieved (Kruskal-Wallis test, two- tailed, $p < 0.05$). Patient 1 significantly modulated a large range of frequency values to move the cursor including 6.6 Hz-26.4 Hz, as well as 38.2 Hz and 52.8 Hz though their control band was only 10-20 Hz. Patient 2 also learned a strategy for modulation covering the large range of 6.6 Hz-16.5 Hz, and 17.5 Hz- 36.2 Hz although their control band was only 12.5-17.5 Hz. In contrast, Patient 3 modulated only 13.3 Hz-33.0 Hz and their control band was 20-30 Hz. For all three patients, higher frequencies (<55 Hz) were not utilized, nor were the lowest frequencies (0.5-5 Hz). Thus patients do not exactly converge on their input beta signal, but they also do not utilize the full spectrum to modulate beta.

4.4 Conclusions

Here we have shown for the first time to our knowledge use of a chronic, fully embedded implantable device for neurofeedback training in human patients. While neurofeedback paradigms using mesoscale and macroscale neural activity have been demonstrated in humans before, these studies were done with to either one-time invasive recordings with epilepsy patients who were undergoing monitoring for seizures [181], paralyzed patients [46], or chronically but in non-human primates [143, 191, 142]. This paradigm demonstrates the feasibility of chronic neurofeedback training in patients at their homes with an invasive, fully embedded neural recording system accessible without the need to visit a hospital or clinic. This study is a proof of principle for future work on the effects of chronic neurofeedback training utilizing similar implanted devices for access to invasive cortical

signals. Implantable neural devices with better power consumption, artifact rejection, streaming latency, and signal-to-noise ratio (SNR) are in development [192]. All these improvements will further enhance neurofeedback learning by reducing concerns about battery consumption, increasing feedback rate and responsiveness of the system to the patient, and improving the SNR of the desired neural features.

Despite the different device configurations and task parameters used in the three patients (table 4.1: beta power extraction method, beta frequency limits, cursor estimation method, stimulation settings, and training session environment), all subjects still performed above chance level. This suggests the cortical beta band signal quality was sufficiently robust to be estimated either by calculating the spectrum from a streamed time-domain signal or a streaming on-chip power estimate. Remarkably, in both cases (streamed time domain or power estimate) the signal was not corrupted by the stimulation signal on the STN lead despite both the sensing and stimulation leads being routed into the same device. Further, despite the slow feedback rate (400 ms) in comparison to typical brain-machine interface cursor control studies ($\leq$ 100 ms), subjects were still able to process the feedback and use it to update their endogenous beta power.

We had subjects modulate cortical sensorimotor beta band power due to its relevance in PD motor symptoms. Patients exhibited modulation to distinct target-specific beta band power levels (4.5D-F), and did so without modulating the entire power spectrum (4.5G-I). Patients' time-to-target improvement was different for different targets (4.4D-F) suggesting that target difficulty must vary from patient to patient. This was likely due to variance in baseline beta power calculations which then translated to an offset bias in the beta-to-cursor map fitting. It is of interest that Patient 3 exhibited significant improvement in performance to the lowest beta target (4.4F). Given that literature suggests PD patients may need to compensate for their excessive beta synchrony throughout the basal ganglia and motor areas through earlier and more drastic desynchronizations prior to movement [163], it is perhaps not surprising that Patient 3 learned desynchronization faster. Potentially, this patient had practice at desynchronizing aberrant beta activity from using it as a compensation during normal motor control, which allowed the patient to exhibit faster improvement to the lower target. Patients 1 and 2 exhibit most improvement in the high beta target (4.4D-E). Patient 2 did find an overt movement strategy that worked for the high target, potentially explaining why improvement was quicker for that target compared to others. Further investigation of synchronization versus desynchronization difficulty can be addressed in longer-term neurofeedback training studies.

While subjects demonstrate they can learn to increase and decrease cortical beta power, modulating it may not alleviate symptoms. It is possible that changing cortical beta power will not affect spiking synchrony at beta frequencies and thus will not affect PD symptoms. It is also possible that the cognitive strategy used to generate or inhibit beta activity is a distinct circuit from the one creating pathological activity. With respect to the neurostimulator device, streaming time domain neural data for many hours requires substantial battery power. Streaming from the power channel requires less power but

still drains the battery [188]. Future systems will use rechargeable batteries eliminating this challenge [192], but current patients implanted without rechargeable batteries may limit their engagement in neurofeedback training to avoid invasive battery replacement surgeries.

While DBS is currently an FDA-approved therapy only for PD, essential tremor, and dystonia, it is being piloted as a therapy for numerous other neurological disorders such as medication-resistant depression, Tourette syndrome, epilepsy, and neuropsychiatric disorders [193, 194, 195, 196]. For many of these disorders, symptom characterization is more challenging than it is in PD. Exploring neural activity in relevant circuits may shed light on signals that can be used as biomarkers for symptom onset, as well as potential targets for neurofeedback therapy. Since neurofeedback has been shown to be effective at influencing neural signals relevant to specific behaviors, it is a promising tool to work in tandem with DBS, with both therapies striving to relieve the patient of pathological neural activity.

4.5 Introduction

4.5.1 Extended Neurofeedback Control in Parkinsonian Patients

The previous sections demonstrate that three PD patients were able to control a one-dimensional cursor with their motor cortical beta oscillations well above chance level after 1-2 hours of practice. Here, we extend this study in one of these patients (Patient 3). We ask first whether extended neurofeedback practice results in performance improvements over 9 days, and second, if the patient performs a sequential neurofeedback- finger tapping task, if we will see any effect of neurofeedback training on finger-tapping behavior. Notably, while this patient is characterized as parkinsonian, their actual diagnosis is Multiple Systems Atrophy (MSA), parkinsonian-subtype. MSA is a neurodegenerative disorder like Parkinson's disease characterized by symptoms that affect the autonomic nervous system and movement. There are two types of of subtypes of MSA – one is the parkinsonian subtype and the other is cerebellar.

4.6 Methods

4.6.1 Settings of the Activa PC + S for Multi-Day Use

As previously described, the Activa PC + S allows for wireless real-time streaming from a fully embedded neural stimulator. Since the device's battery is not rechargeable and real-time streaming applies extra computational demands on the device, care must be taken to ensure that data streaming protocols do not place patients at risk for an early battery change surgery. We computed the number of days of stimulation therapy that

Mode	Current Drain (μA)	Charge Drain in One Day (Coloumbs)
Stim ON	220 μA	19 C
Stim ON + Sensing ON	230 μA	20 C
Stim OFF + Stream Power Channels + Stream Time Domain Channels	2500 μA	216 C
Stim OFF + Stream Power Channels	90 μA	8 C

Table 4.3: Summary of Current Drainage for Different Modes of Activa PC + S (Source: personal communication with Medtronic, PLC).

would be sacrificed by doing a study involving real-time streaming. Estimates of current drain for different device behaviors are summarized in table 4.3.

Using the above table, an estimate for 10-15 hours of practice for the neurofeedback task where time domain is streamed is 90 - 135 C, or the equivalent of 4 - 7 days of conventional open loop stimulation therapy (Stim ON in table 4.3). On the other hand, streaming power channels for 10-15 hours of neurofeedback practice is 3-5 C, or less than a quarter of a day. Thus, doing the task with only the power channel streaming saves substantial battery power.

In this study, only motor cortical power channels from the left hemisphere were streamed while the device was stimulating in a monopolar stimulation configuration. Stimulation was delivered on two STN contacts (+) and the INS case (-) (located behind the pectoralis muscle) at 160 Hz at 3.7V with a pulse width of 60 μsec. Motor cortical beta power (22.5 - 27.5 Hz) was computed on the bipolar signal between cortical strip contacts E10-E11 (same contacts used in study described in previous section). Beta power was computed on-chip and was streamed in realtime during the neurofeedback task.

One important consideration in deciding what bandwidth limits should be used for on-chip power computation is that stimulation artifacts can leak into the power calculation of the desired frequency band [188]. This artifact can overpower the physiological signal, preventing observations of modulation during the task. To eliminate the interference of stimulation in the signal, one can compute the 'folded harmonics' from the stimulation frequency (here, 160 Hz) and the sampling rate of the signal (here, 422 Hz):

```
import numpy as np
artifacts = []
fstim = 160 # Hz
Fs = 422 # Hz
for n in range(1, 18):
    temp = np.abs(n*fstim-np.round(n*fstim/Fs)*Fs)
    artifacts.append(temp)
print 'artifacts: ', np.sort(artifacts)
```

```
artifacts: np.array([ 14, 28, 58, 72, 116,
130, 160, 174, 188, 218, 232, 276, 290, 320,
334, 378, 392])
```

For our stimulation and sampling rate parameters, we find that 28 Hz is a frequency that likely will experience folded stimulation artifacts, so we chose to cutoff our beta band frequency limit off at 27.5 Hz to avoid this leakage.

4.6.2 Task Description

The patient participating in this research performed two types of tasks within the one hour of training completed per day. The first task of the day was always a cued-movement task where an auditory cue was delivered every 6 seconds and at the sound of the cue the patient performed an elbow flexion/extension movement. At the end of sessions sometimes this task was repeated but with a finger-tapping motion instead of with elbow flexion/extension.

The second task was a sequential neurofeedback-finger tapping task, illustrated in 4.6. The task setup was almost identical to the NHP sequential neurofeedback-reaching task described in 2.2.1 but with a finger tapping task instead of reaching task. Finger-tapping was selected because it is part of the Unified Parkinson's Disease Rating Scale (UPDRS, part III) and has been proposed as a movement that can capture elements of speed, amplitude, and rhythm which are taken into account by clinician ratings [197, 198, 65, 199]. The task began the with patient initiating a trial by bringing their cursor (Mario character) to the middle of the screen. Next, either the top (high-beta) or bottom (low-beta) target appeared. Once the cursor entered the cued-target, the target color and text changed, cuing the patient to begin finger tapping. Finger tapping lasted for 6 seconds before the beta cursor and central target re-appeared signaling the start of a new trial.

Baseline beta power streamed from the Activa PC + S fluctuated substantially day-to-day, and even within-day. To account for these fluctuations, the beta-to-cursor map was adjusted each day using the cued-movement task with elbow flexion/extension as calibration, as shown in 4.7. First, a cued-elbow flexion/extension task is performed while beta power estimates are collected. In the first panel of 4.7, clear beta desynchronizations and synchronizations following and prior to the next cue are evident, as expected [189]. A distribution of beta power values from the calibration task is constructed, as shown in black in the middle panel, and the 25th and 50th percentiles are identified and mapped to the bottom and center of the middle target respectively. This process then fully defines the linear beta-to-cursor mapping.

Since motor cortical power from the left hemisphere is used during this task, the patient performed finger-tapping with their right hand. To quantify the kinematics of the tapping, the subject wore an inertial measurement unit (IMU) (SparkFun, MP9250) on the right hand index finger, and a capacitive touch sensor on the right hand thumb

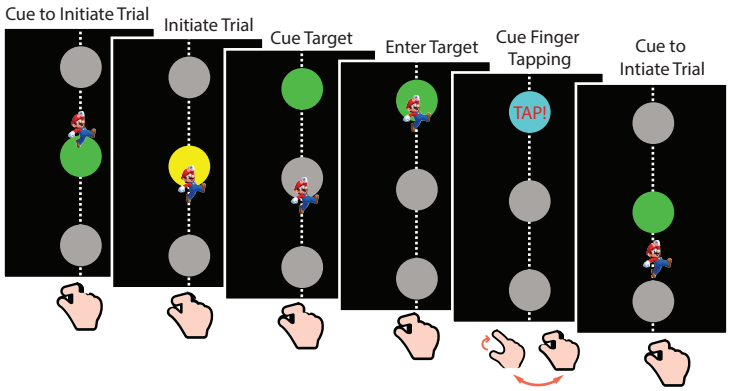

Figure 4.6: Schematic illustrating the sequential neurofeedback-finger tapping task. Trials are initiated by bringing the cursor (Mario character) to the central target. When the cursor enters the central target, a trial has been successfully initiated. Either the top or bottom target is cued. The patient then has 120 seconds to move the cursor to the cued target. If successfully reached, the target turns cyan with the red word 'TAP' displayed on the target, and the cursor disappears. This concomitant color change, 'TAP' cue, and cursor disappearance cues the patient to begin finger tapping for 6 seconds. After 6 seconds, the central target and cursor reappears allowing the patient to begin a new trial.

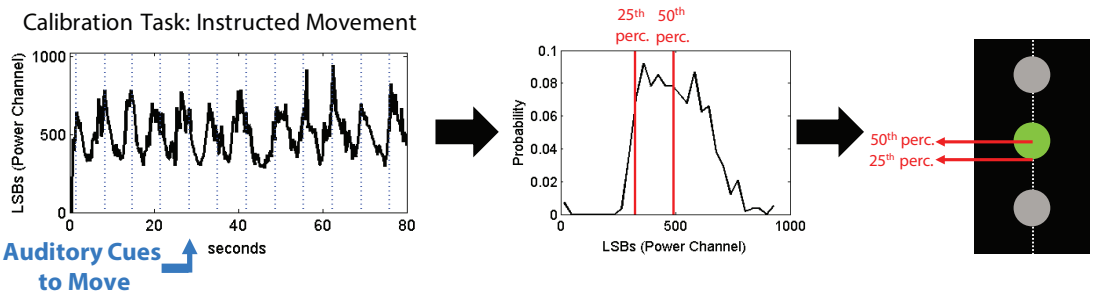

Figure 4.7: Daily calibration method: First, an cued-elbow flexion/extension task is performed. Elbow flexion/extension is performed immediately after auditory cues (vertical blue dotted lines, occurring every 6 seconds). On-chip beta power estimates are streamed while the task is being performed and are plotted in black in units of LSBs (least significant bits). A distribution of beta power values from the calibration task is constructed, as shown in black in the middle panel, and the 25th and 50th percentiles are identified. The 50th percentile is mapped to the center of the middle target, and the 25th percentile is mapped to the bottom edge of the middle target, fully defining the linear beta-to-cursor mapping.

(SparkFun, AT42QT1010). To quantify movements of the left hand, the patient wore an accelerometer on the left hand (SparkFun, ADXL345). Sensors were attached to fingers with custom 3D printed housing and velcro straps as shown in 4.8.

An example of the flow of the task is shown in 4.9. Here, the dotted green line corresponds to the target position (the y-axis is the vertical axis in which the cursor moves). At first the cued target is the bottom target. The patient's cursor (plotted in black) then enters the bottom target, and 6 seconds of finger-tapping ensue (shown in red). The red arrow points out the modulation of one axis of the gyroscope data during the tapping epoch. Once the 6 seconds of tapping are complete, the target moves to the center position, and the patient quickly re-initiates a trial. Next the target moves to the top position, which is quickly acquired followed by another 6 seconds of tapping. The task progresses.

4.7 Results

4.7.1 Neurofeedback Performance Over Days

First we assessed whether the subject was able to generate higher beta power when the high beta target was cued compared to when the low beta target was cued. In 4.10, we show that for almost all blocks on all days, the subject was able to generate a distribution with a higher median for the high beta target compared to the low beta target (right panel, almost all points are greater than zero). In 4.10, the left plot red trace shows a distribution

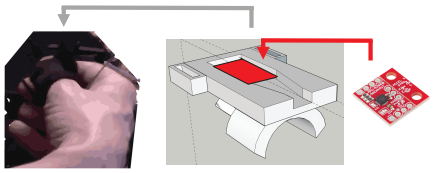

Figure 4.8: Custom-made IMU holder for right hand. On the left is an image of the patient's hand wearing the custom-made IMU, in the middle is a file of custom IMU, and on the right is a schematic of the SparkFun IMU with an arrow indicating how it fits into the casing. The slots on the sides of the IMU holder are for velcro straps to secure the device to the subjects' index finger.

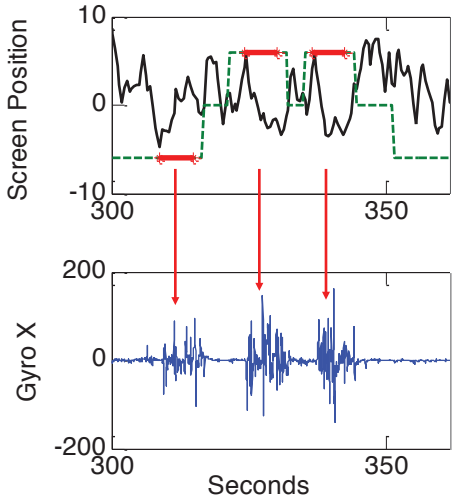

Figure 4.9: Neurofeedback-finger tapping task cursor and gyroscope progression over ~80 seconds. In the top plot, the green dotted line corresponds to beta target position. The black line corresponds to cursor position. The red lines indicate when the cursor and target are removed from the screen because the patient is completing the finger tapping epoch of the task.

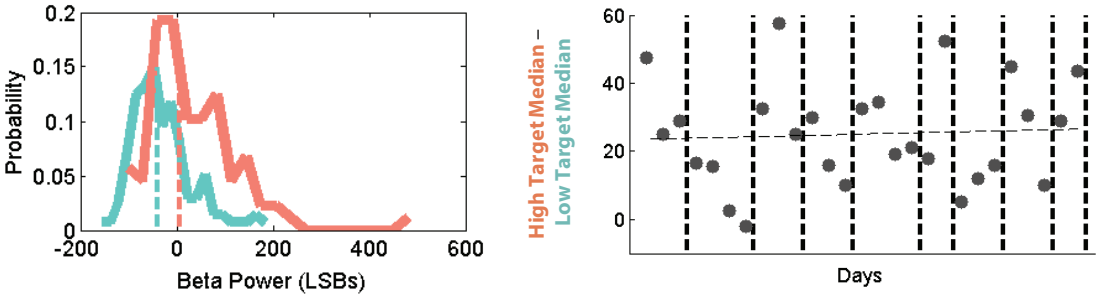

Figure 4.10: On the left is an example of the distribution of beta power for the low beta target and the high beta target for one day. Both distributions have the median of the calibration block (cued movement block) subtracted, hence the negative values of LSBs. The dotted lines point out the median of both distributions. On the right, the median of the high beta target distribution minus the low beta target distribution are plotted for each 10-15 minute block performed over 9 days of training. The dotted vertical lines separate days, and the dotted horizontal line is the linear regression line of all days (not significantly different from zero, $p > 0.05$).

of beta power from one day for times when the high target was on the screen, including unsuccessful trials, and not including times when tapping was occurring. Similarly the green trace shows a distribution of beta power from one day for times when the low target was on the screen. The medians are shown in dotted vertical lines, and the median of the red distribution is higher than the median of the green distribution. For all days and all blocks, the difference between the median of the high and low beta power distributions is calculated, and plotted on the right graph.

When performing a linear regression on the all points (each corresponding to a single 10-15 min block on a single day), there is no clear increase in median difference over days. We also assessed whether the subject got better at reaching the neurofeedback target over days assess by a 'time to target' measure as shown in 4.11. Over days the subject does not exhibit improvements in time to target, though often exhibits within-day improvements in the bottom target time-to-target, and worsening in the top target time-to-target. Overall, these results are consistent with no improvement in time to target.

One observation that may account for the within day changes in time to target are changes in baseline mean beta power. Often on days where the top target became more difficult, the baseline beta power level averaged over minutes had drifted lower. Naturally if the distribution shifts lower, the top targets will become harder.

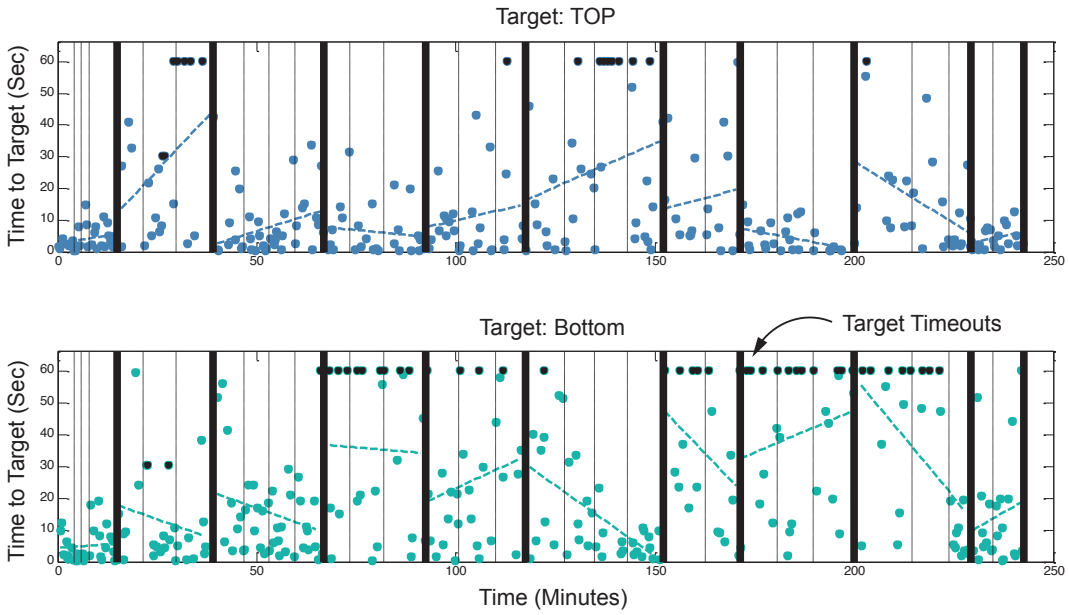

Figure 4.11: Time-to-target for the top and bottom beta neurofeedback targets for all trials over blocks (thin vertical bars) and over days (thick vertical bars). Note that ten days of data are displayed here while 9 days are reported in other figures. Indeed 10 days of data were collected, but the first day was used to select motor cortical contacts, frequency band limits, and used a beta-to-cursor mapping that was slightly different than the subsequent days, so it is only shown here for comparisons of within - day learning. Dotted diagonal lines are within-day regressions between time to target over the course of the day.

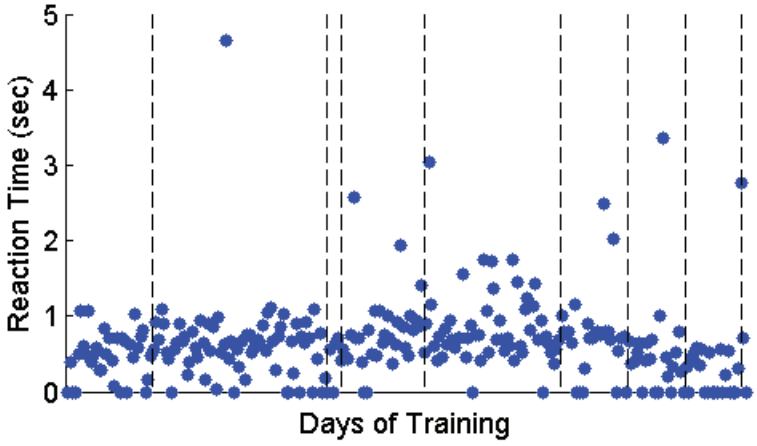

Figure 4.12: Reaction times of all finger-tapping trials. Marked in vertical dashed lines are divisions between days.

4.7.2 Dynamics of Finger Tapping Are Influenced by Neurofeedback

Given that the subject has consistent performance in the neurofeedback task, we next asked if the dynamics of finger-tapping change depending on whether they were preceded by the low beta target or the high beta target. First we assess if there is a significant improvement in tapping reaction-time, signifying that the subject was improving in their ability to react to the tapping cue after completing the neurofeedback task. If there is a learning effect, it does not make sense to include the early days in the analysis since the subject was still learning the contingencies of the task. Plotted in 4.12 are the reaction times of all finger tapping trials across the nine days of training. Some days have very few trials because there was a technical issue in which occasionally the data trigger to start the IMU recording was not received. Some reaction times are very high (> 1.5 sec) because occasionally the patient did not realize that they had achieved the target, or occasionally forgot to start tapping. The patient commented after the 9 days of training that having an auditory cue to indicate the onset of tapping would be helpful. Some reaction times are zero due to anticipation of the tapping cue.

We note that some days exhibit systemically higher RTs than other days (e.g. day 5) shows RTs that have a greater mean than RTs on day 9. On average though, there is no reduction in RTs over days, indicating that there is no systemic trend of the patient learning to initiate taps faster throughout the course of training.

We choose to perform a 2-way unbalanced ANOVA to assess tapping metrics such as RT to assess the difference between taps preceded by the high beta target versus taps preceded by the low beta target. The two factors in the ANOVA are i) low vs. high beta target preceding tap and ii) day of training. Given the high day-to-day variance in RTs,

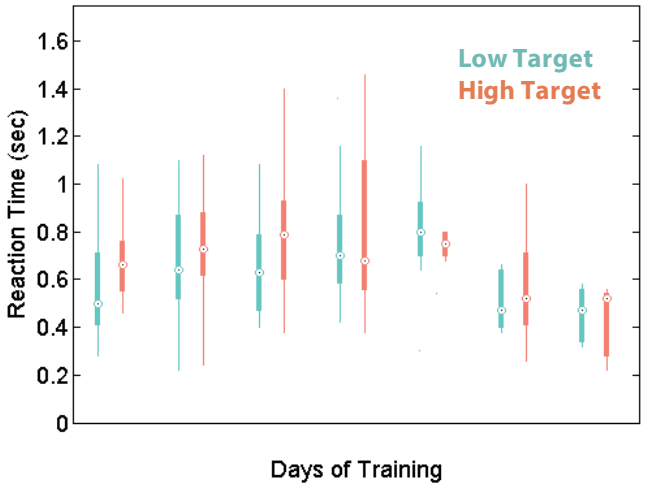

Figure 4.13: RTs ($0.2 < \text{RT} < 1.5$ secs, excluding days 3 and 9), from 4.12plotted by high vs. low target

we wanted to assess within-day differences in high vs. low beta targets. We eliminated trials that had RTs < 0.2 sec, or RTs > 1.5 sec. We also did not analyze days in which there were fewer than 15 trials collected (did not analyze days 3, 9). There was a total of 198 trials analyzed. Plotted in 4.13are boxplots of RTs following low vs. high beta targets by day. On a majority of days, the low beta target precedes faster RTs than the high beta target. The 2-way ANOVA shows an affect of day ($p < 10^{-15}$) and of beta target ($p = 0.0011$).

We also assessed other tapping metrics in addition to tapping RT, shown in 4.4. Tapping Frequency is defined as the dominant frequency when performing an FFT on gyroscope data about the tapping axis. Higher values indicate a faster tapping pattern. Tapping frequency was higher for trials following the low beta target compared to the high beta target. Finger angle speed is defined as the mean rectified angular velocity over the 6 seconds of tapping. A higher finger angle speed indicates a more vigorous tapping sequence with higher positive and negative angular velocities. Finger angle speed was higher for taps following low beta targets compared to higher beta targets, although not significantly. Finally, we assessed whether the number of discrete taps (measured as the number of times the capacitive touch sensor turned from off to on) different following low versus high beta targets. We found that there were more discrete taps following low beta targets compared to high beta targets, although not significant.

In the previously described work in 2using a sequential neurofeedback-reaching task in non-human primates, there was no difference in the kinematics of the arm reaches except for the movement onset time (here, RT). It is surprising, then, to find evidence for more vigorous tapping following low-beta target than high-beta targets. Perhaps, cortical

Table 4.4: Table of 2-way ANOVA test results for tapping metrics

Metric	Tapping following Low vs. High Beta Target	2-way ANOVA, Test for Metric
Tapping Frequency (Hz)	Low > High	p = 0.0460
Finger Angle Speed (rad/sec)	Low > High	p = 0.5213
Number of Discrete Taps Per Trial	Low > High	p = 0.1454

beta oscillations have different dynamics in parkinsonian patients than in healthy NHPs. We analyzed control finger-taps to assess how different cortical beta dynamics were for fast (RT $<$ 3 sec) versus slow (RT $>$ 3 sec) finger-taps. These taps were collected in the same way that the calibration task was – an auditory cue instructed the patient to begin tapping. The patient was instructed to complete two big and fast taps, and then to rest until the next cue. Because the patient performed these taps after the sequential neurofeedback-finger tapping task, they likely more tired than during the calibration phase and thus had a slower reaction time. Thus we were only able to collect 10 example trials where the response to the cue was within 5 seconds. For many of the cues, the patient was still tapping from the previous trial, contaminating the current trial. When plotting the fast vs. slow RT trials during the instructed tapping time, we see differences in cortical dynamics, though admittedly with very few trials.

Though we have very few trials of the controlled finger-tapping task, we still can assess the mean cortical beta dynamics. As shown on the left in 4.14, when aligning to the movement cue fast and slow trials exhibit a beta synchronization in anticipation of the cue, and a desynchronization following the cue. The slow trials, however, do not fully desynchronize, and instead re-synchronize within 0.8 seconds. When assessing the fast versus slow RT trials when aligned to the RT (movement onset time), we see that cortical beta has desynchronized prior to RT in fast trials, but has yet to fully desynchronize prior to the RT in slow trials. Thus, we hypothesize that fast RT trials are associated with a successful, persistent beta desynchronization following the movement cue. On the other hand, we hypothesize that slow RT trials may exhibit a slight beta desynchronization following the movement cue, but pathological beta throughout the cortico-basal-ganglia-thalamic loop may drive M1 to resynchronize, arresting movement onset.

4.8 Conclusions and Future Directions

Future directions of this work will investigate whether the findings that finger tapping RT is faster, tapping frequency is faster, number of taps is greater, and angular speed

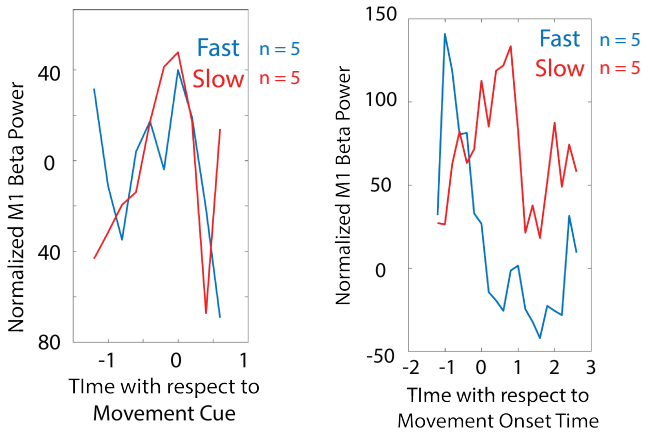

Figure 4.14: Cortical beta dynamics for fast vs. slow RT trials during an instructed finger-tapping task. Cortical beta power is normalized by subtracting the median beta power value from the calibration task.

is greater when finger-tapping follows a low beta target compared to a high beta target holds up in parkinsonian patients as opposed to an MSA patient. It will also further investigate the dynamics of beta power during finger-tapping itself, and if the observed differences in movement kinematics have a neurophysiological signature in beta band.

Technologically, future devices with rechargeable battery [192]will be used in these investigations, to enable time domain-streaming from more than just one channel. In the setup described above, only a couple estimates of on-chip beta power are streamed every 400 ms to reduce battery consumption. If there are changes in the range of the full time-domain signal over the course of recording, or changes in baseline noise level, these will be reflected as changes in beta power when using the on-chip estimate. For example, during the study described above, beta power was also streamed from a bipolar pair of STN electrodes but the baseline level noise due to the stimulation artifact changed on a day-to-day basis and even throughout the course of a session. This noise obscured expected movement-related desynchronizations in beta power [200].

Better filtering and signal conditioning techniques can be used if the full time-domain signal is available. Streaming more channels than a single bipolar pair per electrode strip would also enable better online-computation of beta power, as well as analysis of other brain regions including the basal ganglia during the neurofeedback task. We also hope to be able to make a more engaging and exciting task for patients given the importance of motivation and engagement in learning.

Overall, we hope that the approach described above (sequential neurofeedback-behavioral task) will be used to study not only the motor system and beta oscillations, but also the relationship between other brain areas with unique neural signals and other behaviors.

Chapter 5

Distinct Sources of Neural Variability Drive Neuroprosthetic Control

5.1 Introduction

5.1.1 Modeling Neuroprosthetic Control

Previous chapters focused on beta oscillations in the motor system and their relationship to fast versus slow movements in health and disease. Here, a different experimental paradigm is employed but is still used to study of how population motor activity contributes to movement dynamics. As discussed in chapter 1, closed-loop brain machine interfaces (BMIs) can be thought of as simplified, mini, motor systems. In these systems, unlike our natural motor systems, the experimenter fully observes all neural activity that controls the actuator, knows the exact (and usually linear) transform of how neural activity affects the movement of the actuator, knows the physics of how the actuator moves, and knows exactly what feedback the subject receives. In the natural motor system, experimenters can measure only a subset of the relevant signals for control, do not know the exact non-linear brain to arm mapping, have only an estimation of the physics of how the arm moves, and must contend with sensory feedback that is linked to arm movements [78]. A closed-loop BMI then simplifies many aspects of studying sensorimotor control yet still includes fundamental features of natural motor control. Notably, subjects still must use many neurons to control a lower-dimensional effector, must plan and execute a motor plan, and must use sensory feedback to update their motor plan. Much like the endogenous motor system, there is evidence of subjects learning skillful control of neuroprosthetic effectors [71, 74] and developing an internal model of how to control the effector [80]. This chapter explores principles of how subjects coordinate their population activity to control an effector using both *in silico* simulations of BMI control, and online experiments.

Many closed-loop, continuous-control BMI architectures rely on decoding via a linear readout of noisy population neural activity [169]. However, the neurons responsible for controlling a BMI (BMI output cells) receive inputs from diverse sources. Different inputs likely have dissociable effects on BMI population activity, and contribute to BMI control differently. Indeed, recent work has found that BMI output cells increase shared variability (many cells co-modulating) and decrease private variability (cells modulating independently) as subjects increase the speed and directness of the 2D cursor movements over learning [90]. Further work has demonstrated that BMI decoders that require the production of specific types of correlated modulations are very difficult to learn [91]. Correlated modulation patterns that are generated during BMI control thus represent an informative statistic of population activity, and even predict the learnability of a BMI decoder. What is unclear is first, what types of inputs give rise to correlated versus independent variability in a population. Can a high-dimensional set of tuned inputs with independent noise give rise to shared variability or must inputs already be low-dimensional? Understanding what inputs give rise to shared and private variance in populations will assist in interpreting what variables must be encoded upstream of motor cortical populations. Does controlling a low-dimensional effector necessitate a upstream population activity to first

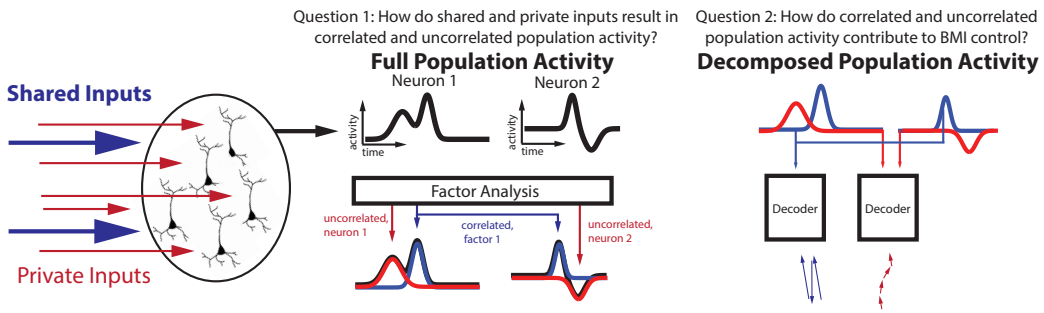

Figure 5.1: A conceptual schematic of the questions addressed in this chapter. First, the population activity of BMI output neurons as a linear combination of private and shared sources. We then simulate the resultant population activity, and apply Factor Analysis to the resultant population activity to assess how different combinations of shared and private input sources affect the amount of correlated and uncorrelated population activity. Second, we assess how the correlated and uncorrelated population activity contribute to online BMI control of a cursor.

become low-dimensional? This work presents a generative model of population neural activity in which a shared input drives many cells simultaneously, and private inputs drive each cell independently. Presented are results of how changes in input parameters of this generative model produce the neural population's observed correlated and uncorrelated variability. Second, how do shared and private variance contribute to neuroprosthetic control? If only shared variance or only private variance is used to control the cursor, how successful is performance? Offline simulations and online experiments are used to test how shared and private variability contribute to BMI control.

5.2 Methods

5.2.1 Terminology

In the following section, we will refer to two input sources: **input shared variability** and **input private variability**. These are defined as the inputs to the population which are simulated in the generative model described below. **Observed shared variability** and **observed private variability** refer to shared and private variability that are parsed by observing the population activity after all the input sources in the generative model have been summed. **Output neurons** refers to the neurons in the simulated population. The terms 'output' and 'observed' should not be confused with the Kalman Filter decoder output. The Kalman Filter decoder output will be referred to as velocity commands, decoded velocities, or predicted velocities.

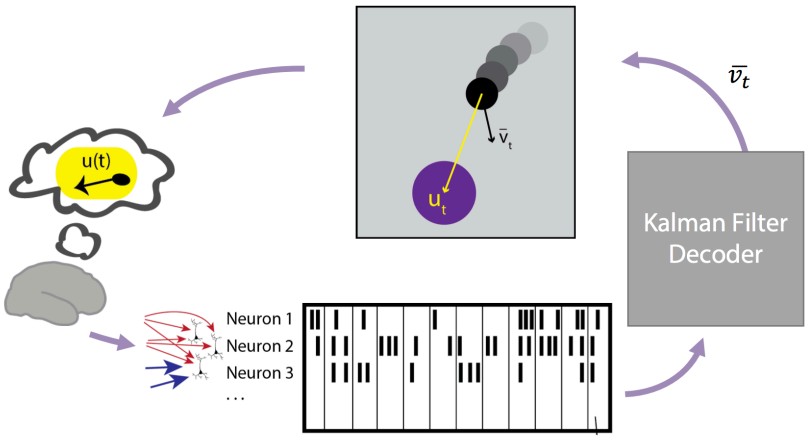

Figure 5.2: Generative model of neuroprosthetic control using a 2D velocity controlled cursor, Kalman Filter decoder, and simulated inputs into a 20-neuron population

5.2.1.1 Generative Model

5.2.1.2 Overview

In order to study how different types of inputs to a population influence the observed shared and private variability observed, we developed a generative model of neural population activity 5.2. Activity for a population of 20 neurons was generated based on the shared (blue arrows) and private (red arrows) population inputs described below and the summed activity of inputs to these neurons were then used as BMI output neurons. The activity of the 20 BMI output neurons was binned in 100 ms bins (illustrated with the raster plot), and used as observations in a Kalman Filter. The decoder predicts intended velocities, $\bar{v}_t$, which is integrated and used to update the position of 2D cursor as is commonly used in cursor BMI experiments [73, 91]. To complete the closed-loop simulation, the cursor state and cued target position were used to infer an intended velocity command, $u(t)$, which would move the cursor closer to achieving and holding within the cued target. This intended velocity command was then used to generate inputs for future population activity as illustrated by the gray brain and its yellow thought bubble.

5.2.1.3 The Factor Analysis Model

Typically Factor Analysis is used to decompose a high-dimensional set of simultaneous observations (y_t) into a summation of i) low dimensional shared patterns, ii) high

dimensional private patterns, and iii) a mean. The Factor Analysis statistical generative model is [201]:

$$y_t = U_{obs} z_{obs,t} + \psi_{obs,t} + \mu_{obs}$$

$$\psi_{obs,t} \sim N(0, \Psi_{obs})$$

$$z_{obs,t} \sim N(0, I)$$

Here, y_t is a vector $\in R^n$ where n is the number of neurons in a population (in these simulations, $n = 20$). y_t represents a vector of spike counts over the course of a 100 ms bin and in real data consists of integer counts. In this simulation, y_t is modeled as Gaussian and can have continuous values. Shared variability, Uz_t can be attributed to the variance of a low-dimensional source ($z_t \in R^m, m < n$) that is mapped through a loading matrix ($U \in R^{nxm}$) to the high-dimensional space of neurons. Private variability (ψ_t) is individual neural variance that is independent of other neurons. That is, Ψ is a diagonal matrix with cross-terms of zeros. Finally, $\mu \in R^n$ is the mean of y_t. This model will be used to parse observed shared variance and observed private variance, but also inspires the simulated structure of input sources to the population. Details on fitting Factor Analysis can be found in A.1.

Note that correlated and uncorrelated variability were extracted with the fit FA model without any knowledge of the generative model parameters, as would be the case in an online experiment. Theoretically this decomposition ought to segment neural activity into exactly the private and shared sources that were used to generate the activity, but given limited data and different levels of noise (untuned activity), the decomposition may converge on a different solution than that of the generative model.

5.2.1.4 Population Inputs

There are four types of inputs to each output BMI neuron; tuned and untuned shared inputs, and tuned and untuned private inputs. Shared inputs represent activity that is generated by m low-dimensional factors, and mapped to each output neuron through a loading matrix U. Private inputs are generated independently for each output neuron. Tuned inputs produce a linear readout of the subjects' intent to reach an instructed target during a two-dimensional velocity-driven cursor BMI task and untuned inputs generate noise with respect to the BMI objective. These inputs are illustrated in more detail 5.3 and 5.4.

Specifically, tuned inputs change according to the desired velocity of the cursor, $u(t)$. Once $u(t)$ is computed by some function F, $u(t) = F(cursor_t, target_t))$, the tuned shared and tuned private inputs are computed based on cosine tuning models to $u(t)$. Tuned

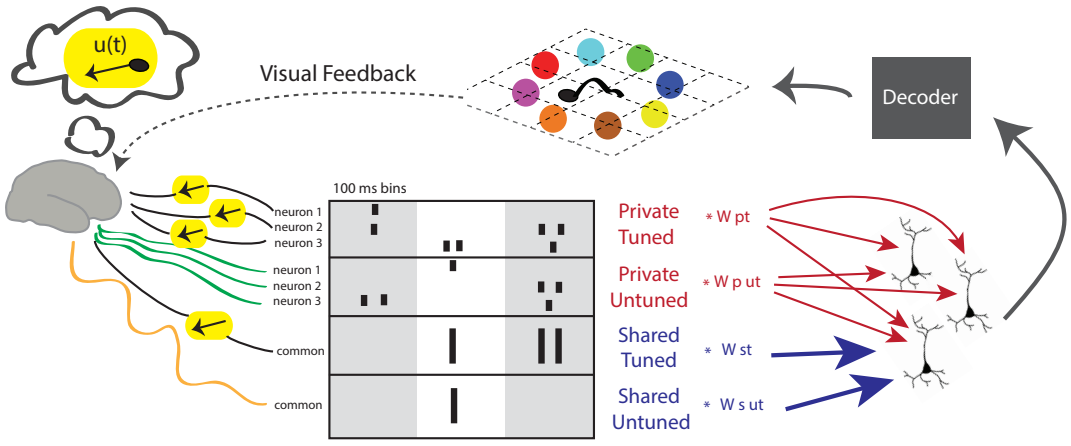

Figure 5.3: Generative model of neuroprosthetic control with four classes of inputs highlighted. The raster plot first shows tuned private activity which is derived from the dot product of $u(t)$ with each output neurons' private tuned cosine tuning model. Next, untuned private activity is sampled from each output neurons' private untuned zero-mean normal distribution and is unrelated to $u(t)$. Tuned shared input is third, consisting of of cosine tuned low-dimensional factors whose activation is determined by cosine tuning to $u(t)$, and then is mapped to the output neurons via $U_{shar,tuned}$. Finally, untuned shared input is generated by each untuned shared factor sampling the standard normal distribution, and is mapped to output neurons via $U_{shar,tuned}$.

shared activity is defined as $z^i_{t,tuned} = u_t \cdot \theta^i_{pref,shar}, 0 < i \leq m$ where $\theta^i_{pref,shar}$ is the unit vector in the preferred velocity direction of tuned factor i. Tuned private is defined as $\psi^j_{t,tuned} = u_t \cdot \theta^j_{pref,priv}, 0 < j \leq n$ where $\theta^j_{pref,priv}$ is the preferred velocity direction of tuned private input to unit j.

Untuned shared activity is defined as $z_{t,untuned} = N(0, I)$ and untuned private activity is defined is $\psi_{t,untuned} = N(0, \Psi_{untuned})$ where $\Psi_{untuned}$ is diagonal with off-diagonal components equal to zero. At each time bin, $z_{t,untuned}$ and $\psi_{t,untuned}$ were drawn from their respective normal distributions. All sources $z_{t,tuned}, z_{t,untuned}, \psi_{t,tuned}, \psi_{t,untuned}$ are weighted before being summed:

$$y_t = \alpha_1 U z_{t,tuned} + \alpha_2 U z_{t,untuned} + \alpha_3 \psi_{t,tuned} + \alpha_4 \psi_{t,untuned} + \mu$$

Finally, to break temporal correlation between the four sources an update rate is included for each source. For each time point, the number of times to sample each source changes according to a Poisson process with rate parameter Λ_{source}= 1. Thus at each time point, for each source, and each factor or neuron, an integer $\lambda^k_{t,s}$ is drawn from a Poisson process with rate parameter Λ_s.Here, $t$corresponds to time bin, s corresponds to source, and k corresponds to factor or neuron index depending on source. Thus, $\lambda^1_{T0,unt.priv.} = 2$ would mean that on bin $T0$, for source untuned private, for neuron 1, the distribution $N(0, \Psi[1, 1])$ would be sampled twice, and sum the two resulting samples to get $\psi^1_{T0,untuned}$. If $\lambda^2_{T1,tun.shar.} = 2$, that would mean that on bin $T1$, tuned factor number two would equal $z^2_{T1,tuned,shar} = 2 * (u_{T1} \cdot \theta^2_{pref,shar})$. Almost always Λ_{source}= 1, with the exception of $\Lambda_{tuned,private}$ which will be varied in results presented below.

To summarize, the parameter set of this generative model include:

- $\Lambda_{tuned,private}, \Lambda_{untuned,private}\Lambda_{tuned,shared}, \Lambda_{untuned,shared}$.
- $U, \Psi_{untuned}, \mu, n, m_{tuned}, m_{untuned}$
- $\theta^j_{pref,priv}, \theta^i_{pref,shar}$ for $1 \leq j \leq n, 1 \leq i \leq m$
- $\alpha_1, ..., \alpha_4$
- $F(cursor_t, target_t)$
- Kalman filter decoder

In the simulations following, Λ_{source} was always equal to 1, unless specifically manipulated. U was set to random values between -1 and 1. $\Psi_{untuned}$ was diagonal and set to values between zero and 1. $\mu$was set to match mean firing rates of real neural data collected during online BMI performance. $n$was set to 20, $m_{tuned} = m_{untuned} = 3$ factors. $\theta^i_{pref,shar}$ and $\theta^j_{pref,priv}$ were set to random unit vectors spanning the 2-D space. The αweights were set between values of 0 and 1. The function mapping cursor state and target position to $u(t)$, $u(t) = F(cursor_t, target_t)$ is defined as:

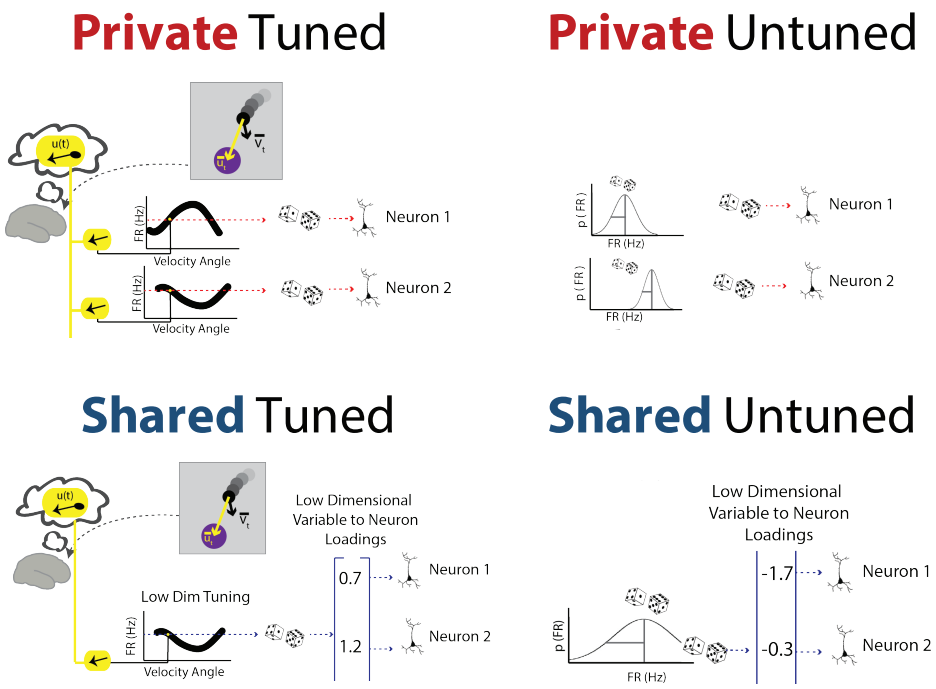

Figure 5.4: A summary of how the four different sources of private and shared, tuning and untuned activity are generated. The yellow circle with black arrow in the middle refers to $u(t)$ interferred from cursor and target position. The graphs for private tuned and shared tuned refer to neuron or factor tuning curves. The graphs for private untuned and shared untuned sources show the probability distributions $N(0, \Psi_{untuned}), N(0, I)$ respectively. Dice refer to the Poisson process used to determine how many times to sample the source.

$$cursor_{pos} \notin target_{rad} : u(t) = 15 * \frac{target_{pos} - cursor_{pos}}{||target_{pos} - cursor_{pos}||}$$

$$cursor_{pos} \in target_{rad} : u(t) = 0.5 * target_{pos} - cursor_{pos}$$

Once all the above parameters were set, samples of data were drawn using actual experimental cursor and target positions to compute $u(t)$ and then each of the individual sources. It was confirmed that for $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4$, on average (across neurons and simulations) each input source contributed equal amounts of variance to the summed population activity. It was also confirmed that the shared tuned and shared untuned low-dimensional variations, z_t, were normally distributed with identity covariance, the private tuned and untuned variations were normally distributed with appropriate variances $\Psi_{i,i}$.

Finally, the algorithm used to generate population activity:

1. Estimate desired velocity, $u(t)$, based on current cursor state and cued target position
2. For each source $z_{t,tuned}, z_{t,untuned}, \psi_{t,tuned}, \psi_{t,untuned}$ compute the number of times to draw from each factor or neurons by sampling from a Poisson process with parameter Λ_{source}.
3. For each source, computed the activations for each neuron or factor using $u(t)$, $\lambda^k_{t,s}.\theta^j_{pref,priv}, \theta^i_{pref,shar}$
4. Sum all sources together: $y_t = \alpha_1 U z_{t,tuned} + \alpha_2 U z_{t,untuned} + \alpha_3 \psi_{t,tuned} + \alpha_4 \psi_{t,untuned} + \mu$
5. Use $y_t \in R^n$as an observation to update the Kalman Filter decoder's cursor velocity prediction.
6. Update the cursor position and velocity.

5.2.2 BMI Task and Decoder

A standard center-out task is used to assess BMI performance. To begin a trial, the cursor is automatically set to the center of the workspace. One of eight peripheral targets, arranged uniformly in a circle of radius 10 cm around the center, is cued. The cursor center must enter the target (radius of 2 cm) must be acquired in less than 10 seconds to avoid a timeout penalty. The cursor is updated every 100 ms using a velocity Kalman Filter decoder. The Kalman Filter models neural activity as a linear combination of x and y cursor velocities. Every 100 ms, a new vector of neural observations is used to update the Kalman Filter's estimate of the subject's (in this case, simulation's) intended velocity by performing first a time update step on the previous estimate of cursor state according to the modeled cursor dynamics, followed by a measurement update step using the neural observations. More details can be found [73].

In order to fit the Kalman Filter decoder for each simulation, a calibration phase was completed to fit a decoder that would drive the BMI cursor from simulated neural activity, as is done in online experiments (described as 'visual feedback' in [73]). In this phase, pre-computed cursor trajectories on a center-out only task were played, and at each time step, cursor velocities were used along with the generative model to create a vector of spike activations resulting from a linear combination of the four different types of input sources. In this calibration phase, neural activity was not used to update the cursor. Using the calibration data, a Kalman filter decoder was fit between neural activity and cursor velocities (previously described as 'visual feedback seeding' in [73]).

5.2.3 Extracting Observed Shared and Observed Private Activity for Online Control

During the calibration ('visual feedback') phase, the generative model creates population neural activity driven by four input sources. Then, to analyze the observed shared and private variability in the resultant population activity, an FA model is fit. Once the model is fit to the calibration neural data, each vector of neural activity at each time point can be decomposed into a vector estimating the shared component and a vector estimating the private component (see A.3.2 for details):

$$y = E(y_{t,shared}|y_t = y) + E(y_{t,private}|y_t = y) + \mu$$

$$E(y_{t,shared}|y_t = y) = UE(z_t|y_t = y)$$

$$E(y_{t,private}|y_t = y) = y - E(y_{t,shared}|y_t = y) - \mu$$

$$E(z_t|y_t = y) = U^{'}(UU^{'} + \Psi)^{-1}(y - \mu)$$

The decomposed activity can then used to drive the BMI cursor to directly simulate the contributions of shared and private activity in BMI performance. Since segmenting neural activity into private and shared activity components reduces its overall variance, a scaling factor was calculated to restore both shared and private components to the same variance level as original activity prior to simulating through the decoder:

$$\beta^i_{shar} = \sqrt{\frac{\sigma^2_{i,full}}{\sigma^2_{i,shar}}}, \beta^i_{priv} = \sqrt{\frac{\sigma^2_{i,full}}{\sigma^2_{i,priv}}}, 1 \leq i \leq n$$

$$\beta_{shar} \in R^{nx1}, \beta_{priv} \in R^{nx1}$$

Here, $\sigma^2_{i,full}$ is the variance of the full neural activity for neuron $i$and $\sigma^2_{i,shar}$, $\sigma^2_{i,priv}$ are the variance of the computed shared and private activity. Thus the shared ($y_{t,shar}$) and private ($y_{t,priv}$) decompositions of neural activity are then:

$$\hat{y}_{t,sc.shared} = \beta_{shar}(UU^{'}(UU^{'} + \Psi)^{-1}(y - \mu)) + \mu$$

$$\hat{y}_{t,sc.private} = \beta_{priv}(y - \mu - UU^{'}(UU^{'} + \Psi)^{-1}(y - \mu)) + \mu$$

We note that this approach of fitting a separate scaling factor for each neuron is not the most principled approach to re-scaling shared and private neural activity. In fact, choosing a different scaling for each neuron makes it such that neurons that are largely driven by private activity and have a very small shared component part with have very small $\sigma^2_{i,shar}$ compared to their $\sigma^2_{i,full}$, resulting in a very large β^i_{shar}. In contrast, units largely driven by shared activity will have β^i_{shar} approximately equal to one. This behavior is undesirable since it causes neurons with low shared activity to contribute heavily to the estimate of $\hat{y}_{t,sc.shared}$. Should this experiment be repeated in either the online or offline context, estimating a single scaling factor for all neurons is advised in order to preserve the proper correlational structure between neurons. Despite this error, we persist. To simulate shared and private contributions to online BMI control, $\hat{y}_{t,sc.shared}$ and $\hat{y}_{t,sc.private}$ are passed into the decoder instead of y_t. All calibration and simulation sessions were run for 3 minutes, sufficient time for a mean of 65.9 trials. Some simulations contained more trials than others due to faster or slower cursor movement.

5.2.4 Online Experiment

The full experimental setup including electrophysiology and task description can be found in chapter 6, as these experiments were included in that data collection. Briefly, a non-human primate who was previously trained in neuroprosthetic control performed a standard centerout 2D cursor task (almost identical to the centerout task described in chapter 2, but now in neuroprosthetic control mode, instead of arm-reaching mode). In this task, the virtual BMI cursor (0.4 cm radius) was set to the middle of the workspace to begin a trial, and one of eight peripheral targets (2.1 cm radius) arranged radially around the center, a distance of 10 cm away, was cued. The subject, using tens of single and multi-units that were being recorded, binned every 100ms, and decoded with a velocity-Kalman Filter, attempted to drive the cursor to the peripheral target. Once at the target, they attempted to hold their cursor within the target for 200 ms, and was delivered an apple juice reward for successful trials. The next trial began when the subject was finished with the reward, or had spent too long attempting a target (timeout time, varying from 15-30 seconds).

In the online experiment, as in the simulations, we also estimate $\hat{y}_{t,sc.shared}$ and $\hat{y}_{t,sc.private}$ and use these values to drive cursor performance. Here, the factor analysis model is not fit from calibration data ('visual feedback block'), but rather from another block when the subject is performing online neuroprosthetic control. The experimental flow involved the first block, calibration block used to seed the Kalman Filter decoder, followed by a second block of short decoder adaptation used to refit the decoder [202]. After these two

blocks, the decoder was fixed for the remainder of the day and the subject performed a third block of 64 trials (8 to each target) using full neural activity to drive the cursor (y_t). A factor analysis model was fit using just data this third block, and subsequent blocks then used estimated $\hat{y}_{t,shared}$ and $\hat{y}_{t,private}$ as inputs to the decoder.

It should be noted that other groups exclusively used $\hat{y}_{t,sc.shared}$ (though parameterized differently and without the scaling factors) to drive the cursor [91], however they do not compare control of $y_t, \hat{y}_{t,sc.shared}$, and $\hat{y}_{t,sc.private}$. Finally, see A.1 for further details on how the model was fit for the simulations and the online experimental results.

5.3 Results

5.3.1 Effect of Source Weights on Observed Correlated Variability

The first question investigated was how variability in the 20-neuron population emerged from inputs by manipulating the αweights. The expectation was that any shared variability input to the population would emerge as observed shared variability once all four sources were summed, but it was unclear how the weighting between tuned and untuned sources may influence the final amount of observed shared variance. We simulated population with the following sets of alpha parameters: $\alpha_{untuned,shared} = 0.2, \alpha_{untuned,private} = [0, 0.25, 0.5]$. Depending on $\alpha_{untuned,private}$ the remaining weights for tuned variance were 0.8, 0.55, or 0.3. A ratio between $\alpha_{tuned,shared}$ and $\alpha_{tuned,private}$ was then swept, where ratio $= \frac{\alpha_{tuned,private}}{\alpha_{tuned,private}+\alpha_{tuned,shared}}$, and ratio sweeps [0, 0.25, 0.5, 0.75, 1.0] for a total of $1 \times 3 \times 5 = 15$ parameter sets tested. For each of these parameter sets, the amount of variance contributing to the input source was compared to the observed variance in the summed population activity. Specifically, during a 'visual-feedback calibration block', the following comparisons were computed:

$$Shared_{input} - Shared_{output} =$$

$$[\sum_{i=1}^{m} var(\alpha_{tuned,shar} * Uz^i_{t,tuned} + \alpha_{untuned,shar} * Uz^i_{t,untuned})^T_{t=0}] - trace(U_{out} * U'_{out})$$

$$Private_{input} - Private_{output} =$$

$$[\sum_{j=1}^{n} var(\alpha_{tuned,priv} * \psi^j_{t,tuned} + \alpha_{untuned,priv} * \psi^j_{t,untuned})^T_{t=0}] - trace(\Psi_{out})$$

If the input private and shared variances are directly reflected in the output shared and private variance, on average $Shared_{input} - Shared_{output} = 0$ and $Private_{input} -$

$Private_{output} = 0$. What happens instead is shown in 5.5. When there is little private tuned input, as on the lefthand-side of the x-axis, private and shared input roughly match private and shared output. In contrast, when there is mostly private tuning and little shared tuning, as on the righthand-side of the x-axis, private input variance is much greater than private output variance. Further, shared input variance is much lower than shared output variance. Thus, when input private activity is heavily tuned, it manifests in the resultant population as shared variance. Adding more private untuned activity (increasing redness in color of each point) dampens this effect. However, the more untuned private activity there is, the less tuned private activity there can be. For example, on the righthand-side of the x-axis, $\alpha_{tuned,private} = 0.8$ when $\alpha_{untun,private} = 0$ (black dots), but $\alpha_{tuned,private} = 0.3$ when $\alpha_{untun,private} = 0.5$ (red dots). Thus, there much less tuned private activity to be interpreted as shared activity when the untuned private activity is higher.

Why would input tuned private activity, that has been designed such that temporal correlations across neurons are broken via the Poisson process, exhibit output correlated shared activity? If we direct attention back to 5.4, the answer comes from the fact that the tuned private activity is being driven by a low-dimensional $u(t)$ signal. Although different neurons may have different tuning curves to the $u(t)$ signal, and although there is a separate Poisson process for each neuron to break temporal correlations, over the course of an entire calibration session the two-dimensional velocities used to compute the private tuned inputs are the same for each neuron. Thus, when modeling the private tuned neural activity, a model that used a two factors and a mapping from the factors to the neural activity fits well. Increasing temporal jitter by decreasing $\Lambda_{tuned,private}$ could be a way for input private, tuned activity to be reflected as private in the output regardless of the fact that the activity is driven by a low-dimensional source. We explore this possibility below.

5.3.2 Effect of varying $\Lambda_{private,tuned}$ on Observed Shared Variability

To investigate how the update rate may influence whether input variability is designated as private or shared in the resultant population, we varied the update rate ($\Lambda_{tuned,private}$) for just the tuned private input source while keeping all other parameters constant ($\alpha_{private,tuned} = \alpha_{private,untuned} = \alpha_{shared,tuned} = \alpha_{shared,untuned} = 0.25$). In the resultant population, we measure how much of the total population variance is designated as a shared by computing a shared-over-total metric:

$$SOT_{out} = \frac{trace(U_{obs}U'_{obs})}{trace(U_{obs}U'_{obs} + \Psi_{obs})}$$

Increases in SOT indicate that an increasing amount of a population's total variability is shared variability.

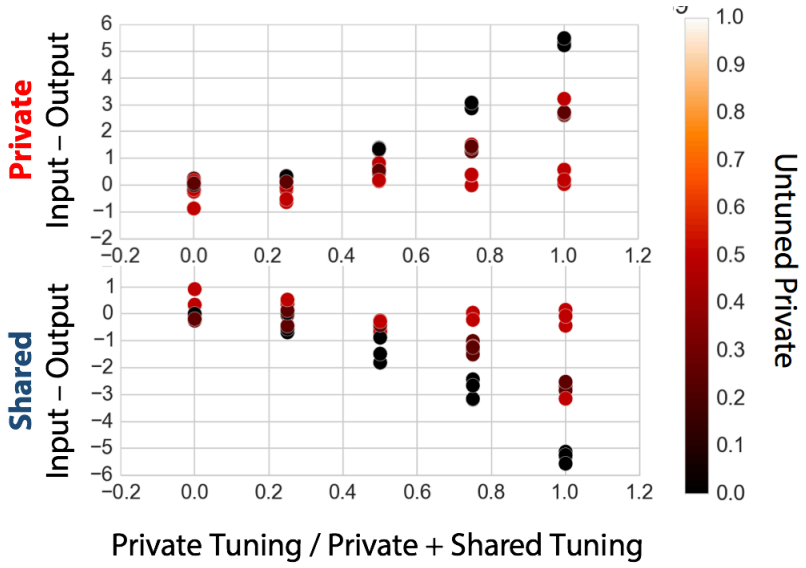

Figure 5.5: Differences in input versus output private and shared variability as a function of $\frac{\alpha_{tuned,private}}{\alpha_{tuned,private}+\alpha_{tuned,shared}}$ and $\alpha_{untuned,private}$.

5.6illustrates that as the update rate increases from 10 events / sec (1 event / 100 ms bin) to more than 100 events / sec, the resultant population exhibits increasing shared variability. This trend is in line with the findings of greater private tuning activity resulting in more shared output activity in 5.5, except that here instead of increasing the weight of private tuning, the weight remains constant but the update rate increases, making the private tuned events more likely to be coincident across neurons. Why SOT increases when the private tuned activity approaches zero is less clear. One possible explanation is that as the rate of private tuned activity decreases, the influence of private tuned activity disappears. Thus, although $\alpha_{private,tuned} = 0.25$, if $\Lambda_{private,tuned} = 0$, then the number of bins with non-zero $\lambda^i_{t,priv.,tuned}$ is very few. If $\alpha_{private,tuned} = 0$ then the contribution of shared sources and untuned private sources becomes: $\alpha_{private,untuned} = \alpha_{shared,tuned} = \alpha_{shared,untuned} = 0.33$, making SOT values of 0.6 - 0.8 understandable for small $\Lambda_{private,tuned}$.

Note that since our Factor Analysis model is fit with neural data binned in 100 msecs, the increase in private tuning event rate above the 10 Hz rate will result in more observable shared activity. For systems with narrower binning, the local minima of shared activity seen in 5.6 will be at the shorter update value.

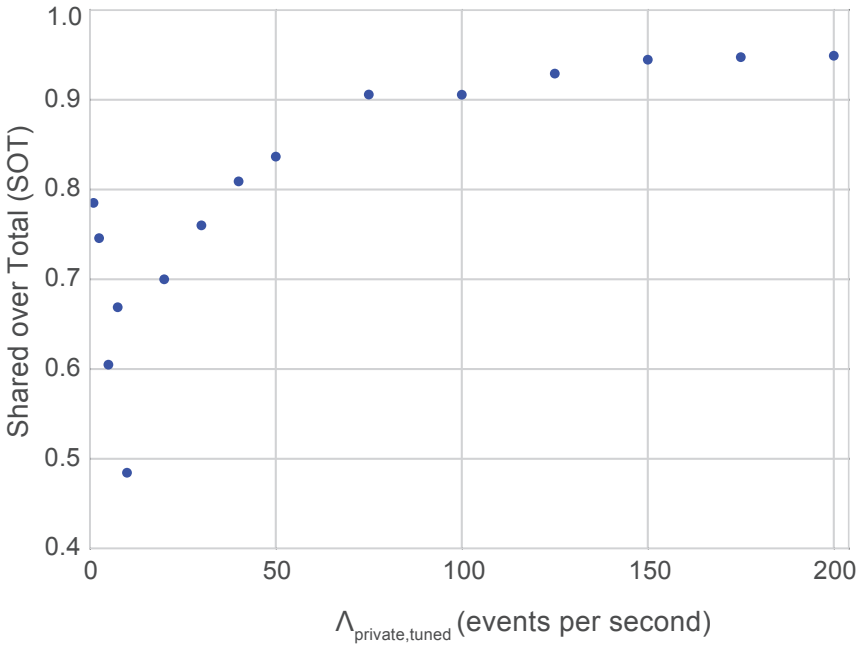

Figure 5.6: Observed shared variability in the population varies as a function of the $\Lambda_{private,tuned}$

5.3.3 Contributions of Observed Shared and Observed Private Activity to Online Control

We next consider how population with different amounts of input private and input shared activity perform in simulated online BMI control. First, cursor speed and accuracy toward the end target are measured in populations containing different $\alpha_{private,untuned}$, $\alpha_{shared,tuned}$, $\alpha_{shared,untuned}$ and then we consider how using observed shared or observed private affects performance.

Using the same populations as studied above in 5.5, now a full closed-loop simulation setup is used to study how populations with different weightings of input sources perform online. In these figures, each point is a single population where its position on the x-axis corresponds to the private tuned to total tuning ratio as before ($\frac{\alpha_{tuned,private}}{\alpha_{tuned,private}+\alpha_{tuned,shared}}$) and the color corresponds to $\alpha_{untuned,private}$. As before, $\alpha_{untuned,shared}$ is set to 0.2.

In the left figure, tuning ratio is plotted versus the path error. Path error computes the total distance between the cursor location and the straight-line path between the center start target and the peripheral target. The longer the path error, the worse the accuracy. The metric plotted was multiplied by -1 such that greater accuracy is at the top of the y-axis. The left figure shows that populations with more private tuning exhibit greater accuracy, and populations with less private tuning exhibit lower accuracy. The right figure shows that populations with no private tuning (leftmost on x-axis) exhibit the fastest cursor speeds. However, populations with a mix of private and shared tuning exhibit lowest speeds, and population with only private tuning exhibit the second highest speeds.

Why is accuracy linear with proportion of private tuning, but speed has a parabolic relationship? Perhaps speed is a better reflection of the observed activity in the resultant population, instead of the input sources, since the resultant population activity is what drives the decoder. Since 5.5 already establishes that relative contributions of input sources of variability do not always dictate the observed sources, plotting speed versus observed shared variability may yield a clearer correlation. In 5.8, speed is plotted versus observed shared over total variability ($SOT_{obs} = \frac{trace(U_{obs} * U'_{obs})}{trace(U_{obs} * U'_{obs} + \Psi_{obs})}$). Although speed does not exactly correlate with SOT_{out} ($R^2 = 0.35$), the relationship is linear. Thus, shared variability in the output population contributes to speed of the cursor. The relationship will not be perfect though, since observed shared variability can vary in its tuning, and more tuned observed shared variability ought to drive the cursor faster.

Thus far, we have shown that:

1. Inputs of private, tuned activity give rise to observed shared variability for specific regimes of $\Lambda_{tuned,private}$

2. Cursor accuracy in simulated online control is best achieved with tuned private inputs

3. Cursor speed is correlated with observed shared variability

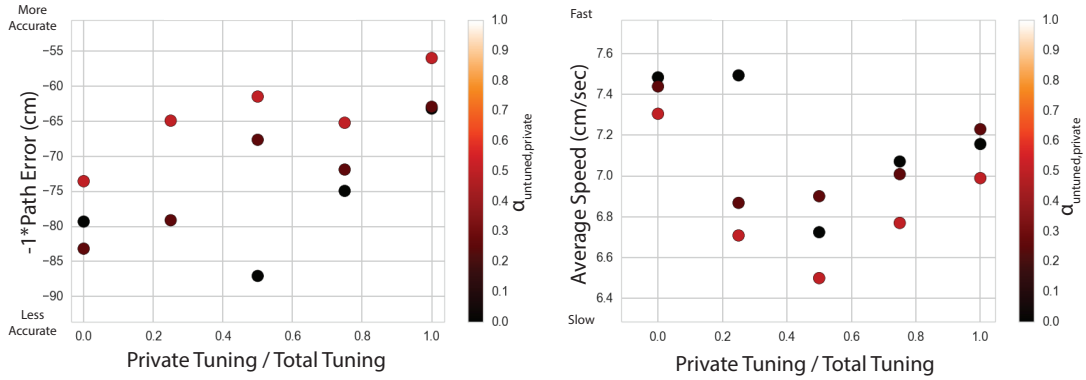

Figure 5.7: Left: Path error in simulated online BMI control as a function of tuned activity ratio. Note that higher on the y-axis is more accurate. Right: Average cursor speed during BMI control as a function of tuning ratio.

Given these findings, we ask whether extracting shared-only activity from the resultant population may improve online simulated BMI control.

5.3.4 Using Observed Shared Activity Online Boosts Speed when Private Tuning is High

So far, it has been shown that private, tuned inputs give rise to observed shared variability for specific regimes of $\Lambda_{tuned,private}$, that cursor accuracy in online control is best achieved with tuned private inputs, and that cursor speed is correlated with observed shared variability. Next, online BMI simulations are performed with scaled, shared activity ($\hat{y}_{t,shared}$) in order to see if extracting observed shared variance and removing observed private variance may serve to give a speed boost.

In simulation, extracting observed shared variance does have the effect of increasing speed, but only for populations dominated by private tuning as shown in 5.9. This is the same figure as 5.7(right), but now the same populations were re-simulated with $\hat{y}_{t,sc.shared}$ as the input to the decoder ($\hat{y}_{t,sc.shared} = \beta_{shar}(U_{obs}U'_{obs}(U_{obs}U'_{obs} + \Psi_{obs})^{-1}(y_t - \mu_{obs})) + \mu$). We hypothesize if (observed shared variance - input shared variance (y axis of 5.5) is plotted against average speed shared - average speed full activity (difference between circles and triangles in 5.9), there would be a clear correlation.

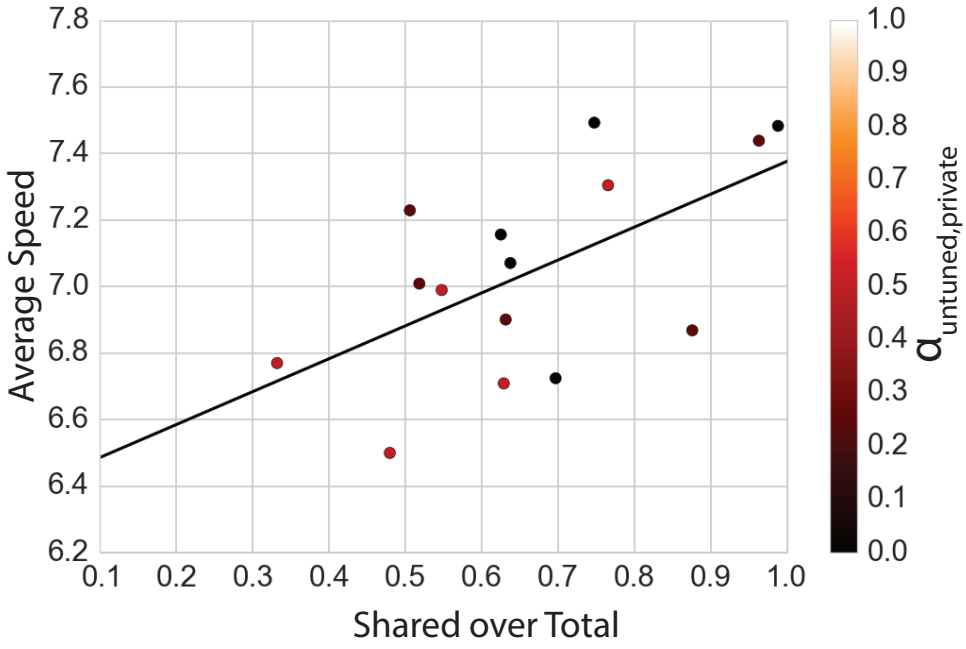

Figure 5.8: Same as 5.7(right) but now x-axis is observed shared over total (SOT_{obs}) activity

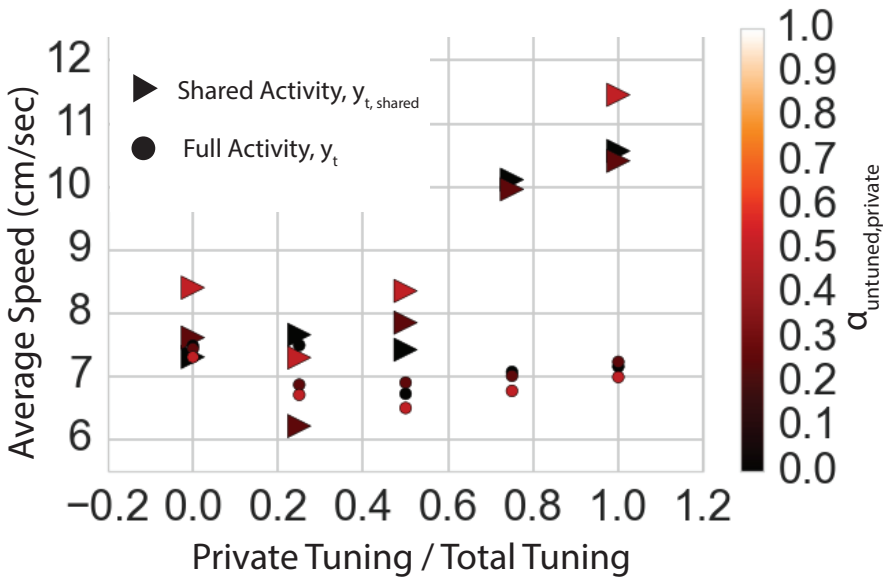

Figure 5.9: For all populations, scaled shared activity ($\hat{y}_{t,shared}$, triangle) was used as observations for the decoder, and the speed of the resultant trajectories were compared to their speed when using the full neural activity (y_t, circle)

5.3.5 Experimental Results Demonstrate Contributions of Distinct Sources of Neural Variability to Online Neuroprosthetic Control

In order to test how using estimated shared variance or estimated private variance online contributes to online performance, we performed a set of catch-trial experiments in one nonhuman-primate. The NHP performed the task described above, but on 30% of trials y_t was replaced with $\hat{y}_{t,sc.hared}$, on 30% of trials y_t was replaced with $\hat{y}_{t,sc.private}$, and the remaining 40% of trials were performed with y_t. The data presented below consists of 1167 trials collected over a few days.

5.3.5.1 Cursor Speed is Faster when Using Scaled Shared Activity

First cursor speed was analyzed since that was the attribute most obviously noticeable to the experimenter when scaled shared activity was used compared to scaled private activity. Plotted in 5.10 is the average cursor speed for trials driven by y_t,$y_{t,sc.private}$, $y_{t,sc.shared}$. When using $y_{t,sc.shared}$, speed is significantly higher than the original y_t, and when using $y_{t,sc.private}$, speed is significantly lower than the original (Significant Kruskal-Wallis test followed by two-tailed Mann-Whitney, number of trials: y_t = 580, $y_{t,sc.shared}$=324, $y_{t,sc.private}$=263, p-values; $y_t vs. y_{t,sc.private}$, p = 1.8e-57, y_t vs. $y_{t,sc.shared}$, p=3.7e-69, $y_{t,sc.shared}$

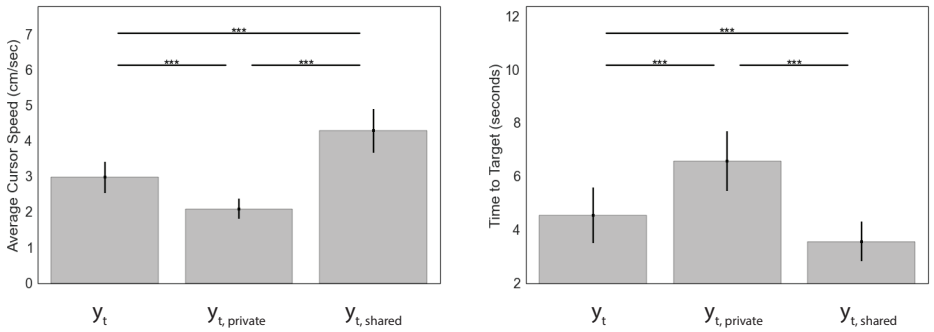

Figure 5.10: Changes in average cursor speed (left) and time to target (right) when using y_t,$y_{t,private}$, $y_{t,shared}$ online to drive the same Kalman Filter decoder.

vs. $y_{t,sc.private}$, p=3.0e-91). This finding matches with the simulation results in the regime of a high ratio of private tuning to shared tuning.

Also assessed was time to target, to see whether the increase in speed was accompanied by a lack of control, as would be evidenced by an increase in time to target for $y_{t,sc.shared}$. Time to target was significantly lower for $y_{t,sc.shared}$ than y_t, and was significantly higher for $y_{t,sc.private}$ than y_t. (Significant Kruskal-Wallis test followed by two-tailed Mann-Whitney, number of trials: y_t = 580, $y_{t,sc.shared}$=324, $y_{t,sc.private}$=263, p-values; $y_t vs. y_{t,sc.private}$, p=6.5e-45, $y_t vs. y_{t,sc.shared}$, p=5.3e-20, $y_{t,sc.shared} vs. y_{t,sc.private}$, p=3.0e-65). Thus, using scaled shared variance online did not make the cursor so fast that it was uncontrollable.

5.3.5.2 Average Cursor Path Error is Higher when Using Scaled Shared Activity

Finally, we investigate how cursor accuracy changes for trials driven by y_t, $y_{t,sc.private}$, $y_{t,sc.shared}$. The above analysis already shows that cursor accuracy is not so low when using $y_{t,sc.shared}$ that the time to target is negatively affected. In fact, time to target is significantly lower (faster) when using $y_{t,sc.shared}$ compared to either $y_{t,sc.private}$ or y_t. Plotted in 5.11 are the average and total path errors. The average path error for $y_{t,sc.shared}$ is significantly higher than the path error for y_t, $y_{t,sc.private}$ (Significant Kruskal-Wallis test followed by two-tailed Mann-Whitney, number of trials: y_t = 580, $y_{t,sc.shared}$=324, $y_{t,sc.private}$=263, p-values, $y_t vs. y_{t,sc.shared}$, p=3.5e-04, $y_{t,sc.shared} vs. y_{t,sc.private}$,p=3.4e-06), and the total path error for $y_{t,sc.private}$ is significantly higher than the total path error for y_t, $y_{t,sc.shared}$ (Significant Kruskal-Wallis test followed by two-tailed Mann-Whitney, number of trials: y_t = 580, $y_{t,sc.shared}$=324, $y_{t,sc.private}$=263, p-values, $y_t vs. y_{t,sc.private}$, p=2.8e-07, $y_{t,sc.shared} vs. y_{t,sc.private}$, p=1.9e-08).

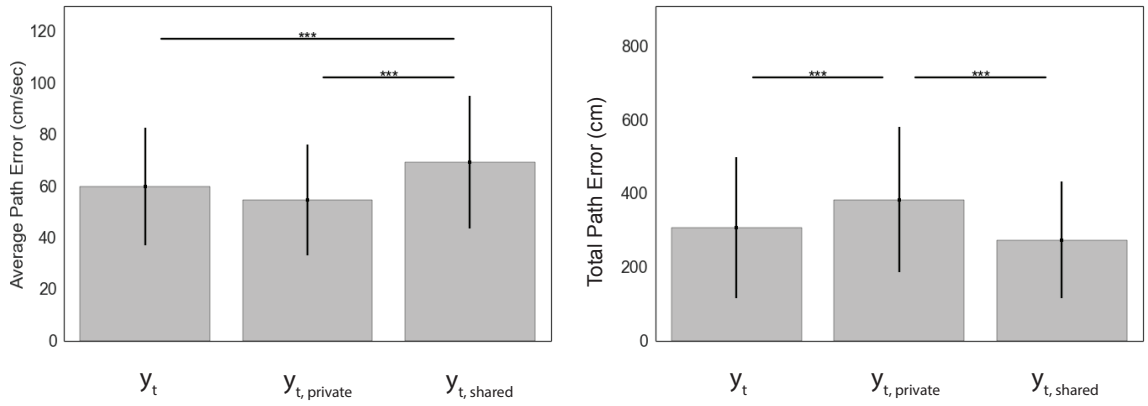

Figure 5.11: Changes in average path error (left) and total path error (right) when using y_t,$y_{t,sc.private}$, $y_{t,sc.shared}$ online to drive the same Kalman Filter decoder.

5.4 Conclusions

We first simulated inputs to a population of neurons to parse how inputs can influence to presence of observed shared variability in a population. We found that tuned, private inputs can generate observed shared variability, since they are themselves driven by a low-dimensional signal, $u(t)$. We also found that increasing the update rate $\Lambda_{tuned,private}$ can result in more observed shared variability, since each neurons' private encodings of low-dimensional $u(t)$ are more likely to be coincident. Similarly, decreasing the update rate below the BMI rate will also serve to increase observed shared variability, since the tuned private source contributes less and less to the total population activity. This relationship between private inputs and observed shared variability suggests a mechanism by which activity that is private may become shared. Given observations of private variability dominating observed population activity shown in the beginning of learning a BMI and shared variability dominating in late learning [90], perhaps an analogous increase in update rate drives this change in structure.

BMI populations with greater observed shared variance exhibit greater speed. This makes sense in our model since the populations with the greatest observed shared variance are the ones with the least untuned private inputs. Thus, all shared variance is from untuned shared variance, and tuned shared and private variance. Since untuned shared variance is the same in all of the above populations, increases in speed can be attributed to the tuned private and tuned shared variance that become observed shared variance. Since the Kalman filter decoder is a linear decoder, when inputs encoding the same $u(t)$ are synchronous as they are when they constitute shared variance, the decoder has more observations of the same $u(t)$ signal. In our model when inputs are asynchronous as in tuned private inputs, some neurons will encode $u(t)$ while the other will have no activity, giving the decoder fewer observations of the $u(t)$ signal within a single timepoint. Note that in general, just because a population has shared variance doesn't necessarily mean

that this variance is encoding a useful variable, it just happens to be true in our case because we did not vary the untuned shared input.

Observed private activity still drives the cursor towards the target, but along a slower trajectory more accurate (on average) trajectory. While it has been postulated that a low-dimensional shared representation of neural activity may be sufficient for driving proficient BMI control and even that learning may be limited to this space [91], we show by decomposing neural variance into shared and private sources in an online experiment, that each has its own contribution to performance.

A number of unanswered questions remain. Why is private activity used online more accurate in the online experiment? If all tuned private activity appears as observed shared variability, there should be little tuned information left in the observed private activity. It is possible that the regime most analogous to the online results is a private tuned ratio of approximately 0.5, indicating that some private tuned activity will appear as observed output but not all. In this regime, might there be conflicts between private tuning and shared tuning? Might a single neuron's shared variability component be tuned to a different direction as their private component? Perhaps separating observed private and observed shared sources could improve overall information in the population. Another related question is how the dimensionality of the shared variance changes depending on whether it is driven by private tuned inputs or shared tuned inputs, or a mix of both. Perhaps the dimensionality of the shared variance when using private tuned inputs is higher. Establishing this could improve our ability to figure out which regime real data resides in.

Another simplification in the model is the assumption that subjects generate a representation of $u(t)$ where they always point at the target. As highlighted in [203], subjects can only aim exactly where they want to go if they have a uniform distribution of tuning curves about the angular axis. In the case that they don't, they must learn to re-aim depending on the tuning of the units chosen as 'output units'. Thus, this model could be improved if the function used to generate $u(t)$ was fit based on actual subject behavior. Doing so could allow for better comparisons of how a simulated population compares to the a real neural population during the same behavior, and used to explore principles of how real neural populations operate.

Finally, the generative model of population activity models neural activity using cosine tuning to an intended direction of movement, a representational model of how motor areas would encode movement objectives. Although grounded in decades of electrophysiological findings [29], this generative model is very simple, and does not capture recent findings that neurons change their tuning properties based on a variety of task parameters [204, 205, 206]. However even if many independent tuning curves do not present a useful model of motor encoding, the result that independent, noisy observations of a low-dimensional quantity can appear as correlated activity is still an important principle. At the most basic processing level, a high dimensional number of primary sensory neurons samples the three dimensional world and give rise to inputs that drive motor areas, so somewhere in this processing is a collapse of high dimensional private observations to a lower-dimensional

shared representation. While this model is a simple abstraction, it yields a number of basic principles that guide our investigation into studying motor population computation during actions sequences in the next chapter.

Chapter 6

Distinct Action Sequences Performed During Neuroprosthetic Control Reveal High-Dimensional Neural Dynamics

6.1 Introduction

How does motor cortex generate action sequences which constitute our ability to reach for the newspaper to read it ourselves one morning, and then reach for the newspaper and give it to our friend the next? The initial newspaper-directed part of these hand reaches will resemble each other kinematically, but are these identical portions generated by the same motor cortical neural activity patterns? With a focus on primary motor cortex (M1), we ask how do neural activity patterns for a given action change depending on the sequence in which the actions fall? Early work in understanding motor cortical coding of motor sequences reports changes in M1 neural firing rates [207] and M1 neural correlations [208] during planning phases for the same motor actions that occur in different planned sequences. These results suggest that motor cortex encodes information about sequence context or future movement [209, 172, 210]. However, we cannot be certain that all aspects of motor actions are in fact identical when performed in one sequence versus another. Perhaps a seemingly identical motor action measured kinematically in fact requires more co-contraction of the trunk or stabilizing muscles when performed in one sequence compared to another. Thus, finding neural activity patterns that are different for the same kinematic action performed in different sequences may just be evidence of different muscular activations for the same kinematic action. For example, in speech production, the same sound may be produced with different muscle activations depending on the sequence it falls due to the dynamics of the muscles and jaw [211]. Another challenge in interpreting differences in neural firing or correlations is that others have found consistent M1 firing patterns for a given action regardless of sequence [212]. One approach to simplifying the problem of studying natural movements and assessing whether the observed changes in neural activity contribute substantially to behavior is to use a brain-machine interface (BMI) where a population of motor cortical neurons drive the velocity commands to a virtual 2D cursor. In this setup, the experimenter knows exactly which neurons contribute to movements of the cursor, and know the exact mapping between neural activity and kinematics of the cursor behavior, as discussed in chapter 1. Further, the experimenter can decompose neural activity patterns to assess how different parts contribute to cursor control.

Another complication with studying neural encoding of action sequences is that even if kinematics and muscular patterns are identical when performing the same action in different sequences, the issue of which coordinate frame M1 represents has remained largely unresolved [30, 29]. Thus, studying neural encoding of action requires the experimenter to first select their favorite coordinate frame for M1, a selection that is not obvious. A different approach to studying the neural correlates of action is leave the question of which movement parameters M1 represents and instead focus on identifying principles of how M1 generates movement. The redundancy of neural solutions to generate the same EMG output may explain how different neural firing patterns generate the same movement depending on the sequence it falls in. However, understanding the principles of why certain solutions are used for one context versus another remain unclear. Some have

embarked on studying properties of M1 population activity and generated proposals of how M1 generates some types of behaviors, though few have focused on action sequences directly. For example, it has demonstrated that M1 population activity occupies a low-dimensional subspace, and subjects have difficulty generating neural activity that departs from the space even when doing so would allow for task-related rewards [91]. However, how neural activity traverses the low dimensional subspace to generate action sequences remains unclear. Further, why the space tends to be higher-dimensional than required by the task is also unknown. Another group has proposed a theory of movement generation which models neural activity during very well-practiced, fast arm reaches, as observations from a low-dimensional linear dynamical system, with a low-dimensional population state evolving according to a consistent lawful linear dynamical process [33, 34]. When subjects initiate their reaches, the low dimensional population state starts at a reach-specific initial position and consistent temporal dynamics evolve its state to produce a temporal pattern that has the correct frequency content to produce EMG activity [33]. This model captures substantial neural variance, and explains rotational characteristics of population activity that other representational models do not capture. Whether a consistent dynamical system can also explain complex behaviors such as feedback corrected reaches, or reaches requiring the subject to hold for an extended period of time before reaching is beginning to be explored [35, 36, 37].

Here, we investigate how M1 generates action sequences by using a brain-machine interface. In this preparation, we record from tens of M1 single and multi-units and linearly map the activity of these neural activity patterns to the velocity of a 2D cursor on a screen [213]. In different tasks, the subject must generate distinct cursor trajectories, allowing us to precisely study the neural activity patterns that generate different action sequences. The BMI preparation gives us an exact assay of the behavioral variables (velocities sent to the cursor) and neural-to-behavioral mapping (decoder used in the BMI), in contrast to natural reaching experiments where behavior is more challenging to rigorously capture and the neural to behavioral mapping is unknown and nonlinear. Within the BMI preparation, we first ask how neural activity used to generate individual behavioral actions differs depending on which action sequence the behavior is embedded. We assess if the neural activity patterns for different action sequences occupy a consistent subspace, as predicted by [91]. Given that the subspace found by Sadtler is higher dimensional than the task dimensions, there are multiple ways in which subjects may generate within-subspace neural activity which have the same behavioral output. Perhaps they use this redundancy to have flexibility in how they generate the same behavior depending on which action sequence it falls within. We also assess the mean firing rates of each neuron for a given behavioral action. One hypothesis for how subjects perform BMI control is that they infer through experience the tuning parameters of the neuron-to-cursor mapping in the decoder [71, 73] and then reliably generate each neural firing rate corresponding to the direction they want the cursor to move. Again, there is substantial redundancy in the many neuron to two-dimensional cursor mapping, so we assess whether the same behavioral actions (cursor velocities) are generated by different mean firing rates depending on

action sequence. These analyses reveal if in a well-controlled BMI task mean firing rates and covariation patterns differ for generation of identical behavioral actions depending on action sequence. These analyses are the beginning steps to understanding principles of how M1 leverages its redundancy to generate activity patterns.

Finally, we investigate if the differences in M1 neural activity for the same behavioral action can be explained by M1 neural activity abiding by temporal rules that can be modeled with a linear dynamical system. Indeed, using a linear dynamical system fit from arm movements has already been shown to capture substantial neural activity in an online BMI task, yielding improved BMI performance compared to models that do not model the temporal rules of neural activity [214, 215]. However, we don't ask if this model is the best way of representing neural activity for online BMI performance, but rather if it can distill principles of computation in M1 that yield different neural activity for identical behaviors within action sequences. We investigate whether the inclusion of these temporal rules explains our subspace difference and mean firing rate difference observations. We conclude proposing that the high-dimensionality and temporal laws in M1 give rise to a flexible generator of action sequences.

6.2 Methods

6.2.1 Surgery, electrophysiology, and experimental setup

Two male rhesus macaques (Macaca mulatta, RRID: NCBITaxon:9544) were chronically implanted with arrays of 128 Teflon-coated tungsten microwire electrodes (35 μm in diameter, 500 μm separation between microwires, 16 × 8 configuration, 6.5 mm length, Innovative Neurophysiology, Durham, NC) in the left upper arm area of primary motor cortex (M1) and posterior dorsal premotor cortex (PMd). Localization of target areas was performed using stereotactic coordinates from a neuroanatomical atlas of the rhesus brain [216]. LFP activity was recorded at 1 kHz using either the 128-channel Multichannel Acquisition Processor (Plexon, Inc., Dallas, TX) (Monkey J) or the 256-channel Omniplex D Neural Acquisition System (Plexon, Inc.) (Monkey G). Single-unit and multi-unit activity were sorted online, after setting channel thresholds. Thresholds were set at the beginning of each session based on 1–2 min of neural activity recorded as the animal sat quietly (i.e. not performing a behavioral task). For all monkeys and tasks in this study, visual feedback of the BMI output was shown by a circular cursor on the task screen. Prior to this study, Monkeys G and J were trained at reaching tasks and spike-based brain-machine interface (BMI) cursor tasks for 1 year. All procedures were conducted in compliance with the NIH Guide for the Care and Use of Laboratory Animals and were approved by the University of California, Berkeley Institutional Animal Care and Use Committee.

6.2.2 Tasks

Two BMI tasks are performed – a centerout task and an obstacle avoidance task. In the centerout task, subjects control a 2D cursor to move from the center target to one of eight cued peripheral targets distributed radially around a 6.5 cm (Monkey J) or a 10 cm (Monkey G) radius circle. For Monkey J, trials were initiated by entering the center target and holding for a variable time. The go cue after the hold period was indicated by the center target changing color and the peripheral target illuminating, cuing a reach to that target. For Monkey G trials were initiated automatically by resetting the cursor to the center position, and illuminating a peripheral target. A liquid reward was provided after the cursor successfully reached the target with a peripheral hold time of 200 ms.

In the obstacle avoidance task, Monkeys G and J performed different tasks. Monkey G performed an obstacle avoidance task with a very similar structure to the center-out task. The only difference was that a square obstacle (side length 2 or 3 cm) would appear in between the center cursor reset position and peripheral target. If the cursor center entered the obstacle, the trial would end in an error and subjects would repeat the trial. Monkey J's obstacle-avoidance task required a point-to-point movement between an initial (not necessarily center) target and a peripheral target. On arrival at the initial target, an ellipsoid obstacle appeared on the screen. If the cursor entered the obstacle at any time during the movement to the peripheral target, an error resulted and the trial was repeated. Target positions and obstacle sizes and positions were selected to vary the amount of obstruction, radius of curvature around the obstacles, and spatial locations of targets. Trials were constructed to include no obstruction, partial obstruction with low-curvature, full obstruction with a long distance between targets and full obstruction with a short distance between targets thus requiring a high curvature.

We analyzed 9 days of data from Monkey G and 4 days of data from Monkey J where on each day, monkeys performed both the centerout and obstacle tasks with the same BMI decoder. Only successful trials were analyzed.

6.2.3 Decoding

In both tasks, subjects must draw upon their action repertoire to generate trajectories that bring their cursor to the cued peripheral target. In both the Kalman Filter decoder (Monkey G) and the Point Process Filter decoder (Monkey J), the cursor position and velocity are modeled with a state space with position and velocity variables with a dynamics process:

$$x_t = Ax_{t-1} + w_{t-1}, w_{t-1} \sim N(0, W)$$

$$n_t \sim p(n_t|x_t)$$

where x_t is a 5 x 1 vector corresponding to $x_{pos}, y_{pos}, x_{vel}, y_{vel}$, and an offset variable. Note that x_t always refers to the full 5 x 1 vector, whereas x_{pos}, x_{vel} each refer to specific cursor

variables. Additionally, n_t are binned spike counts (100 ms bins for Monkey G, 5ms bins for Monkey J).

For the Kalman Filter, the observations process $n_t \sim p(n_t|x_t)$ is:

$$n_t \sim N(Cx_t, Q)$$

$$n_t = Cx_t + q_t, q_t \sim N(0, Q)$$

Here, C is an $nx5$ matrix where each row represents the preferred direction and magnitude of each neuron in a cosine tuning model. In these experiments, neural activity is modeled as a linear function of velocity, so $C[:, 1, 2] = 0$. At each time bin, an estimate of x_t is computed from the previous estimate of $x(\hat{x_{t-1|t-1}})$ through the time update process $(\hat{x_{t|t-1}} = E(x_t|\hat{x_{t-1|t-1}}))$, and then a refined estimate of is computed using the observed spike counts, (the measurement update, $\hat{x_{t|t}} = E(x_t|\hat{x_{t|t-1}}, n_t)$). For the Kalman Filter, these two update steps can be summarized with the follow expression:

$$\hat{x_{t|t}} = F\hat{x_{t-1|t-1}} + K_t n_t$$

where K_t is the Kalman Gain and takes into account $C, Q, \hat{P_{t|t}}$ where $\hat{P_{t|t}} = E((\hat{x_{t|t}} - x_t)(\hat{x_{t|t}} - x_t)^T)$. K_t approaches steady-state in BMI tasks within seconds ($K_t \rightarrow K$, [217]). Further details about velocity Kalman Filter decoding are available [73, 218].

For the Point Process Filter, the observations process $n_t \sim p(n_t|x_t)$ is:

$$p(n_t|x_t) = \prod_{i=1}^{n} (\lambda_i(t|x_t, \phi_i)\Delta)^{n_t^i} exp(-\lambda_i(t|x_t, \phi_i)\Delta)$$

where $\lambda_i(t|x_t, \phi_i)$ is each output units' instantaneous firing rate modeled as a log-linear function of x_t:

$$\lambda_i(t|x_t, \phi_i) = exp(\beta_i + \widetilde{\alpha_i} x_t)$$

where $\tilde{\alpha_i}$ is tuning model parameter that only includes the velocity terms: $\tilde{\alpha_i} = [0, 0, \alpha_{xvel}, \alpha_{yvel}, 1]^T$. Thus, $\beta_i, \tilde{\alpha_i}$ are the parameters that model how an output units' firing rate (λ_i) modulates with velocity, and the expression for $p(n_t|x_t)$ describes the Poisson Point Process model that takes into account the velocity-modulated λ_i. The Point Process Filter decoder also has a time-update and measurement update steps that may be summarized with the following expression:

$$\hat{x_{t|t}} = F\hat{x_{t-1|t-1}} + \hat{P_{t|t}} \sum_{i=1}^{n} \tilde{\alpha}_i (n_t^i - \lambda(t|\hat{x_{t|t-1}}, \phi_i)\Delta)$$

Further details about Point Process Filter decoding can be found in [219, 220, 221].

6.2.4 Analysis

6.2.4.1 Repertoire Similarity

The repertoire similarity metric used in Figure 1G computes what percentage of the 32 sections have differences in distributions that are more than half of the mean of their distributions:

$$similarity = 1 - \frac{\sum_{b=1}^{32} |p_{co}(b) - p_{obs}(b)| > 0.5 * [p_{co}(b) + p_{obs}(b)]}{32}$$

where $p_{co}(b)$ refers to the empirical probability of observing a velocity command that falls in section b.

6.2.4.2 Subspace Overlap Metric

The subspace overlap metric is a uni-directional metric that computes what fraction of variance from one shared space is captured by another shared space. For example, if one has two factor analysis models (reminder: $k$is dimensionality of z, U is the loading matrix between $z$to x, Ψ is a diagonal matrix characterizing private variance and $\mu$is the population mean):

$$FA_1 = \{k_1, U_1, \Psi_1, \mu_1\}, FA_2 = \{k_2, U_2, \Psi_2, \mu_2\}$$

The subspace overlap from between the two first requires finding the main shared variance (UU^T_{main}, see A.2) and projection matrix to main shared spaced (P_{main}):

$$overlap_{1\to2} = \frac{trace(P_{main,2} UU^T_{main,1} P^T_{main,2})}{trace(UU^T_{main,1})}$$

Note that $overlap_{1\to2} \neq overlap_{2\to1}$. When computing this metric for within-task vs. across-task overlap as in 6.3.1.1, both directionalities were considered.

6.2.4.3 Comparisons of mean firing rate for a given command

In 6.3.2G, we compare the mean firing rate of each neuron for a given velocity command across both across tasks as well as within task. We followed the following procedure for each comparison outlined in 6.3.2D:

1. Across Task vs. Within Task:
    (a) First, for each day, spike counts and neural push velocity vectors from 100ms bins from successful trials were accumulated.
    (b) Spike counts and their corresponding neural push vectors from the 100ms bins from each task were randomly assigned to subset 1 or subset 2 (yielding CO_subset1, CO_subset2, OBS_subset1, OBS_subset2).

(c) Within each subset, neural spike counts and corresponding neural push vectors were further discretized by assigning the neural push vectors to their corresponding velocity sections.

(d) For each subset and each velocity section, a mean neural firing rate vector was computed. If fewer than 15 bins were assigned to a particular velocity section, the mean neural firing rate vector was populated with NaNs.

(e) To compute the across-task differences in mean firing rates, the absolute difference in mean firing rate between i) CO_subset1 and OBS_subset1 ii) CO_subset2 and OBS_subset2 was computed for each velocity section. These calculations yielded two n x 1 vectors per velocity section. Each of these vectors was then summed across neurons, yielding 2 'across-task-population-difference-estimates' for each velocity section. If any of the subsets did not have enough data to estimate a mean firing rate, then the population difference estimate that used that subset was not used.

(f) To compute the within-task difference, the same procedure described above in (e) was followed, except instead of taking the difference between i) CO_subset1 and OBS_subset1 and ii) CO_subset2 and OBS_subset2, now differences were taken between i) CO_subset1 and CO_subset2 and ii) OBS_subset1 and OBS_subset2. This procedure also yielded two 'within-task-population-difference-estimates'.

(g) Each velocity section then had two across-task-population-difference-estimates and two within-task-population-difference-estimates. The absolute difference between the first across-task-population-difference-estimate and first within-task-population-difference-estimate was used to compute one across-vs-within-task-population-difference-estimate. This was repeated for the second within and across task population-difference-estimate.

(h) At the end, for each day, 2 x n_section estimates (or fewer, if some sections has too few samples to estimate) were obtained to estimate the across-vs-within-task-population-difference-estimate. These values constituted Bar 1 in 6.3.2G.

2. Across Task Close vs. Across Task Far:

(a) This procedure was very similar to the Across Task vs. Within Task procedure. Instead of creating two CO subsets and two OBS subsets using the full task data, instead two CO and OBS subsets were created based on time in session. On each day, the 16 trials of the CO task session closest in time to the OBS task session was identified and used to create CO_close (light green in 6.3.2E). Then, 16 trials of the CO task session furthest in time to the OBS task session was identified and used to create CO_far (dark green in 6.3.2E). Similar sections were made for the obstacle task (16 trials closest to CO session used

to make OBS_close – light blue in 6.3.2E, 16 trials furthest from CO session used to make OBS_far—dark blue in 6.3.2E).

(b) For each subset and velocity section, a mean population firing rate was computed using data from that subset.

(c) To compute the close and far across-task-population-difference-estimate, the absolute differences between i) OBS_close and CO_close and ii) OBS_far and CO_far were computed for each velocity section, respectively. After summing each vector across neurons, this yields one close and far across-task-population-difference-estimate for each velocity section.

(d) To compute the close-vs-far-across-task-population-difference-estimate, the absolute difference was taken between the close and far across-task-population-difference-estimate for each velocity section, yielding for each day n_section estimates (or fewer, if some sections had too few samples to estimate).

3. Within Task Close vs. Within Task Far:

(a) Same as above, but with subsets chosen to be i) CO_subset,close1 as the 16 trials before the midpoint of the CO task, ii) CO_subset, close2 as the 16 trials after the midpoint of the CO task, and iii) OBS_subset, close1 iv) OBS_subset, close2, v) CO_subset, far1 as the first 16 trials of the CO task, vi) CO_subset, far2 as the last 16 trials of the CO task, vii) OBS_subset, far1, and viii) OBS_subset, far2. Here, initial differences were taken between: i. CO_subset, close1 and CO_subset, close2 ii. OBS_subset, close1 and OBS_subset, close2 iii. CO_subset, far1 and CO_subset, far2 iv. OBS_subset, far1 and OBS_subset, far2 And then final differences were taken between: CO_close vs. CO_far and OBS_close vs. OBS_far.

6.2.4.4 Data and Performance Metrics for Online Subspace Testing

For the tests of how using the shared space from one task influences performance on the other task (6.3.1.1E, F), data from 9 days with Monkey G were used (overlapping, but not identical to the 9 days used for the main centerout vs. obstacle comparisons). For these tests, the metrics used for comparison between online performance with within-task vs. across-task subspaces was normalized time to target and normalized path length:

BMI Performance Metrics Time to target was computed as the amount of time it took to leave the center and arrive at the peripheral target for successful trials. Faster times to target indicate better performance with the decoder. One issue in comparing within vs. across task shared spaces for online control is that there are four permutations we would ideally aim to test on a single day: CO_space-CO_task, OBS_space-CO_task, CO_space-OBS_task, OBS_space-OBS_task. Further, to compute CO_space and OBS_space,

an initial block of performance was needed on respective tasks. Finally, we also used decoders that were either seeded with the CO task or OBS task. All of the above conditions resulted in an inability to perform all decoder-task-subspace permutations on a single day. Unfortunately, comparing BMI performance across days in this animal yielded highly variable results depending on unit quality for that particular day, motivation level, and as we find later – whether the CO or OBS task was used as the seed decoder. Thus, we compute normalized performance metrics for easier comparisons across days. On each day, regardless of whether the CO or OBS task is used to seed the decoder, and regardless of which shared space is tested on which task, there is always a block in which the decoder is used to perform the task tested on that day without any shared space manipulations. Thus, we normalized all time to target metrics using that block. For example, assume on day 1 we want to test how a shared space fit from the CO task performs when the subject is trying to do the OBS task. We could compute time to target for all of the CO_space-OBS_task successful trials. Then, for each trial to target i, we subtract the mean time to target measured on that same day, with that same decoder, performed on the same task (here, OBS), with full activity n_t without any shared space manipulations, to target i. This way, any biases in the decoder related to speed or direction do not influence our ability to detect differences in using one shared space compared to another.

The same procedure was used for path length. First raw path length was computed as the sum of the distances the cursor travel from time point to time point over the course of a successful trial to target $i$during a specific task. Then, normalized path length was computed by subtracting the mean path length during baseline performance on the same task, with same decoder, without shared space manipulations, to the same target from raw path length.

Rewards per minute were computed by estimating the number of rewards obtained within 2 minute windows during task performance. Thus if a particular block was performed for 11 minutes, 5 values were obtained for 'rewards per minute' for that block. No normalization was performed on the rewards per minute estimates presented in 6.3.2.

Tuning Mismatch Metric For each BMI output unit, we compared the mismatch in what the decoder tuning curve model to a model estimated from online performance. For each block of online BMI performance, we regressed the cursor velocities against the neural activity patterns using linear regression, to yield an estimate of how the subject was 'using' the BMI output neuron i in online control, v^i_{task}. We then took the difference between the decoder-fit preferred vector direction $C(i, 3:4)$ (see 6.2.3 for details about decoder) and preferred vector direction fit from the task v^i_{task}, $dv^i = C(i, 3:4) - v^i_{task}$, and summed d^i_vover all neurons $i$to yield a single DV for each block. The angle of $DV$was compared across blocks in 6.3.2.

6.2.4.5 Linear Dynamical System

The linear dynamical system that was fit to our data used very similar methods to [214]. Here we fit the following model to our neural data:

$$z_t = Az_{t-1} + w_t, w_t \sim N(0, W)$$

$$n_t = Cz_t + q_t, q_t \sim N(0, Q)$$

If the above model looks similar to the Kalman Filter Decoder, it's because it's exactly the same. The Kalman Filter Decoder assumes that kinematics of the cursor (x_t) follow a dynamics process (first equation), and that neural activity (n_t) is a linear readout of the cursor kinematics. In contrast, the way this model is used in the work below, is the model the dynamics of the neural activity itself. Thus, $z_t$is a low-dimensional latent factor, must like it was in FA. However, now instead of $E(z_{t-1}z_t^T) = E(z_t z_{t+1}^T) = 0$ as in FA, now there are temporal correlations modeled in z_t. Thus, the activation of the latent variables cannot jump around from time point to time point, but must follow some temporal rules governed by A.

Fitting In all analyses, we use a 15-state linear dynamical system. Thus, $z_t \in R^{15}, A, W \in R^{15x15}$. To fit an LDS (and all the models the LDS is compared to in 6.1, 6.4) for a particular task's data, 20% of trials from the CO and OBS task from a particular day is held out as training data ($N_{test,CO}$ and $N_{test,OBS}$). The remaining 80% of trials were used as training data. Since temporal structure is vital to maintain in fitting the LDS, each 'datapoint' was a single trial instead of just a bin of data. Each trial gets one turn being in the test dataset, and 4 turns being in the training dataset.

For the test and training data sets, trials were organized in structures and included a full 1 second of data before their actual start. An LDS was fit using expectation-maximization [222]implemented in python package (https://github.com/mattjj/pylds) with the following relevant parameters:

- number of states: 15
- number of EM iterations: 30
- initialization: from Factor Analysis fit to data, as in [214]
  - $C_0 = U$
  - $Q_0 = \Psi$
  - Let $z_1 = z_{1:\hat{T}-1}, z_2 = z_{\hat{2}:T}$
  - $A_0 = (z_1^T z_1)^{-1}(z_1^T z_2)$
  - $W_0 = cov(z_2 - A_0 z_1)$

Prediction, Filtering, Smoothing Once the LDS for a particular fold was fit, it was then used to make predictions on held-out data. We review the following terminology:

- Prediction: $E(z_t|n_{1:t-1}, z_0, P_0)$
- Filtering (e.g. what is done during Kalman Filter Decoding): $E(z_t|n_{1:t}, z_0, P_0)$
- Smoothing: $E(z_t|n_{1:T}, z_0, P_0)$

In order to make any of the above predictions, we first must estimate z_0 and P_0 (which is $E(z_0^T z_0)$). In order to do this, we use the whole trial estimate in order to estimate the tenth bin (1 sec after start of trial), $E(z_{10}, P_{10}|n_{1:T}, z_0 = 0, P_0 = I)$. Since we data for our training and testing data starting one second before the actual start of the trial, we can then use $E(z_{10}, P_{10})$ as the initial starting point for our actual trial estimates.

Armed with values of z_0, P_0 for our actual trial starts, we can now predict, filter, and smooth to our hearts' content according to [222].

Dynamics Ratio In order to determine how much the dynamics process actually contributes to estimates of $z_t$compared to the observations process, we compute a "dynamics ratio" for each bin in each trial, as is done [214]. For each bin, the norm of the dynamics update (or time update), the norm of the innovations update (or measurement update), and ratio between the two are computed:

$$dyn = |x_{t|\hat{t}-1} - x_{t-\hat{1}|t-1}|$$

$$innov = |K_t n_t - C x_{t|\hat{t}-1}|$$

$$ratio = \frac{dyn}{dyn + innov}$$

Estimating Subspace Overlap Within an LDS For each LDS fit, there is C matrix fit which defines the mapping from $z_t$to n_t. This C defines a low-dimensional plane (specifically 15 dimensional) within R^n, that is analogous to the shared space from FA. We can see how similar these planes are by defining the LDS-subspace overlap metric:

1. For each LDS, concatenate all $\hat{z}_t$ into an pxTmatrix Z.
2. Compute main shared variance (UU^T_{main}) and the normalized projection matrix (P_{main}) of the following co-variance matrix: CZZ^TC^T (see **??**). Note that in the LDS, there is no requirement that $z_t \sim N(0, I)$, so we compute the empirical co-variance of $z_t$instead of assuming it is the identity (as done in FA).
3. Compute overlap from $LDS_1 \rightarrow LDS_2$:

$$overlap_{LDS_1 \rightarrow LDS_2} = \frac{trace(P_{main,2} UU^T_{main,1} P^T_{main,2})}{trace(UU^T_{main,1})}$$

6.3 Results

Action sequences are comprised of sequence of smaller units, which in motor control have been termed primitives [223], submovements [224], fragments [225], or strokes [226]. While each concept has its own properties, they all describe a complete set of basis movements that are used as building blocks to create the full set of more complex movements. We term these set of basis movements as the action repertoire. Drawing from and combining basis movements from the action repertoire yields action sequences for solving goal-directed tasks in different environments (6.3A).

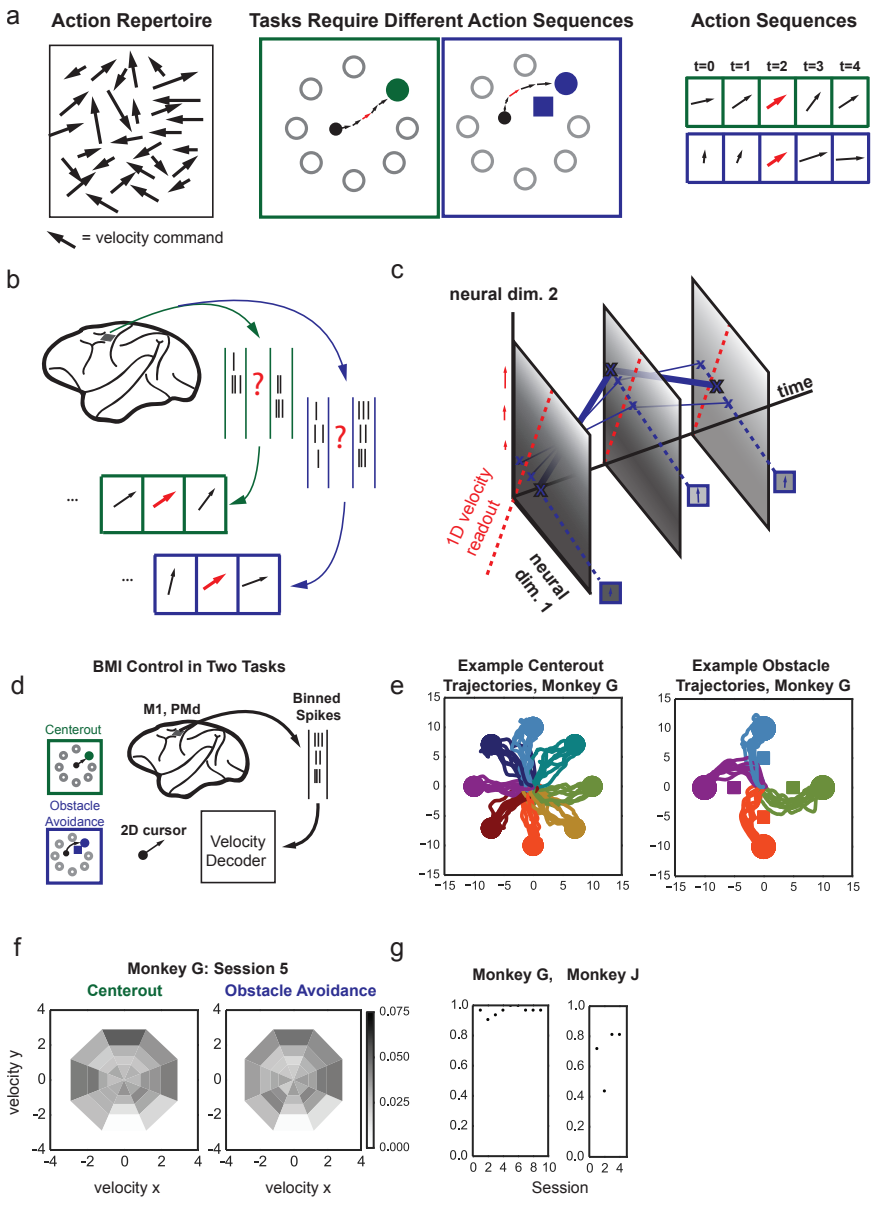

Figure 6.3.0 Action sequences in distinct BMI tasks. a) Left: A schematic illustrating the repertoire of action commands, which in the case of BMI, consist of individual velocity commands of varying directions and amplitudes, Middle: Two distinct action sequences, required due to differences in task demands, contain at least one of the same action (red arrow). Right: Red arrow is embedded in sequences with different histories and futures. b) We ask, how do the neural activity patterns that generate the red arrow in the first, green action sequence differ from the neural activity patterns that generate the red arrow in the second, blue action sequence? c) Illustrated with a one dimensional

velocity readout space (red dotted line) and two neural dimensions (x and y dimensions of each gray plane). The axis along which the three planes lie is a time axis illustrating the evolution of neural activity over time. The blue dotted lines are perpendicular to the red dotted line on all three planes (representing sequential time points). Neural activity patterns that lie on the blue line are decoded identically (arrow in blue box), and thus the three trajectories plotted exhibit differnet neural activity but are decoded identically. d) The brain machine interface loop used in this investigation. Single and multi unit activity is decoded from primary motor cortex (M1) and dorsal premotor cortex (PMd), and binned into 100 ms (Monkey G) or 5 ms (Monkey J) bins. Then a velocity decoder is used to convert binned spikes into predicted cursor velocities. These predicted velocities update the position of the cursor, which is driven either in a straight centerout task, or a curved obstacle avoidance task. e) Example straight, centerout trajectories, or curved obstacle trajectories for Monkey G. f) Distribution of velocity commands over an entire centerout session and entire obstacle session. Gray scale refers to how frequently a particular command was used (darker indicates more frequent use). g) Similarity of velocity command distributions between the centerout and obstacle task sessions within day (9 days for Monkey G, 4 days for Monkey J).

Here, we study the neural generation of the basis movements for BMI control, velocity commands, and specifically ask how neural encoding of a single velocity command changes depending on the sequence in which the command is embedded (6.3B). In both natural motor control and in BMI control, there are many more neurons than number of behavioral dimensions to control. For example, if we were to map the activity of two neurons onto a BMI with a 1 dimensional velocity vector (6.3C, dashed red line), at each discrete time point (pictured as grey planes), there are many neural activity patterns that can yield the same behavioral output. In 6.3C, three different neural activity trajectories through time are pictured that all yield the same behavioral output (pictured as boxed black velocity arrows). Since variations orthogonal to the velocity readout line do not affect behavior, subjects have flexibility in how they generate desired movements.

Here, we designed two BMI tasks – a centerout task and an obstacle avoidance task. In the centerout task, subjects control a 2D cursor to move from the center of the workspace to a cued peripheral target. When reaching the target, they are required to hold within the target for 200 ms. In the obstacle avoidance task, subjects must perform the same centerout movements but now must avoid a square obstacle placed exactly between the straight line connecting the center and the peripheral target (Monkey G), or must perform target to target trajectories with an ellipse obstacle in the path (Monkey J). We analyzed 9 days of data from Monkey G and 4 days of data from Monkey J where on each day, monkeys performed both the centerout and obstacle tasks with the same BMI decoder (6.3D). Only successful trials were analyzed.

In both tasks, subjects must draw upon their action repertoire to generate trajectories that bring their cursor to the cued peripheral target. In both the Kalman Filter decoder

(Monkey G) and the Point Process Filter decoder (Monkey J), the cursor position and velocity are modeled as states with a dynamics process. At each time bin, an estimate of cursor state is computed from the dynamics process propagating the prior state estimate, and then a refined estimate of cursor state computed using the observed spike counts (see 6.2.3 for details). Both processes can be summarized as:

$$x_t = Fx_{t-1}\hat{|}n_{t-1} + G(n_t)$$

where x_t is cursor position and n_t is binned neural spike counts. In the case of the Kalman Filter, $G(n_t) = Kn_t$ where K is the Kalman Gain. In the case of the Point Process Filter, $G(n_t) = \hat{P_{t|t}}C^T n_t$ where C is a matrix $\in R^{nx5}$ with i row of C equal to $\tilde{\alpha_i}$ (see 6.2.3).

In the following analyses, we define the basis movements of BMI control as the contribution of the observations to the update of cursor state, or $G(n_t)$. We choose not to define the actual cursor velocity as the basis movement because the subject does not control the dynamics process, and so the only way for them to update the state of the cursor is through generating neural observations that will direct the cursor. There is evidence that subjects develop an internal model of the cursor dynamics during BMI control [80], making it likely that their neural commands are issued to the BMI with full knowledge of how it will update control of the cursor. Since we are interested in how subjects generate control, we define the basis movement for a particular bin as the velocity command $G(n_t)$.

As evidenced from example cursor trajectories in 6.3E, subjects succeed at generating successful trajectories in both tasks. Cursor trajectories in the centerout task are straight whereas ones in the obstacle avoidance task are more curved. In subsequent analyses, we pool over all targets and trials during performance of a given task to ask questions about subjects select neural commands to generate desired action sequences. Thus, we first ask how similar the distribution of velocity commands is across the two tasks. Each velocity command is binned into one of eight angular bins and one of four magnitude bins as shown in 6.3F for an example session from Monkey G. For each animal and each session, the similarity of distribution of commands (see 6.2.4.1) is close to 1, indicating very similar distributions of commands (6.3G). Thus, subjects draw from similar distributions of velocity commands across tasks.

6.3.1 Manifold Similarity

6.3.1.1 Task-Specific Comparison

Given that subjects draw from very similar distributions of velocity commands, we next ask whether neural activity used to control the BMI in different tasks lies within the same manifold parameterized by a low dimensional set of latent variables. Recent work demonstrates that subjects' selections of neural activity to control a 2D cursor BMI are well described by combinations of low-dimensional latent variable activations [91]. It is suggested that creating new cursor trajectories would involve using the original activation

patterns but strung together in a different order [227, 91]. If subjects are overall generating the same velocity commands in the centerout and obstacle tasks (6.3G), and are generating each of these commands with the same latent variables activations, then we expect the manifold fit from neural activity during the centerout task versus the obstacle task to be nearly identical (6.3.1.1A). On the other hand, if performing the different cursor sequences required in the different tasks involves producing velocity commands with distinct latent variable activations, the manifold fit from different tasks may differ (6.3.1.1B).

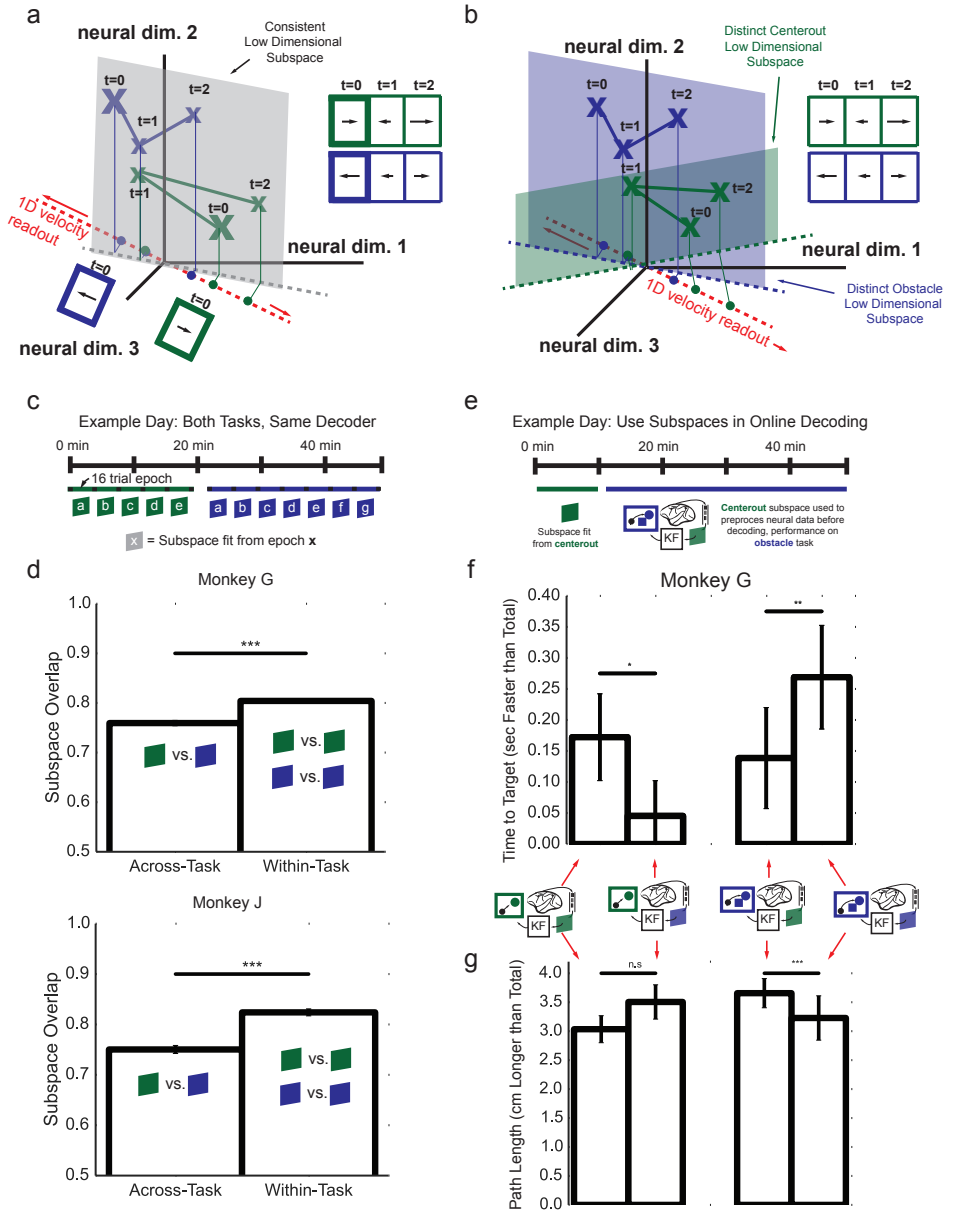

Figure 6.3.1.1 Manifold Similarity. a) Example of where the neural activity commands for both centerout (green) and obstacle (blue) action sequences are drawn from the same manifold. The red dotted line again represents a one dimensional readout space, and the boxes with arrows illustrate the decoded action sequence. b) Example of where neural activity commands for the same centerout and obstacle action sequences as shown in (a) are now drawn from different task specific manifolds. c) Illustration of typical session where first the centerout task is performed for about 20 min (green line), followed by performance of the obstacle task for another 20 minutes (blue line). On some sessions this order is reversed. Within these blocks, consecutive 16 trial epochs are defined and subspaces are fit from the data in the epoch. These subspaces are illustrated by the green and blue parallelograms. d) Comparison of across task vs. within task subspace overlaps for Monkey G (top) and Monkey J (bottom). e) Example of typical session in which and estimated subspace is used in online decoding. First, a normal 64 trial epoch is performed (illustrated by green line), in this case on the centerout task. A subspace is fit from the data in that block. Then the subspace is used in online decoding during performance in a second block, in this case, on the obstacle task. The BMI loop shows the subspace (green) and task (blue) used in online testing of subspaces. f) Normalized time to target and g) normalized path length for different combinations of tasks used to fit subspaces and tasks performed.

In order to assess similarity of the task-specific manifolds, trials from each task on each day were split into 16-trial epochs (6.3.1.1C). Spike counts from neurons that were used to control the BMI were binned at 100 ms. Trials to all targets were concatenated to yield an nxT array, where n is the number of neurons, and T is the total number of bins (sum of bins over all 16 trials). Next, Factor Analysis (FA) was applied to estimate a low-dimensional manifold for each 16-trial epoch [201]. FA decomposes the activity of each 100ms bin of neural spike counts (vector) into a summation of i) a mean rate $\mu \in R^n$, ii) private signal $\psi_t \sim N(0, \Psi)$ where Ψ has diagonal covariance, and 3) shared signals Uz_t where z_t is a low-dimensional latent variable $\in R^k, k < n, U \in R^{nxk}$. The low-dimensional manifold is defined as the column space of U. Once an FA model was fit for each 16-trial epoch (see Appendix A.1 for methods), we compared within task versus across task manifolds. Each manifold was compared to every other manifold within its own task on that same day, and every other manifold in the different task on that same day. Manifolds were compared using a subspace overlap metric, comprising of an estimate of how much variance from Manifold A is captured in Manifold B, and vis versa (see 6.2.4.2). Both estimates contributed to the overall estimate of within-task versus across-task subspace overlap. In both monkeys, within task overlap was significantly higher than across task overlap (Monkey G, Kruskal Wallis test statistic = 56.10, p = 6.89e-14, n_within = 2194, n_across = 1046, Monkey K Kruskal Wallis test statistic = 59.07, p = 1.52e-14, n_within = 432, n_across = 452). Thus, surprisingly, tasks occupied significantly different subspaces.

6.3.1.2 Online Test of Task-Specific Comparison

Though the centerout and obstacle subspaces are significantly different, we sought to assess exactly how different a subspace overlap of 0.05 – 0.1 is in more concrete metrics by using the task-specific manifolds as a neural pre-processing step in the online BMI loop. This approach estimates n_t^{shared}, or the component of each neural spike count vector that is due to the low-dimensional manifold, Uz_t, and mean, μ, (removes the component due to ψ_t). Since some neurons in are not well explained by the low-dimensional factors, their had a much lower variance than the original n_t. This resulting in much slower movement of the cursor. In order to remedy this, a scaling factor for each neuron i was derived:

$$\beta^i = \sqrt{\frac{\sigma^2_{i,full}}{\sigma^2_{i,shar}}}$$

where $\sigma^2_{i,full}$ is the variance of n_t^i during the manifold-fitting block and $\sigma^i_{i,shar}$ is the variance of $U(i,:)z_t$ during the manifold-fitting block. Here we call β the vector of all $n\beta^i$s. Note that β defined here is the same as β_{shar} from chapter 5.

Thus, the final estimate of manifold-specific activity that was used as an input into the decoder (instead of n_t directly), was:

$$\hat{n}_{t,shared} = \beta(UU^{'}(UU^{'} + \Psi)^{-1}(n_t - \mu)) + \mu$$

Then $\hat{n}_{t,shared}$ is used as the observation in the online BMI instead of n_t. If the centerout and obstacle task manifolds are very different in axes that have non-zero projections onto the decoder axis (e.g. 6.3.1.1B), then $\hat{n}_{t,shared}$ will be very different depending on whether the manifold if fit on the CO or OBS task, resulting in distinct cursor velocities. Thus, if the manifolds are different, then there should be a proportional difference in task performance when using task A manifold versus task B manifold.

On each day (Monkey G only, data from 9 days not exactly overlapping days used in main comparison), a decoder was fit and fixed. Subjects proceeded to perform either the centerout or obstacle task with this decoder, constituting a manifold-fitting block. Their neural data during this session was used to fit a task-specific manifold using FA (6.3.1.1E). Finally, this manifold was used to estimate $\hat{n}_{t,shared}$ which entered the BMI decoder.

We assess performance for each combination of task and manifold (centerout manifold-centerout task, obstacle manifold- centerout task, centerout manifold-obstacle task, obstacle manifold-obstacle task, indicated by color coded BMI loops between 6.3.1.1F,G). Since decoders had day-to-day variability in their cursor speeds, we normalize performance metrics in each session by the mean of the metric in the manifold-fitting block performed with the same decoder. 6.3.1.1F and 6.3.1.1G show the normalized time to target and path length respectively. When subjects are using a BMI loop with a mismatch in the task performed and the task used to train the manifold, time to target performance suffers (Monkey G, Kruskal Wallis test, normalized time to target: CO task-CO manifold vs.

CO task-OBS manifold, statistic = 6.09, p = 0.0136, n co task-co manifold = 443, n co task-obs manifold = 548, OBS task-CO manifold vs. OBS task-OBS manifold, statistic = 6.76, p = 0.00933, n obs task-co manifold = 368, n obs task-obs manifold = 452) and path length of cursor trajectories becomes longer for the obstacle task, but is unaffected in the centerout task(Monkey G, Kruskal Wallis test, normalized path length: CO task-CO manifold vs. CO task-OBS manifold, statistic = .021, p = 0.885, n co task-co manifold = 443, n co task-obs manifold = 548, OBS task-CO manifold vs. OBS task-OBS manifold, statistic = 18.13, p = 2.06e-05, n obs task-co manifold = 368, n obs task-obs manifold = 452).

Thus, the differences in the centerout and obstacle manifolds described in 6.3.1.1D are in a task-relevant dimension, and result in worse performance than when the task used to train the manifold and the task performed match. These online results substantiate the findings that across task manifolds are different, and are different in a task-relevant dimension.

6.3.2 Tuning Differences

One hypothesis of how subjects generate sequences of movement is that they string together basis movements from their action repertoire, where each basis movement is consistently generated by a specific neural activity pattern. In BMI control, one theory of how subjects learn to generate skillful control is that they learn how the decoder maps the activity of each neuron into a velocity command, and learn to produce neural activity patterns that will push the cursor in the desired direction [71, 73]. Perhaps these consistent activity patterns for basis movements can be characterized by activations of the low-dimensional factors spanning the intrinsic manifold [227]. However, given the finding in 6.3.1.1 of distinct across-task low-dimensional spaces, we re-assess the initial assumption that subjects use a consistent neural activity pattern to generate specific basis movements. One possibility is that subjects can have different across-task subspaces, but on average use the same neural command to produce basis movements (6.3.2A). Pictured here is a 3 dimensional neural space, with a 1 dimensional readout space (red dotted line). The green and blue planes represent the intrinsic manifolds for the centerout and obstacle tasks respectively. The intersection of these planes with the red plane represents the mean neural activity pattern for a given velocity command (the planes are assumed to be symmetrically distributed in the axis normal to the red plane for ease of visualization). Here, the green and blue planes are different, but the mean neural activity pattern for a given velocity readout command is the same. Alternatively, it is possible that the different blue and green subspaces also have different mean neural activity patterns for the same velocity command (6.3.2B).

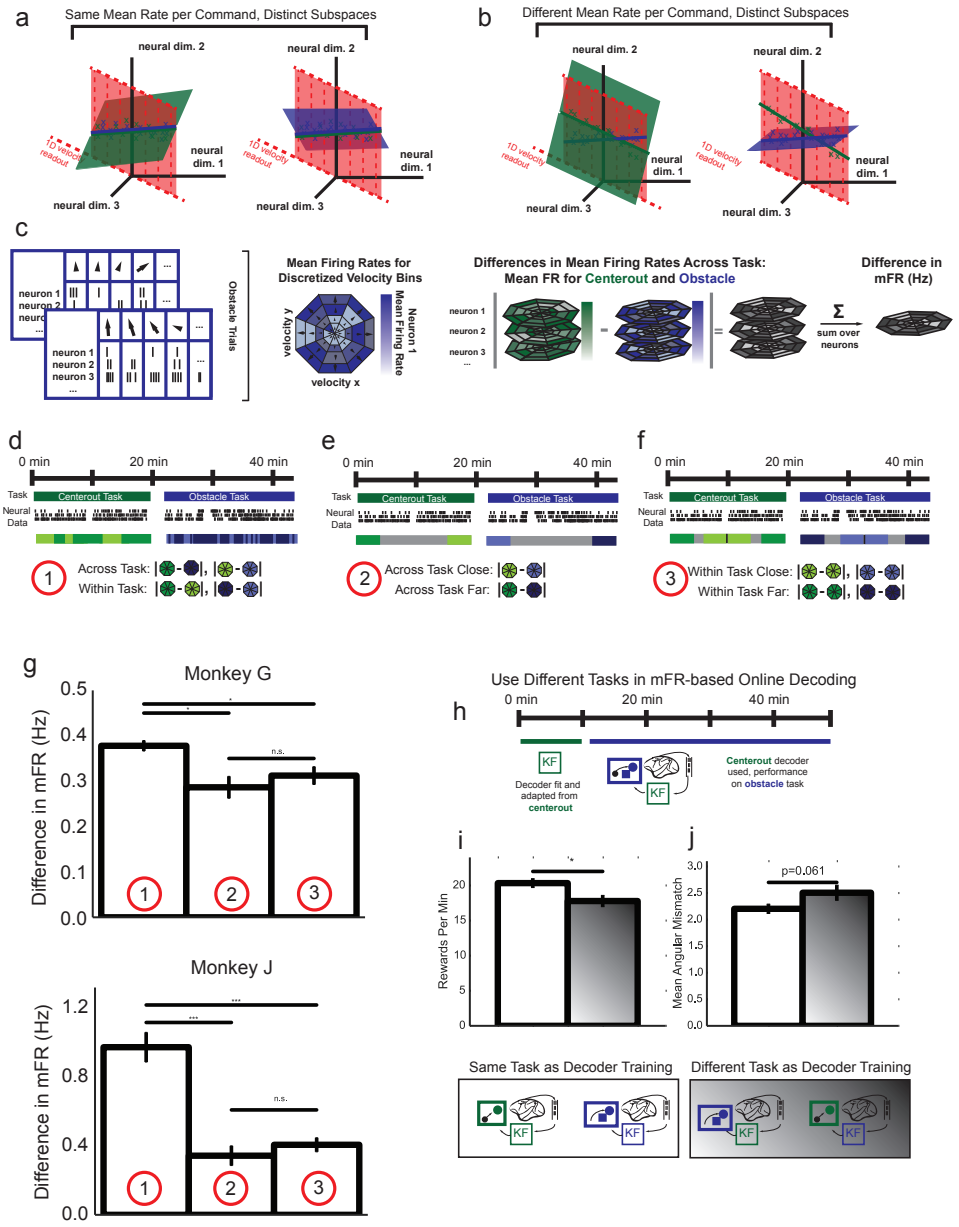

Figure 6.3.2 Tuning Differences. a) Illustration of how mean firing rate for a given command can be consistent across task, but the subspaces can be different. Here, the red dotted line is again a 1D readout space. The red plane illustrates the intersection between the red dotted line and the green and blue planes. In both illustrations, the blue and green planes have the same mean firing rate for a given command, illustrated by the solid blue and green lines. b) Illustration of how mean firing rate for a given command can be different with different subspaces, same format as (a). c) Illustration of analysis flow to assess tuning differences. d) Illustration of main within task vs. across task comparison.

Two subsets of CO (light and dark green) and OBS (light and dark blue) data that are compared. e) Control for drift in firing rates over time, comparison of across task data that is close in time versus across task data that is far in time. f) Control for drift in firing rates over time, comparison of within task data that is close versus within task data that is far. g) Differnences in mean firing rates for the comparisons outlined in d, e, f. h) Online test to illustrate how differences in tuning impact online decoding. In the timeline, the green bar represents seeding and adapting a decoder during the centerout task. The blue line represents assessing the decoder performance during the obstacle task. i) Comparison of rewards per minute between blocks where decoder fitting task matches the performed task (white bar) and blocks where decoder fitting task does not match the performed task (gray bar). j) Same as i) but comparing angular mismatch between decoder's assumed preferred direction for a given output unit and inferred preferred direction from task for a given output unit.

We first use a non-parametric method to assess consistency of neural activity patterns for a given velocity command across tasks. We build up an estimate of the mean firing rate for each velocity command for each task as follows (6.3.2C). For each 100 ms bin with a successful trial for a particular task, we extract the two-dimensional velocity command ($v_t \in R^2$). Each velocity command is then discretized into one of 32 sections, as shown in 6.3.2C. The sections split the radial space into 8 divisions and magnitude space into 4 divisions for a total of 32 sections. After all velocity commands have been assigned to their sections, the corresponding neural activity patterns that were used to generate the velocity commands are collected and averaged to yield a mean firing rate for each neuron, for each velocity atom section (blue octagon in 6.3.2C). Then, we compare across task mean firing rates by taking absolute differences in mean firing rate for each neuron in each section, yielding many gray octagons (one per output unit in the decoder). Finally, we sum over all neurons to yield a final octagon describing the absolute vector difference between mean firing rates for each command across tasks. Sections with fewer than 15 instances in a task were excluded from analysis to prevent low counts from influencing results.

A number of comparisons were made to assess across-task versus within task differences in mean firing rate. 6.3.2D describes the first comparison. For each task, each 100 ms bin was randomly assigned to subset 1 or subset 2 (light vs. dark colors). This division allows us to compare across task differences (dark green vs. dark blue and light green vs. light blue, were randomly chosen instead of dark green vs. light blue and dark blue vs. light green) compared to within task differences (dark green vs. light green, dark blue vs. light blue). If there are greater across task differences than within task differences, we expect that the differences in across task vs. within task will be much greater than zero (bar labeled with '1' in 6.3.2G).

To assess whether differences in across-task vs. within-task mean firing rate were due to a slow drift in firing rates over the day, or due to consistent adaptation over the

day, we also compared near vs. close sections both across task (6.3.2E), and within task (6.3.2F). If slow drift in mean firing account for the across-task vs. within-task mean differences observed in bar 1, we would expect that across-task-close would show a much smaller difference than across-task-far (6.3.2E), resulting in a large difference in 6.3.2G, bar 2. Similarly, we would expect that within-task-close differences would be much smaller within-task-far (6.3.2F), again resulting in a large difference in 6.3.2G, bar 3. However, in both animals, bar 1 is significantly larger than bars 2, 3 (Monkey G, Kruskal Wallis test for differences across three bars statistic = 16.72, p = 2.34e-04, bar 1 n = 540, bar 2 n = 65, bar 3 n = 159, Mann Whitney Test for Multiple Comparisons with Bonferroni Correction, bar 1 vs. bar 2 p 5.69e-3, bar 1 vs. bar 3 p = 6.52e-04, Monkey J, Kruskal Wallis test for differences across three bars statistic = 33.04, p = 6.69e-08, bar 1 n = 132, bar 2 n = 23, bar 3 n = 64, Mann Whitney Test for Multiple Comparisons with Bonferroni Correction, bar 1 vs. bar 2 p = 7.37e-5, bar 1 vs. bar 3 p = 8.74e-7). This finding indicates that there are differences in mean firing rate for a given command across-task, that are not present within-task, and that are not accounted for by drifts or adaptation over the course of the day. Thus, the cartoon illustrated in 6.3.2B is an accurate depiction of our current knowledge of how the intrinsic manifold and mean neural activity for a given velocity command differ across tasks.

Similar to our previous manifold analysis in 6.3.1.1, we were curious how big of a mean firing rate difference of 0.4 Hz (Monkey G) or 0.9 Hz (Monkey J) is in more concrete metrics. Also similarly to 6.3.1.1, we were able to test how much these differences matter in an online experiment. Since the Kalman filter decoder used with Monkey G assumes that neurons have an underlying linear tuning model, if a decoder is trained on the centerout task then the fit linear tuning model will be different than if the decoder was trained on the obstacle task. If a subject is using a decoder with the incorrect underlying linear tuning model, then the velocity command computed from the decoder will not match the subjects' intended velocity command possibly resulting in errors [80]. This experiment is similar to assessing how inaccurate a human is when driving a car that always goes in slightly the wrong direction than intended. If the degree of wrongness is high perhaps in both the left and right direction, driving performance will be heavily effected, whereas if it is only slight, it may not be noticed in short bouts.

To conduct this experiment, we seed a decoder using visual feedback of a specific task, and perform decoder adaptation on this seed as the subject performs that same task. Next, we use this decoder on both the centerout task and the obstacle task (6.3.2H). Finally, we assess performance metrics when Monkey G used a decoder trained on the same task that he performing, versus when using a decoder trained on the opposite task. 6.3.2I, left shows that Monkey G has significantly higher rewards per minute when the decoder he is using is trained on the same task as the one being performed (white bar) versus when the decoder is trained on the opposite task as the one being performed (grey bar) (Monkey G, Kruskal Wallis test, W = 5.90, p = 0.01516, n_same = 92, n_diff = 60). Further, when comparing the angle difference (angle mismatch) between an estimate of each neurons' tuning during online task performance and in the decoder, same-task-

decoder-training mismatch was smaller than different-task-decoder-training (Monkey G, Kruskal Wallis test, W = 3.52, p = 0.061). Thus, Monkey G exhibits a decrement in task performance when using decoders trained on the incorrect task possibly because of an angular mismatch in the decoders' assumed preferred direction of neurons and his own model of that neurons' preferred direction.

6.3.3 Expanded Tuning Models

Our previous analysis demonstrated that subjects exhibit different mean firing rates for producing the same velocity command, and that such differences could have a detrimental role when training a decoder on one task and using it on the next. In this cursor task, the decoder uses neural firing rates, usually 20-50 dimensional, at a given time bin to produce a two-dimensional velocity command. With such a large redundancy, subjects may internalize a more complex model of how the cursor responds to neural activity instead of the simplest velocity tuning curve model that the decoder assumes. Perhaps the subject represents lagged or leading velocities and positions, generating different neural activity patterns for the same instantaneous velocity command that occurs in a different sequence (e.g. as in 6.3A, center). We first investigate whether greater across-task mean firing rate differences for a specific velocity section are predicted by greater differences in what trajectory the velocity command is embedded within. Specifically, for a given velocity section (6.3.3A, red arrow), we scan through each task and mark the occurrence of when this velocity occurs. Then, the velocity commands 400 ms before and 400 ms after our command of interest are extracted to yield a 'command-aligned action segment' (6.3.3A). Action segments are collected for each task and averaged across segments to yield a task-specific mean command-aligned action segment which we term the command-aligned behavioral PSTH. An example of a command-aligned behavioral PSTH is shown for Monkey G (6.3.3B). Finally, we ask whether commands with greater across-task differences in their x and y velocity behavioral PSTHs (from -400 ms to 400 ms) exhibit greater differences in the mean firing rates for the specific command (at 0 ms). 6.3.3C shows the correlation between across-task differences in behavioral PSTHs (y axis) and across task difference in mean neural firing rate (x axis), where each point corresponds to an action segment on a specific day. Both animals exhibit significant correlations between differences in their mean neural firing rates at 0ms and the differences in their behavioral PSTHs (Two-sided t-test for significant slope in linear regression, Monkey G, p = 2.39e-21, n=283, Monkey J, p = 5.27e-19, n = 84). Thus, the sequence in which the velocity command is embedded influences the neural activity patterns produced. This finding supports the hypothesis that subjects have a representation of the cursor that extends beyond the instantaneous command they are producing, not unlike findings that representation of movement are better characterized by temporal segments than instantaneous movement parameters.

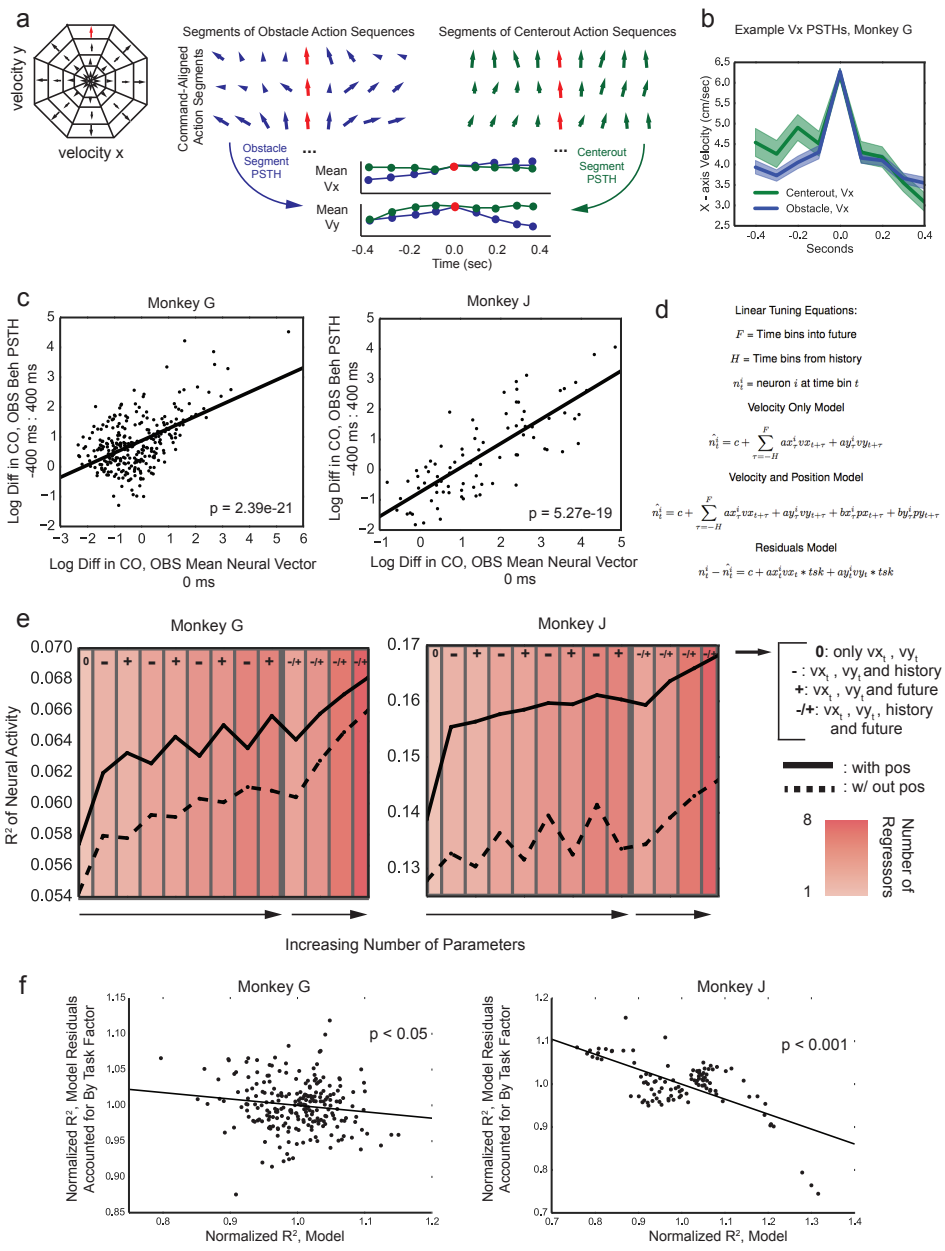

Figure 6.3.3 Expanded Tuning Models. a) Schematic of non parametric method used to assess the mean action sequence in which velocity commands are embedded for CO and OBS tasks. Briefly, for a particular task in a session, all examples of a particular velocity command are aggregated, along with the 400 ms before and after the command. These 900 ms are averaged, yielding a mean action sequence for the task. b) Example x velocity action sequence. c) Correlation between log of the difference in CO and OBS mean neural activity patterns (at 0ms) for a given velocity command, and log of the difference in the mean CO and OBS action sequence from a history of 400ms to future of 400ms

for the same velocity command. Each point is a velocity command from a particular session. d) Expanded parametric linear tuning models where neural activity is modeled as a linear readout of velocity or velocity and position. The residuals model explains the residuals of the original model as a linear combination of task specific velocity commands. e) Illustration of velocity tuning models (dotted line) and position and velocity tuning models (solid line) with different history or future lags. After the solid black line, both history and future lags are included. f) Correlation between normalized of each of the models and normalized of the residual models. As the models explain more variance, the apparent task specific velocity commands is reduced.

We next investigate which cursor parameters (position, velocity) and time lags account most for neural activity differences across tasks by using linear models. First, the neural activity at each time bin is modeled as a linear combination of cursor positions and neural push velocities at various lags, as formalized in 6.3.3D. These models use data from both tasks. The amount of neural variance that each model can explain is plotted in 6.3.3E. At the top of the figure is a symbol for each model illustrating if it includes lagged cursor movement terms (-), leading terms (+), neither (0), or both (+/-). The solid line illustrates models that have position and velocity terms, and the dashed line illustrates models that only have velocity terms. The color intensity of each column refers to how many regressors are used in the velocity version of each model (corresponding to the inclusion of terms from increasingly far from 0 lags). The amount of neural variance accounted for increases with the addition of position terms (solid line is greater than dashed line), and with the inclusion of more regressors.

Since we are specifically interested in accounting for neural firing pattern differences in the centerout and obstacle tasks that were reported in 6.3.3C, for each linear model in 6.3.3E we assess how much of the remaining, unexplained neural variance can be accounted for by a task-specific neural firing pattern (residual model, 6.3.3D). If a model does not account well for the differences in firing rates across tasks, then the residual model which has a binary task factor (tsk = 0 for the centerout task, 1 for the obstacle task), ought to account for a large amount of neural variance. On the other hand, if a model does account for the across-task differences in neural firing rates, then the residual model will account for little neural variance.

In 6.3.3F, we plot the normalized R2 for each model from 6.3.3E versus the normalized R2 of the residual model for each model. In this plot, one point is one model fit on one day of data, and the normalized R2 value is the R2 for that point divided by the mean R2 across all models for that particular day. The significant negative correlation in both animals demonstrates that as the models in 6.3.3E begin to include more cursor movement parameters, the amount of neural variance accounted for increases, and the across-task differences in firing rates are reduced (Two-sided t-test test for non-zero slope, Monkey G $p = 0.0317$, $n = 234$, Monkey J $p = 2.69e\text{-}14$, $n = 104$). In other words, the difference in the centerout and obstacle tasks is merely that the same actions are strung together in

different sequences, so modeling neural variance as a function of an increasing number of velocity lags and leads essentially models neural variance as a function of task.

At this point, we have the notion that subjects may have a complex internal model of how the neural-to-cursor mapping works. As we add more cursor kinematic parameters to our model the amount of neural variance explained improves. However, it must be challenging for subjects to keep track of so many previous and anticipated future cursor kinematics. What if subjects have a very simple model of how the cursor works, but there are guidelines on how their neural activity can evolve from time point to time point? Thus, subjects may maintain a simple model of cursor kinematics, but due to the temporal guidelines on their neural population activity, in different situations they generate activity differently. We explore this possibility next.

6.3.4 Constraints on Neural Activity

There are numerous models that could describe how neural activity evolves from one time point to the next. Possibly, neural activity must evolve smoothly, making a low pass filter applied to individual neurons an accurate description. Maybe neural activity only changes at specific frequencies, so employing a bandpass filter could describe those changes. Maybe there is additive noise in neural activity, preventing the above models from achieving accurate descriptions of neural activity, so employing a Weiner filter would best describe the evolution of activity. Maybe neural activity changes only occur in a low-dimensional space, so employing dimensionality reduction methods prior to modeling temporal dynamics is advisable. We do not assess the validity of every possible model in the following work. Instead, we draw on previous literature and our findings thus far to narrow the set of models we consider. First, we and others find that high dimensional neural data recorded from motor cortex during performance of BMI tasks is well described by low dimensional neural modes [90, 228, 91]. Thus, we choose to use models that reduce the high dimensionality of neural activity to a lower dimensional population state. Prior work has used a linear dynamical system (LDS) formulation:

$$z_t = Az_{t-1} + w_t, w_t \sim N(0, W)$$

$$n_t = Cz_t + q_t, q_t \sim N(0, Q)$$

where z_t is low-dimensional and represents the 'population state', and n_t is a vector of neural data as usual. Note, that there is no cursor kinematics information in this model, unlike the linear tuning model. When this model is trained on arm movements, and used online to make predictions (by linearly mapping z_t to a two-dimensional cursor state), it exhibited superior decoding compared to Weiner filters, Kinematic Kalman Filters, and low pass filters [214]. Thus, given its power in explaining online BMI neural data and our inclination that there may be temporal dynamics in our data we also consider how well the same LDS with 15 states fits our data.

First, how well can the LDS model account for the neural data compared to the full linear tuning model (model from 6.3.3E with largest number of regressors, including position)? For each day and each task, a 80% of data is used to estimate a linear tuning model and a LDS model. Performance on 20% of the held-out task data is then assessed. Since the linear tuning model is acausal, incorporating cursor kinematic information from future points, we compared the performance of the linear tuning model to the performance of smoothed estimates of observations from the LDS. For both animals, the LDS significantly outperforms the linear tuning model (6.1A Kruskal Wallis test , Monkey G: statistic = 25.95, p = 3.51e-07, n = 18 per model, Monkey J: statistic = 11.29, p = 7.78e-4, n = 8 per model). Thus, modeling neural activity as a linear function of a previous low-dimensional neural activity state without any cursor kinematics movements more successfully predicts neural activity than using a linear tuning model with tens of terms devoted to capturing cursor movements.

We next consider how well the LDS fits neural activity compared to FA, a model that reduce dimensionality but does not model temporal dynamics and a Gaussian low pass filter, a model that incorporates temporal dynamics but not dimensionality reduction. For a fair comparison, we only use the filtered state estimate of the LDS to estimate the filtered observation estimate, instead of the smoothed state estimate we used in 6.1A. Since there are many neural signals that do not contribute the decoder substantially and are well-predicted by a low-pass filter, we also ask how much of the task-specific neural activity the LDS captures (6.1B). We define task-specific neural activity as predicted neural activity passed through the decoder to yield predicted cursor velocity commands. In both animal, the LDS explains significantly more variance in cursor velocity commands than FA or the low-pass filter (Monkey G, Kruskal Wallis test for differences across three models = 27.90, p = 8.72e-07, n = 18 for all models, Mann Whitney Test for Multiple Comparisons with Bonferroni Correction, LDS vs. FA p = 1.57e-05, LDS vs. LPF p = 2.25 e-06, Monkey J, Kruskal Wallis test for differences across three models = 15.74, p = 3.82e-04, n = 8 for all models, Mann Whitney Test for Multiple Comparisons with Bonferroni Correction, LDS vs. FA p = 7.41e-03 , LDS vs. LPF p = 9.39e-04).

In an extreme case where there are little temporal dynamics (A matrix is zero), the LDS will approximate an FA model. We confirm how much the dynamics are being used by a metric termed the 'dynamics ratio' [214], or a ratio of the magnitude of the time-update step to the sum of the magnitudes of the time-updates step and the measurement-update step. Previous work fitting an LDS to arm-movement data has reported dynamics ratio values of 0.35-0.50 [214]. For both animals, the LDS dynamics ratio is 0.45-0.6, indicating that in the LDS, linear dynamics contributed substantially to the state estimate. Thus, the LDS is a model that captures substantial neural variance compared to the tuning curve model, an FA model, and in Monkey J, a low pass filter with a Gaussian window of 300 ms 6.1,6.2. In Monkey G, we suspect that due to a number of poorly isolated units included in the decoder that do not contribute to cursor movement substantially, a low pass filter was a good model due to the activity of these units. If we instead assess decoded velocity commands by the LDS, FA, and low pass filter, which isolates the components

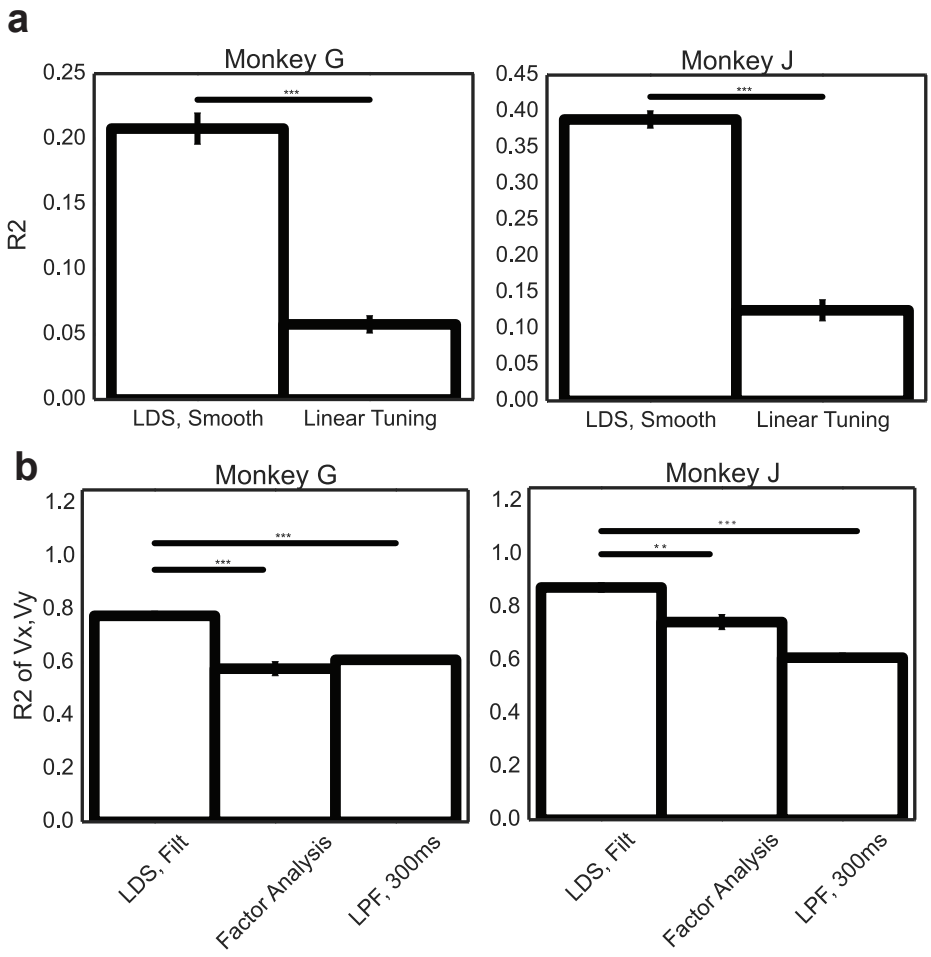

Figure 6.1: Linear Dynamics explain most variance. a) Comparison between the linear dynamical system model and the full linear tuning model, both of which use observation of past and future neural (LDS) or kinematic (linear tuning) activity to estimate the neural activity at the current time point. b) Filtered estimate of decoded neural activity from LDS compared to the FA model, and a simple 300 ms gaussian kernal low pass filter. Neural activity is estimated for each of these models and then the neural push from these estimates is compared to the true neural push.

of neural activity that contribute to movement, we find that the LDS outperforms the others 6.1. Finally, in the LDS, the linear dynamics are substantially used.

6.3.5 Generalized Linear Dynamics

In this work, we have hypothesized that subjects are using a consistent model to control the BMI cursor in both tasks, and we have sought to uncover a framework that captures how neural activity is generated for one task that also generalizes to the other task. Thus far, we report that subjects do not use a consistent intrinsic manifold to generate activity in different tasks, nor do they generate the same velocity commands consistently across tasks. Given how well an LDS fits the neural data compared to other models, how general are the linear dynamics across tasks? For each model trained on a subset of one task's data, its performance on held-out data is compared to performance on data from the non-trained task shown in the bottom row of 6.2 (Kruskal Wallis test, Monkey G, n = 18 per model, test statistic = 7.752, p = 5.37e-03, Monkey J, n = 8 per model, test statistic = 0.172) . For both the FA model and LDS model, the ability of the model to predict within-task held out neural data (left bar) is compared to predictions on across task neural data (right bar). We also compare how much of the task-relevant neural activity is captured in 6.3. In both animals, the difference between within task and across task estimates of velocity commands predicted by the LDS is significantly smaller than for FA (Kruskal Wallis test, Monkey G, n = 18 per model, test statistic = 23.43, p= 1.29e-06, Monkey J, n = 8 per model, test statistic = 11.29, p = 7.78e-4). Indeed the linear dynamics and observation model from one task generalize better to the other.

Within the population state-space of the LDS, do the different tasks exhibit distinct subspaces that they occupy? If they do, then perhaps our findings from 6.3.1.1 are consistent with temporal dynamics influencing which how neural activity patterns evolve. Since different action sequences require distinct ordering of velocity commands, and the linear dynamics influence how neural activity patterns can evolve through the population state space, perhaps the different action sequences are initialized at distinct points in the state space that can take advantage of the different linear dynamics to guide activity patterns through the space in a desired sequence. In this framework, subspace overlap difference across are to be expected. We performed a similar analysis as in 6.3.1.1D, but now instead of estimating the subspace from an FA model, we use the LDS model to compute the subspace for a 16-trial epoch (see 6.2.4.5). Indeed, when comparing LDS subspace overlaps, they are lower across-task versus within task (6.4B).

We confirm if our intuition that the subspace overlap differences are due to the neural activity patterns following consistent linear dynamics through different population state spaces. Since we hypothesize that neural activity patterns are following linear dynamics to generate different action sequences, a selected velocity command from a given action sequence ought to be produced with neural activity that has a high dynamics ratio within the sequence it falls in. On the other hand, if we interchanged that velocity command and corresponding neural activity pattern with an identical velocity command (vel_new)

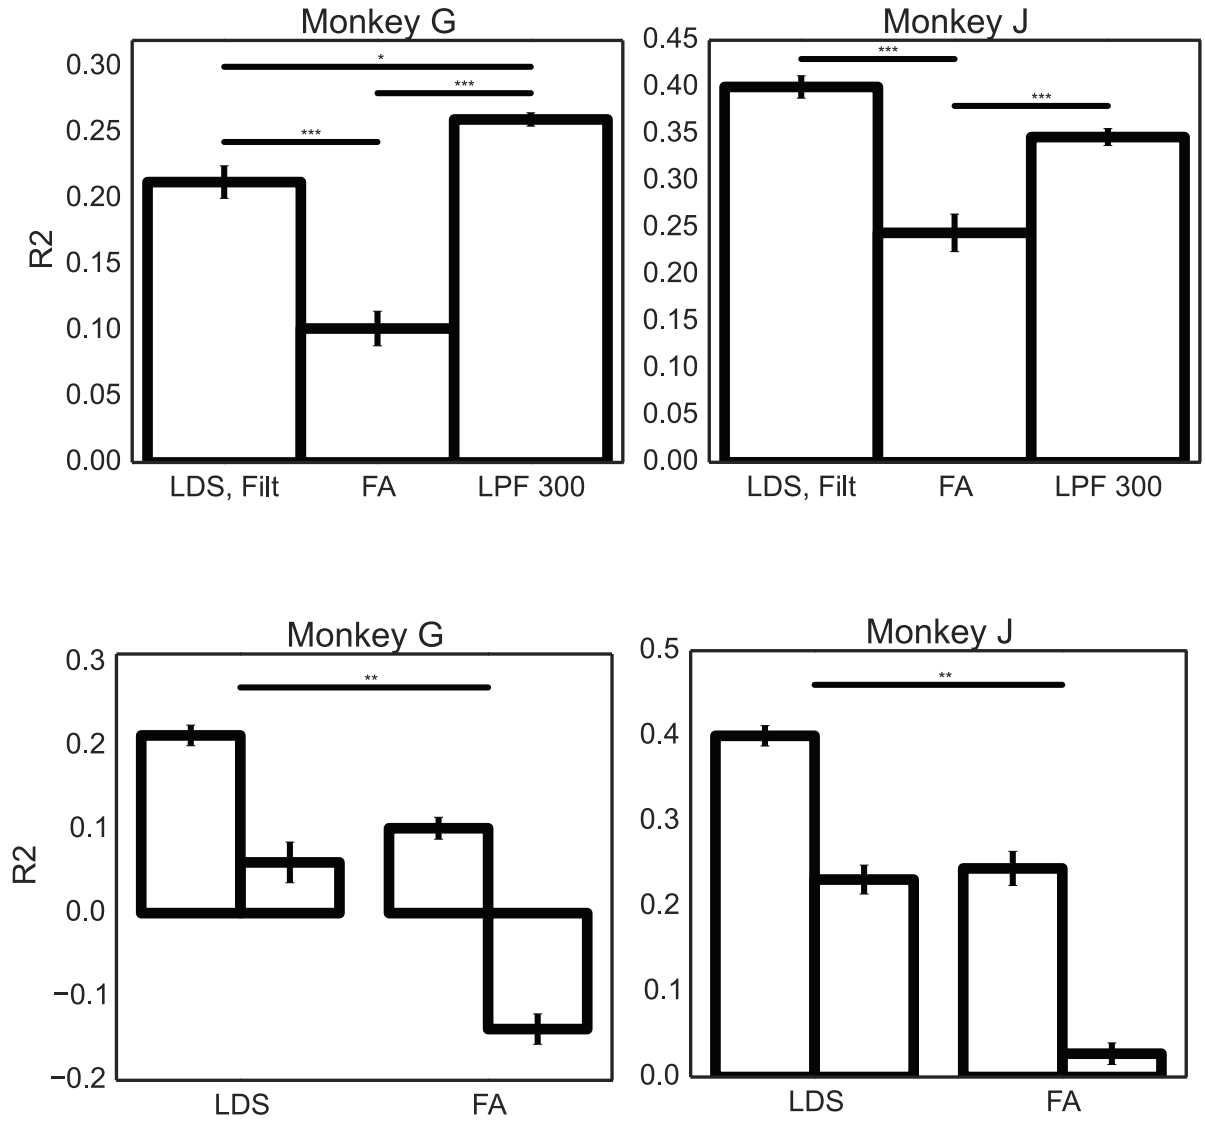

Figure 6.2: a) Comparisons of the neural activity estimates from the filtered LDS, FA, and low pass filter. b) Comparisons of neural activity predictions on within-task held-out data (left bar) or across-task data (right bar) for the LDS and FA models. The LDS generalizes much better across-task than the FA model does.

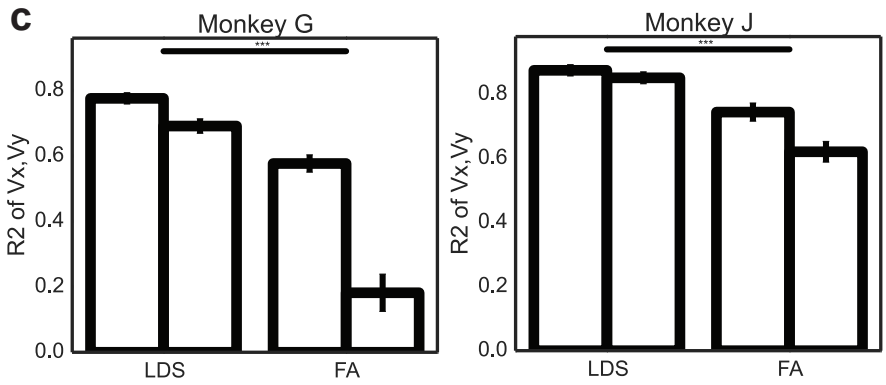

Figure 6.3: Comparisons of generalizability of the LDS and FA models based on accuracy of decoded velocities from estimated neural activity for the held-out within task data (left bar) versus across task data (right bar). The LDS generalizes much better across-task than the FA model does.

and its corresponding neural activity pattern (n_new) from another action sequence, the dynamics ratio of this interchanged command will likely be lower. In this analysis, vel_new can come from an action sequence that is directed toward the same target in the same task, a different target in the same task, or from the other task. 6.4C shows that the dynamics ratio is significantly decreasing as the action sequence become less similar to the original (Monkey G: Cuzick's test for significant ordering of groups, z = 6.01, p = 1.86e-09, decreasing order , Monkey J: Cuzick's test for significant ordering of groups, z = 4.59, p = 4.53e-06, decreasing order). Thus substituting in an identical velocity command from increasingly different action sequences would make the original sequence have decreasing dynamics ratio. Thus, identical velocity commands are produced differently in such a way that keeps the dynamics ratio high in respective action sequences.

6.3.5.1 High Dimensional Linear Dynamics Enable Complex Action Sequences

Finally, we assess how the dimensionality of the model fit with Factor Analysis compares to the dimensionality of the Linear Dynamical Model. One prediction is that since the LDS models have temporal constraints, in order for the model to capture the same variance as the Factor Analysis model, it would need to have a higher dimensionality. Indeed, 6.5shows that to be true.

6.4 Conclusions

Above we have shown that for two different BMI tasks, subjects draw neural activity commands from different low dimensional manifolds. These subspaces differ in a task-relevant manner – using the subspace fit on the opposite task results in task performance

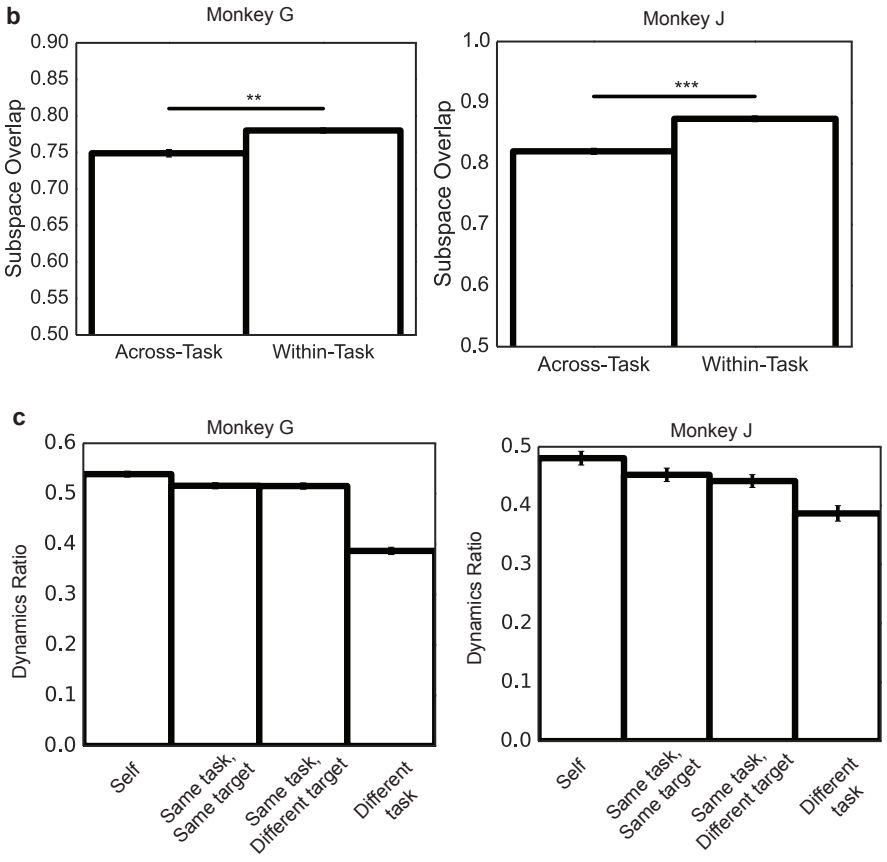

Figure 6.4: b) Main shared subspaces overlap in the LDS for epochs of 16 trials (as in 6.3.1.1D). c) Dynamics ratio when i) using velocity commands from own action sequence, ii) replacing velocity commands with ones from same task and target, iii) replacing velocity commands with ones from same task but different target, or iv) replacing velocity commands with ones from different task.

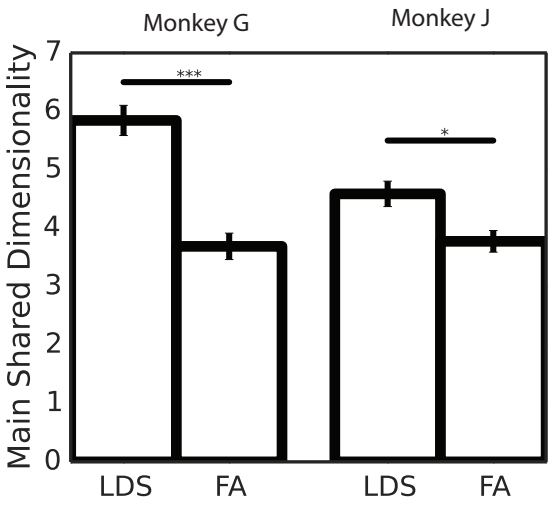

Figure 6.5: Main shared dimensionality of LDS vs. FA models.

decrements. These different tasks also use, on average, different neural activity patterns for generating the same command. When linear decoders are fit on one task, they result in performance decrements when used to perform the other task. Using more complex models of tuning that include parameters capturing tuning to the past and future of the cursor position and velocity improves tuning model performance. However when compared to models that don't have an explicit encoding model (e.g. FA or LDS), the variance captured is low. How do subjects represent the activity of the cursor, then? What rules govern how neural activity is generated at different time bins?

We turned to modeling temporal relationships between consecutive bins of neural activity using an LDS. We find that the LDS explains more variance than FA or the tuning curve models. We also find that the LDS generalizes across task better than FA (fit on one task, test on other), meaning that the dynamics process fit with one task must be general in the second task. This is surprising given that the subspace overlap across tasks is still significantly different in the LDS. Thus, even though fitting data with one task may not fully span the space of the LDS, the dynamics that are defined by the A matrix are still accurate even in parts of the space that the training data does not span. Thus, the A matrix defines a set of temporal rules that generalizes across tasks. We confirm that individual action sequences follow these dynamical rules by comparing how 'dynamical' they would be if we were to substitute out velocity commands and corresponding neural activity that produced the command for the same velocity command from a different action sequence going to the i) same target during the same task, ii) different target during the same task, iii) different task. The finding that these substitutions progressively make the action sequence less dynamical confirms that neural activity does indeed abide by

dynamical rules, justifying the tuning and subspace differences previously noted. Below we discuss the implications of this work for theories of motor control and representation.

6.4.1 Low Dimensional Manifolds for Computation

Although there is no encoding model explicitly built into how a low dimensional manifold would contribute to movement, it has been proposed that low-dimensional neural activity could be responsible for driving movements efficiently [227]. Indeed other groups have found tuning of low-dimensional factors to the cursor in a 2D BMI task [91], suggesting that if tuning

We find, however, that subspaces are slightly different for different tasks, despite the decoder and effector being identical. Perhaps using a higher dimensional manifold is necessary as more tasks are performed, however, this suggests that the low dimensional manifold is not sufficient for explaining how neural activity is generated from a specific effector.

6.4.2 Representation of Movement

In motor tasks, it has been found that tuning curves can drift over time during performance of the same, well-learned task [229, 230, 231, 232], or can change as the task requirements change [233, 234]. Similarly, we demonstrate in the linear tuning curve analyses, that even in a very simple motor task with tens of neurons, a linear decoder, and two output dimensions, that in order to successfully drive the cursor during different tasks subjects do not exhibit a consistent representation of cursor movement. Previous work defending representation of movement variables in motor cortex has discussed many possible reasons for this drift in representation. First, tuning curves could be non-linear instead of linear with movement, thus different distributions of neural-kinematic pairs may result in different linear tuning curves. We find that the kinematics are quite similar from task to task though (6.3G). We also do a non-parametric tuning analysis and find greater across-task versus within-task differences in neural activity patterns for a given velocity command (6.3.2). Second, tuning curves could differ due to unfit co-variates. We do sweep position and velocity co-variates at various lags. Perhaps though, there are uncaptured co-variates such as changes in posture, or physical movement that differ from task to task. Third, subjects could exhibit different types of aiming for different tasks. As found in [203], there are different aiming strategies for different targets, and correlating neural activity with the inferred aimed direction yields more accurate estimates of tuning than correlating activity with the cursor direction. Perhaps in our task, subjects aim differently depending on whether there is an obstacle present or not, making our tuning curves based on neural push seem inaccurate across task. Fourth, it is possible that noisy plasticity drives changes in tuning curves over time, in a random walk fashion [229]. If this were the case though, we would expect to see greater differences in within and across-task comparisons between near and far epochs in 6.3.2. Finally, perhaps the

subjects are still learning the task and adjusting their strategy online. This seems unlikely since performance on either task does not improve over days, or within the course of a single day. Overall, given the evidence we present demonstrating differences in tuning for different tasks, it seems unlikely that these differences could support a representational view of motor cortical population activity.

6.4.3 High-Dimensional Linear Dynamics

In contrast to tuning models and manifolds, incorporating linear dynamics significantly improves the explained neural variance. Indeed, linear dynamical systems have been used to describe population neural activity before, and have even been used online in BMI control and demonstrated to exhibit better performance than typical encoding models [214]. In [214], the LDS is fit from arm reaching data and used online in a BMI task where monkeys continue to use their arm during the task. This demonstrates that during movement, compared to other models that do not take into account temporal dynamics, an LDS is the best estimate of the population state.

In this work we show that not only does an LDS exhibit the most explanatory power of neural activity compared to tuning models or manifold, but also that it generalizes best across different tasks with the same effector. It is notable that this is true even though the different tasks explore slightly different subspaces within the LDS state space. Thus, the dynamics captured by the A matrix are accurate even for portions of state space that are not as well explored.

What exactly does the A matrix capture? Why are there temporal dynamics? Given that neurons are embedded in a recurrently connected network, it may be challenging for them to instantaneously change their rate of firing from one moment to the next. Further, there may be low-dimensional modes of activity that reflect tighter connectivity that also exhibit temporal constraints. However, in the manner in which the LDS is fit in our results above, the A matrix has the challenging task of not only capturing intrinsic neural dynamics, but also of capturing how neural dynamics may change in response to inputs. This is less of a problem in feedforward tasks where input into the network likely does not arrive in time to update the neural activity pattern, but becomes important to disambiguate when assessing tasks that rely on feedback. In other words, in the dynamics equation, there is no control input that could capture task information that is being fed into the network. This is true in [215, 214] as well, and may be the reason why task complexity and effector dimensionality are thought to increase the dimensionality of the linear dynamical system needed to capture neural activity [235]. Instead of thinking of the A matrix as being responsible for capturing all task-related modulations of a low-dimensional population state, if the dynamics process was instead replaced by :

$$z_t = Az_{t-1} + Bu_t + w_t, w_t \sim N(0, W)$$

where u_t is a task-related input (cursor position and velocity with lags), then the A

matrix would be exclusively responsible for explaining intrinsic neural dynamics. A should then not necessarily change according to task, but rather should be a constant reflection of the temporal constraints imposed by the network. Of course A may vary depending on which data is used to fit it, but this is an estimation concern.

Further, when action sequences are generated by neural activity patterns that are not well explained by dynamics (e.g. low dynamics ratio), is this because of intrinsic neural noise? Or is it because of a large input that causes the neural population state to deviate largely from its initial trajectory? In future work we seek to explore how the innovations updates to the LDS are either used to re-direct action sequences correctly, or could just be errors.

Chapter 7

Conclusion

Motor actions constitute the way in which we interact with the world. The topic of this thesis focuses on understanding principles of how the population of our millions of motor system neurons coordinate during different types of behavior.

7.1 Summary of Contributions

7.1.1 Population Activity and Behavior Reflected by Motor Beta Band Oscillations

In chapters 2 and 3 we asked what the neural signatures are of withholding movement. We developed a novel sequential neurofeedback-motor behavior task to specifically test the relationship between population level beta band oscillations in motor cortex, hypothesized to be involved in maintenance of current status, and movement onset. It was discovered that neurofeedback was an effective tool to manipulate beta band oscillations, and confirmed in three non-human primates that beta band oscillations were tightly correlated with movement onset time, but not other movement parameters like initial speed, initial acceleration, or maximum speed. It was also confirmed that it was beta band oscillations indeed that accounted most strongly for this effect. Thus, beta band oscillations were shown to be more tightly linked to movement onset than could be claimed by correlational studies.

Given that beta band oscillations were related to movement onset, and previous work shows populations of single and multi-units also related to movement onset, we sought to connect the three measurements (beta band oscillations, population level activity, and movement onset) by characterizing how beta band oscillations reflected population activity in different behavioral states. Population level single and multi unit activity was analyzed recorded from the same motor cortical brain region during the neurofeedback and natural motor reaching tasks. We found that beta band oscillations reflect similar "non-moving" population states during the neurofeedback task and during natural motor reaching tasks. Thus, although the neurofeedback task is a new, non-natural behavior, the underlying population exhibited similar firing patterns during neurofeedback-induced beta oscillations as during naturally occurring beta band oscillations during motor control.

Thus, we were able to i) develop a method to study the link between local field potential activity and behavior, ii) show a tight correlation between a specific local field potential activity pattern and behavior, iii) demonstrate that the underlying neuronal population activity exhibits similar characteristics when the local field potential activity pattern occur naturally or via neurofeedback. Taken together, this study links signatures in local field potential signals with different underlying population states, both of which are indicative of withholding movement. Future studies could i) use similar methods to more rigorously test correlations between local field potential activity and behaviors of interest, and ii) use beta band oscillations as a biomarker of movement being withheld or the underlying population activity signifying withholding of movement.

7.1.2 Demonstration of Neurofeedback Control in Parkinsonian Patients

In addition to BMIs being useful scientific tools, in chapter 4 we investigated whether BMIs may be useful as possible therapeutic systems. In three parkinsonian patients, we show in a proof of concept study that they can exhibit neurofeedback control over their motor cortical beta band oscillations, which had not yet been shown. Notably, the system used to perform the neurofeedback study was a wireless, fully embedded neural recording and stimulation system, the Activa PC + S. This fully embedded, wireless capability allowed us to perform some of the neurofeedback sessions in the comfort of patients' own homes. Further, software used to perform this closed loop study has been made publicly available. This project thus has brought neurofeedback and BMI control out of the lab, and hopefully future studies will continue to trend towards at-home care for increased patient convenience.

We also perform one extended at-home study where we assess how motor cortical beta band activity is linked to finger tapping movements. This study provides pilot data pointing to beta desynchronization being a useful tool for patients to practice to facilitate faster movement onset, and faster and more vigorous finger tapping. Future studies could leverage our findings that parkinsonian patients can learn to volitionally control their neural activity, and investigate how the activity in other brain structures is effected by changing motor cortical beta activity. Further, future studies will continue at-home studies and better characterize across-brain beta band activity during finger tapping movement following neurofeedback to see how beta band dynamics change over the course of movement following neurofeedback.

7.1.3 Establish how Motor Population Shared Variability May Emerge, and how Shared and Private Variability Contribute to Online Control

Beta band oscillations are only one statistic of population neural activity, so we next turned our attention to populations of single and multi-units that coordinate to perform 2D cursor control movements with a BMI, representing a simplified, mini, motor system. In chapter 5, we simulate and test in experiments how different decompositions of population activity give rise to cursor movements. First we develop a generative model of BMI output unit population activity that allows us to ask questions about how correlated vs. uncorrelated and tuned vs. untuned inputs to a population become summed together. We find that tuned, uncorrelated inputs can give rise to observed correlated activity when summed in the population, and that fast, tuned uncorrelated inputs can give rise to even more correlated activity when summed in the population. This simulation yields possible mechanisms in which uncorrelated observations can become summed and yield correlated activity downstream.

Further, we see how population correlated versus uncorrelated activity gives rise to online BMI control in simulation and in a real online catch trial experiment. We find that on average, decomposing neural population signals into correlated and uncorrelated parts and sending them into the decoder individually yields faster and more accurate BMI performance respectively. Many groups discard uncorrelated population activity in their BMI experiments, but understanding exactly how uncorrelated activity supports performance requires future investigation.

7.1.4 Demonstration of Neural Dynamics in a Feedback Control Task

Finally, in chapter 6 we build on our investigation of how correlated and uncorrelated activity supports cursor trajectories by studying how populations of single and multi-units coordinate to construct different 2D cursor action sequences in a BMI. Here we show that the correlations in neural activity differ when an expert subject performs straight-line centerout task compared to a curved-line obstacle task. Further, even the mean neural activity patterns for a given velocity command sent to the cursor differ across the two tasks. We find that the neural data is best explained by a linear dynamical system (LDS), which models neural activity as emissions from a lower-dimensional latent state that abides by lawful temporal dynamics. The LDS fit on one task also fits neural activity from the other task well, unlike other models. Further, the individual segments of an action sequence are shown to be selected to best follow the dynamics of the low dimensional latent state, yielding an understanding of why mean firing rates for a given velocity command may differ across task, and why correlation patterns may differ.

Future work will investigate the nature of the temporal dynamics modeled, how much they reflect the task being performed versus the network the neurons are in, and how deviations from the lawful temporal dynamics contribute to BMI performance.

7.2 Future Directions

7.2.1 Studying Local Field Potential signals with Neurofeedback

The origin and behavioral significance of specific frequencies of local field potential signals in the brain is debated. Developing methods to perturb these signals would shed light into how the underlying neural firing patterns reflected in these signals contribute to behavior. Non-invasive stimulation methods are one approach, but using neurofeedback is a stimulation-free (and artifact-free!) approach. Further, if the LFP signal of interest is implicated in disease or injury, neurofeedback may be a promising therapy. Beta oscillations remain a signal of interest due to their exacerbation in Parkinson's disease. Going forward, understanding how beta oscillations are transmitted throughout the basal ganglia, thalamus, and motor cortical areas both in healthy conditions and PD condi-

tions will help clarify how they may reflect impaired computation. New chronic recording systems are being developed to allow just those investigations to take place [192].

7.2.2 Studying Sensorimotor Control with BMIs

Given how many principles of neural computation remain to be understood about simple, closed-loop BMI systems, continuing to study these simple systems is appropriate. The field of systems motor neuroscience is starting to shift away from representational views of motor control towards generative ones, and with that comes a need for new principles of motor cortical computation. In chapter 6 we begin to develop a linear dynamical system to model neural activity during 2D cursor control. Extending this model to include feedback control terms will shed light on what aspects of neural activity are due to true intrinsic neural dynamics, and which are due to inputs to the system such as visual feedback of the cursor. Further studies may then be able to use this framework to decompose intrinsic versus extrinsic components of neural activity, and begin to understand what relationship the extrinsic inputs have with behavior, as is already starting to be done [236]. Perhaps extrinsic inputs represent visual or sensory feedback, strategy shifts, or even mood changes.

7.2.3 Conclusion

BMIs used for science investigations offer new ways to perturb, study, and theorize about our complex, distributed motor system. This thesis makes contributions towards exploring the previously complex-to-perturb local field potential features, towards studying the origins of correlated and uncorrelated population activity and their effects on behavior, and towards testing models of how motor cortical population activity may coordinate the generate action sequences. Further work leveraging the potential of closed-loop BMIs as scientific tools provide promising new paths of scientific and clinical exploration.

Bibliography

[1] Patrick Haggard. Human volition: towards a neuroscience of will. *Nature reviews. Neuroscience*, 9(12):934–946, December 2008.

[2] G Rizzolatti, G Luppino, and M Matelli. The organization of the cortical motor system: new concepts. *Electroencephalography and Clinical Neurophysiology*, 106(4):283–296, April 1998.

[3] R. P. Dum and P. L. Strick. The origin of corticospinal projections from the premotor areas in the frontal lobe. *Journal of Neuroscience*, 11(3):667–689, March 1991.

[4] Michael S. A. Graziano. New Insights into Motor Cortex. *Neuron*, 71(3):387–388, August 2011.

[5] Eric R. Kandel, James H. Schwartz, Thomas M. Jessell, Steven A. Siegelbaum, and A. J. Hudspeth, editors. *Principles of Neural Science, Fifth Edition*. McGraw-Hill Education / Medical, New York, 5th edition edition, October 2012.

[6] C. S. Sherrington. Flexion-reflex of the limb, crossed extension-reflex, and reflex stepping and standing. *The Journal of Physiology*, 40(1-2):28–121, April 1910.

[7] Yifat Prut and Eberhard E. Fetz. Primate spinal interneurons show pre-movement instructed delay activity. *Nature*, 401(6753):590–594, October 1999.

[8] E. E. Fetz, S. I. Perlmutter, Y. Prut, K. Seki, and S. Votaw. Roles of primate spinal interneurons in preparation and execution of voluntary hand movement. *Brain Research. Brain Research Reviews*, 40(1-3):53–65, October 2002.

[9] Eugene Braunwald, Anthony S. Fauci, Dennis L. Kasper, et al. *Harrison's Principles of Internal Medicine: Textbook, Self-Assessment and Board Review*. McGraw-Hill Professional Publishing, February 2001. Google-Books-ID: BcOhPwAACAAJ.

[10] Sylvain Lavoie and Trevor Drew. Discharge characteristics of neurons in the red nucleus during voluntary gait modifications: a comparison with the motor cortex. *Journal of Neurophysiology*, 88(4):1791–1814, October 2002.

[11] C. Nicholas Riddle, Steve A. Edgley, and Stuart N. Baker. Direct and Indirect Connections with Upper Limb Motoneurons from the Primate Reticulospinal Tract. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29(15):4993, April 2009.

[12] C. Nicholas Riddle and Stuart N. Baker. Convergence of pyramidal and medial brain stem descending pathways onto macaque cervical spinal interneurons. *Journal of Neurophysiology*, 103(5):2821–2832, May 2010.

[13] A. V. Poliakov and M. H. Schieber. Limited functional grouping of neurons in the motor cortex hand area during individuated finger movements: A cluster analysis. *Journal of Neurophysiology*, 82(6):3488–3505, December 1999.

[14] Michael Graziano. The organization of behavioral repertoire in motor cortex. *Annual Review of Neuroscience*, 29:105–134, 2006.

[15] J. M. Macpherson, C. Marangoz, T. S. Miles, and M. Wiesendanger. Microstimulation of the supplementary motor area (SMA) in the awake monkey. *Experimental Brain Research*, 45(3):410–416, 1982.

[16] Parashkev Nachev, Henrietta Wydell, Kevin O'Neill, Masud Husain, and Christopher Kennard. The role of the pre-supplementary motor area in the control of action. *Neuroimage*, 36(3-3):T155–T163, 2007.

[17] N Picard and P L Strick. Motor areas of the medial wall: a review of their location and functional activation. *Cerebral cortex (New York, N.Y.: 1991)*, 6(3):342–353, June 1996.

[18] G. Luppino, M. Matelli, R. M. Camarda, V. Gallese, and G. Rizzolatti. Multiple representations of body movements in mesial area 6 and the adjacent cingulate cortex: an intracortical microstimulation study in the macaque monkey. *The Journal of Comparative Neurology*, 311(4):463–482, September 1991.

[19] G. E. Alexander, M. R. DeLong, and P. L. Strick. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, 9:357–381, 1986.

[20] Amiram Grinvald and Rina Hildesheim. VSDI: a new era in functional imaging of cortical dynamics. *Nature Reviews Neuroscience*, 5(11):874, November 2004.

[21] György Buzsáki, Costas A. Anastassiou, and Christof Koch. The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6):407–420, June 2012.

[22] Henrik Lindén, Tom Tetzlaff, Tobias C. Potjans, et al. Modeling the Spatial Reach of the LFP. *Neuron*, 72(5):859–872, August 2011.

[23] Supratim Ray and John H. R. Maunsell. Different Origins of Gamma Rhythm and High-Gamma Activity in Macaque Visual Cortex. *PLoS Biol*, 9(4):e1000610, April 2011.

[24] Erin L. Rich and Joni D. Wallis. Spatiotemporal dynamics of information encoding revealed in orbitofrontal high-gamma. *Nature Communications*, 8(1):1139, October 2017.

[25] Henrik Lindén, Klas H. Pettersen, and Gaute T. Einevoll. Intrinsic dendritic filtering gives low-pass power spectra of local field potentials. *Journal of Computational Neuroscience*, 29(3):423–444, December 2010.

[26] Hernan Gonzalo Rey, Carlos Pedreira, and Rodrigo Quian Quiroga. Past, present and future of spike sorting techniques. *Brain Research Bulletin*, 119(Pt B):106–117, October 2015.

[27] Cynthia A Chestek, Vikash Gilja, Paul Nuyujukian, et al. Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. *J. Neural Eng.*, 8(4):045005, August 2011.

[28] Andrew B. Schwartz, X. Tracy Cui, Douglas J. Weber, and Daniel W. Moran. Brain-Controlled Interfaces: Movement Restoration with Neural Prosthetics. *Neuron*, 52(1):205–220, October 2006.

[29] Stephen H. Scott. Inconvenient Truths about neural processing in primary motor cortex. *The Journal of Physiology*, 586(5):1217–1224, March 2008.

[30] John F. Kalaska. From intention to action: motor cortex and the control of reaching movements. *Advances in Experimental Medicine and Biology*, 629:139–178, 2009.

[31] Stephen H. Scott. Inconvenient Truths about neural processing in primary motor cortex. *The Journal of Physiology*, 586(5):1217–1224, March 2008.

[32] Eberhard E. Fetz. Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and Brain Sciences*, 15(4):679–690, December 1992.

[33] Mark M. Churchland, John P. Cunningham, Matthew T. Kaufman, et al. Neural population dynamics during reaching. *Nature*, advance online publication, June 2012.

[34] Krishna V. Shenoy, Maneesh Sahani, and Mark M. Churchland. Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience*, 36(1):337–359, 2013.

[35] K. Cora Ames, Stephen I. Ryu, and Krishna V. Shenoy. Neural dynamics of reaching following incorrect or absent motor preparation. *Neuron*, 81(2):438–451, January 2014.

[36] Biljana Petreska, M. Yu Byron, John P. Cunningham, et al. Dynamical segmentation of single trials from population neural data. In *Advances in neural information processing systems*, pages 756–764, 2011.

[37] Sergey D. Stavisky, Jonathan C. Kao, Stephen I. Ryu, and Krishna V. Shenoy. Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions. *Neuron*, 2017.

[38] Dongjin Seo, Ryan M. Neely, Konlin Shen, et al. Wireless Recording in the Peripheral Nervous System with Ultrasonic Neural Dust. *Neuron*, 91(3):529–539, August 2016.

[39] Mijail D. Serruya, Nicholas G. Hatsopoulos, Liam Paninski, Matthew R. Fellows, and John P. Donoghue. Instant neural control of a movement signal. *Nature*, 416(6877):141–142, March 2002.

[40] Dawn M. Taylor, Stephen I. Helms Tillery, and Andrew B. Schwartz. Direct Cortical Control of 3D Neuroprosthetic Devices. *Science*, 296(5574):1829–1832, June 2002.

[41] Jose M. Carmena, Mikhail A. Lebedev, Roy E. Crist, et al. Learning to Control a Brain–Machine Interface for Reaching and Grasping by Primates. *PLOS Biol*, 1(2):e42, October 2003.

[42] Leigh R. Hochberg, Mijail D. Serruya, Gerhard M. Friehs, et al. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442(7099):164–171, July 2006.

[43] J. R. Millan, F. Renkens, J. Mourino, and W. Gerstner. Noninvasive brain-actuated control of a mobile robot by human EEG. *IEEE Transactions on Biomedical Engineering*, 51(6):1026–1033, June 2004.

[44] Mariska Vansteensel, Elmar Pels, Martin Bleichner, et al. *Fully Implanted Brain–Computer Interface in a Locked-In Patient with ALS*, volume 375. November 2016. DOI: 10.1056/NEJMoa1608085.

[45] Chethan Pandarinath, Paul Nuyujukian, Christine H. Blabe, et al. High performance communication by people with paralysis using an intracortical brain-computer interface. *eLife*, 6:e18554, February 2017.

[46] Leigh R. Hochberg, Daniel Bacher, Beata Jarosiewicz, et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375, May 2012.

[47] Jennifer L Collinger, Brian Wodlinger, John E Downey, et al. 7 degree-of-freedom neuroprosthetic control by an individual with tetraplegia. *Lancet*, 381(9866):557–564, February 2013.

[48] Ander Ramos-Murguialday, Doris Broetz, Massimiliano Rea, et al. Brain–machine interface in chronic stroke rehabilitation: A controlled study. *Annals of Neurology*, 74(1):100–108, July 2013.

[49] S. R. Soekadar, M. Witkowski, C. Gómez, et al. Hybrid EEG/EOG-based brain/neural hand exoskeleton restores fully independent daily living activities after quadriplegia. *Science Robotics*, 1(1):eaag3296, December 2016.

[50] Chet T. Moritz, Eberhard E. Fetz, and Steve I. Perlmutter. Direct control of paralysed muscles by cortical neurons. *Nature*, 456(7222):nature07418, October 2008.

[51] Yukio Nishimura, Steve I. Perlmutter, and Eberhard E. Fetz. Restoration of upper limb movement via artificial corticospinal and musculospinal connections in a monkey with spinal cord injury. *Frontiers in Neural Circuits*, 7:57, 2013.

[52] C. Ethier, E. R. Oby, L. E. Miller, and M. J. Bauman. Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature*, 485(7398):nature10987, April 2012.

[53] Francesco Tenore, Ander Ramos, Amir Fahmy, et al. Towards the control of individual fingers of a prosthetic hand using surface EMG signals. *Conference proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*, 2007:6146–6149, 2007.

[54] M. Jiang, R. Wang, J. Wang, and D. Jin. A Method of Recognizing Finger Motion Using Wavelet Transform of Surface EMG Signal. *Conference proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*, 3:2672–2674, 2005.

[55] Ramses E. Alcaide-Aguirre, David C. Morgenroth, and Daniel P. Ferris. Motor control and learning with lower-limb myoelectric control in amputees. *Journal of Rehabilitation Research and Development*, 50(5):687–698, 2013.

[56] Simon Little, Alex Pogosyan, Spencer Neal, et al. Adaptive Deep Brain Stimulation In Advanced Parkinson Disease. *Annals of Neurology*, 74(3):449–457, September 2013.

[57] Boris Rosin, Maya Slovik, Rea Mitelman, et al. Closed-Loop Deep Brain Stimulation Is Superior in Ameliorating Parkinsonism. *Neuron*, 72(2):370–384, October 2011.

[58] Christianne N. Heck, David King-Stephens, Andrew D. Massey, et al. Two-year seizure reduction in adults with medically intractable partial onset epilepsy treated with responsive neurostimulation: final results of the RNS System Pivotal trial. *Epilepsia*, 55(3):432–441, March 2014.

[59] Nikolaus Wenger, Eduardo Martin Moraud, Stanisa Raspopovic, et al. Closed-loop neuromodulation of spinal sensorimotor circuits controls refined locomotion after complete spinal cord injury. *Science Translational Medicine*, 6(255):255ra133–255ra133, September 2014.

[60] David M. Schultz, Lynn Webster, Peter Kosek, et al. Sensor-driven position-adaptive spinal cord stimulation for chronic pain. *Pain Physician*, 15(1):1–12, February 2012.

[61] Paul Boon, Kristl Vonck, Kenou van Rijckevorsel, et al. A prospective, multicenter study of cardiac-based seizure detection to activate vagus nerve stimulation. *Seizure*, 32:52–61, November 2015.

[62] Brian J. Wenzel, Joseph W. Boggs, Kenneth J. Gustafson, and Warren M. Grill. Closed loop electrical control of urinary continence. *The Journal of Urology*, 175(4):1559–1563, April 2006.

[63] Federico Cirett Galán and Carole R. Beal. EEG Estimates of Engagement and Cognitive Workload Predict Math Problem Solving Outcomes. In *User Modeling, Adaptation, and Personalization*, Lecture Notes in Computer Science, pages 51–62. Springer, Berlin, Heidelberg, July 2012.

[64] E. W. Anderson, K. C. Potter, L. E. Matzen, et al. A User Study of Visualization Effectiveness Using EEG and Cognitive Load. *Computer Graphics Forum*, 30(3):791–800, June 2011.

[65] Leena Subramanian, John V. Hindle, Stephen Johnston, et al. Real-Time Functional Magnetic Resonance Imaging Neurofeedback for Treatment of Parkinson's Disease. *The Journal of Neuroscience*, 31(45):16309–16317, November 2011.

[66] Holger Gevensleben, Birgit Holl, Björn Albrecht, et al. Distinct EEG effects related to neurofeedback training in children with ADHD: a randomized controlled trial. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, 74(2):149–157, November 2009.

[67] Mark P. Jensen, Kevin J. Gertz, Amy E. Kupper, et al. Steps toward developing an EEG biofeedback treatment for chronic pain. *Applied Psychophysiology and Biofeedback*, 38(2):101–108, June 2013.

[68] James Sulzer, Ranganatha Sitaram, Maria Laura Blefari, et al. Neurofeedback-mediated self-regulation of the dopaminergic midbrain. *NeuroImage*, 75C:176–184, March 2013.

[69] Ander Ramos-Murguialday, Doris Broetz, Massimiliano Rea, et al. Brain-Machine-Interface in Chronic Stroke Rehabilitation: A Controlled Study. *Annals of neurology*, 74(1):100–108, July 2013.

[70] Eberhard E. Fetz. Operant Conditioning of Cortical Unit Activity. *Science*, 163(3870):955–958, February 1969.

[71] Karunesh Ganguly and Jose M. Carmena. Emergence of a Stable Cortical Map for Neuroprosthetic Control. *PLoS Biol*, 7(7):e1000153, July 2009.

[72] Aaron C. Koralek, Xin Jin, John D. Long Ii, Rui M. Costa, and Jose M. Carmena. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature*, 483(7389):331–335, March 2012.

[73] Amy L. Orsborn, Helene G. Moorman, Simon A. Overduin, et al. Closed-Loop Decoder Adaptation Shapes Neural Plasticity for Skillful Neuroprosthetic Control. *Neuron*, 82(6):1380–1393, June 2014.

[74] Beata Jarosiewicz, Steven M. Chase, George W. Fraser, et al. 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(49):19486–19491, December 2008.

[75] Robert J Schafer and Tirin Moore. Selective attention from voluntary control of neurons in prefrontal cortex. *Science (New York, N.Y.)*, 332(6037):1568–1571, June 2011.

[76] Frank Scharnowski, Chloe Hutton, Oliver Josephs, Nikolaus Weiskopf, and Geraint Rees. Improving Visual Perception through Neurofeedback. *The Journal of Neuroscience*, 32(49):17830–17841, December 2012.

[77] S. J. Johnston, S. G. Boehm, D. Healy, R. Goebel, and D. E. J. Linden. Neurofeedback: A promising tool for the self-regulation of emotion networks. *NeuroImage*, 49(1):1066–1072, January 2010.

[78] Matthew D. Golub, Steven M. Chase, Aaron P. Batista, and Byron M. Yu. Brain–computer interfaces for dissecting cognitive processes underlying sensorimotor control. *Current opinion in neurobiology*, 37:53, April 2016.

[79] Amy L Orsborn and Bijan Pesaran. Parsing learning in networks using brain–machine interfaces. *Current Opinion in Neurobiology*, 46(Supplement C):76–83, October 2017.

[80] Matthew D. Golub, Byron M. Yu, and Steven M. Chase. Internal models for interpreting neural population activity during sensorimotor control. *eLife*, 4:e10015, December 2015.

[81] Reza Shadmehr, Maurice A. Smith, and John W. Krakauer. Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, 33:89–108, 2010.

[82] Vikash Gilja, Paul Nuyujukian, Cindy A. Chestek, et al. A high-performance neural prosthesis enabled by control algorithm design. *Nature Neuroscience*, 15(12):1752–1757, December 2012.

[83] Matthew T. Kaufman, Mark M. Churchland, Stephen I. Ryu, and Krishna V. Shenoy. Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience*, 17(3):440–448, March 2014.

[84] Andreas K Engel and Pascal Fries. Beta-band oscillations — signalling the status quo? *Current Opinion in Neurobiology*, 20(2):156–165, April 2010.

[85] Coralie de Hemptinne, Elena S Ryapolova-Webb, Ellen L Air, et al. Exaggerated phase-amplitude coupling in the primary motor cortex in Parkinson disease. *Proceedings of the National Academy of Sciences of the United States of America*, 110(12):4780–4785, March 2013.

[86] Coralie de Hemptinne, Nicole C. Swann, Jill L. Ostrem, et al. Therapeutic deep brain stimulation reduces cortical phase-amplitude coupling in Parkinson's disease. *Nature Neuroscience*, advance online publication, April 2015.

[87] Nicole C. Swann, Coralie de Hemptinne, Adam R. Aron, et al. Elevated synchrony in Parkinson disease detected with electroencephalography. *Annals of Neurology*, 78(5):742–750, November 2015.

[88] Byron M. Yu, John P. Cunningham, Gopal Santhanam, et al. Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity. *Journal of Neurophysiology*, 102(1):614–635, July 2009.

[89] Emilio Bizzi and Vincent C. K. Cheung. The neural origin of muscle synergies. *Frontiers in Computational Neuroscience*, 7, April 2013.

[90] Vivek R. Athalye, Karunesh Ganguly, Rui M. Costa, and Jose M. Carmena. Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control. *Neuron*, 93(4):955–970.e5, February 2017.

[91] Patrick T. Sadtler, Kristin M. Quick, Matthew D. Golub, et al. Neural constraints on learning. *Nature*, 512(7515):423–426, August 2014.

[92] Gerwin Schalk and Eric C. Leuthardt. Brain-computer interfaces using electrocorticographic signals. *IEEE reviews in biomedical engineering*, 4:140–154, 2011.

[93] Tomislav Milekovic, Beata Jarosiewicz, Anish A. Sarma, Leigh R. Hochberg, and John P. Donoghue. Increases in beta frequency band LFP activity mark low engagement of motor cortex in voluntary movement intentions in people with longstanding tetraplegia. In *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)*, November 2013.

[94] G.V. Ranade, K. Ganguly, and J. Carmena. LFP beta power predicts cursor stationarity in BMI task. In *4th International IEEE/EMBS Conference on Neural Engineering, 2009. NER '09*, pages 482–485, April 2009.

[95] V.N. Murthy and E.E. Fetz. Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *Journal of Neurophysiology*, 76(6):3968–3982, 1996.

[96] Michael Denker, Sébastien Roux, Marc Timme, Alexa Riehle, and Sonja Grün. Phase synchronization between LFP and spiking activity in motor cortex during movement preparation. *Neurocomputing*, 70(10–12):2096–2101, June 2007.

[97] Jacob Reimer and Nicholas G. Hatsopoulos. Periodicity and Evoked Responses in Motor Cortex. *The Journal of Neuroscience*, 30(34):11506–11515, August 2010.

[98] Ryan T Canolty, Karunesh Ganguly, and Jose M Carmena. Task-dependent changes in cross-level coupling between single neurons and oscillatory activity in multiscale networks. *PLoS computational biology*, 8(12):e1002809, 2012.

[99] Alexa Riehle, Sonja Grün, Markus Diesmann, and Ad Aertsen. Spike Synchronization and Rate Modulation Differentially Involved in Motor Cortical Function. *Science*, 278(5345):1950–1953, December 1997.

[100] Flavio Fröhlich and David A. McCormick. Endogenous Electric Fields May Guide Neocortical Network Activity. *Neuron*, 67(1):129–143, July 2010.

[101] Gyorgy Buzsáki. *Rhythms of the Brain*. Oxford University Press, August 2006.

[102] M.A Whittington, R.D Traub, N Kopell, B Ermentrout, and E.H Buhl. Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *International Journal of Psychophysiology*, 38(3):315–336, December 2000.

[103] A. Jackson, R. L. Spinks, T. C. B. Freeman, D. M. Wolpert, and R. N. Lemon. Rhythm generation in monkey motor cortex explored using pyramidal tract stimulation. *The Journal of Physiology*, 541(3):685–699, June 2002.

[104] Anita K. Roopun, Steven J. Middleton, Mark O. Cunningham, et al. A beta2-frequency (20-30 Hz) oscillation in nonsynaptic networks of somatosensory cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 103(42):15646–15650, October 2006.

[105] Anita K. Roopun, Mark A. Kramer, Lucy M. Carracedo, et al. Temporal Interactions between Cortical Rhythms. *Frontiers in Neuroscience*, 2(2):145–154, December 2008.

[106] S. N. Baker, E. Olivier, and R. N. Lemon. Coherent oscillations in monkey motor cortex and hand muscle EMG show task-dependent modulation. *The Journal of Physiology*, 501 (Pt 1):225–241, May 1997.

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

[108] B A Conway, D M Halliday, S F Farmer, et al. Synchronization between motor cortex and spinal motoneuronal pool during the performance of a maintained motor task in man. *The Journal of Physiology*, 489(Pt 3):917–924, December 1995.

[109] Georgopoulos. Motor cortex: A changing perspective. *Experimental Brain Research*, Suppl.(22):175–183, 1992.

[110] N. Kopell, G. B. Ermentrout, M. A. Whittington, and R. D. Traub. Gamma rhythms and beta rhythms have different synchronization properties. *Proceedings of the National Academy of Sciences of the United States of America*, 97(4):1867–1872, February 2000.

[111] Toru Tsujimoto, Tatsuya Mima, Hideki Shimazu, and Yoshikazu Isomura. Directional organization of sensorimotor oscillatory activity related to the electromyogram in the monkey. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, 120(6):1168–1173, June 2009.

[112] Claire L. Witham and Stuart N. Baker. Network oscillations and intrinsic spiking rhythmicity do not covary in monkey sensorimotor areas. *The Journal of Physiology*, 580(3):801–814, May 2007.

[113] Stuart N. Baker. Oscillatory interactions between sensorimotor cortex and the periphery. *Current Opinion in Neurobiology*, 17(6):649–655, December 2007.

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

[115] Doug Rubino, Kay A Robbins, and Nicholas G Hatsopoulos. Propagating waves mediate information transfer in the motor cortex. *Nat. Neurosci.*, 9(12):1549–1557, 2006.

[116] Stuart N. Baker, Matthew Chiu, and Eberhard E. Fetz. Afferent Encoding of Central Oscillations in the Monkey Arm. *Journal of Neurophysiology*, 95(6):3904–3910, June 2006.

[117] Elodie Lalo, Thomas Gilbertson, Louise Doyle, et al. Phasic increases in cortical beta activity are associated with alterations in sensory processing in the human. *Experimental brain research*, 177(1):137–145, February 2007.

[118] W. Jiang, Y. Lamarre, and C. E. Chapman. Modulation of cutaneous cortical evoked potentials during isometric and isotonic contractions in the monkey. *Brain Research*, 536(1–2):69–78, December 1990.

[119] W. A. Mackay. Synchronized neuronal oscillations and their role in motor processes. *Trends in Cognitive Sciences*, 1(5):176–183, August 1997.

[120] Kazuhiko Seki and Eberhard E. Fetz. Gating of Sensory Input at Spinal and Cortical Levels during Preparation and Execution of Voluntary Movement. *The Journal of Neuroscience*, 32(3):890–902, January 2012.

[121] Daofen Chen and Eberhard E. Fetz. Characteristic membrane potential trajectories in primate sensorimotor cortex neurons recorded in vivo. *Journal of Neurophysiology*, 94(4):2713–2725, October 2005.

[122] Timothy J. Buschman, Eric L. Denovellis, Cinira Diogo, Daniel Bullock, and Earl K. Miller. Synchronous Oscillatory Neural Ensembles for Rules in the Prefrontal Cortex. *Neuron*, 76(4):838–846, November 2012.

[123] Timothy J. Buschman and Earl K. Miller. Top-Down Versus Bottom-Up Control of Attention in the Prefrontal and Posterior Parietal Cortices. *Science*, 315(5820):1860–1862, March 2007.

[124] Bram-Ernst Verhoef, Rufin Vogels, and Peter Janssen. Synchronization between the end stages of the dorsal and the ventral visual stream. *Journal of Neurophysiology*, 105(5):2030–2042, May 2011.

[125] Jung H. Lee, Miles A. Whittington, and Nancy J. Kopell. Top-down beta rhythms support selective attention via interlaminar interaction: a model. *PLoS computational biology*, 9(8):e1003164, 2013.

[126] Eberhard E. Fetz. Volitional Control of Cortical Oscillations and Synchrony. *Neuron*, 77(2):216–218, January 2013.

[127] J. N. Sanes and J. P. Donoghue. Oscillations in local field potentials of the primate motor cortex during voluntary movement. *Proceedings of the National Academy of Sciences*, 90(10):4470–4474, May 1993.

[128] Maryam Saleh, Jacob Reimer, Richard Penn, Catherine L. Ojakangas, and Nicholas G. Hatsopoulos. Fast and Slow Oscillations in Human Primary Motor Cortex Predict Oncoming Behaviorally Relevant Cues. *Neuron*, 65(4):461–471, February 2010.

[129] Ryan T. Canolty, Karunesh Ganguly, Steven W. Kennerley, et al. Oscillatory phase coupling coordinates anatomically dispersed functional cell assemblies. *Proceedings of the National Academy of Sciences*, 107(40), September 2010.

[130] Aaron C. Koralek, Rui M. Costa, and Jose M. Carmena. Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron*, 79(5):865–872, September 2013.

[131] Tanuj Gulati, Dhakshin S. Ramanathan, Chelsea C. Wong, and Karunesh Ganguly. Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nature Neuroscience*, 17(8):1107–1113, August 2014.

[132] W. Singer. Consciousness and the binding problem. *Annals of the New York Academy of Sciences*, 929:123–146, April 2001.

[133] Simal Ozen, Anton Sirota, Mariano A. Belluscio, et al. Transcranial electric stimulation entrains cortical neuronal populations in rats. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(34):11476–11485, August 2010.

[134] Subramaniam Venkatraman, Jose M. Carmena, Subramaniam Venkatraman, and Jose M. Carmena. Behavioral modulation of stimulus-evoked oscillations in barrel cortex of alert rats. *Frontiers in Integrative Neuroscience*, 3:10, 2009.

[135] Flavio Fröhlich. Endogenous and exogenous electric fields as modifiers of brain activity: rational design of noninvasive brain stimulation with transcranial alternating current stimulation. *Dialogues in Clinical Neuroscience*, 16(1):93–102, March 2014.

[136] Lisa Marshall, Halla Helgadóttir, Matthias Mölle, and Jan Born. Boosting slow oscillations during sleep potentiates memory. *Nature*, 444(7119):610–613, November 2006.

[137] Claudia Wach, Vanessa Krause, Vera Moliadze, et al. The effect of 10 Hz transcranial alternating current stimulation (tACS) on corticomuscular coherence. *Frontiers in Human Neuroscience*, 7:511, 2013.

[138] Nick J. Davis, Simon P. Tomlinson, and Helen M. Morgan. The Role of Beta-Frequency Neural Oscillations in Motor Control. *Journal of Neuroscience*, 32(2):403–404, January 2012.

[139] Matteo Feurra, Giovanni Bianco, Emiliano Santarnecchi, et al. Frequency-Dependent Tuning of the Human Motor System Induced by Transcranial Oscillatory Potentials. *The Journal of Neuroscience*, 31(34):12165–12170, August 2011.

[140] Alek Pogosyan, Louise Doyle Gaynor, Alexandre Eusebio, and Peter Brown. Boosting Cortical Activity at Beta-Band Frequencies Slows Movement in Humans. *Current Biology*, 19(19):1637–1641, October 2009.

[141] Karen A. Moxon and Guglielmo Foffani. Brain-Machine Interfaces beyond Neuroprosthetics. *Neuron*, 86(1):55–67, April 2015.

[142] Ben Engelhard, Nofar Ozeri, Zvi Israel, Hagai Bergman, and Eilon Vaadia. Inducing Gamma Oscillations and Precise Spike Synchrony by Operant Conditioning via Brain-Machine Interface. *Neuron*, 77(2):361–375, January 2013.

[143] Kelvin So, Siddharth Dangi, Amy L. Orsborn, Michael C. Gastpar, and Jose M. Carmena. Subject-specific modulation of local field potential spectral power during brain–machine interface control in primates. *Journal of Neural Engineering*, 11(2):026002, April 2014.

[144] Dennis J. McFarland, William A. Sarnacki, and Jonathan R. Wolpaw. Effects of training pre-movement sensorimotor rhythms on behavioral performance. *Journal of Neural Engineering*, 12(6):066021, 2015.

[145] B. Babadi and E.N. Brown. A Review of Multitaper Spectral Analysis. *IEEE Transactions on Biomedical Engineering*, 61(5):1555–1564, May 2014.

[146] J Cuzick. A Wilcoxon-type test for trend. *Stat. Med.*, 4(1):87–90, January 1985.

[147] Raed A. Joundi, Ned Jenkinson, John-Stuart Brittain, Tipu Z. Aziz, and Peter Brown. Driving Oscillatory Activity in the Human Cortex Enhances Motor Performance. *Current Biology*, 22(5):403–407, March 2012.

[148] Richard Courtemanche, Naotaka Fujii, and Ann M. Graybiel. Synchronous, Focally Modulated β-Band Oscillations Characterize Local Field Potential Activity in the Striatum of Awake Behaving Monkeys. *Journal of Neuroscience*, 23(37):11741–11752, December 2003.

[149] Jan-Mathijs Schoffelen, Robert Oostenveld, and Pascal Fries. Neuronal coherence as a mechanism of effective corticospinal interaction. *Science (New York, N.Y.)*, 308(5718):111–113, April 2005.

[150] Mikael Lundqvist, Jonas Rose, Pawel Herman, et al. Gamma and Beta Bursts Underlie Working Memory. *Neuron*, 90(1):152–164, April 2016.

[151] Noa Fogelson, Alek Pogosyan, Andrea A. Kühn, et al. Reciprocal interactions between oscillatory activities of different frequencies in the subthalamic region of patients with Parkinson's disease. *European Journal of Neuroscience*, 22(1):257–266, July 2005.

[152] Thomas Gilbertson, Elodie Lalo, Louise Doyle, et al. Existing motor state is favored at the expense of new movement during 13-35 Hz oscillatory synchrony in the human corticospinal system. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 25(34):7771–7779, August 2005.

[153] J. P. Donoghue, J. N. Sanes, N. G. Hatsopoulos, and G. Gaál. Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *Journal of Neurophysiology*, 79(1):159–173, January 1998.

[154] M M McCarthy, C Moore-Kochlacs, X Gu, et al. Striatal origin of the pathologic beta oscillations in Parkinson's disease. *Proceedings of the National Academy of Sciences of the United States of America*, 108(28):11620–11625, July 2011.

[155] Jonathan Cannon, Michelle M. McCarthy, Shane Lee, et al. Neurosystems: brain rhythms and cognitive processing. *The European Journal of Neuroscience*, 39(5):705–719, March 2014.

[156] Eduardo E. Benarroch. Effects of acetylcholine in the striatum Recent insights and therapeutic implications. *Neurology*, 79(3):274–281, July 2012.

[157] Yoland Smith, Dinesh Raju, Bijli Nanda, et al. The thalamostriatal systems: Anatomical and functional organization in normal and parkinsonian states. *Brain Research Bulletin*, 78(2–3):60–68, February 2009.

[158] Yoland Smith, Dinesh V. Raju, Jean-Francois Pare, and Mamadou Sidibe. The thalamostriatal system: a highly specific network of the basal ganglia circuitry. *Trends in Neurosciences*, 27(9):520–527, September 2004.

[159] Daniel K Leventhal, Gregory J Gage, Robert Schmidt, et al. Basal ganglia beta oscillations accompany cue utilization. *Neuron*, 73(3):523–536, February 2012.

[160] N. Yamawaki, I. M. Stanford, S. D. Hall, and G. L. Woodhall. Pharmacologically induced and stimulus evoked rhythmic neuronal oscillatory activity in the primary motor cortex in vitro. *Neuroscience*, 151(2):386–395, January 2008.

[161] Renato N. Watanabe and Andre F. Kohn. Fast Oscillatory Commands from the Motor Cortex Can Be Decoded by the Spinal Cord for Force Control. *The Journal of*

Neuroscience: The Official Journal of the Society for Neuroscience, 35(40):13687–13697, October 2015.

[162] Soroush Zaghi, Mariana Acar, Brittney Hultgren, Paulo S. Boggio, and Felipe Fregni. Noninvasive brain stimulation with low-intensity electrical currents: putative mechanisms of action for direct and alternating current stimulation. *The Neuroscientist: A Review Journal Bringing Neurobiology, Neurology and Psychiatry*, 16(3):285–307, June 2010.

[163] Nathan C. Rowland, Coralie De Hemptinne, Nicole C. Swann, et al. Task-related activity in sensorimotor cortex in Parkinson's disease and essential tremor: changes in beta and gamma bands. *Frontiers in Human Neuroscience*, 9:512, 2015.

[164] Pascal Fries. Rhythms for Cognition: Communication through Coherence. *Neuron*, 88(1):220–235, October 2015.

[165] N. Kopell, M. A. Whittington, and M. A. Kramer. Neuronal assembly dynamics in the beta1 frequency range permits short-term memory. *Proceedings of the National Academy of Sciences of the United States of America*, 108(9):3779–3784, March 2011.

[166] Matthew D. Best, Aaron J. Suminski, Kazutaka Takahashi, Kevin A. Brown, and Nicholas G. Hatsopoulos. Spatio-Temporal Patterning in Primary Motor Cortex at Movement Onset. *Cerebral Cortex*, page bhv327, January 2016.

[167] Afsheen Afshar, Gopal Santhanam, Byron M. Yu, et al. Single-trial neural correlates of arm movement preparation. *Neuron*, 71(3):555–564, August 2011.

[168] Mark M. Churchland and Krishna V. Shenoy. Delay of Movement Caused by Disruption of Cortical Preparatory Activity. *Journal of Neurophysiology*, 97(1):348–359, January 2007.

[169] A. P. Georgopoulos, J. F. Kalaska, R. Caminiti, and J. T. Massey. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 2(11):1527–1537, November 1982.

[170] Timothy P. Lillicrap and Stephen H. Scott. Preference Distributions of Primary Motor Cortex Neurons Reflect Control Solutions Optimized for Limb Biomechanics. *Neuron*, 77(1):168–179, January 2013.

[171] Mark M. Churchland, Gopal Santhanam, and Krishna V. Shenoy. Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *Journal of Neurophysiology*, 96(6):3130–3146, December 2006.

[172] D. W. Moran and A. B. Schwartz. Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology*, 82(5):2676–2692, November 1999.

[173] Charidimos Tzagarakis, Nuri F. Ince, Arthur C. Leuthold, and Giuseppe Pellizzer. Beta-Band Activity during Motor Planning Reflects Response Uncertainty. *Journal of Neuroscience*, 30(34):11270–11277, August 2010.

[174] Nicholas Hatsopoulos, Jignesh Joshi, and John G. O'Leary. Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles. *Journal of Neurophysiology*, 92(2):1165–1174, August 2004.

[175] Maxwell A. Sherman, Shane Lee, Robert Law, et al. Neural mechanisms of transient neocortical beta rhythms: Converging evidence from humans, computational modeling, monkeys, and mice. *Proceedings of the National Academy of Sciences of the United States of America*, 113(33):E4885–4894, August 2016.

[176] Gabriel Tan, John Thornby, D. Corydon Hammond, et al. Meta-analysis of EEG biofeedback in treating epilepsy. *Clinical EEG and neuroscience*, 40(3):173–179, July 2009.

[177] Simon Little and Peter Brown. The functional role of beta oscillations in Parkinson's disease. *Parkinsonism & Related Disorders*, 20, Supplement 1:S44–S48, January 2014.

[178] Joshua A. Goldberg, Thomas Boraud, Sharon Maraton, et al. Enhanced synchrony among primary motor cortex neurons in the 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine primate model of Parkinson's disease. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 22(11):4639–4653, June 2002.

[179] Benjamin Pasquereau and Robert S. Turner. Primary Motor Cortex of the Parkinsonian Monkey: Differential Effects on the Spontaneous Activity of Pyramidal Tract-Type Neurons. *Cerebral Cortex*, 21(6):1362–1378, June 2011.

[180] Nicole C. Swann, Coralie de Hemptinne, Svjetlana Miocinovic, et al. Gamma Oscillations in the Hyperkinetic State Detected with Chronic Human Brain Recordings in Parkinson's Disease. *Journal of Neuroscience*, 36(24):6445–6458, June 2016.

[181] G Schalk, K J Miller, N R Anderson, et al. Two-dimensional movement control using electrocorticographic signals in humans. *Journal of neural engineering*, 5(1):75–84, March 2008.

[182] A.-T. Avestruz, W. Santa, D. Carlson, et al. A 5 W/Channel Spectral Analysis IC for Chronic Bidirectional Brain #x2013;Machine Interfaces. *IEEE Journal of Solid-State Circuits*, 43(12):3006–3024, December 2008.

[183] Scott Stanslaski, Peng Cong, Dave Carlson, et al. An implantable bi-directional brain-machine interface system for chronic neuroprosthesis research. *Conference proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*, 2009:5494–5497, 2009.

[184] J. Herron, T. Denison, and H.J. Chizeck. Closed-loop DBS with movement intention. In *2015 7th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 844–847, April 2015.

[185] Elena Ryapolova-Webb, Pedram Afshar, Scott Stanslaski, et al. Chronic cortical and electromyographic recordings from a fully implantable device: preclinical experience in a nonhuman primate. *Journal of Neural Engineering*, 11(1):016009, February 2014.

[186] Philip A. Starr, Chadwick W. Christine, Philip V. Theodosopoulos, et al. Implantation of deep brain stimulators into the subthalamic nucleus: technical approach and magnetic resonance imaging-verified lead locations. *Journal of Neurosurgery*, 97(2):370–387, August 2002.

[187] Kiarash Shahlaie, Paul S. Larson, and Philip A. Starr. Intraoperative computed tomography for deep brain stimulation surgery: technique and accuracy assessment. *Neurosurgery*, 68(1 Suppl Operative):114–124; discussion 124, March 2011.

[188] P. Khanna, S. Stanslaski, Yizi Xiao, et al. Enabling closed-loop neurostimulation research with downloadable firmware upgrades. In *2015 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 1–6, October 2015.

[189] N. E. Crone, D. L. Miglioretti, B. Gordon, et al. Functional mapping of human sensorimotor cortex with electrocorticographic spectral analysis. I. Alpha and beta event-related desynchronization. *Brain: A Journal of Neurology*, 121 (Pt 12):2271–2299, December 1998.

[190] Wei Wu, Yun Gao, Elie Bienenstock, John P. Donoghue, and Michael J. Black. Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Computation*, 18(1):80–118, January 2006.

[191] P. Khanna and J.M. Carmena. Changes in reaching reaction times due to volitional modulation of beta oscillations. In *2015 7th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 340–343, April 2015.

[192] D. Bourget, H. Bink, S. Stanslaski, et al. An implantable, rechargeable neuromodulation research tool using a distributed interface and algorithm architecture. In *2015 7th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 61–65, April 2015.

[193] Tejas Sankar, Nir Lipsman, and Andres M. Lozano. Deep Brain Stimulation for Disorders of Memory and Cognition. *Neurotherapeutics*, 11(3):527–534, July 2014.

[194] Takashi Morishita, Sarah M. Fayad, Masa-aki Higuchi, Kelsey A. Nestor, and Kelly D. Foote. Deep Brain Stimulation for Treatment-resistant Depression: Systematic Review of Clinical Outcomes. *Neurotherapeutics*, 11(3):475–484, July 2014.

[195] Nealen G. Laxpati, Willard S. Kasoff, and Robert E. Gross. Deep Brain Stimulation for the Treatment of Epilepsy: Circuits, Targets, and Trials. *Neurotherapeutics*, 11(3):508–526, July 2014.

[196] Mark K. Lyons. Deep Brain Stimulation: Current and Future Clinical Applications. *Mayo Clinic Proceedings*, 86(7):662–672, July 2011.

[197] Dustin A. Heldman, Alberto J. Espay, Peter A. LeWitt, and Joseph P. Giuffrida. Clinician versus machine: reliability and responsiveness of motor endpoints in Parkinson's disease. *Parkinsonism & Related Disorders*, 20(6):590–595, June 2014.

[198] Ana Lisa Taylor Tavares, Gregory S. X. E. Jefferis, Mandy Koop, et al. Quantitative measurements of alternating finger tapping in Parkinson's disease correlate with UPDRS motor disability and reveal the improvement in fine motor control from medication and deep brain stimulation. *Movement Disorders: Official Journal of the Movement Disorder Society*, 20(10):1286–1298, October 2005.

[199] Dustin A. Heldman, Joseph P. Giuffrida, Robert Chen, et al. The modified bradykinesia rating scale for Parkinson's disease: reliability and comparison with kinematic measures. *Movement Disorders: Official Journal of the Movement Disorder Society*, 26(10):1859–1863, August 2011.

[200] Michael Cassidy, Paolo Mazzone, Antonio Oliviero, et al. Movement-related changes in synchronization in the human basal ganglia. *Brain: A Journal of Neurology*, 125(Pt 6):1235–1246, June 2002.

[201] B. S. Everitt. Factor analysis. In *An Introduction to Latent Variable Models*, Monographs on Statistics and Applied Probability, pages 13–31. Springer Netherlands, 1984. DOI: 10.1007/978-94-009-5564-6_2.

[202] Siddharth Dangi, Suraj Gowda, Helene G. Moorman, et al. Continuous Closed-Loop Decoder Adaptation with a Recursive Maximum Likelihood Algorithm Allows for Rapid Performance Acquisition in Brain-Machine Interfaces. *Neural Computation*, 26(9):1811–1839, June 2014.

[203] Steven M. Chase, Andrew B. Schwartz, and Robert E. Kass. Latent inputs improve estimates of neural encoding in motor cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(41):13873–13882, October 2010.

[204] Mark M. Churchland and Krishna V. Shenoy. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *Journal of Neurophysiology*, 97(6):4235–4257, June 2007.

[205] Shinji Kakei, Donna S. Hoffman, and Peter L. Strick. Muscle and Movement Representations in the Primary Motor Cortex. *Science*, 285(5436):2136–2139, September 1999.

[206] R. Caminiti, P. B. Johnson, C. Galli, S. Ferraina, and Y. Burnod. Making arm movements within different parts of space: the premotor and motor cortical representation of a coordinate system for reaching to visual targets. *Journal of Neuroscience*, 11(5):1182–1197, May 1991.

[207] Xiaofeng Lu and James Ashe. Anticipatory Activity in Primary Motor Cortex Codes Memorized Movement Sequences. *Neuron*, 45(6):967–973, March 2005.

[208] Nicholas G. Hatsopoulos, Liam Paninski, and John P. Donoghue. Sequential movement representations based on correlated neuronal activity. *Experimental Brain Research*, 149(4):478–486, April 2003.

[209] Nicholas G. Hatsopoulos, Qingqing Xu, and Yali Amit. Encoding of Movement Fragments in the Motor Cortex. *Journal of Neuroscience*, 27(19):5105–5114, May 2007.

[210] L. Paninski. Superlinear Population Encoding of Dynamic Hand Trajectory in Primary Motor Cortex. *Journal of Neuroscience*, 24(39):8551–8561, September 2004.

[211] D. J. Ostry, P. L. Gribble, and V. L. Gracco. Coarticulation of jaw movements in speech production: is context sensitivity in speech kinematics centrally planned? *Journal of Neuroscience*, 16(4):1570–1579, February 1996.

[212] H. Mushiake, M. Inase, and J. Tanji. Neuronal activity in the primate premotor, supplementary, and precentral motor cortex during visually guided and internally determined sequential movements. *Journal of Neurophysiology*, 66(3):705–718, September 1991.

[213] Andrea M. Green and John F. Kalaska. Learning to move machines with the mind. *Trends in Neurosciences*, 34(2):61–75, February 2011.

[214] Jonathan C. Kao, Paul Nuyujukian, Stephen I. Ryu, et al. Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nature Communications*, 6:ncomms8759, July 2015.

[215] Jonathan C. Kao, Stephen I. Ryu, and Krishna V. Shenoy. Leveraging neural dynamics to extend functional lifetime of brain-machine interfaces. *Scientific Reports*, 7(1):7395, August 2017.

[216] George Paxinos, Xu-Feng Huang, and Arthur W Toga. The Rhesus Monkey Brain in Stereotaxic Coordinates. April 2013.

[217] W. Q. Malik, W. Truccolo, E. N. Brown, and L. R. Hochberg. Efficient Decoding With Steady-State Kalman Filter in Neural Interface Systems. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(1):25–34, February 2011.

[218] H. G. Moorman, S. Gowda, and J. M. Carmena. Control of Redundant Kinematic Degrees of Freedom in a Closed-Loop Brain-Machine Interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(6):750–760, June 2017.

[219] M.M. Shanechi, G.W. Wornell, Z.M. Williams, and E.N. Brown. Feedback-Controlled Parallel Point Process Filter for Estimation of Goal-Directed Movements From Neural Signals. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 21(1):129–140, January 2013.

[220] Maryam M. Shanechi, Amy L. Orsborn, Helene G. Moorman, et al. Rapid control and feedback rates enhance neuroprosthetic control. *Nature Communications*, 8:ncomms13825, January 2017.

[221] Maryam M. Shanechi, Amy L. Orsborn, and Jose M. Carmena. Robust Brain-Machine Interface Design Using Optimal Feedback Control Modeling and Adaptive Point Process Filtering. *PLoS computational biology*, 12(4):e1004730, April 2016.

[222] Zoubin Ghahramani and Geoffrey E. Hinton. Parameter Estimation for Linear Dynamical Systems. Technical report, 1996.

[223] Peter Pastor, Mrinal Kalakrishnan, Franziska Meier, et al. From dynamic movement primitives to associative skill memories. *Robotics and Autonomous Systems*, 61(4):351–361, April 2013.

[224] T. E. Milner. A model for the generation of movements requiring endpoint precision. *Neuroscience*, 49(2):487–496, July 1992.

[225] A. S. F. Leyton and C. S. Sherrington. Observations on the Excitable Cortex of the Chimpanzee, Orang-Utan, and Gorilla. *Quarterly Journal of Experimental Physiology*, 11(2):135–222, July 1917.

[226] Morasso, P. and Mussa Ivaldi, F. A. Trajectory Formation and Handwriting: A Computational Model. *Biological cybernetics*, 45:131–142, 1982.

[227] Juan A. Gallego, Matthew G. Perich, Lee E. Miller, and Sara A. Solla. Neural Manifolds for the Control of Movement. *Neuron*, 94(5):978–984, June 2017.

[228] John P. Cunningham and Byron M. Yu. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509, November 2014.

[229] Uri Rokni, Andrew G. Richardson, Emilio Bizzi, and H. Sebastian Seung. Motor learning with unstable neural representations. *Neuron*, 54(4):653–666, May 2007.

[230] S. P. Wise, S. L. Moody, K. J. Blomstrom, and A. R. Mitz. Changes in motor cortical activity during visuomotor adaptation. *Experimental Brain Research*, 121(3):285–299, August 1998.

[231] Andrew J. Peters, Simon X. Chen, and Takaki Komiyama. Emergence of reproducible spatiotemporal activity during motor learning. *Nature*, 510(7504):263–267, June 2014.

[232] William A. Liberti, Jeffrey E. Markowitz, L. Nathan Perkins, et al. Unstable neurons underlie a stable learned behavior. *Nature Neuroscience*, 19(12):1665–1671, December 2016.

[233] Rony Paz and Eilon Vaadia. Learning-Induced Improvement in Encoding and Decoding of Specific Movement Directions by Neurons in the Primary Motor Cortex. *PLoS Biol*, 2(2):e45, February 2004.

[234] C. S. Li, C. Padoa-Schioppa, and E. Bizzi. Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron*, 30(2):593–607, May 2001.

[235] Peiran Gao, Eric Trautmann, Byron M. Yu, et al. A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv*, page 214262, November 2017.

[236] Chethan Pandarinath, Daniel J. O'Shea, Jasmine Collins, et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *bioRxiv*, page 152884, June 2017.

[237] Zoubin Ghahramani and Geoffrey E. Hinton. The EM Algorithm for Mixtures of Factor Analyzers. Technical report, 1997.

Appendix A

Factor Analysis

A.1 Fitting Factor Analysis (FA)

A.1.1 Overview of Factor Analysis (FA)

This section will outline the Factor Analysis (FA) model and the manner in which we fit Factor Analysis models for data analysis in chapters 5 and 6. FA is a method that can be applied as a dimensionality reduction tool, much like principal component analysis (PCA), probabilistic principal component analysis (PPCA), or independent component analysis (ICA). Each method is different. PCA aims to reduce the dimensionality of a high-dimensional dataset with a low-dimensional set of orthogonal axes that capture the maximum variance. These axes, or loading vectors, turn out to be equivalent to the eigenvectors of $X^T X$where X is a Txn matrix of n-dimensional data observed T times after subtracting the mean $\mu(\mu \in R^n$ from each entry of $X)$.

PPCA is similar to PCA, but can be described by a probabilistic generative model:

$$n_t = Uz_t + \psi_t + \mu$$

$$z_t \sim N(0, I)$$

$$\psi_t \sim N(0, \sigma^2 I)$$

Each observation n_t is described as a linear sum of a low-dimensional latent variable z_t multiplied by a constant loading matrix U, plus a noise term ψ_t. In PPCA, ψ_t is draw from a nomal distribution with covariance equal to a diagonal matrix $\sigma^2 I$, such that each entry of the diagonal is identical. In contrast, FA models ψ_t as drawn from a normal distribution with covariance equal to diagonal matrix Ψ where each entry of the diagonal can differ. In both PPCA and FA, model parameters are fit using expectation-maximization [237].

In all work described in chapter 5 and chapter 6, an FA model is estimated from neural data using the python sklearn implementation of FA (http://scikit-learn.org/stable/modules/generated/sklearn.decomposition.FactorAnalysis.html). The following workflow was used to estimate the FA model:

1. Choose FA dataset. In chapter 5 it would be the visual feedback calibration session for the simulated online shared BMI control, or the resultant population activity to parse the observed shared vs. private neural variability. In chapter 6 it would be the 16-trial epochs of online BMI performance for the subspace-overlap analysis, or the 64 trial block of online BMI data for the online shared BMI control.

2. Aggregate the binned neural data from all correct trials from training dataset and concatenate them into a matrix $N \in R^{Txn}$ where T is the number of bins across all trials and $n$is the number of output neurons.

3. Randomly select 90% of points to be used for training FA, yielding $N_{train} \in R^{(0.9*T)xn}$ and $N_{test} \in R^{(0.1*T)xn}$

4. For a given k(dimensionality of z) initialize and fit an FA model:

   (a) Initialize FA with $z \in R^k$

   (b) Fit FA with N_{train} . The number of times a fit is attempted is 3, max number of EM iterations is 1000, stopping tolerance for EM is 0.01 (all defaults for sklearn.decomposition.FactorAnalysis method).

5. Compute the loglikelihood of N_{test} using the fit FA model. Store it.

6. Repeat steps 4, 5 for different values of k(usually $1 \leq k \leq 10$).

7. Repeat steps 3, 4, 5, 6 for different random subselections of N_{test}, N_{train} (usually 5 times).

8. Plot mean loglikelihood for FA models for each k(averaged over different iteraitons, different subselections of N_{test}, N_{train}).

9. Select k with maximum log likelihood.

10. Fit a new FA model with dimensionality of $z$equal to k, and with all data N. This is the final FA model.

A.2 Main Shared vs. Shared Variability

Chapter 6 refers to "Main Shared Variability" instead of just "Shared Variability". Main shared variability is extracted after fitting FA. The shared signals were ordered by their variance by diagonalizing via the singular value decomposition (SVD) [88]:

$$u, s, v = np.linalg.svd(UU^T)$$

The top p singular values that were needed to describe > 0.9 of the total variance were used, and the remainder were set to zero:

$$s_{reduced} = [s[1:p], zeros(1, k-p)]$$

$$UU^T_{main} = u * s_{reduced} * v$$

Also required in the computation of the main shared subspace overlap is the projection matrix on to the main shared space. This constitutes a matrix that takes a point and projects it onto the subspace defined by UU^T_{main}:

$$s_{reduced,proj} = [ones(1,p), zeros(1,k-p)]$$

$$P_{main} = u * s_{reduced,proj} * v$$

Throughout the chapter "main shared" variance refers to the use of UU^T_{main} instead of UU^T. Further, dimensionality of "main shared" variance refers to p instead of k, though note that $p \leq k$. Finally, the projection onto the main shared space is defined as P_{main}. These methods are also reviewed in the supplement of [90].

A.3 Estimating z_t given n_t

In chapter 5 and chapter 6, individual bins of population activity are decomposed into private and shared components. Below we review how this decomposition emerges:

A.3.1 Conditional Distribution of Multivariate Gaussians

Consider a multivariate Gaussian:

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} \sim N(\begin{bmatrix} \mu_x \\ \mu_y \end{bmatrix}, \begin{bmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{xy}^T & \Sigma_{yy} \end{bmatrix})$$

Now say we want to estimate the conditional distribution for y_t given a particular value x of x_t: $p(y_t|x_t = x)$. Because of the awesome rules of multivariate Gaussians, the conditional distribution is also Gaussian: $p(y_t|x_t = x) = N(\mu_{y|x}, \Sigma_{y|x})$:

$$\mu_{y|x} = \mu_y + \Sigma_{yx}\Sigma_{yy}^{-1}(x - \mu_x)$$

$$\Sigma_{y|x} = \Sigma_{yy} - \Sigma_{yx}\Sigma_{xx}^{-1}\Sigma_{xy}$$

A.3.2 Factor Analysis Model Estimates

Now, turning our attention to the factor analysis model:

$$\begin{bmatrix} n_t \\ z_t \end{bmatrix} \sim N(\begin{bmatrix} \mu \\ 0 \end{bmatrix}, \begin{bmatrix} \Sigma_{nn} & \Sigma_{nz} \\ \Sigma_{nz}^T & \Sigma_{zz} \end{bmatrix})$$

We compute:

$$\Sigma_{zz} = E((z_t - 0)(z_t - 0)^T) = E(z_t z_t^T) = I$$

$$\Sigma_{nn} = E((n_t - \mu)(n_t - \mu)^T = E((Uz_t + \psi_t + \mu - \mu)(Uz_t + \psi_t + \mu - \mu)^T)$$

$$= E(Uz_t z_t^T U^T + Uz_t\psi_t^T + \psi_t z_t^T U^T + \psi_t\psi_t^T) = UU^T + \Psi$$

$$\Sigma_{nz} = E((n_t - \mu)(z_t - 0)^T) = E((Uz_t + \psi_t + \mu - \mu)z_t^T) = E(Uz_t z_t^T + \psi_t z^T) = UE(z_t z_t^T) = UI = U$$

Thus:

$$\begin{bmatrix} n_t \\ z_t \end{bmatrix} \sim N(\begin{bmatrix} \mu \\ 0 \end{bmatrix}, \begin{bmatrix} UU^T + \Psi & U \\ U^T & I \end{bmatrix})$$

Finally, using the formula from above, the conditional distribution of $p(z_t|n_t = n) = N(\mu_{z|n}, \Sigma_{z|n})$ can be computed:

$$\mu_{z|n} = \mu_z + \Sigma_{zn}\Sigma_{zz}^{-1}(n - \mu)$$

$$\mu_{z|n} = U^T(UU^T + \Psi)^{-1}(n - \mu)$$

$$\Sigma_{z|n} = I - U^T(UU^T + \Psi)^{-1}U$$

Thus, the shared estimate of $n_{t,shared}$ is $E(n_t|\hat{z}_t, n_t = n) = UE(z_t|n_t = n) = UU^T(UU^T + \Psi)^{-1}(n - \mu)$

And the private estimate is $n_t - E(n_t|\hat{z}_t, n_t = n)$.