UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Neural Correlates of Learning of Brain-Machine Interfaces

Permalink
https://escholarship.org/uc/item/6zf6q4bc

Author
You, Albert King

Publication Date
2020

Peer reviewed|Thesis/dissertation

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

Neural Correlates of Learning of Brain-Machine Interfaces

by

Albert King You

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Joint Doctor of Philosophy
with the 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 Anca Dragan
Professor Christoph E. Schreiner

Spring 2020

Neural Correlates of Learning of Brain-Machine Interfaces

Copyright 2020
by
Albert King You

Abstract

Neural Correlates of Learning of Brain-Machine Interfaces

by

Albert King You

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

University of California, Berkeley

Professor Jose M. Carmena, Chair

The brain has an incredible capacity to learn how to control various effectors, ranging from those endogenous the body to those that are artificially implanted. Moreover, changes to these effectors can be readily adapted by the brain, albeit on varying timescales depending on the amount of perturbation. For example, many studies have shown adaptation to force perturbations or visual-motor rotations to occur within minutes, attributing the adaptation to inputs from the cerebellum. This flexible adaptation of the brain makes the use of brain-machine interfaces (BMIs) an attractive rehabilitative option for a variety of motor pathologies such as stroke and amyotrophic lateral sclerosis. BMIs allow users to interact with their environments by using signals recorded directly from the brain and transforming them into actions taken on by an external effector such as a computer cursor or robotic arm.

Past work has shown neurons to adapt to decoders over time, suggesting BMIs as an exciting paradigm for neural rehabilitation. Modern decoding methods involve using closed-loop decoder adaptation (CLDA) allowing for decoder weights to converge quickly, and users to gain high levels of control with very little training. Furthermore, these weights can be updated over learning in the event there are changes in the neural population. While great for assistive devices, it is unclear how these decoders impact neural adaptation over time. To answer this question, we conducted four studies

examining the interplay between decoder adaptation and changes in neural activity over learning.

Previous studies in the field have shown neural activity to fire in increasingly correlated patterns, converging onto low-dimensional spaces as performance with a BMI improves, indicating consolidation of coordinated firing patterns. However, these studies were conducted with static decoders. To examine how neurons adapt alongside CLDA, we used nonhuman primate models to observe how neurons in motor cortex adapt when decoder weights change in a BMI task. We found that there is a concomitant exploration-exploitation strategy that occurs where changes in decoder weights yield increases in exploratory neural activity, which in turn consolidates over many days. These results are remarkably paralleled by those in the psychophysics literature, suggesting similar mechanisms of adaptation between the two regimes. Furthermore, we found that initial levels of decoder adaptation bear no impact on the rate of neural adaptation. That is, decoders resulting in better initial performance of a BMI did not affect how quickly neurons adapted over multiple days.

However, we found that there are differences in adaptation between neurons used as decoder inputs compared to other neurons in the local supporting network. While neurons that were delegated as decoder inputs explored and eventually collapse onto a low-dimensional space, neurons in the supporting local network had comparatively less exploration and consolidation of activity. Strikingly though, these support neurons show task-relevant behavior throughout learning. In combination, these results suggest that neural adaptation to BMIs is driven by network solutions that aim to improve performance through exploring new patterns rather than suppressing preexisting ones. Rather than improving control of the BMI by decreasing noise, motor cortex seems to boost the signal quality instead. Together, the results in this body of work highlight mechanisms of neural adaptation to neuroprosthetic devices that mirror those of natural motor adaptation. Exploiting these mechanisms may lead to better methods of utilizing BMIs for neural rehabilitation.

To Mom and Dad

Contents

Contents ii

Chapter 1 Introduction 1

1.1 The Brain-Machine Interface 1

1.2 Using Brain-Machine Interfaces to Study Neural Circuits 3

1.3 Neural Adaptation to BMIs 4

1.4 BMIs as 2-Learner Systems 4

1.5 Chapter Previews 5

1.5.1 Flexible Neural Variance Facilitates Skill Learning 5

1.5.2 The Rate of Neuroprosthetic Learning is Invariant to Initial Levels of Closed-Loop Decoder Adaptation 5

1.5.3 Local Network Adaptation Underlying Neuroprosthetic Learning 6

1.5.4 Neural Correlates of Control of a Kinematically Redundant BMI 6

1.6 Neural Correlates of Adaptation to Brain-Machine Interfaces 6

Chapter 2 Flexible Neural Variance Facilitates Skill Learning 7

2.1 Abstract 7

2.2 Introduction 8

2.3 Results 9

2.3.1 Factor Analysis Predicts Flexible Modulation of Consolidated Patterns 11

2.3.2 Neural Variance Flexibly Changes 13

2.3.4 Neural Exploration Adapts Consolidated Patterns 15

2.3.5 Modulation Depth Facilitates Changes in Neural Variance 18
2.4 Discussion 19
2.5 Methods 21
2.5.1 Surgical Procedures and Electrophysiology 21
2.5.2 BMI Task 22
2.5.3 Kalman Filter 23
2.5.4 Factor Analysis 24
2.5.5 Preferred Direction and Modulation Depth 28
2.5.6 Statistical Analyses and Testing 28
Chapter 3 The Rate of Neuroprosthetic Learning is Invariant to Initial Levels of Closed-Loop Decoder Adaptation 29
3.1 Abstract 29
3.2 Introduction 30
3.3 Results 31
3.3.1 Initial decoder assistance level does not impact neuroprosthetic learning rates 33
3.3.2 Neural adaptation underlying BMI learning with different levels of initial CLDA 35
3.3.3 Coordinated neural plasticity over training is independent of initial CLDA 38
3.3.4 Neural "reweighting" following decoder changes 40
3.4 Discussion 41
3.5 Methods 42
3.5.1 Electrophysiology 42
3.5.2 Behavior 43
3.5.3 BMI implementation 43
3.5.4 Closed-loop decoder adaptation (CLDA) and SmoothBatch 43
3.5.5 Analysis 44
3.5.6 Factor Analysis 45
3.5.7 General linear regression model with categorical covariates and statistical tests 46

Chapter 4 Local-Network Adaptation Underlying Abstract Skill Learning 47
4.1 Abstract 47
4.2 Introduction....48
4.3 Results....49
4.4 Discussion....57
4.5 Methods....59
4.5.1 Animal Subjects 59
4.5.2 Electrophysiology 59
4.5.3 Experimental Setup and Behavioral Training 60
4.5.4 BMI Tasks....60
4.5.5 Preprocessing Pipeline....60
4.5.6 Factor Analysis 61
Chapter 5 Neural Correlates of Control of a Kinematically Redundant BMI 63
5.1 Abstract 63
5.2 Introduction....64
5.3 Results....64
5.3.1 Endpoint Trajectories Become More Stereotyped Over Learning....64
5.3.2 State-Space Linearization 67
5.3.2 Neurons Form a Stable Map to a Converging Decoder 71
5.4 Discussion....73
5.5 Methods....75
5.5.1 Experimental Setup....75
5.5.2 Analysis Pipeline 76
5.5.3 Factor Analysis 76
Chapter 6 Discussion 77
6.1 Neural adaptation in 2-learner systems 77
6.2 Mechanisms of neural adaptation 78
6.3 Open questions 79

6.4 Dissertation Contributions and Publications 80
Bibliography 82

Acknowledgments

It goes without saying that my PhD could not have been possible without the support from those around me as well as those who have come before me. My time here in grad school has been the most transformative years of my life thus far, both in research and as a person. I have made lifelong friends in the process with whom I owe this milestone to.

First, thanks to all the Carmenites from years past. I am standing on the shoulders of giants, and not a day goes by where I forget that. Suraj, you were the first person I worked with in lab, and I cherish those moments we had working with Gromit. I wish we could have overlapped more in the lab, but am so, so, so grateful that I still get to discuss results with you and ask for help years after you've graduated. Thank you also to Karunesh for the most beautiful dataset that we still analyze a decade later in addition to serving on my Qualifying Exam committee, as well as helpful discussions for my projects.

Next, thank you to Sam, Preeya, Nuria, and Ryan for being so welcoming when I first started. You guys set the bar so high; I knew right away I would never be able to fill those shoes. Sam, without your guidance, I would not have made it past Year 1. Thanks for believing in me at my lowest. Preeya, thank you so much for being such a fantastic mentor all these years. You're unfathomably intelligent, hardworking, and patient – all traits I can only hope to emulate. Nuria, lab is always a little brighter with you around (stop leaving us!). We still need to find something to play on Steam. Also, thanks for joining the lab a couple months before me, thus absolving me of any senior lab member duties, hehe. Vivek! I wish you came back from Portugal sooner. You were always the catalyst for intensely funny conversations and your extroversion is contagious. I felt so much more open after meeting you. Thank you also for being such a great mentor, both for analysis methods, which my PhD largely follows, but also in life. I miss working with you.

Tanner and David, I don't know what I'd do without you guys. I'm so happy we were able to meet before even getting to Berkeley. Thanks for being my rubber duckies during my bouts of complaining and stress. I'm so humbled to be able to call myself your cohort- and lab-mates. You've been such great friends both in and out of lab. Tanner, thanks for always including me in Thanksgiving and Christmas gatherings. Being away from home can be lonely sometimes; you have no idea how much those moments have meant to me. David, you're the best late-night work partner. I always enjoyed working at night knowing that you were around if I needed a high-five.

Ellen, thanks for being one of my closest friends and snack buddies in the lab (despite you rejecting my FB friend request your first year). I've enjoyed being able to commiserate with you about everything. More importantly, I'm thankful that we were able to work on an awesome project together. You're so easy to work with and I'm so proud of how much you've grown since joining the lab. I wish you much success with everything you do!

Maki, it's hard to imagine going back to a time without you here. The sheer amount of work you do for the lab is astounding; how you manage to keep track of everything baffles me. Aside from work, I feel like we became fast friends and I'll miss :yifangtime:. Thank you also for introducing me to bouldering (with Ema). I definitely found a new lifelong hobby because of you.

Finally, Jose, you're the best PI and mentor anyone could hope for. I know I wasn't easy my first years - thank you for not giving up on me. I feel proud of how our relationship has evolved over the years and I'm lucky to be able to call myself a Carmenite. Thank you for all the guidance and support you've given me. I don't know any other PI who would be willing to meet on Christmas day. I remember the first time I explicitly came to you for help: "How do I ask a scientific question?". You gave a reassuring, yet oddly unhelpful answer at the time. You said, "You just need to keep struggling and eventually you'll figure it out. It's a process and there isn't anything really more I can say that can help. But you'll figure it out." Years later, I think about that moment and I finally understand what you meant.

To the rest of Carmena lab, thank you for the friendships, discussions, and memories. Manu, it has been such a pleasure working with you on projects. Thank you for everything you've done. I hope you continue to work hard-play hard with whatever you choose to do after college. Qbert and Gabby, I know you'll continue to shine in grad school. Proud of all of you. Of course, none of my work would be possible without the hardest working members of the lab - our furry friends. Thank you, Paco, Rico, Jeeves, Sebastian, Cartman, and Gromit for such beautiful data we still analyze years later. Thank you also to Gromit and Luigi for being my first monkeys and for being so patient during my training process. Thank you to Yoda and Obi-Wan for being the cutest little babies to train. And thank you to Wallace for being such a hard worker and being Gromit's friend. Finally, Keller, I love you and will always miss you.

My PhD experience was successful not only because of the Carmena lab family, but also because of the friendships I made through various hobbies and clubs I found at Cal. Thanks to the Soggies for letting me play Dota with you and BATS and Axel for introducing me to badminton. I've made lifelong friends because of both. Romain, Atish, and Maring - I feel so lucky to call you guys my best friends. To Romain, thanks for

always being there for me. Love you, brother. Maring, thanks for picking me up from the airport at 6 am multiple times. We should get McDonald's at 7 am again sometime. Atish, thanks for being a great mentor and friend. You sure do make Stanford look good. I'll miss getting Saul's with you and Sophia. To BEAST, Kristin, and Cohort 2015, thank you for the amazing support and friendships.

Lastly, thank you to my family. Kaitlin, you've been my rock. Thank you so much for everything you've done for me and with me. I've vented so many times to you. I'm so grateful to have you in my life (and Churro and Tiny!). Mom and Dad, thank you for raising me to be who I am. I'm so proud to be your son. This dissertation is for you. Thank you.

Chapter 1

Introduction

Nearly one in six people in the world suffer from neurological disorders including strokes, Parkinson's disease, mood disorders, traumatic brain injuries, and more [1]. Many of these disorders affect patients' motor functions. In stroke, for example, it is common for patients to suffer from hemiparesis following a stroke due to lesions in motor cortex along with subcortical regions. Other disorders such as amyotrophic lateral sclerosis (ALS) affect the nerves innervating muscles, weakening physical function. In many of these cases (along with other cases such as amputees), restoring motor function is a current technological challenge. While the efferent connections may be disrupted in these patients, often times the afferent pathways are still intact, meaning that movement of the limbs can still be sensed even though those movements cannot be self-generated. Thus, upstream anatomical structures are promising target locations for developing therapeutic and prosthetic devices.

1.1 The Brain-Machine Interface

Brain-machine interfaces (BMIs) are a class of neuroprosthetic device that aim to allow users to interact with their external environment directly with signals recorded from the brain [2]–[8]. Such a device allows users to control external effectors such as robotic arms, wheelchairs, or computer cursors so that movement of endogenous effectors is not required. Typically, these devices act in a closed-loop system where users receive feedback on the current state of the effector they are controlling and can modulate their neural activity to correct for any errors. An example is shown in Figure 1.

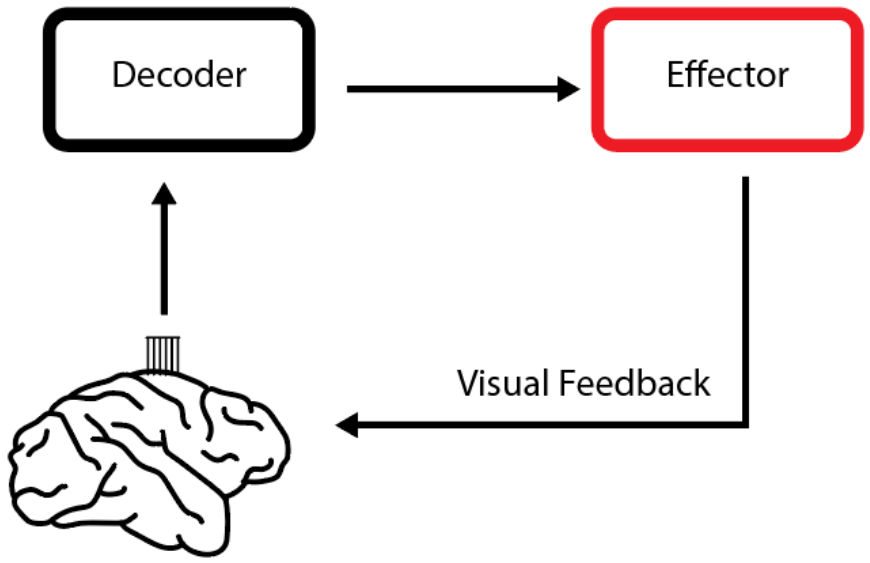


Figure 1 Closed-loop BMI. A typical BMI paradigm involves recording neural activity and translating that activity into movements of an effector (e.g. artificial arm) using a decoder. Many BMI experiments rely on visual feedback to close the loop. However, stimulation can also be given via an encoder to give a sense of tactility or proprioception.

Various effectors can be used depending on the application of the BMI. For example, an artificial arm would make sense for an amputee while a keyboard or mouse might be important for an ALS patient. For the most part, the loop is closed via visual feedback. However, other sensory feedback modalities are also possible, including touch and sound [9]–[11]. Furthermore, an additional component can be added to the loop above; an encoder can be used to translate effector states (e.g. position or velocity) into neurostimulation patterns [12]. These paradigms allow for closed-loop stimulation BMI paradigms, where the effector is a neural state a subject is explicitly controlling.

Modern decoders tend to be varying flavors of Kalman filters though other methods such as Wiener filters and support vector machines are sometimes used [3], [6], [13]–[19]. These decoders are typically linear models that are fast to run online and easy to interpret. More complex decoders using recurrent neural networks are also under investigation [20] but will not be a main focus of this body of work.

Finally, the neural signals recorded depends greatly on the application for the BMI. For example, control of external effectors that require movements involve signals recorded from motor cortex. However, it has been shown that other brain areas can be surrogates for decoding, such as primary visual cortex coupled with auditory feedback [9]–[11]. The type of signal can also depend on many factors. For instance, noninvasive techniques, such

as using electroencephalography (EEG) are popular with human experiments since they require no surgical procedure [21]–[23]. However, these signals have relatively low signal to noise ratio and cannot interrogate deeper brain structures. Other techniques involve more invasive procedures, often requiring a surgical technique to lower electrodes intracortically. These electrodes penetrate the dura and offer fantastic signal to noise ratio. Furthermore, they allow experimenters to choose between using single neuron spiking activity or local field potentials for decoding. In all, while the concept of a BMI is concise, many considerations need to be made to maximize the efficiency and efficacy of these devices; these define the motivation for this body of work.

1.2 Using Brain-Machine Interfaces to Study Neural Circuits

While many subfields have attempted to study how our brains work on a network level, one challenge that remains is the ability to determine what the roles of specific cells are. For example, psychophysics experiments, particularly those involving populations with variable pathologies have shed light on which neural structures are important for certain tasks. Other fields have taken a more molecular approach, attempting to map out the entire connectome. These types of experiments have expanded neuroscience's understanding of general principles of how the brain functions but offer very little in the scope of neural learning and neural adaptation.

BMIs are a fantastic paradigm to model neural learning because we have the ability to specify which cells are important for control by delegating them as inputs into the decoder. We can define *de novo* effectors and track how neurons adapt as control improves, thereby allowing us to correlate changes in neural activity with behavioral outputs. These experiments are often done as proxies to natural motor adaptation, particularly when recorded units are chosen in motor cortex and related areas [5], [6], [24]–[28].

Of course, one limitation with BMIs is that we are still unable to observe what happens to the rest of the network. While we can select which cells act as inputs for effector control, those cells are not isolated in the brain. The supporting network structure may also adapt, and it is unclear what upstream and downstream structures are doing during this process. Groups have begun more audacious experiments, such as using wide-field calcium imaging coupled with BMI to track how thousands of neurons adapt [9]. However, the same problem remains – there are still structures hidden during these learning processes that we cannot observe.

1.3 Neural Adaptation to BMIs

Past work has demonstrated varying modes of neural adaptation over BMI learning. Original studies showed changes in the tuning directions of neurons as task performance improved [5], [6]. Furthermore, it has been shown that the depth of the tuning direction curves (i.e. modulation depth) increases over learning, particularly in direct neurons [6], [15], [16], [29]. Perhaps most remarkably, these changes were reversible and specific to the decoder [6], [29], indicating BMI learning as an association to a cortical state that can rapidly change [29], [30].

More recent work has looked specifically at how neurons adapt on what we refer to as "low-dimensional" neural spaces or manifolds [4], [26]–[28], [31], [32]. Neurons tend to fire in concert and therefore, we can typically characterize a population of n neurons' firing activity in k dimensions, where $k < n$. Note that this is true for both BMI and natural motor control, and depending on the task, k can sometimes be an order of magnitude smaller than the full neural space [33]. Work specific to the BMI field has shown neurons to collapse onto these low-dimensional spaces as performance improves, indicating consolidation of learned patterns [26]. Intuitively, as more neural variance is captured in these low-dimensional spaces, the indication is that the population activity can be approximated with fewer dimensions, hence consolidation. Furthermore, neurons have been shown to adapt more readily to control spaces (i.e. decoders) that lie within the same manifold as the neural activity and require more time to adapt to control spaces that diverge from those manifolds [27], [32]. Given enough time, those low-dimensional spaces can rotate, causing an increased alignment between the neural space and the control space [26], [32]. While these results provide valuable insight on driving forces behind neural adaptation in certain BMI contexts, it is unclear how they apply to BMI tasks where the decoder periodically changes, albeit within manifold.

1.4 BMIs as 2-Learner Systems

One framework to pose the questions of interest here is to consider BMIs as "2-learner" systems. That is, allowing the decoder and the brain to co-adapt as performance improves [16]. The idea of a 2-learner system is powerful because it loosens the constraints on certain features of the decoder. For example, neurons routinely drift and drop out over many weeks and months of recording. A static decoder would have a hard time dealing with these changes, leading to errors in the effector the subject is attempting to control. By refitting the decoder weights periodically, we can get a closer estimate of the neural intentions at any given point in time.

These decoder refittings are performed via closed-loop decoder adaptation (CLDA) [15], [34], [35]. There are multiple methods of adapting these decoder weights, but they all act in similar fashions. CLDA aims to estimate decoder weights based on limited neural activity. The algorithms can then run online for 10-15 minutes and converge while the subjects are attempting to perform the task. Over time, the weights change until a stable decoder is found. Note however, that this is not what we refer to as a 2-learner system in this context since we do not expect neurons to adapt on such short timescales. Rather, we refer to 2-learner systems as those where decoders change on the order of days rather than minutes. That is, a decoder may be used and held static for a series of days before a new decoder is fit. However, using CLDA allows us to reseed these decoder weights much faster and more accurately than using more traditional methods. While previous work has shown neurons to form neuroprosthetic maps in 2-learner contexts [36], little is known about how neurons might consolidate on low-dimensional manifolds and what the driving forces behind those changes might be.

1.5 Chapter Previews

1.5.1 Flexible Neural Variance Facilitates Skill Learning

We begin in Chapter 2 by examining how neurons adapt on varying timescales over learning given periodic changes in decoder weights. In particular, we use factor analysis to parse out the neural variance in terms of correlated population activity and neural variance unique to each unit. We find that neurons increase population-level coordination over learning, indicating consolidation of firing patterns. However, changes in decoder weights cause shorter timescale increases in neural variance without decreasing correlated patterns. Our results indicate a concomitant exploration-exploitation strategy underlying neuroprosthetic learning.

1.5.2 The Rate of Neuroprosthetic Learning is Invariant to Initial Levels of Closed-Loop Decoder Adaptation

In Chapter 3, we investigate more closely how neurons adapt alongside decoder changes. We find that initial decoder performance does not affect adaptation rates of neurons and consolidation rates are similar regardless of initial BMI performance. These results highlight the use of well-fit decoders from the onset of learning to improve effector control. In a clinical setting, this would allow for BMIs to be used as both assistive and rehabilitative devices.

1.5.3 Local Network Adaptation Underlying Neuroprosthetic Learning

While many studies have examined what happens to the decoder input neurons over learning, much is still left to be understood about other neurons in the network. We tackle this question by analyzing other locally recorded neurons during BMI learning (indirect neurons) and find that while both direct and indirect neurons adapt, they do so differently. Direct neurons tend to explore the neural space more and consolidate over learning while indirect neurons suppress activity. Notably, both subpopulations remain task relevant throughout learning. These results suggest that neurons adapt coordinated firing rate patterns to boost the biological signal-to-noise ratio for control.

1.5.4 Neural Correlates of Control of a Kinematically Redundant BMI

Lastly, we ask how neurons adapt in a more complex task involving controlling a kinematically redundant effector in 2D space. We find that animals are able to approximate control of the effector with degrees of freedom, in part due to the decoder architecture. However, results suggest that the brain estimates a state-space linearization. Despite the complexity of the task, control is relatively low-dimensional and we find that neurons consolidate similarly to other kinematically deterministic tasks.

1.6 Neural Correlates of Adaptation to Brain-Machine Interfaces

The studies in this work highlight the brain's amazing ability to adapt to various types of decoders and showcase its ability to learn abstract skills. Our results show various mechanisms that allow the brain, especially in cortex, to undergo the changes we observe, and offer explanations for how the brain derives network solutions that are robust to changes in the decoder. The parallels between our results and those found in natural motor control suggests a common neural strategy that allows neural adaptation to abstract effectors. These results show the propensity of the brain to efficiently find solutions for control, shaping how we can both utilize BMIs to study natural motor control as well as leverage neural adaptation mechanisms to construct future BMIs.

Chapter 2

Flexible Neural Variance Facilitates Skill Learning

2.1 Abstract

One hallmark of natural motor control is the brain's ability to adapt to perturbations ranging from temporary visual-motor rotations to paresis caused by stroke. These adaptations require modifications of learned neural patterns that can span the time-course of minutes to months. Previous work with brain-machine interfaces (BMI) has shown that over learning, neurons consolidate firing activity onto low-dimensional neural subspaces, and additional studies have shown that neurons require longer timescales to adapt to task perturbations that require neural activity outside of these subspaces. However, it is unclear how the motor cortex adapts alongside task changes that do not require modifications of the existing neural subspace over learning. To answer this question, five nonhuman primates were used in three BMI experiments, which allowed us to track how specific populations of neurons changed firing patterns as task performance improved. In each experiment, neural activity was transformed into cursor kinematics using decoding algorithms that were periodically readapted based on natural arm movements or visual feedback. We found that decoder changes caused neurons to increase exploratory-like patterns on within-day timescales without hindering previously consolidated patterns regardless of task performance. The flexible modulation of these exploratory patterns in contrast to relatively stable consolidated activity suggests a simultaneous exploration-exploitation strategy that adapts existing neural patterns during learning.

2.2 Introduction

Our brains have an intrinsic ability to adapt to perturbations. However, the degree to which we can adapt depends on the similarity to previously learned behaviors and the time allowed for adaptation [37]–[40]. Brain-machine interfaces (BMI) allow us to interrogate specific populations of neurons and examine how they adapt during learning [3], [4], [41]. BMI decoders translate neural activity ranging from firing rates to field potentials into actions taken by an end effector (e.g. computer cursor or robotic arm). It has been shown in multiple studies that the brain adapts and forms stereotyped firing patterns over learning [5]–[7], [10], [16], [42]–[44]. We refer to these stereotyped patterns as neuroprosthetic maps [6], [16], [29]. More recent work has shown how the brain coordinates the activity of neurons over learning and collapse upon a low-dimensional neural subspace as performance improves when given a fixed decoder [26], [32].

In addition, similar results regarding neuroprosthetic map formation have been shown in BMI experiments where decoders are periodically readapted to fit changes in the neural population [16]. We refer to these paradigms as "2-learner systems." In these experiments, closed-loop decoder adaptation (CLDA) was used to refit decoder weights to updated neural intentions at the start of the day, allowing for quick improvements in BMI performance [14]–[17], [35]. However, while CLDA incrementally fits the overall control space closer to the neural subspace, the decoder weights of individual neurons could be affected due to changes in the ensemble of neurons, the variance of neural activity, or sources of noise. Thus, new decoder weights inherently introduce slight perturbations to individual neurons even if the overall population activity is better modeled. In these experiments, task performance continued to improve alongside these decoder refits. Thus, we asked how neurons adapt given slight decoder changes throughout learning. Furthermore, what is the timescale of this adaptation? Refitting a new decoder happens on the timescale of minutes. Thereafter, the decoder is fixed for the duration of the day and sometimes held fixed for multiple days. How do neurons respond to relatively fast changes in decoder weights and still maintain the ability to consolidate firing patterns over longer multi-day timescales?

Since CLDA fits to the neural intentions of the subject, one can assume that any changes in weights in the decoder are mainly rotations within the low-dimensional neural subspace, sometimes referred to as the intrinsic manifold. As reported in recent studies, such perturbations are quickly learned within-day whereas out-of-manifold perturbations are more difficult for neurons to adapt indicating resistance to changes in the neural subspace [27], [28], [32]. Thus, neurons may be adapting to periodically refit decoders as they would to fixed decoders since the control spaces are not significantly perturbed by new decoder weights. However, these decoder weight changes introduce errors on

individual neurons such as misalignment of preferred tuning directions, potentially creating slight out-of-manifold perturbations. In order to adapt to these errors, we hypothesize that neurons increase their neural firing rate variance while maintaining relatively stable, coordinated population-level firing patterns on short, within-day timescales.

To observe the effects of perturbations on the consolidation of neural firing patterns throughout learning, we analyzed data from three BMI experiments performed on a total of five male rhesus macaques. These experiments refit decoders at various frequencies, ranging from repeated decoder swaps within-day to performing CLDA at the beginning of each day (Figure 2). These differences allowed us to gain increased insight on how neurons adapt on short versus long-term timescales. As expected, our results show that neurons increased neural variance immediately after decoder-refitting. However, the variance dropped as the decoder was held fixed, thereby increasing the fraction of neural variance that was aligned with the low-dimensional manifold. Importantly, these manifolds did not change within day nor across many days. Rather, neural variance was slowly collapsed onto a relatively stable neural subspace over learning. These flexible changes in neural variance alongside stable low-dimensional neural spaces from the onset of learning suggest a concomitant exploration and exploitation strategy during motor learning, where new neural patterns are formed by adapting preexisting patterns.

2.3 Results

In this study we analyzed data from three previous studies [6], [16], [17] involving five monkey subjects. In each experiment, decoders were periodically changed and monkeys were trained over the time-course of days thus allowing us to ask how neural consolidation occurs alongside changes in decoder weights over learning. Furthermore, all three experiments consisted of center-out tasks in 2D space whereby monkeys were instructed to move the effectors to eight peripheral targets (Figure 2B). In Experiment 1 (Figure 2C, green), two biomimetic decoders were used each day. The animal (Monkey P) had previous experience with one decoder (DecoderOLD) but was naïve to the second (DecoderNEW). While performance with the old decoder was stable and high across all four days, control with the new decoder improved across days (Figure 2C, green).

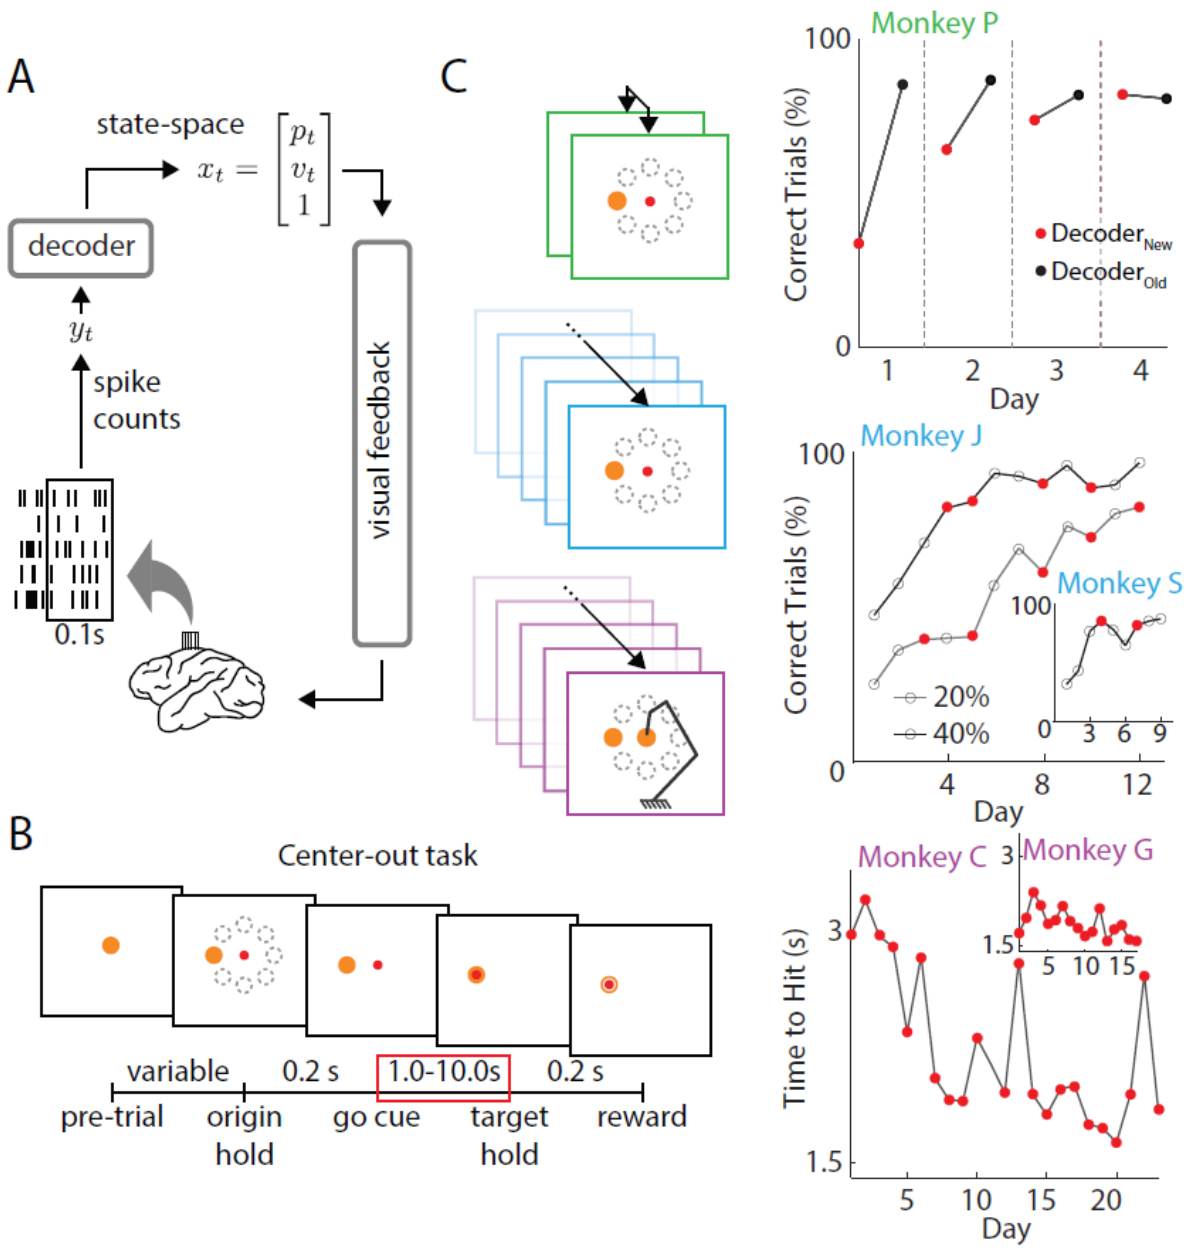


Figure 2 Task design and performance. (A) Closed-loop BMI. Neural signals were recorded from motor cortices in the brain of rhesus macaques. Spikes were recorded and decoded to update a state-space vector, moving the effector in the animals' visual field (adapted from Moorman and Gowda et al., 2017) [17]. (B) In all three experiments, monkeys were instructed to perform a center-out task, moving the effector from a center target to one of eight randomly selected peripheral targets. Successful trials were rewarded with juice. Only rewarded trials during movement periods (red box) were analyzed in this study (adapted from Moorman and Gowda et al., 2017) [17]. (C) Three experiments were analyzed in this study. *Green*, two decoders were used each day [6]. After proficient control of one decoder (black), a new decoder was introduced at the beginning of each successive

day (red). After a period each day, decoders were swapped back to the previously learned decoder. Percent accuracies were calculated as the ratio of successful trials to number of self-initiated trials. *Blue*, decoders were re-adapted periodically across training and neural populations occasionally changed from day-to-day [16]. Monkeys were trained to perform a center-out task with decoders of varying amounts of adaptation (CLDA) as defined by the initial performance of each time series (i.e. 20% and 40%). Monkey J was trained on two time series. After 12 days of 20% CLDA, he was switched to a 40% decoder. Monkey S remained on a 20% decoder for the entirety of the experiment. Red dots indicate decoder re-adaptation. *Purple*, a virtual arm with 4 degrees of freedom (DOF) was controlled with a new decoder each day [17]. Monkeys trained to move the effector in 2D space and improved reach times over learning (Monkey C, $R^2 = 0.431$, p = 6.6e-4; Monkey G, $R^2 = 0.298$, p = 0.0233). CLDA was performed each day and percent accuracy was saturated from Day 1 for both animals (˜80% correct trials across all days).

In Experiment 2 (Figure 2C, blue), both monkeys (J and S) learned to control the BMI with suboptimal decoders as indicated by their initial performance on Day 1 (Figure 2B). Closed-loop decoder adaptation (CLDA) was used to intermittently refit decoder weights over the time-course of learning (Figure 2C, middle, red dots). Note that while Monkey J had two time series of data collected, the 40% CLDA condition (indicating 40% performance accuracy after CLDA training on the first day) occurred the day after the 20% CLDA condition was completed.

Finally, in Experiment 3 (Figure 2C, purple), two monkeys (C and G) moved a kinematically redundant 4 degree-of freedom (DOF) virtual arm to the same set of targets as in the previous experiments. CLDA was performed for roughly 10 minutes at the beginning of each day leading to high performance from the first day. However, decoders weights were still changed from day to day; correlations of these weights are shown in Figures 9 and 10. While the percent correct was saturated from the first day (˜80% for both animals), their reach times continued to decrease across training indicating skill improvement (Figure 2C, purple).

2.3.1 Factor Analysis Predicts Flexible Modulation of Consolidated Patterns

As it has been previously reported [26], [27], [32], [41], factor analysis (FA) provides an intuitive model for how neurons coordinate firing patterns. In short, FA analyzes population-level firing trends and decomposes each neuron's firing activity into covarying (shared) and uncorrelated (private) components (Figure 3A). Importantly, since we were interested in how neural populations adapted over learning, only units that were stably recorded across the entire timeseries for each experiment were used. Geometrically, this shared space (or often referred to as an intrinsic manifold) refers to a low-dimensional

neural subspace where neurons fire in correlated ways (Figure 3B). Subsequently, the private variance (PV) captures the variance of firing rates of a neuron that is independent from other neurons and has been suggested to be a correlate of exploratory patterns [26].

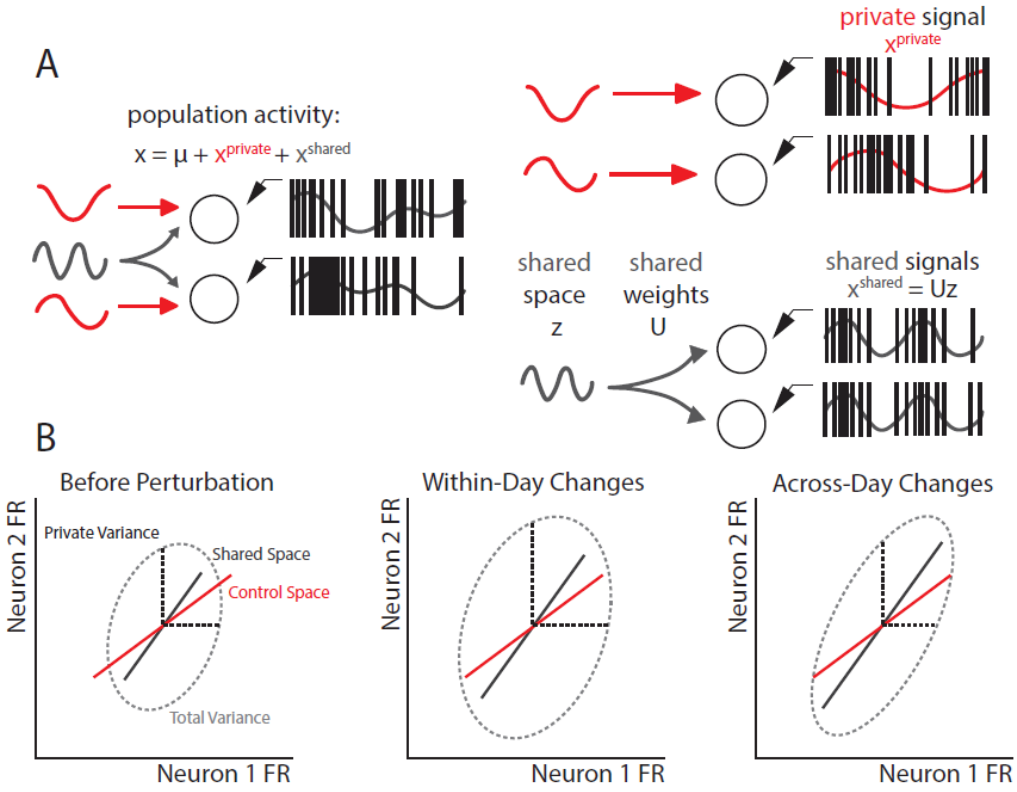


Figure 3 Factor analysis model predictions. (A) Factor Analysis (FA) model. Neural signals can be demeaned (μ) and decomposed into covarying signals (shared) and uncorrelated signals (private). Intuitively, shared signals indicate a low-dimensional neural space where neurons fire in coordinated ways [26]. (B) Predictions of neural subspace changes in response to decoder re-adaptation. Neural firing patterns between two neurons are shown in the dashed ellipses. Private variance (black, dashed) explains the variance that is independent from other neurons. Before decoder re-adaptation, there is an intrinsic shared space (gray). Since CLDA fits to neural intentions, we expect the decoder plane (control space, red) to be fairly aligned with the neural subspace (shared space, gray). Lengths of lines indicate amount of variance captured by each component. Immediately following decoder re-adaptation, we expect no changes in the angle of the shared space. Rather, we expect an increase in the overall neural variance in order to better capture the control space. Note that this increases the shared variance, but not the ratio of shared to total variance. Over multiple days of learning with the same (or similar) decoder, we expect consolidation patterns to occur with increased shared activity (longer gray line) and decreased private activity.

Previous work in the field has focused primarily on the effects of BMI learning on shared spaces and it has been shown that decoders that align with the neural shared space (i.e. intrinsic manifold) can be learned more readily within a day whereas decoders that are explicitly out of alignment with the intrinsic manifold cannot be as quickly learned [27]. Since our experiments rely on biomimetic movements and CLDA to fit effector kinematics to the neural intentions of the animal, we would expect the decoders used in these experiments to be roughly within the intrinsic manifold and therefore we would expect stable alignment of the shared space (Figure 3B). Hypothetically, if a decoder was perfectly fit to the neural intentions of an animal, we would expect no changes in either the angle or the size of the shared space. However, due to the high dimensionality of the neural populations recorded during these BMI tasks, decoder weights change slightly upon a decoder refit. Thus, we hypothesize that since the intrinsic manifold is relatively stable on short timescales, neurons must increase their overall firing rate variance (including both PV and shared variance) to accommodate perturbations in the control space (i.e. decoder weights) (Figure 3B, middle). Subsequently, decoders that are less aligned with the intrinsic manifold would cause a larger increase in overall neural variance.

Additionally, work by Athalye et al., 2017 showed that neurons increased their proportion of shared variance to total variance (SOT) over learning with a fixed decoder, indicating an emergence of coordinated or consolidated neural activity. While CLDA is used to periodically adapt the decoder weights, animals are able to learn the task over the time course of days [16]. Together, this leads to the hypothesis that SOT increases even in the context of a 2-learner system (Figure 3B, right). Furthermore, since task performance is comparable before and after CLDA, some neural adaptation (e.g. changes in total variance) must occur within the same day (Figure 3B, middle). In summary, we expect the neural variance of BMI input neurons to increase as a function of decoder unfamiliarity on a short timescale while shared structures are preserved and consolidated over a longer timescale.

2.3.2 Neural Variance Flexibly Changes

Past work has suggested that PV is linked with exploratory neural patterns that are less goal-potent than shared activity but sufficient for achieving target hits [26]. Since decoder refitting inherently causes some small perturbations in the decoder readout, we would expect exploratory patterns to be engaged to accommodate increases in error. Congruent with past results, we found that the PV increased each time the monkeys controlled the BMI with a new decoder (Figure 4 A, C, E). In Experiment 1, training with DecoderNEW led to higher PV in each of the four days compared to training with DecoderOLD, which was already well-learned (t-test, $p = 1.3e\text{-}7$). In Experiment 2, the PV was calculated each day the decoder was refit and all days following the perturbation (Figure 4C). The PV

consistently increased the day following a decoder refit for both monkeys, though only Monkey J had enough perturbations to test for significance (t-test; Monkey J, p = 0.0334). Lastly, in Experiment 3, where a new decoder was refit each day, the magnitude of PV change was correlated with the amount of change in decoder weights (Figure 4E). Decoders that were similar to ones previously learned resulted in smaller PV changes (correlation; Monkey C, R = -0.70, p = 2e-4; Monkey G, R = -0.63, p = 6.3e-3). Implicit in this result is that PV decreases when more familiar (or fixed) decoders are used. We found that indeed PV decreased when the task switched to using decoders that were more familiar (Figure 5A) or fixed for a period of days (Figure 5B) echoing results from Athalye et al., 2017.

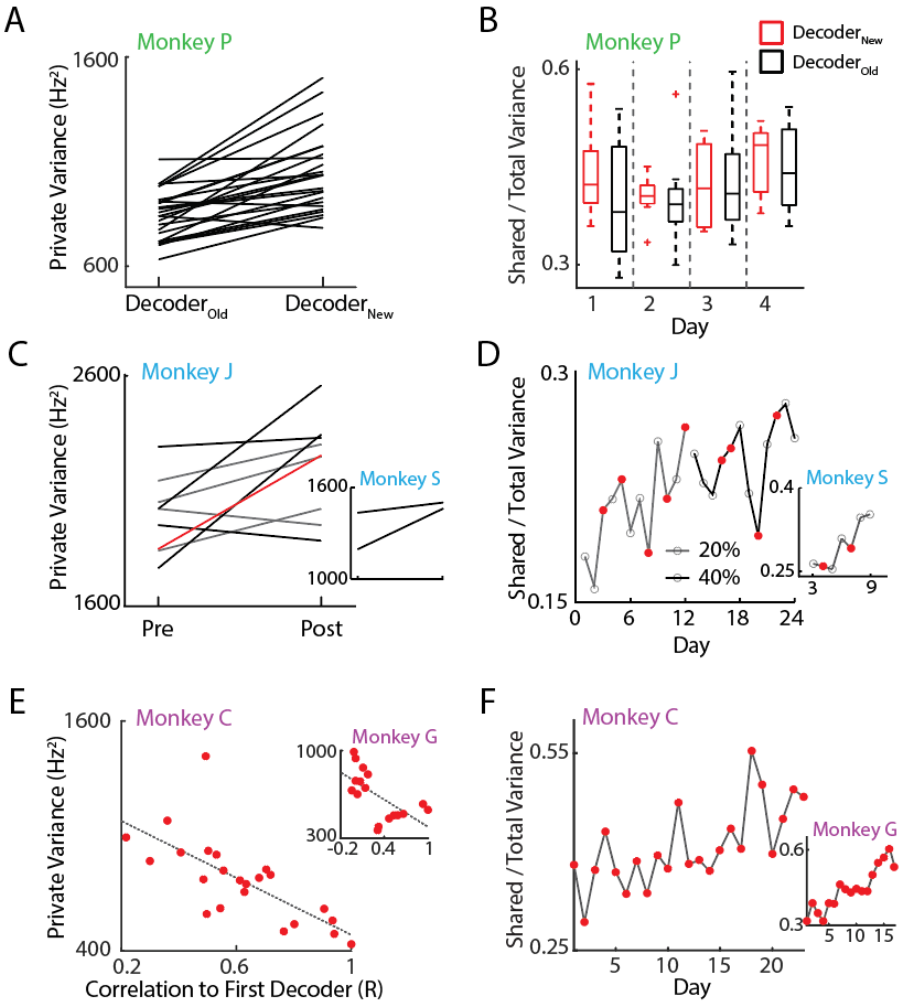


Figure 4 Exploratory and consolidated patterns. (A) Private variance (PV) changed within day. Using the new decoder increased the PV at the beginning of each day (t-test; p = 1.5e-6). (B) Consolidated patterns did not change within day. The shared-to-total variance ratio (SOT) was stable each day and did not change as a result of switching decoders. (C) PV increases when decoders are re-adapted. Across both 20% and 40% time

series, PV increased when decoders were refit with CLDA. The red line indicates the transition from the 20% experiment to the 40% experiment. (Wilcoxon rank-sum test; Monkey J, p = 0.0315; Monkey S n.s., only two decoder re-adaptations were performed). (D) SOT increased over both time series across days (linear regression; Monkey J, R2 = 0.46, p = 2.8e-4; Monkey S, R = 0.81, p = 6.1e-3). (E) PV was higher in cases where decoder weights were more different from the first day. (correlation; Monkey C, R = -0.70, p = 2e-4; Monkey G, R = -0.63, p = 6.3e-3). (F) SOT increased over days with re-adapted decoders each day. (linear regression; Monkey C, R2 = 0.435, p = 3e-4; Monkey G, R2 = 0.829, p = 3.9e-7).

Remarkably, these changes in PV were exhibited on the timescale of minutes. Private variance decreased within the same day when an unfamiliar decoder was replaced with a well-learned decoder (Figure 5A). Conversely, switching back to the unfamiliar decoder the following day increased the PV (Figure 4A). These results show the flexibility of neural variance, which may be required when adapting to new tasks or decoders.

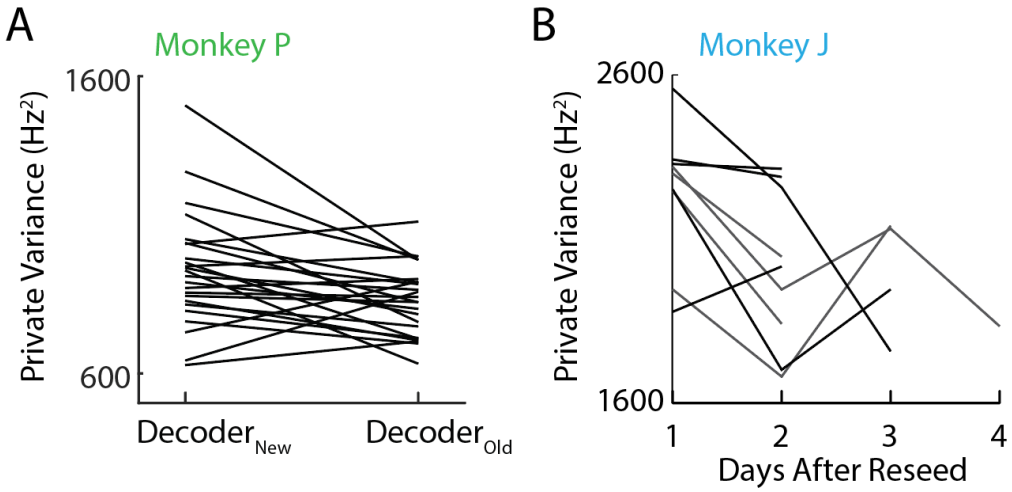


Figure 5 Private variance decreases over learning with a fixed decoder. (A) Private variance decreased for Monkey P when the task switched from using the unfamiliar decoder to the familiar decoder. Each line indicates changes in PV for a single target for each day (t-test, p = 0.008). (B) When decoder weights were constant in Experiment 2, PV decreased (Wilcoxon rank-sum test, p = 0.0315).

2.3.4 Neural Exploration Adapts Consolidated Patterns

We asked how neural consolidation emerged in a 2-learner system despite frequent changes in neural firing patterns due to changes in decoder weights. To track the proportion of neural activity that fired in coordinated patterns, we normalized the shared variance by the total neural variance (SOT). Thus, increases in SOT act as a measure for how covaried

neurons' firing activities are and can be thought of as the degree of consolidation of neural patterns [26]. We found that while PV was modulated due to decoder changes, the SOT remained stable on within-day timescales in Experiment 1 where decoders were switched in the middle of the day (Figure 4B). Together, these results indicate a scaling of the total variance (including shared variance) when a perturbation arises as predicted in Figure 3B (middle). In addition, SOT increased in experiments where a single decoder was used each day (Figure 4 D, F) (regression, Monkey J, $R^2 = 0.46$, p = 2.8e-4; Monkey S, $R^2 = 0.81$, p = 6.1e-3; Monkey C, $R^2 = 0.435$, p = 3e-4; Monkey G, $R^2 = 0.829$, p = 3.9e-7). The decoupling of changes in PV and SOT over different timescales suggests that neural exploration and exploitation can simultaneously occur, and increased exploration does not occur at the expense of consolidated patterns. In contrast, increases in SOT over the time-course of many days despite flexible changes in PV may suggest that exploratory patterns consolidate over time and that neurons are adapting learned behaviors rather than generating completely novel patterns.

Furthermore, we asked if the shared spaces were stable across different decoders or if poor performance with a new decoder would lead to adaptation to a different low-dimensional space. To answer this question, we computed the "shared alignment" between the neural subspaces used to control the two decoders in Experiment 1 (Figure 6A). We found that the shared spaces used to control the task were nearly identical despite large differences in task performance (Figure 6B). Similarly, we compared the shared alignment between adjacent days and across days for Experiments 2 and 3 and found that shared spaces were stable on short and long timescales (Figure 6 B-D). In all cases, the shared alignment was high in adjacent blocks or days (~0.8 across all animals), suggesting that neurons rely heavily on preexisting patterns regardless of task performance.

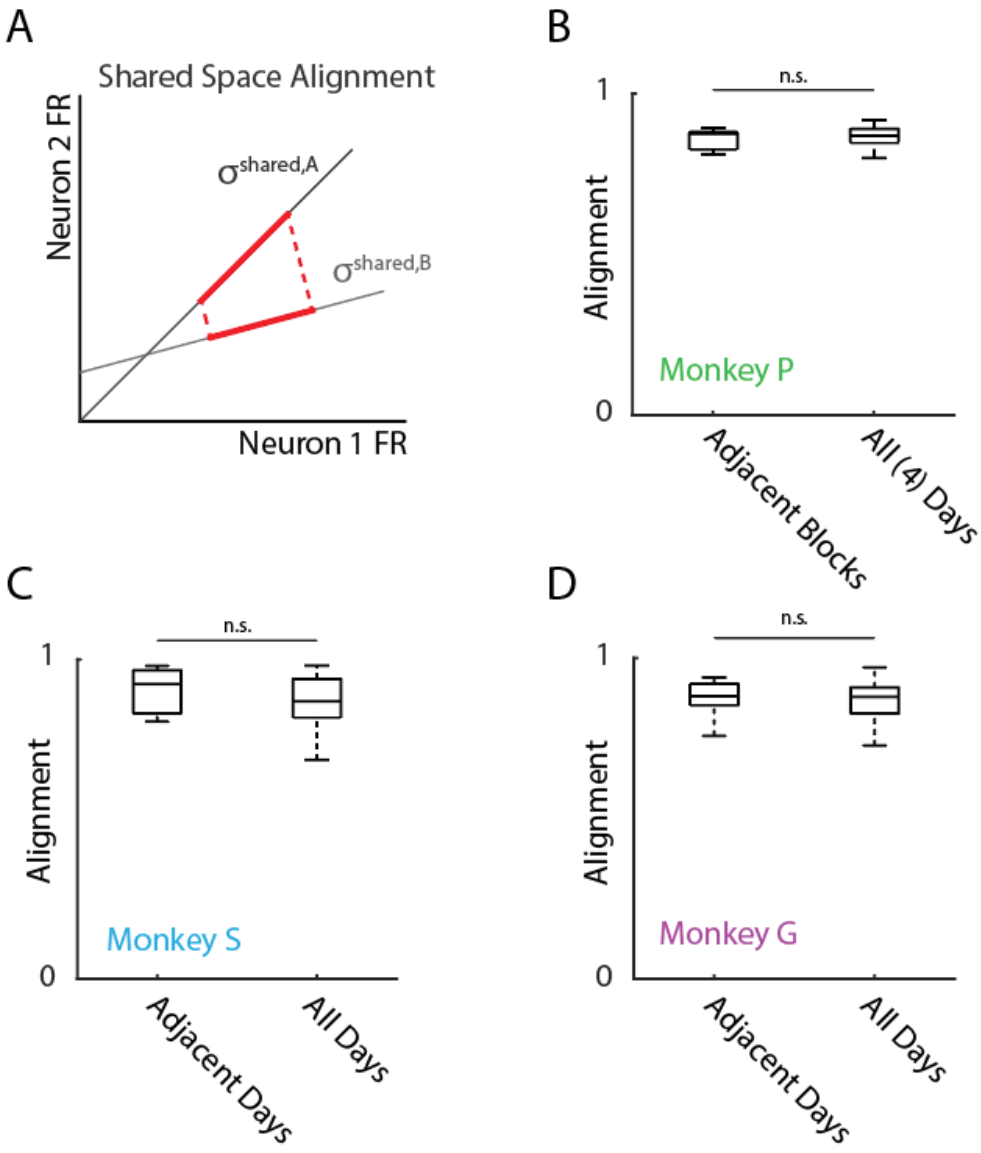


Figure 6 Shared spaces are resilient to perturbations. (A) The similarity (alignment) between shared spaces was calculated by projecting the shared space of one experimental block onto another. The shared space alignment was defined as the amount of variance captured by this projection. (B) Alignment was calculated between adjacent blocks in Experiment 1 as well as all pairwise comparisons among all blocks across the four days. The shared alignment was high between decoder swaps indicating a stable space immediately following a decoder change. This shared alignment was not significantly different between adjacent days and over the time course of four days ($p > 0.05$, Kolmogorov-Smirnov test). (C) Same as B except over days rather than individual blocks for Experiment 2. Shared alignment was not significantly different between short and long time periods ($p > 0.05$, Kolmogorov-Smirnov test). (D) Same as C but for Experiment 3 ($p > 0.05$).

2.3.5 Modulation Depth Facilitates Changes in Neural Variance

Finally, we asked how neurons are able to flexibly modulate their neural variance on short, within-day timescales. Intuitively, our results imply that the dynamic range of firing rates across all neurons must also be flexible to account for changes in neural variance. That is, solely increasing the mean firing rate of neurons does not necessarily increase the variance. Rather, we would expect some modulation of firing rates. Past literature has shown preferred direction tuning in BMI neurons. The dynamic range of a neuron's tuning curve is often referred to as the modulation depth (MD) of the neuron [5], [6], [16], [29]. Past work has shown MD to change over many days of BMI learning. However, given the flexibility of the neural variance, we hypothesized that MD could also change on a shorter timescale.

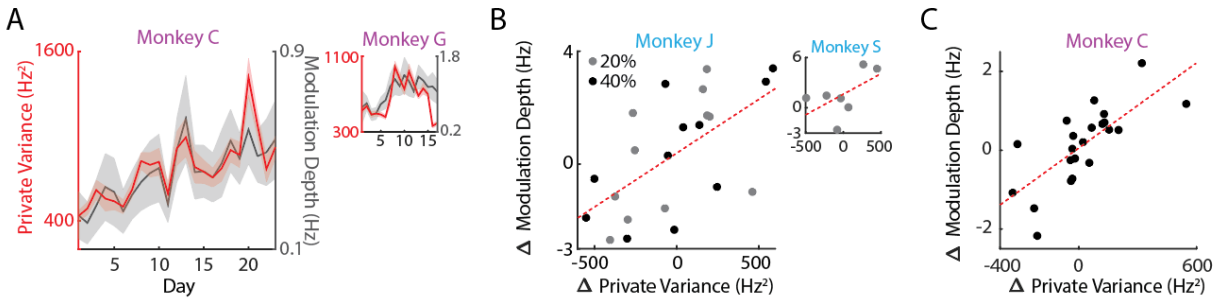


Figure 7 Modulation depth may facilitate PV changes. (A) PV fluctuations are proportional to changes in MD (correlation; Monkey C, R = 0.86, p < 1e-5; Monkey G, R = 0.63, p = 7e-3). Lines indicate mean values of PV and MD; shaded regions show s.e.m. across targets for each day. While the changes in PV and MD were not significantly correlated for Monkey G in (C), the overall trends are proportional. (B) Changes in PV were closely linked with changes in modulation depth (MD). Each dot represents one decoder change. Regression was conducted on combined data across the 24 days for Monkey J (correlation; R = 0.63, p = 0.0013). Correlation was not significant for Monkey S due to too few data but still shows same trends (correlation, R = 0.569, p = 0.18). (C) Same as in (B) for Monkey C in Experiment 3 (correlation; R = 0.73, p = 1e-4). Results were not significant for Monkey G due to large fluctuations in channels used for recording from day-to-day.

We calculated the preferred tuning models of each neuron by fitting a cosine tuning curve to the mean firing rates corresponding to cursor velocity directions [6], [16], [29], [45], [46]. The MD of each neuron was then determined by taking the peak-to-peak amplitude of the cosine model. As predicted, we found that the MD was tightly coupled to the PV (Figure 7A) (correlation; Monkey C, R = 0.86, p < 1e-5; Monkey G, R = 0.63, p = 7e-3). Furthermore, the changes in MD were coupled to the changes in PV and was

also able to change on a within-day timescale (Figure 7 B, C) (correlation; Monkey J, R = 0.63, p = 0.0013; Monkey C, R = 0.73, p = 1e-4). Notably, regression lines very closely intersected the origin indicating that MD does not inherently increase due to task learning but could be due to modulation of PV. These results suggest that neurons are able to quickly increase the firing rates specifically in their preferred directions in response to perturbations in the control of the BMI. This increase in MD can be thought of as a "stretching" of the tuning curve or an amplitude gain that increases the variance of firing rate in the preferred direction without necessarily affecting the overall mean firing rate of the neuron.

2.4 Discussion

Previous work has shown BMI performance to improve over days alongside decoder re-adaptations [16], [32]. Here, we asked how neurons changed firing patterns over learning to allow for these performance improvements despite changing decoder weights. Our results support our hypothesis that motor cortex flexibly modulates neural variance when errors arise (due to decoder readaptation) without eroding preexisting coordinated patterns useful for generating goal-potent actions. Thus, at least in a 2-learner context, BMI learning involves leveraging exploratory patterns to adapt learned neural patterns rather than generating new ones.

Across all three experiments and five monkeys, the shared over total variance ratio (SOT) was remarkably stable across learning while private variance (PV) reliably fluctuated whenever a different decoder was introduced (Figure 4). Moreover, changes in PV were dependent on how different the decoders were from previously learned control spaces (Figure 4E). This was particularly noticeable in Experiment 1 where PV drastically changed within the same day while SOT remained unchanged within and across days. Notably, SOT did increase over longer timescales (i.e. Experiments 2 and 3) where perturbations were smaller. Experiment 1 used two separate decoders, with a new one introduced after achieving expert control with the first decoder. Furthermore, decoder perturbations happened each of the four days, each day ending with the previously learned decoder. The frequency and magnitude of these perturbations may have hindered consolidation effects observed in the other two experiments. Nonetheless, the SOT was stable within and across days and did not decrease due to repeated perturbations.

Also as predicted, we observed no changes in the low-dimensional neural subspace both within-day and across days (Figure 6 B-D). To quantify any changes, we calculated the "shared alignment" which represented the magnitude of the projection of one manifold onto another. In all three experiments, the alignment between any two adjacent blocks (or days for Experiments 2 and 3) were roughly 0.8-0.9 indicating similar subspaces on

short timescales. These alignments were also similar across days suggesting minimal rotation of the manifold even over the time-course of learning.

Our results also suggest that modulation depth (MD) may play a critical role in facilitating fast changes in neural variance. While past studies have repeatedly shown changes in MD over learning, it has been unclear what the exact role of this may be. We found that the MD was tightly correlated with PV. More specifically, changes in PV were closely tracked by directly proportional changes in MD (Figure 7). Importantly, fluctuations in both were fast (compared to neural consolidation rates) in both increasing and decreasing directions (Figure 3, 6, 7). Intuitively, increasing MD increases the dynamic range of neural firing rates, allowing for higher total neural variance. In contrast to past studies showing overall changes in MD over learning [6], [16], [29], these fast changes in modulation depth indicate that the ability to fluctuate firing rates may be an inherent property of motor cortex rather than a learned behavior.

Lastly, our findings reflecting flexible changes without greatly affecting established coordinated patterns is compatible with that of natural motor adaptation. Our results suggest that neuroprosthetic learning and motor adaptation may have similar underlying mechanisms. Studies have shown the importance of subcortical structures and the cerebellum in adaption to motor perturbations and BMI tasks [10], [37], [47]. Furthermore, past studies provide evidence on the role of the basal ganglia for refining cortical patterns during BMI on shorter and longer timescales [4], [30], [48], [49]. In a motor adaptation task then, it is possible these structures are turning the "variability knob" in motor cortex. If we consider the production of motor actions as a combination of consolidated activity from the motor system with some exploratory variability [50], [51], then our results suggest that PV may be a neural correlate of the latter. Furthermore, changes in neural variance might reflect an increased gain in exploratory variability without directly affecting learned motor behaviors. In such a scenario, neuroprosthetic learning – particularly with closed-loop decoder adaptation (CLDA) – may be encouraging changes from subcortical structures we see manifested as changes in motor cortex.

Our results are complementary with those from previous studies [26]–[28], [32]. It has been shown that intrinsic manifolds are relatively stable and thus are less able to adapt to decoders that are explicitly out of that manifold [27]. While learning out-of-manifold decoders is difficult within-day, it has also been shown that these shared spaces can rotate over many days of training [26], [32]. In this study, we show how neurons respond to changes in decoder weights on both timescales before task performance is saturated. Our results regarding stable shared spaces corroborate these past studies but we highlight the importance of the neural variance (and specifically the PV) during learning. While we echo past results that neural activity preferentially stays in a constant neural subspace during learning, we found that neurons are also able to quickly increase

their private variance when perturbations are introduced. This simultaneous modulation of neural variance whilst maintaining a robust manifold suggests a neural learning strategy that is resilient to perturbations.

2.5 Methods

2.5.1 Surgical Procedures and Electrophysiology

All procedures for each of the three experiments 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.

2.5.1.1 Experiment 1

One adult male rhesus macaque was chronically implanted with 64-channel Teflon-coated tungsten microelectrode arrays. Implants were in the arm regions of primary motor cortex (M1) and dorsal premotor cortex (PMd) in the left hemisphere. There was an additional implant in the arm area of M1 in the right hemisphere, totaling three implants. Unit activity was recorded using a MAP system (Plexon Inc., Dallas, TX) and sorted with an online spike-sorting application (Sort Client; Plexon). Units stable across days were used for decoding [6].

2.5.1.2 Experiment 2

Two adult male rhesus macaques were chronically implanted with 128-channel microwire electrode arrays. Implants were bilateral, targeting the arm areas of M1. Neural activity was recorded using the same MAP system as in Experiment 1. Multi-unit activity was sorted using Sort Client for Monkey S and channel-level activity was used for Monkey J [16].

2.5.1.3 Experiment 3

Two adult male rhesus macaques were chronically implanted with microwire electrode arrays with 128 channels in each hemisphere, targeting arm areas of M1 and PMd in Monkey G and primary sensory cortex (S1) and M1 in Monkey C. Neural activity was filtered and recorded using an OmniPlex system (Plexon Inc., Dallas, TX). Single and multi-unit activity was sorted with a template-matching algorithm (PlexControl) [17].

2.5.2 BMI Task

In all three of the experiments we analyzed, monkeys performed a center-out BMI task. Subjects were instructed to enter and hold in the center target to initiate a trial. Upon hold completion, one of eight peripheral targets would appear. Subjects then moved the cursor towards the target after this go cue. Trials were successful if the cursor entered and stayed in the peripheral target for a short hold time (~250 ms) for each experiment. Decoders were trained differently in the three experiments (outlined below), but in all cases, subjects were familiar with the task structure before attempting control with BMI. Therefore, subjects' performance improvement can be attributed to BMI learning rather than task learning. Once decoder weights were held fixed each day, task control was achieved by modulating neural activity, which was recorded and fed through the decoder to produce cursor movements. In order to observe neural correlates of coordinated control, we analyzed movement portions (i.e. leaving center target to entering peripheral target) of successful trials for each experiment (red box, Figure 2B).

2.5.2.1 Experiment 1

In Experiment 1, two macaques were instructed to complete a BMI center-out task as described above. After achieving proficient control with an "old" decoder, a new biomimetic decoder was introduced. This decoder was trained on Day 1 (Figure 2C). While many weights were similar between the two decoders, many of the weights were significantly different [6]. Each day, both decoders were used for BMI control in a series of two blocks. To track proficiency with each decoder, the percent of correct trials was calculated for each block for each day. For each decoder, a linear filter (Wiener filter) was used to fit neural activity to kinematic activity in the elbow and shoulder joints during a manual reaching task using a BKIN KINARM exoskeleton robot. Decoder outputs were then translated into Cartesian space by using a Jacobian matrix. For more details, we redirect the reader to Ganguly and Carmena, 2009.

The Wiener filter model assumes that the cursor velocity is a linear combination of neural activity across small differences in time. In other words, the output of the decoder at time t is determined by a weighted sum of the neural activity from $t-k$ for some set of k representing the time lags. More specifically:

$$y(t) = \boldsymbol{b} + \sum_{k=0}^{M} \boldsymbol{a}(k) * x(t-k) + \epsilon(t)$$

where ϵ refers to any residual errors and M represents the number of time lags used (in Experiment 1, this was 10). The parameters $\boldsymbol{a}$ and $\boldsymbol{b}$ were determined to optimally fit the model [6], [52].

2.5.2.2 Experiment 2

The second experiment we analyzed was originally conducted by Orsborn et al in 2014. To summarize, prior to BMI control, both subjects were trained to conduct the manual control version of the center-out task using the KINARM. An initial Kalman filter (KF, outlined below in 2.5.3) decoder was fit on day 1 on these manual reaches. Closed-loop decoder adaptation (CLDA) was then used to better fit decoder weights to the neural intentions [15]. To briefly summarize, CLDA updates decoder weights by inferring intentions of the subject based on task goals. Decoder weights were updated through a weighted average of old and new decoder weights. CLDA was run at the beginning of day 1 to achieve some level of control (~20% accuracy). Over days, decoder weights were occasionally refit with CLDA, sometimes in conjunction with neuron swaps to recover performance. After 12 days, the procedure was repeated for Monkey J with initial performance elevated to 40% accuracy. For further details regarding CLDA, we redirect readers to the methods outlined in Orsborn et al., 2012.

2.5.2.3 Experiment 3

In Experiment 3, two monkeys (Monkeys C and G) controlled a four degree-of-freedom (DOF) kinematic chain in a center-out task. Subjects were first trained to perform a cursor center-out task using manual control with a KINARM exoskeleton as in previous studies, then were transitioned to BMI control of a 2D cursor before learning to perform the full 4 DOF task with BMI control. Movements of the kinematic chain were constrained to two dimensions thereby creating a kinematically redundant system. A KF decoder fit neural activity to four individual joint angle velocities. Visual feedback of the kinematic chain moving from target to target was used to seed the KF decoder and CLDA was used to update neural intentions each day. Unlike Experiment 2, CLDA was run until the subjects' performance saturated each day (~85% accuracy). Since the percent correct was saturated and stable across days, task improvement was tracked by decreasing reach times. We refer readers to Moorman et al., 2017 for more experimental details.

2.5.3 Kalman Filter

Here, we briefly describe the Kalman filter model for BMI decoding used in Experiments 2 and 3. For more details, we refer the reader to Orsborn et al., 2014 and Moorman et al., 2017. In general, the state-space (e.g. kinematic space) is represented by $x_t = [p_t\ v_t\ 1]$

where p_t represents the positions of the effector and v_t represents the velocities of the effector. In Experiment 2, the state-space contains the horizontal and vertical positions and velocities. In Experiment 3, the state-space refers to the joint angles and angular velocities of the 4 DOF kinematic chain. The Kalman Filter model assumes that x_t varies as $x_{t+1} = \mathbf{A}x_t + w_t$ where w_t is a noise term ($w_t \sim N(0, \mathbf{W})$). A advances the state using the current positions and velocities and W is set up so as to allow the joint velocities to evolve independently:

$$\mathbf{A} = \begin{bmatrix} \mathbf{I} & \Delta t\mathbf{I} & 0 \\ \mathbf{0} & 0.8\mathbf{I} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \mathbf{W} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & 0 \\ \mathbf{0} & 0.1\mathbf{I} & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

where $\mathbf{I}, \mathbf{0}$ are identity and 0 matrices, respectively, and Δt is 100 ms. Note that in these models, $\mathbf{A}$ contains a scaling factor (0.8) in order to decay the velocities over time.
The observations y_t represent the neural activity in the past 100 ms. The KF model of neural firing assumes that $y_t = \mathbf{C}x_t + q_t$, $q_t \sim N(0, \boldsymbol{Q})$. The KF estimates x_t, the state variable, from the previous t observations $\{y_0, \ldots, y_t\}$ and produces the state estimate x_t and confidence $\mathbf{P}_t$. The KF uses the last estimate x_{t-1} and advances that belief to produce $x_{t|t-1} = Ax_{t-1}$, and then updates the prediction when a new observation becomes available: $x_t = x_{t|t-1} + K_t(y_t - Cx_{t|t-1})$. The Kalman gain $\mathbf{K}_t$ is then determined by the model parameters $\{\mathbf{A}, \mathbf{C}, \mathbf{W}, \mathbf{Q}\}$ and $\mathbf{P}_t$. Since our model parameters are fixed, $\mathbf{K}_t$ converges to a steady-state $\mathbf{K}$, and we estimate the kinematics to follow the simpler equation:

$$x_t = (\mathbf{I} - \mathbf{KC})\mathbf{A}x_{t-1} + \mathbf{K}y_t$$

2.5.4 Factor Analysis

Factor analysis (FA) was used in this study to approximate low-dimensional manifolds; here we provide an outline of the method [26]. FA decomposes the neural activity into three components: the mean firing rates (μ) of each neuron, the population-level coordinated firing activity (x^{shared}), and the remaining component corresponding to each neuron's uncorrelated activity to the rest of the population ($x^{private}$). That is,

$$x = \mu + x^{private} + x^{shared}$$

These latter two components have covariance matrices $\Sigma^{private}$ and Σ^{shared} with dimensions $N \times N$, where N is the number of neurons. The combination of these two parts yields the total variance,

$$\Sigma^{total} = \Sigma^{shared} + \Sigma^{private}$$

The dimensionality was estimated by using cross-validated log-likelihood to determine the number of factors that would best describe held-out data such that ~90% of the shared variance could be captured [26].

Importantly, only units stable across all days for each experiment were used for these analyses. Since we were interested in how the same ensemble of units adapted over learning, including all units used for each day would contaminate across-day results.

2.5.4.1 Private Variance

We define private variance (PV) as an $N \times N$ covariance matrix where the diagonals represent independent variances for each of the N neurons. In Experiment 1, the PV was calculated for each target and each block. Average PV values were found by averaging across the 8 targets for each block. Similarly, PV was calculated for each day (and one day after) a decoder was refit in Experiment 2. This gave us PV values before and after refitting and a t-test was used to determine differences in PV. Furthermore, to track how PV changed over days with stable decoders, PV was calculated for days without decoder changes as well. In Experiment 3, since CLDA was run each day of the experiment and percent accuracy on Day 1 was already as good as the last day (i.e. proficient control), we instead examined the correlation of the decoder for each day with the decoder on the first day as a metric for the difference between decoders. We then plotted this correlation against the average private variance over all neurons for that day (Figure 4E). Note that this correlation was shown for stable units in Figure 4E; Supplemental Figure 8A shows this correlation calculated using stable channels for Monkey G, which include channels that were used in the BMI task but were either physiologically different across recording days or contained no physiological signal.

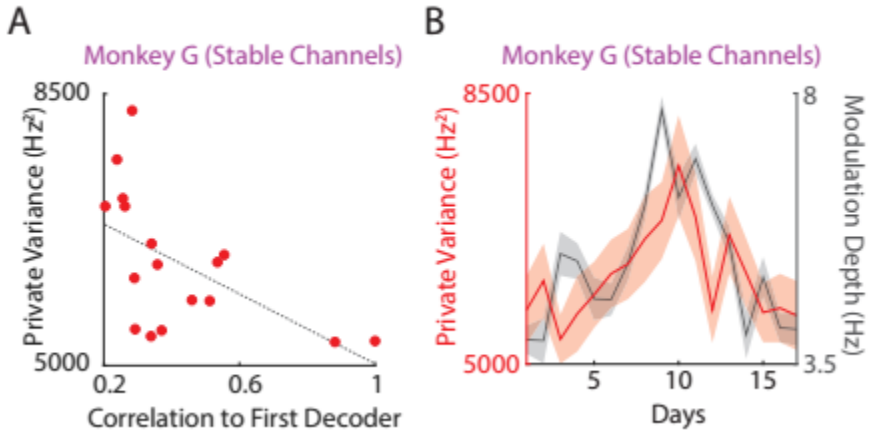


Figure 8 Private variance modulation in stable channels for Monkey G. (A) Same as Figure 4E but with stable channels in Monkey G (correlation; R = -0.57, p = 0.0169). (B) Same as Figure 4A inset but with stable channels in Monkey G (correlation; R = 0.61, p = 0.0088).

2.5.4.2 Shared / Total Variance

We defined the shared/total variance ratio (SOT) as

$$\text{SOT} = \frac{\text{trace}(\Sigma^{shared})}{\text{trace}(\Sigma^{total})}.$$

Intuitively, this allowed us to measure the proportion of neural variance that was coordinated. In order to track how neurons in the population changed coordination patterns over learning, we only considered units that were stable across the entirety of the experiments for these analyses, determined by waveform and inter-spike interval statistics. In all experiments, the SOT was calculated for each block/day for all stable units. In Experiment 3, Monkey G had several "unstable" units throughout the experiment which were not considered in these analyses. In addition, a variety of noisy channels were included as decoder inputs. Figures 3 and 4 show results with stable units for Monkey G, however the results still hold using all stable channels (Figure 9, 10).

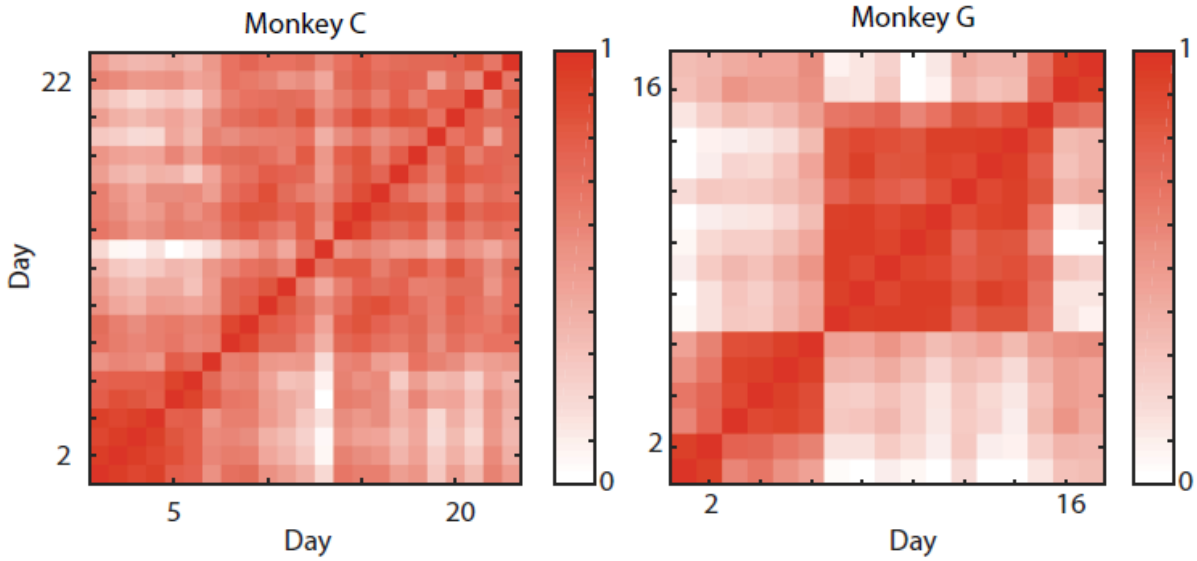


Figure 9 Decoder correlations. Decoder observation model weights were correlated between pairwise days for Monkey C (left) and Monkey G (right). Only stable units were used in this plot for Monkey G. Decoder correlations across all stable channels for Monkey G are shown in Figure 10.

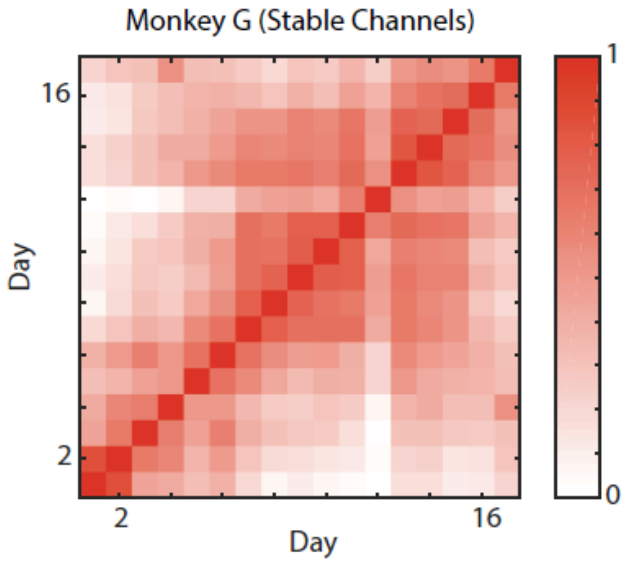


Figure 10 Decoder correlations for stable channels in Monkey G. Same as Figure 9, but for all stably recorded channels in Monkey G.

2.5.4.3 Shared Space Alignment

In order to determine how similar manifolds were between experimental blocks (or days), we calculated the "alignment" between shared spaces [26]. Geometrically, it shows how much of one subspace projects onto a second as a fraction ranging from 0 to 1. Moreover, a perfectly aligned subspace results in an alignment value of 1 while orthogonal subspaces have an alignment of 0. To calculate the shared space alignment, we find the projection of the shared variance of block (or day) A onto the shared space of block (or day) B.

$$\text{shared alignment} = \frac{\text{trace}\left(P_{U^B} \Sigma^{A,\text{shared}} P_{U^B}^{T}\right)}{\text{trace}\left(\Sigma^{A,\text{shared}}\right)},$$

where $P_U^B \in R^{N \times N}$ is an projection matrix onto the shared subspace of block B ($\text{col}(U^B)$), and $\Sigma^{A,shared}$ is the shared variance of block A.

To measure how similar shared spaces were between decoder changes, we first found the shared alignment between pairwise days. In Experiment 1, comparisons were then made between all decoder block changes and all other pairwise comparisons. Similarly, to test if shared spaces were aligned on short and longer timescales in Experiments 2 and 3, shared alignments between adjacent days were compared against all other pairwise combinations of days. For each experiment, a Kolmogorov-Smirnov test was used to determine if the distributions of shared alignment values were significantly different.

2.5.5 Preferred Direction and Modulation Depth

We also looked at the correlation of modulation depth and private variance. Cursor kinematics were separated into eight 45-degree bins and corresponding firing rates were used to determine the preferred tuning direction (PD) for each neuron by fitting the firing rate f to a cosine function with mean firing rate μ:

$$f = [B_1\ B_2\ B_3] \times \begin{bmatrix} \mu \\ \sin(\theta) \\ \cos(\theta) \end{bmatrix}$$

where θ corresponds to cursor movement angle and $\mathbf{B}$ are coefficients estimated through linear regression. PD was calculated as $\mathrm{PD} = \arctan(B_2/B_3)$, resolved to the correct quadrant. Modulation depth (MD) was then defined as the peak-to-peak amplitude of this curve. Intuitively, it shows how specific neurons fire in their PD and is calculated as

$$\mathrm{MD} = \sqrt{B_1^2 + B_2^2}$$

To track how MD and PV changed together over learning, we correlated the changes in MD (i.e. derivative) with the changes in PV across days. As before, Figure 7A (inset) shows the results when considering only stable units for Monkey G to remain consistent with the other animals. However, the results are consistent if all stable channels are used as well (Figure 8B).

2.5.6 Statistical Analyses and Testing

We used t-tests to reject null hypotheses with a confidence of 5% and used linear regressions and correlations to fit trends. Kolmogorov-Smirnov tests were used to test whether distributions were the same. Wilcoxon rank-sum tests were used to compare private variance distribution differences. Where reported, error bars represent standard error of the mean over 8 targets. All analyses were done in MATLAB by writing custom scripts.

Chapter 3

The Rate of Neuroprosthetic Learning is Invariant to Initial Levels of Closed-Loop Decoder Adaptation

3.1 Abstract

Reliable control of neuroprosthetic devices often involves neural adaptation over learning. Advancements in decoding neural signals with closed-loop decoder adaptation (CLDA) algorithms have reliably enhanced subjects' initial performance levels while still allowing for neural adaptation to occur, thereby creating a "2-learner" paradigm. While it has been shown that the brain can adapt alongside decoder adaptation, it is unclear how varying levels of decoder adaptation might affect the rates at which neurons adapt. Here, we examined the neural adaptation of rhesus monkeys in a brain-machine interface (BMI) task where various levels of decoder adaptation were used. Our results show that while the initial performance levels were a function of decoder adaptation, the rate of neural adaptation remained constant. Furthermore, the rates of adaptation were correlated with rates of behavioral improvement. Together, these results show that the rate of neuroprosthetic learning is independent from the initial levels of decoder adaptation.

This work was in collaboration with Bing Liu.

3.2 Introduction

Brain-machine interfaces (BMI) are neuroprosthetic devices that decode recorded neural activity into actions performed by an effector such as a robotic arm or computer cursor. As indicated by the name, BMIs have two main components that interplay with each other – the brain and the machine (e.g. a computer that handles the decoding). Many studies have focused on how the brain adapts to a static, fixed decoder [5], [6], [29], [32]. It has been shown that through training, neurons change their preferred tuning directions as well as their modulation depths (i.e. the amplitude of their preferred tuning curves) [6], [16], [29]. These changes have been shown to stabilize over learning as the control accuracy saturates. Additionally, groups have shown that de novo control of a BMI is possible with an accurate decoder. These complementary studies thus pique questions about how adaptation between the brain and the machine might occur.

Orsborn et al., 2014 showed that indeed the two systems can adapt together as performance improves and thus frames BMIs as a "2-learner" system [16]. To adapt the decoder alongside any changes in neural firing patterns following neuron dropouts, closed-loop decoder adaptation (CLDA) was used to periodically refit the weights of the decoder to the existing neural population. Specifically, it was shown that changes in modulation depth (MD) and preferred tuning directions (PD) were affected by the decoder weights. Decoder weights that more accurately reflected the natural firing properties of the neurons caused smaller changes in neural adaptation while inaccurate decoding was linked to larger changes in MD or PD. Implicit in these results then is that low levels of CLDA encourage large amounts of neural adaptation while high levels of CLDA require less adaptation from the brain. That is, there is a balance between neural and decoder adaptation.

Interestingly, more recent work looking at how neurons adapt on low-dimensional neural subspaces has suggested that neural activity consolidates onto relatively low-dimensional spaces as skill with a BMI improves [26], [27], [31], [32]. It is unclear how differences in CLDA level might affect consolidation patterns or the rate of consolidation over learning. One hypothesis is that with an accurate decoder, performance is good from learning onset, and therefore no consolidation is required. However, it is also possible that neural consolidation occurs as a result of skill refinement and is not driven by decoder performance.

To address these questions directly, we investigated the learning-related behavior and neural adaptation over training with different initial levels of CLDA. In particular, we investigated how the different levels of initial decoder adaptation influence the rate of improvement of learning to control a prosthetic device. Our results suggest that the rate of this neuroprosthetic learning is invariant to the initial assist levels of decoder

adaptation, and the neural adaptation underlying learning, not the assist of decoder adaptation, is essential for learning skillful neuroprosthetic control.

3.3 Results

To test the effects of initial decoder adaptation level on both behavioral and neural changes over learning, closed-loop decoder adaptation (CLDA) was used to periodically adapt the decoders as behavior improved. Multiple training series were conducted for each of two rhesus macaques. Decoder adaptation level was quantified by the percentage of correct trials achieved on Day 1 for each of the training series. That is, 20% CLDA would indicate that 20% of trials attempted were successful immediately following decoder adaption. Monkey J was tested with 20%, 40%, and 80% CLDA conditions while Monkey S was trained on 20% and 60% CLDA conditions. Following the first decoder adaptation, CLDA was briefly run at the start of days where there was a loss or swap of neurons in the recorded population in order to restore performance. In all cases, after CLDA was run at the start of the day, the decoder was then fixed for 1-3 days until CLDA was run again (Figure 11C).

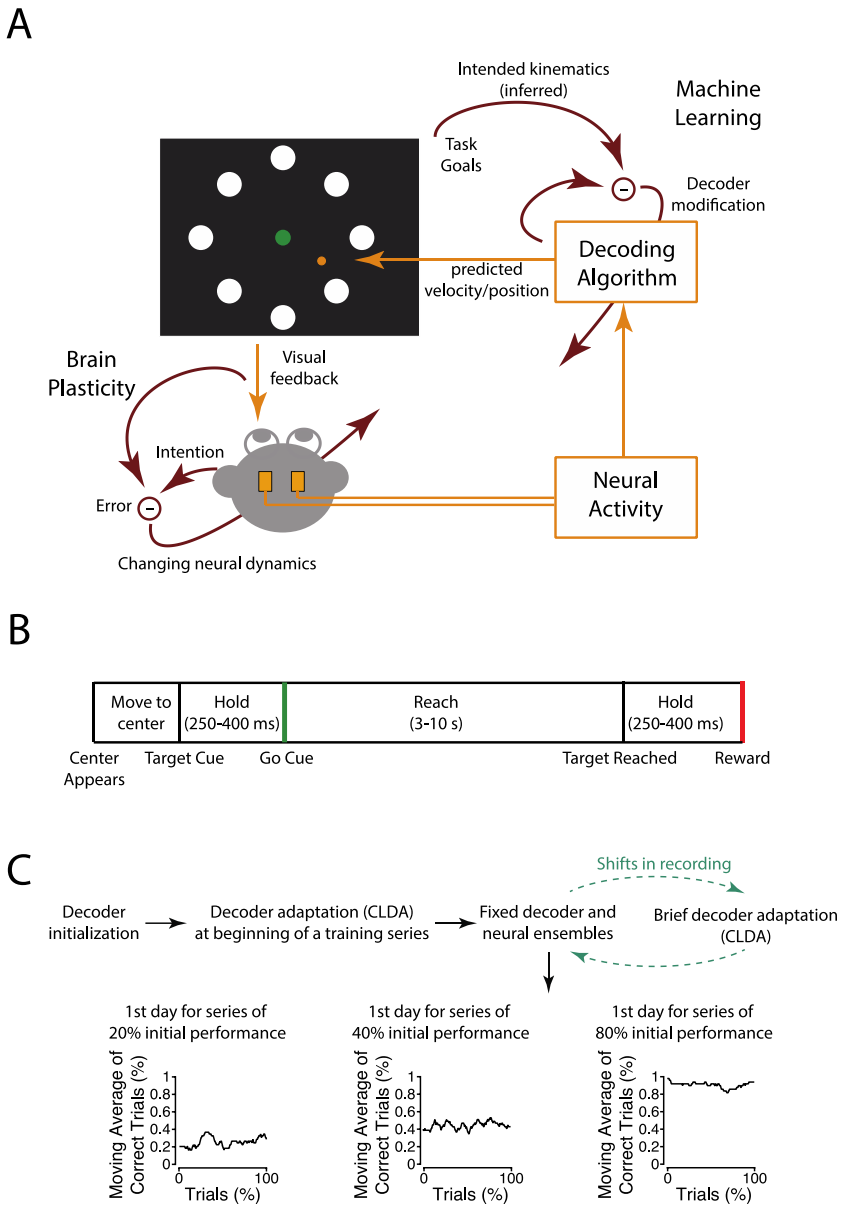


Figure 11. Experimental paradigm. (A) Experimental setup and control loops for a two-learner BMI paradigm. The close loop BMI is indicated by a yellow circle. The CLDA and the neural activity adaptation are indicated in red, which represent the two learners. During the experiment, nonhuman primate subjects were trained to perform a 2D center-out task with brain control. During the task, the subject's arm was confined to the primate chair. (B) The timeline of the center-out behavioral task. (C) Paradigm for experiments with different levels of initial decoder adaptation. Experiments were performed with different degrees of initial performance improvement (also the amount of CLDA) on the first day of the training series (for the three series of experiments in this paper, the initial CLDA improved the percentage correction of the first day training up to around 20%, 40%, and 80%, respectively). Then the decoder was fixed in the later training sessions of the series. CLDA was used briefly if recorded neurons shifted or if performance dropped.

3.3.1 Initial decoder assistance level does not impact neuroprosthetic learning rates

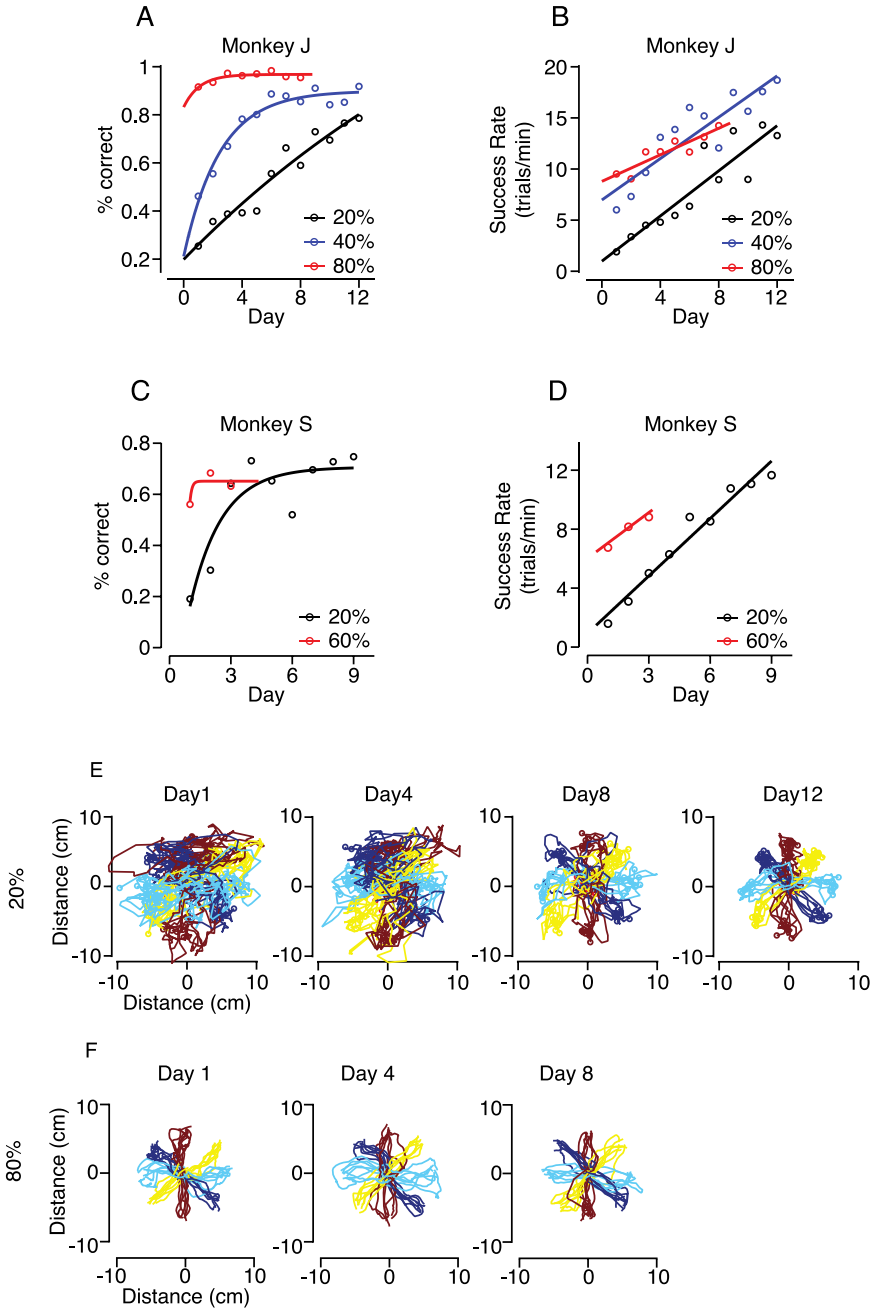


Figure 12. Rate of neuroprosthetic learning is invariant with initial levels of CLDA. (A) BMI task performance quantified by percentage of correct tasks with three different levels of initial decoder adaptation conditions. Different colors indicate different initial decoder adaptation levels. Thick lines are exponential fitting of data (The R^2 of

linear fittings for each data sets are 0.94, 0.95, and 0.75 for three conditions. Training days are 12, 12, and 8 days for different series of experiments). (B) BMI task performance quantified by success rate. Thick lines are linear regression (R^2 of fittings are 0.83, 0.79, and 0.83, and slopes are 1.1, 1.0, and 0.65 for different initial CLDA levels). (C-D) are the same behavioral matrixes from another subject, the training days are 9 and 3 days for 20% and 60% initial level of CLDA conditions. For the percentage correct, thick lines are exponential fitting of data, R^2 are 0.81 and 0.80. For success rate, slopes of linear fittings are 1.22 and 0.97 with R^2 of 0.96 and 0.96. (E) Example of brain control motion trajectories in different days of the 20% initial performance training series; different colors indicate different aiming targets. (F) Example of brain control motion trajectories on different days of the 40% initial performance training series, at the same training days as in H; different colors indicate different aiming targets.

While changing the amount of initial CLDA levels for each time series affected the initial percent accuracy, the success rates increased similarly for all conditions. Higher levels of initial CLDA allowed the monkeys to achieve a higher initial percent accuracy without requiring massive behavioral exploration (Figure 12 E, F). However, despite saturation in percent accuracy in late learning, the success rate continued to improve, indicating further skill refinement. These success rates improved linearly and were independent from initial decoder conditions during the training period (Figure 12 B, D) (linear regression; Monkey J: $R^2 = 0.87$; Monkey S: $R^2 = 0.96$). The slopes of these success rates were not significantly different among the various conditions (ANOVA; Monkey J: $F = 1.12$, $p = 0.34$; Monkey S: $F = 0.26$, $p = 0.63$). These non-significant changes in rates of learning suggest that varying levels of decoder adaptation impact the initial performance of the BMI but potentially bear little impact on neural adaptation.

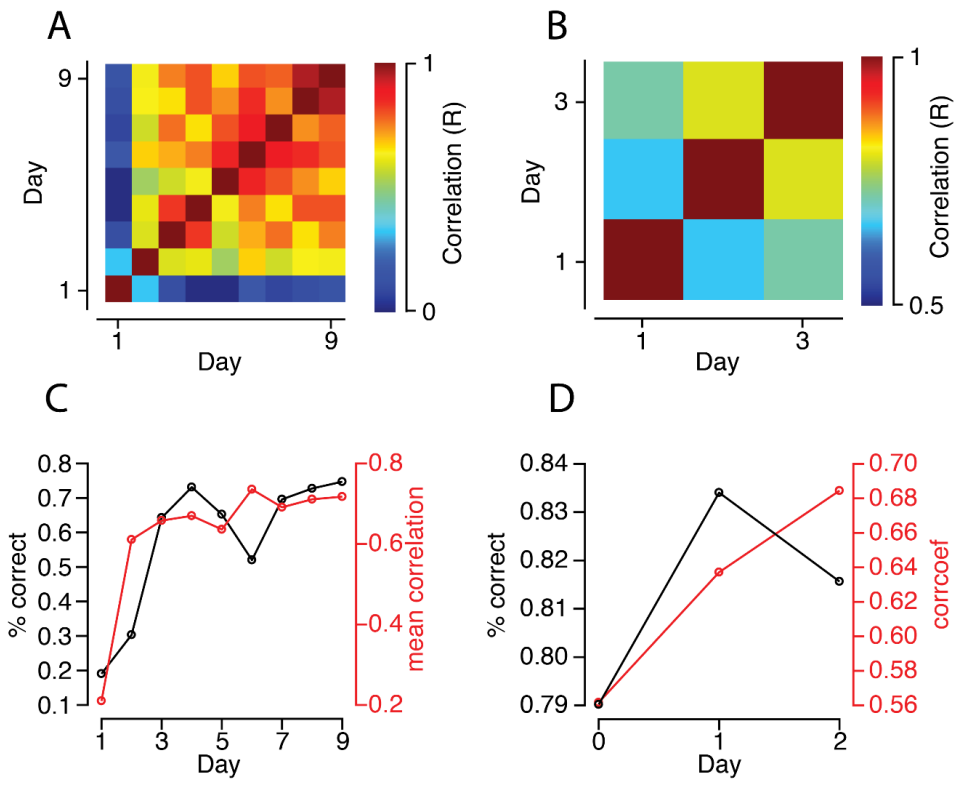


Figure 13 Stable cortical maps consolidated over training with different levels of initial CLDA for monkey S. (A-D) Same data as in Figure S1, but for subject S at the 20% and 60% initial CLDA conditions.

3.3.2 Neural adaptation underlying BMI learning with different levels of initial CLDA

To examine the neural correlates of skill learning, we first related the neural firing rates with the movement directions of the cursor as a tuning curve (Figure 14 A-C). We then tracked changes in the preferred tuning directions (PD) and modulation depths (MD) for each neuron over training days for each condition. Intuitively, the MD shows the dynamic range of the tuning curves, indicating the magnitude of preference for a given movement direction [6], [16], [29]. Significant changes in PD were observed when lower levels of decoder adaption were used (Figure 14D). Furthermore, on a population-level, the PDs consolidated over learning (Figure 13, 15). Consistent with past studies, the acquisition of this consolidated neuroprosthetic map was tracked closely by the percent accuracy in the BMI task [5], [16] and was expedited by the initial levels of CLDA.

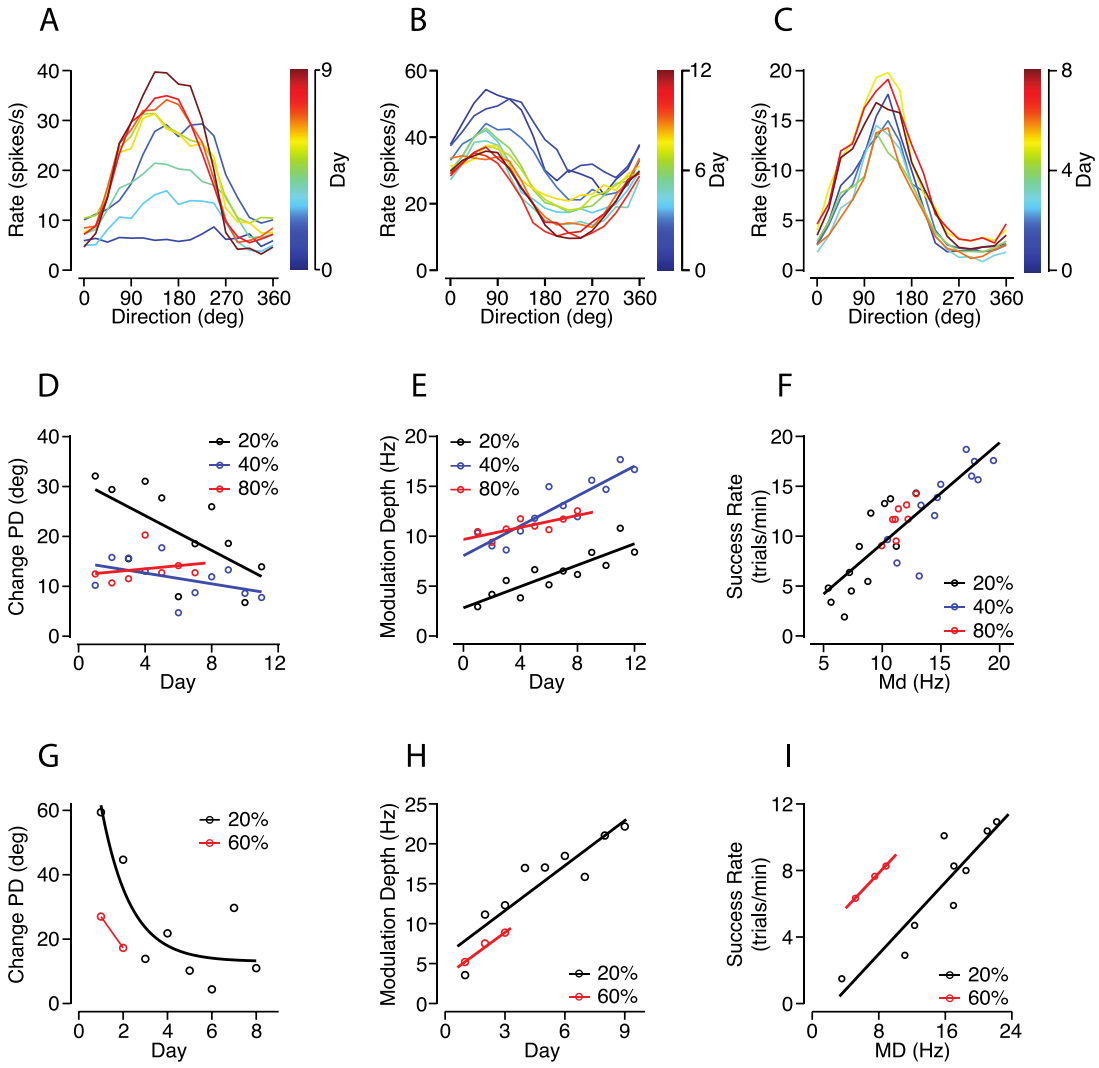


Figure 14. Neural adaptation that underlies neuroprosthetic learning in different initial CLDA levels. (A-C) Neural tuning to motion directions from example units over training. (Neurons in A are from subject S, 20% initial CLDA experiment; and in B and C are from subject J, 40% and 80% initial CLDA condition experiments. Different colors indicate different days during the training session). (D) Changes of preferred directions between consecutive training days across all stable BMI neurons for subject J. The early PD change is significantly larger than later PD change for 20% and 40% initial CLDA experiment with p=0.04, and 0.02, and not significantly different for the 80% CLDA experiment, p=0.56, Wilcoxon signed rank test. Lines are linear fittings of data, slopes of fittings are -1.7, -0.54, and 0.32 with R^2 of 0.40, 0.20, and 0.05 for 20%, 40%, and 80% conditions. (E) Changes of modulation depths for each series of training for subject S (main) and J (inset). (For subject J, early MD is significantly smaller than the MD in late learning for all three CLDA initial conditions, p=1.9e-04, 4.0e-04, and 0.04 for 20%, 40%, and 80% data, Wilcoxon signed rank test. Lines are the linear fittings of MD data, the

slopes are 0.54, 0.67, and 0.25, with R^2 of 0.73, 0.71, and 0.46). (F) Neural firing modulation depth increase correlated with the success rate improvement. Different colors indicate different initial decoder adaptation levels for subject J. The black line is the linear regression of data for all three conditions (r=0.85, p=9.5e-10, n=32). (G-I) are the same data as in D-F, but for monkey S. (G) Changes of neurons' preferred direction. Black line is exponential fitting of data with 20% initial CLDA, R^2 is 0.91 (For the 20% experiment, the early and later parts of preferred direction are significant different, p=0.03, signed rank=46, n=48, one tail Wilcoxon signed rank test). (H) Development of MD across training, the early MD is significantly smaller than later MD (p=7.3e-04, signed rank=2, n=48, one tail Wilcoxon signed rank test). The thick lines are linear fitting of data, slopes of fittings are 1.88 and 1.84, with R^2 of 0.81 and 0.98. (I) Correlation between the success rate and the MD, the thick lines are linear regression. The correlation coefficients are 0.9 and 0.99, with p values of 8.0e-4 and 0.04 for low (n=9) and high initial (n=3) CLDA conditions, respectively.

Interestingly, the neural MD increased significantly over learning in all conditions. Like the success rate, the rates of these increases were independent of the initial decoder adaptation levels (Figure 14 E, H) (linear regression; Monkey J: $R^2 = 0.89$; Monkey S: $R^2 = 0.87$). Furthermore, slopes of the MD increases were not significantly different between conditions (ANOVA; Monkey J: $F = 1.65$, $p = 0.21$; Monkey S: $F = 0.11$, $p = 0.75$). These increased MDs were also highly correlated with behavioral success rates (Figure 14F) (correlation; Monkey J: 0.85, p = 9.5e-10), indicating the neural basis of BMI learning in different CLDA conditions. Thus, while the acquisition of the neuroprosthetic map may be affected by the initial levels of decoder adaptation, the improvement in proficiency of skillful control is correlated with changes in the neural MD rather than the initial decoder adaptation levels.

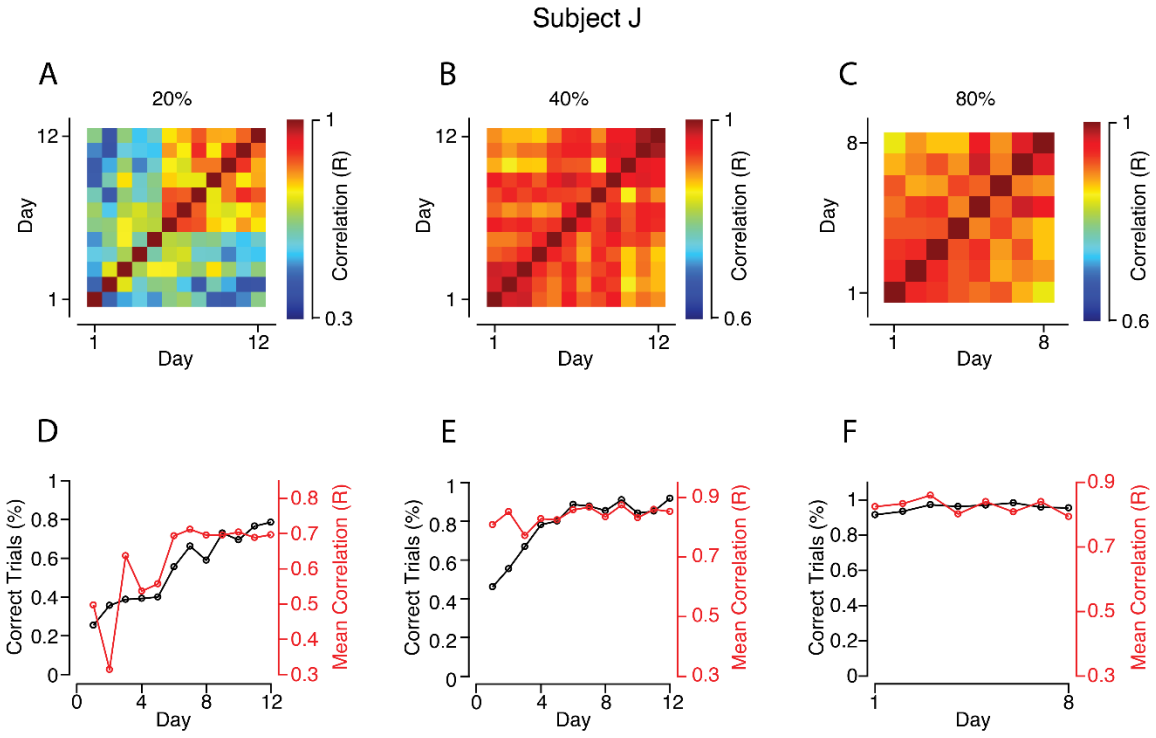


Figure 15 Stable cortical maps consolidated over training with different levels of initial CLDA for monkey J. (A-C) Color maps of pairwise correlation of neural ensemble tuning maps across different sessions in the series with different initial performances (A, 20%; B, 40%; C, 80%) for subject J. (D-F) Mean correlation of neural ensemble tuning map between training days overlaid with the percentage of correct trials across training for all different initial performance conditions.

3.3.3 Coordinated neural plasticity over training is independent of initial CLDA

Past work has shown that coordinated neural patterns emerge through neuroprosthetic learning [26]. However, it is unclear how such coordinated neural responses may emerge in the face of varying initial performance that derive from different levels of CLDA. We conducted factor analysis (FA) on the neural activity across all targets pooled together on a 100-ms decoder timescale for each condition to examine how coordinated patterns might emerge differently for the different CLDA conditions. The proportion of the neural variance underlying the entire population relative to the overall neural variance (SOT) was calculated as a correlate of consolidation. As with MD, we found that SOT increased over training in all conditions (linear regression; Monkey J: $R^2 = 0.81$; Monkey S: $R^2 = 0.56$) , and the rates of change in SOT were not significantly different between the different initial CLDA conditions (Figure 16 A, D) (ANOVA; Monkey J: $F = 2.51$, $p = 0.1$; Monkey S: $F = 0.11$, $p = 0.75$). The SOT was also highly correlated with the success rate

(correlation; Monkey J: r = 0.68, p = 2.24e-5) (Figure 16 B, E) as well as MD (correlation; Monkey J: r = 0.83, p = 3.52e-9) (Figure 16 C, F). Higher initial levels of CLDA did not affect initial levels of coordination in the recorded populations (Figure 14G). These results suggest that neural consolidation rates are invariant to initial decoder adaptation levels. Moreover, the consolidation was linked to changes in success rate, and not hindered by increased percent accuracy through CLDA.

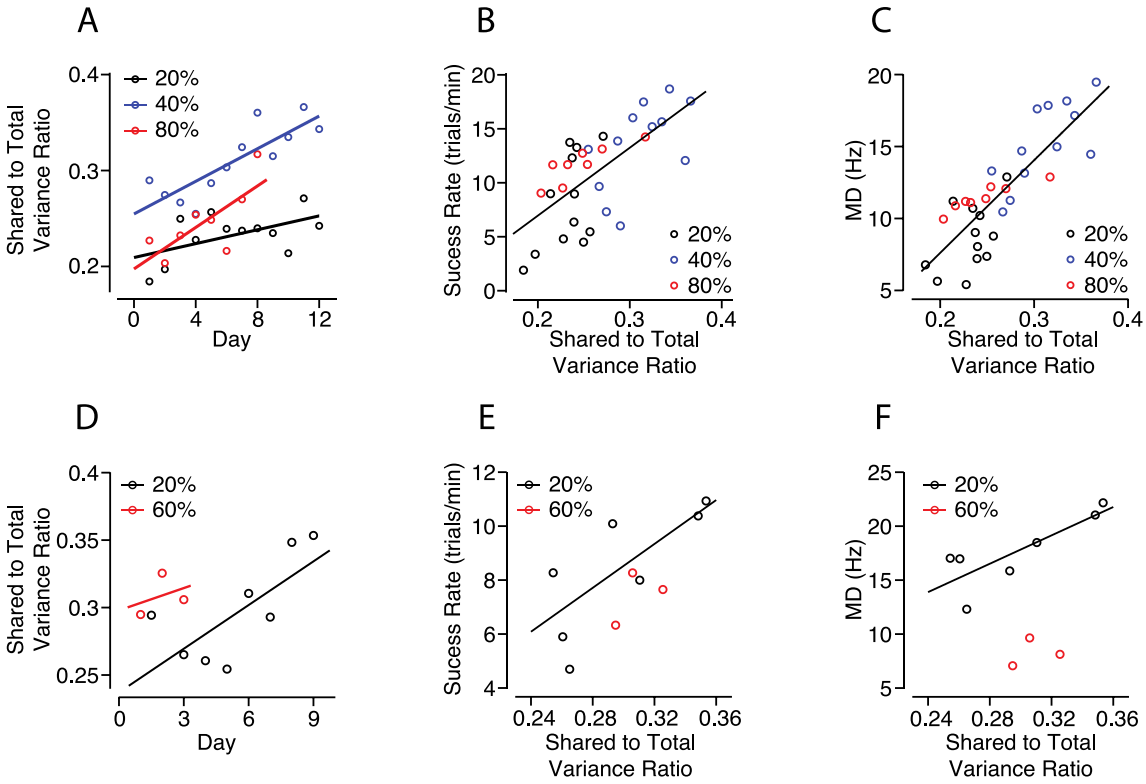


Figure 16. The developments of coordinated neural activities are independent of initial CLDA levels. (A) Ratio of shared variance over total variance developed over training for different initial levels of decoder adaptation for subject J (Lines are the linear fitting of data, the slopes of the fittings are 3.6e-3, 8.5e-3, and 10.8e-3, and the R^2 for the fittings are 0.29, 0.70, and 0.55). (B) Relationship of shared variance with the success rate (black line is the linear regression of all data from three different conditions. Correlation coefficient is 0.68, p=2.24e-5, n=32; Correlation coefficients for each condition are 0.56, 0.58, and 0.85 with p values of 0.059, 0.048, and 0.007, n=12, 12, and 8). (C) Relationship of shared variance with the modulation depth (the black line is the linear fitting for data from all conditions. Correlation coefficient is 0.83, p=3.5e-9, n=32; for each initial CLDA condition the correlation coefficients are 0.51, 0.72, and 0.94 with p values of 0.089, 0.008, and 0.0006, n=12, 12, and 8). (D-F) are the same data, but for monkey S. (D) Ratio of shared variance over total variance developed over training for different initial levels of decoder adaptation for subject S (the slopes of fittings are 0.011 and 0.006, R^2 of fittings are 0.52 and 0.13, n=8 and 3, for 20% and 60% conditions). (E) Correlation between the shared to total variance ratio and the success rate, the black line is the linear regression

of the 20% data. Correlation coefficient is 0.77, p=0.04, n=7. The 60% data were not fit because there were only three data points. (F) For subject S, the black line is the linear regression of the 20% data. The R^2 of fitting is 0.61, correlation coefficient is 0.82, p=0.02, n=7. The 60% data were not fit because there were only three data points.

3.3.4 Neural "reweighting" following decoder changes

While neurons adapt over learning alongside changes in decoder weights, we wanted to know how the brain leverages learned neural patterns to control re-adapted decoders. Note that because the 80% CLDA condition for Monkey J was conducted one month following the 20% and 40% conditions, which were done in a continuous 24-day period, we only focused the following analyses on the 20% and 40% CLDA conditions where 10 out of 16 neurons used for control were retained.

When plotted side-by-side, the SOT increased continuously when transitioning from the 20% to 40% condition despite changes in both the percentage correct and success rates (Figure 12 A-D). Looking at the 10 stable units, we found that the preferred direction tuning was stable across these decoder changes (Figure 17A). However, the correlations between unit MDs changed after decoder changes (Figure 17D). Together, these results suggest a "reweighting" mechanism where the consolidated patterns and tuning directions are more resilient to decoder changes, but control is adapted and improved by changing the firing rates in the preferred direction on shorter timescales.

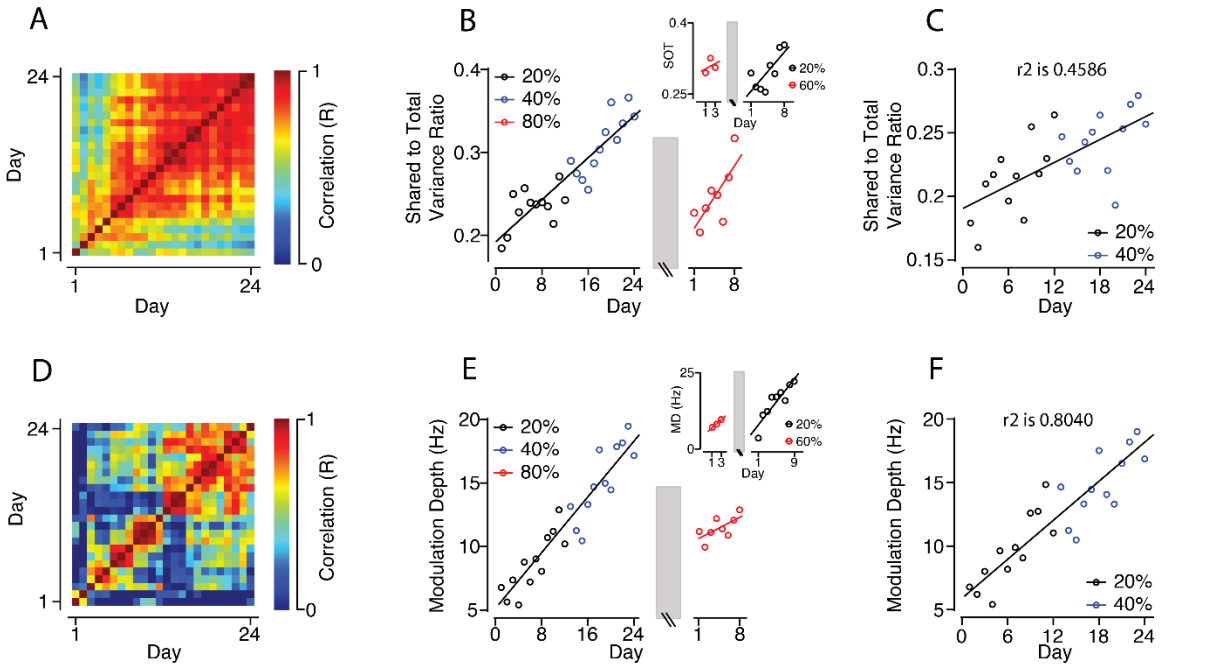


Figure 17 Neural adaptation across 20% and 40% experiments in monkey J. (A) Color maps of pairwise correlation tuning maps between different sessions from a subset of stable units ensemble across 20% and 40% training series (n=10 units, from subject J). (B) Considering the training of 20% and 40% series of experiments as one training, SOT across all units (16 units). Red dots are SOT for 80% series experiment. The gray block indicates there is a long time gap between two series of training. Straight lines are linear fitting. Slope of fittings are 0.006, and 0.011, R^2 of fittings are 0.81 and 0.55. For SOT of subject S, R^2 of linear fittings are 0.13 and 0.57 for 60% and 20% initial CLDA conditions. (C) Shared to total of neural covariance (the black line is linear fitting of data, R^2=0.46) over training series across 20%, 40% stable units (n=10, from subject J). (D) Color map of the pairwise correlation of MD between different sessions from the stable units ensemble across 20% and 40% training series (n=10 units, from subject J). (E) MD across the stable units, the straight line is linear fitting. Slope of the fits are 0.56, and 0.25. The R^2 of fitting is 0.89 and 0.46. For subject S in the subset, R^2 of fittings are 0.99 and 0.81. (F) Mean neural modulation depth (Linear fitting, R^2=0.80) over training series across 20%, 40% stable units (n=10, from subject J).

3.4 Discussion

Brain-machine interfaces are exciting in clinical contexts not only as potential assistive devices, but because they may offer rehabilitation to users as well. Multiple studies have shown that the brain adapts as control of a BMI improves. Translationally, this could mean using BMIs to engage neural changes in motor cortex as well as other cortical and subcortical regions. Orsborn et al., 2014 showed that changes in neural modulation depths and preferred tuning directions were shaped by decoder weights. Furthermore, the degree of neural adaptation, measured by mean MD and changes in PD, was dependent on the task percent correct. Based on these past findings, we thus wanted to know if high initial decoder performance would hinder how the brain learns and consolidates new patterns. While the magnitudes of neural adaptation were affected by decoder adaptation levels, we found that the rate of changes were unaffected by initial decoder adaptation level. Our results suggest that CLDA simultaneously allows for improved initial performance of a BMI while still facilitating neural adaptation.

Our results also implied the importance of longer time of training to establish stable, robust BMI control specific circuits in order to own the BMI, which could not be achieved through CLDA. Also, other designs of CLDA algorithm and training protocol can be investigated that will further improve the proficiency of learning. With brief CLDA in the middle of a training series, we observed significant improvement of success rate, MD, and SOT through an interactive decoder and neural co-adaptation. However, in this 'two-learner' system, the initial CLDA only adapted to intuitive tunings of arbitrarily recorded neurons, which potentially could be sub-optimal for the decoding. Thus, it still

not clear what is the optimal method to initialize the decoder, if the initial brain searching is necessary for getting an optimal decoder, and even what is the neural principal of optimization for this decoder design. All of these needs further investigation.

At the population level, changes in neural firing patterns have been characterized by changes low-dimensional neural subspaces [26], [27], [31], [32]. Studies have shown neurons to collapse their activity onto low-dimensional spaces over learning, suggesting consolidation of firing rate patterns as performance improves. We found that while initial BMI performance was dependent on CLDA level, the rates of neural consolidation were invariant. Higher levels of CLDA indicate better decoding of each individual unit (i.e. high-dimensional). Thus, our results indicate that while CLDA may be better fitting to each unit, skillful control of a BMI is marked by population-level consolidation onto low-dimensional spaces.

Synergistically combing neural adaptation and machine learning is an important concern in the context of translational BMIs. Our results show that neural adaptation occurs at similar rates regardless of decoder adaptation level as subjects gain skillful control of a BMI. Rather, the initial decoding accelerates the control that can be achieved on the first day of practice. Thus, we have shown that is beneficial to have high levels of CLDA from learning onset in order to immediately act as an assistive device while still having the potential to be a neural-rehabilitative device.

3.5 Methods

3.5.1 Electrophysiology

Two male non-human primates (rhesus macaque, Macaca mulatta) were chronically implanted with 128 channels electrodes arrays in the arm region of the primary motor cortex (M1). 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 at Berkeley Institutional Animal Care and Use Committee. Neural activity was recorded using the MAP system (Plexon) with 128 channels. Channel-level activity was separated into individual neural spikes using an online spike sorting application (Sort Client, Plexon) by setting a threshold of 5.5 times the standard deviation of the mean signal amplitude [15], [16], [53]. At the beginning of each day, this threshold was acquired while the subject sat quietly without performing any behavioral task. A neuron template was used for recording after the channel threshold was defined.

3.5.2 Behavior

Each monkey was trained to perform a self-paced 2D center-out reaching task to eight targets [15], [16]. To initiate a trial, subjects moved the cursor to the center target, and held it for 250-400 ms. After the go cue (color changes in the center target), the subjects moved the cursor to the indicated peripheral target, presented in a pseudo-random order. A successful trial required achieving the indicated target and holding for 250-400 ms. Successful trials were rewarded with juice. Reaching time was 3-10 seconds. Failing to reach the target in that time or leaving the target early resulted in trial abortion, and failed trials were repeated.

3.5.3 BMI implementation

In a BMI experiment, the cursor movement was controlled in real time by neural activity using a position-velocity Kalman Filter (KF) decoder. The KF assumes two linear state models:

$$x_t = Ax_{t-1} + w_t$$
$$y_t = Cx_t + q_t$$

where x_t and y_t are the current cursor state and neuron activities, w_t and q_t are noise terms. In the cursor state evolution model, the cursor state is defined as:

$$x_t = [p_t^{hor},\ \ p_t^{ver},\ \ v_t^{hor},\ \ v_t^{ver},\ \ 1]$$

where p and v are the position and velocity of cursor movement, and hor and ver indicate the horizontal and vertical direction of movement. The 1 constant is used to compensate the nonzero mean observation, such as the baseline firing rates. In the observation model, y_t is the neural activity from BMI units, which were updated in non-overlap 100 ms bins [15], [16], [35].

Initial decoder was seeded mainly using a visual feedback protocol where the subject passively watches a cursor movement as in a center-out task. After the decoder is seeded, continuous CLDA was performed using the SmoothBatch algorithm [14]–[16]. In each training series, CLDA was used at the beginning of training to improve the performance to a specific level. After the decoder adaptation, the subject performed BMI using the fixed decoder in the training series, or a briefly adapted decoder if the neural activity shifted (such as lost neurons in the BMI ensemble).

3.5.4 Closed-loop decoder adaptation (CLDA) and SmoothBatch

CLDA is a technique that leverages the machine-learning algorithm to continuously modify the decoder parameters in order to match the neural activity pattern during closed-loop BMI operation. The design of the CLDA algorithm involves not only which and how the parameters adapt, but also the rate of updating the adapted parameters [14], [15], [35], [54]. In this study, an algorithm 'SmoothBatch' was used, which updates the decoder parameters at the intermediate time scale (1-2min). In SmoothBatch, the subject's intended kinematics was inferred with the knowledge of the task goal [15], [16], [35]. This inferred kinematics and the online neural activities were collected in the batch-based period to estimate the update of the motion kinematics. The estimation was then weighted with the previous parameters using an exponential averaging window to yield the updated decoder parameters.

To investigate how the levels of decoder adaptation affects BMI prosthetic leaning, in this study, three series of experiments were performed with different initial levels of CLDA. For the first training day of a series, the termination of using CLDA was based on the estimation of subject behavioral performance, which are around 20%, 40%, and 80% of correct trials. Then the decoder was fixed for each training series unless there is neuron dropping or performance dropping in the middle of training. For the 20% and 40% initial CLDA experiments, the training series was 12 days, and for the 80% initial CLDA experiment, the training series was 8 days.

3.5.5 Analysis

Behavior performance was quantified using the percentage of completed trials amount the initialed trails (percentage correct), and the number of trials that finished in one minute (success rate).

The neuron directional tuning was estimated by calculating the mean firing rate related to the BMI controlled cursor motion directions. Preferred directions (PD) of neurons were calculated by modeling firing rates to the corresponding directions of cursor movements using a cosine-tuning model. The modulation depth (MD) was calculated as the specificity of preference of a neuron's firing activity based on the same tuning model. Directions are binned in segments of 22.5°, which gives 16 direction bins. We assessed the goodness of fitting using the R^2 only tuning curves with R^2 larger than 0.3 were further analyzed.

With learning, the PD of neurons changed and stabilized in all conditions (Fig. 3D), forming a stable neuroprosthetic map [6], [16], [29], which is quantified by the correlation of direction tunings at the neuron ensemble level (Figure 14 G-I). The mean correlation coefficient of the ensemble tuning map between training days tracked closely

the percentage accuracy of brain control behavioral results in all conditions, which echoes the previous results that the formation of a stable map underlies the neuroprosthetic learning (Figure 14 J-L).

To compare the behavioral performance and the neural firing across the learning, the training series was divided into early and late periods. Specifically, for the 20% and 40% experiment, the results from the first and last three days were defined as early and late learning periods. For the 80% initial CLDA experiment, the first and last two days were defined as early and late. For subject S, we compared the first two days and the last two days as the early and late parts of training.

3.5.6 Factor Analysis

To address how the neural ensemble coordinately consolidates over neural prosthetic learning, factor analysis was used to decompose the population activity into private and shared signals [26]. If N units were recorded in the experiment, the neural responses $x \in \mathrm{R}^{\mathrm{N}}$, were modeled as the sum of the mean responses μ, private signal ε, and the shared signal $z \in \mathrm{R}^{\mathrm{K}}$, where $\mathrm{K} < \mathrm{N}$. Then the neural response can be modeled as:

$$x = \mu + \varepsilon + \Lambda z$$

the transformation matrix $\Lambda \in \mathrm{R}^{\mathrm{N}*\mathrm{K}}$ projects the neural response vector x onto a lower dimension variable z, where $\mathrm{z} \sim \mathrm{N}(0, \mathrm{I})$. In our analysis, the $x^{\mathrm{shared}} = \Lambda z$, and the $x^{\mathrm{private}} = \varepsilon$, where $\varepsilon \sim \mathrm{N}(0, \psi)$. The z and the private signal ε are independent. Thus, the neural population responses can be presented by:

$$x = \mu + x^{\mathrm{private}} + x^{\mathrm{shared}}$$

where x^{private} and x^{share} are independent. If we denote $x^{\mathrm{private}} \sim \mathrm{N}(0, \Sigma^{\mathrm{private}})$ and $x^{\mathrm{private}} \sim \mathrm{N}(0, \Sigma^{\mathrm{shared}})$, then

$$\Sigma^{\mathrm{private}} = \psi$$
$$\Sigma^{\mathrm{shared}} = \Lambda\Lambda^{\mathrm{T}}$$
$$\Sigma^{\mathrm{total}} = \Sigma^{\mathrm{private}} + \Sigma^{\mathrm{shared}}$$

The FA model was fitted using the data pooled from all trials with different targets. Model parameters were fitted by log-likelihood maximization of the EM algorithm for a shared dimensionality of $\mathrm{K} < \mathrm{N}$. We analyzed the eigenvalues of the shared variance, and the share dimensionality is defined as how many dimensions are needed to explain 90% of the variance [26].

3.5.7 General linear regression model with categorical covariates and statistical tests

To investigate the difference of slopes of learning curves for different conditions, we fitted a linear regression model with categorical covariates using a dummy variable, which represents the different initial levels of CLDA. To investigate the influence of the initial CLDA levels to the slope of learning, a learning rate-condition interaction term was included to the model. Then to test the significance between the slopes of learning is equivalent to test if the coefficients of interaction terms are significant from 0, which can be derived from the ANOVA test.

Chapter 4

Local-Network Adaptation Underlying Abstract Skill Learning

4.1 Abstract

The emergence of coordinated neural dynamics in neurons controlling a brain-machine interface (BMI) has been shown to occur with neuroprosthetic learning. However, it is unclear if these changes in coordinated firing activity are common to all neurons in motor cortex or if they are unique to the neurons with direct control. To answer this question, we used a BMI paradigm with two macaque monkeys to examine how coordination of activity in neurons that act as decoder inputs (direct subpopulation) compares to that of the rest of the recorded population (indirect subpopulation) as performance improves. We find that direct neurons increase variance in their firing rates more than indirect neurons to explore the neural space, forming coordinated patterns over time. Remarkably, both subpopulations were found to be comparably correlated with task variables throughout learning. Our results suggest that BMI control is improved by increasing activity and exploration of direct neurons, rather than suppressing activity of indirect neurons.

This work was in collaboration with Ellen L. Zippi.

4.2 Introduction

Brain-machine interfaces (BMIs) allow users to interact with their external environment by transforming neural signals into actions taken by effectors, such as computer cursors or prosthetic arms. Because volitional movement of the BMI is desired, primary motor cortex (M1) is a commonly used region for the selection of neural signals to use as input to the BMI as its activity is correlated with voluntary limb movements [5], [45], [55], [56]. Previous studies have demonstrated adaptations in neural firing patterns in motor cortex as subjects learn to control BMIs. These neuronal adaptations have been shown to be stable across time, readily recalled, and resistant to interference of learning control with a second decoder [6]. Furthermore, subjects were able to volitionally control the BMIs without movements of the natural arm over learning despite using decoders that were initially trained on natural arm movements. This highlights the ability for neurons in M1 to dissociate from the rest of the motor circuit [55] and suggests that the animals learn an abstract solution to BMI control that can be reentered and gradually refined [4].

It has also been shown that neurons directly controlling the BMI undergo consolidation of coordinated activity patterns to acquire skillful control of the BMI [4], [26], [27]. As BMI control improves, covariance of neural activity between neurons increases, suggesting an increase in correlated activity in the population. These changes are paired with decreases in neural dimensionality. Notably, these low-dimensional subspaces have been shown to rotate in neural space, indicating changes in coordinated firing patterns of neurons. Additional studies have demonstrated that changes in these low-dimensional structures, sometimes referred to as manifolds, take many days and BMI tasks where the decoder is better aligned to the starting manifold are easier for subjects to learn [27], [28], [32].

While these studies focused on learning in the population of neurons whose neural activity was used as input to the BMI decoder, less is known about how other cells in the same brain region whose activity is not directly required for BMI control adapt over time. Here, we ask how the rest of motor cortex adapts as the brain learns to control a BMI. Specifically, we compare the subset of motor cortical neurons that were assigned a causal role during BMI control ("direct neurons") to all other recorded motor cortical neurons that were recorded, but not used as direct input to the BMI ("indirect neurons"). Note that while indirect neurons have also been used to refer to cells outside of motor cortex, we define indirect neurons specifically in the scope of motor cortex in this study. To understand how neuroplasticity occurs in the brain during abstract learning, we focus on

changes that occur between subpopulations in a local network. Comparing the changes in direct and indirect subpopulations that emerge over learning may help us understand how a network solution is generated.

Past work has shown that both direct and indirect neurons within motor cortex change preferred tuning directions over learning [29]. However, compared to direct neurons, indirect neurons were shown to undergo smaller changes in modulation depth, defined as the peak-to-peak range of the cosine tuning curve [29]. The stability of tuning properties of the indirect population suggests that the indirect neurons play an active support role in learning BMI control, possibly by shaping activity of the whole network. Alternatively, indirect activity may have a negative role in learning BMI control, such as introducing noise, making credit assignment more difficult and the relative reduction in modulation may serve to limit interference by this population [3]. A recent study found that changes in modulation depth correspond to changes in neural variance and exploration in the direct population as control of a BMI is learned. Subsequently, another possible explanation for the differential modulation of neuronal populations based on their causal link to the BMI is that there is a differential amount of neural exploration. Furthermore, it is unclear whether these differences lead to differences in coordinated firing. In this study, we observe how direct and indirect neurons adapt over learning to understand the driving forces behind neural adaptation in the context of BMI learning.

4.3 Results

Two rhesus macaques were chronically implanted with bilateral microelectrode arrays in primary motor and dorsal premotor cortices, with electrodes from a single hemisphere used for BMI control and subsequent analyses. The monkeys learned to perform a two-dimensional, self-initiated center-out BMI task, in which they drove a cursor under neural control to one of eight randomly selected peripheral targets for a juice reward (Figure 18A). The next trial was initiated by driving the cursor back to the center target.

To control the BMI, spiking activity from a selected subpopulation of neurons was decoded into cursor positions (direct subpopulation). The remaining neurons were recorded but did not directly control the BMI (indirect subpopulation). We analyzed population-level activity changes in both the direct and indirect subpopulations over the course of learning. Changes over learning were analyzed over ‘‘training epochs,’’ with each epoch consisting of a constant number of trials.

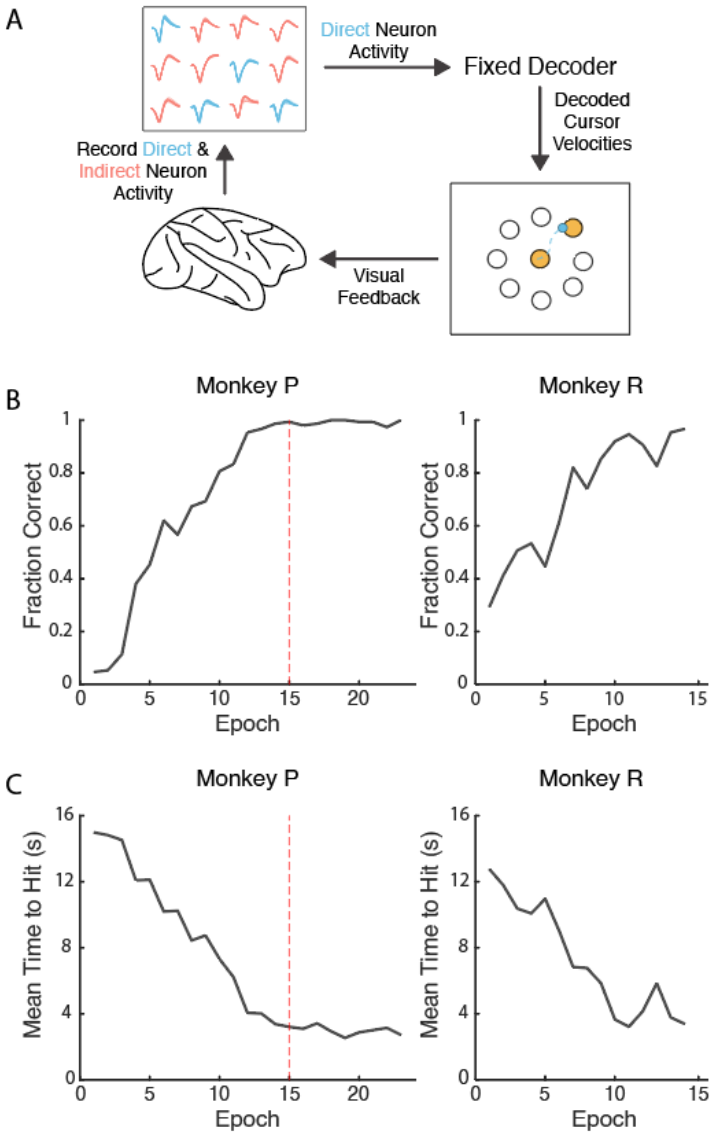


Figure 18 Experimental setup and behavioral performance. (A) Experimental setup from Ganguly & Carmena, 2009. (B) Performance improves over the first 15 training epochs for Monkey P and Monkey R. Each training epoch consists of 150 initiated trials. Fraction of initiated trials that were successful increased over training epochs. (C) The time to reach a target decreased over training epochs.

To estimate the relative amount of neural exploration over learning, we calculated the variance in neural firing rate and normalized it by the mean variance in early learning (first seven training epochs). There was a significant increase in this relative variance in late learning (last seven training epochs) in the direct subpopulations for both animals, suggesting an increase in neural exploration by the direct subpopulation (Figure 19 A).

Factor analysis (FA) was conducted to separate correlated and uncorrelated patterns from the neural activity of the entire stable, recorded population (direct and indirect subpopulations combined). As a measure of coordinated population-level activity, we considered the proportion of the neural variance that was captured by these correlated patterns (i.e. shared variance) and thus calculated the shared-over-total variance (SOT) ratio [26]. The SOT ratio increased from early to late learning, signifying an emergence of coordinated neural firing patterns across the entire population (Figure 19 B).

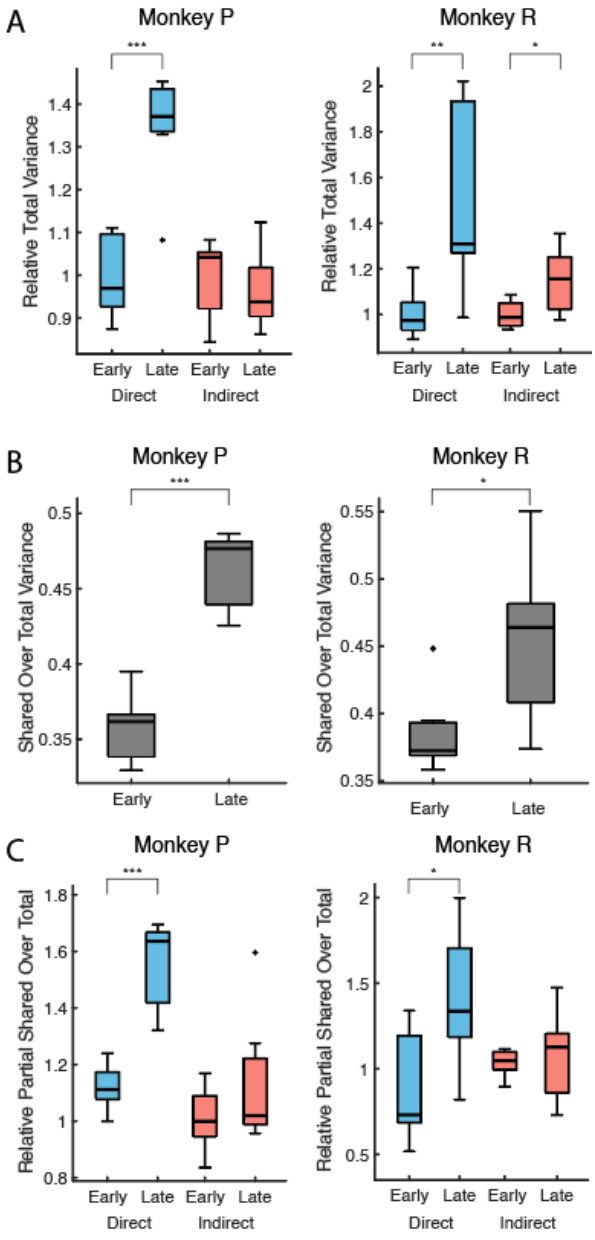


Figure 19 Neural activity consolidates with learning. (A) Direct neurons exhibit a larger relative increase in neural variance than indirect neurons (Unpaired t-test; Monkey

P: direct p = 8.63e-5, indirect p = 0.523; Monkey R: direct p = 0.007, indirect p = 0.019). (B) The shared-over-total variance (SOT) ratio increases between early and late learning (Unpaired t-test; Monkey P, p = 2.22e-6; Monkey R, p = 0.020), indicating an overall consolidation of neural activity in the entire recorded population (direct and indirect). (C) Respective contributions of direct and indirect sub-populations to the SOT ratio in early and late learning relative to contributions in early learning (Unpaired t-test; Monkey P: direct p = 2.99e-5, indirect p = 0.219; Monkey R: direct p = 0.022, indirect p = 0.624).

In order to determine if this increase in SOT ratio was driven by an increase in coordinated patterns across the entire population or within a subpopulation, we considered the partial shared-over-total variance (pSOT) ratio. Intuitively, the pSOT ratio asks how much of the overall neural consolidation was driven by one subpopulation versus the other. We see that, relative to early learning, there was a significant increase in pSOTdirect in late learning but not in pSOTindirect (Figure 19 C), signifying that the increase in coordination of population activity seen across the whole recorded population was driven by the increase in coordination of population activity within the direct subpopulation.

To characterize how this exploration manifested geometrically in neural space, we quantified these changes by calculating the shared alignment pairwise between each training epoch's shared covariance matrix for each subpopulation (Figure 19 A) [26]. The shared alignment compares how much of the shared space of one epoch projects onto the shared space of another epoch. Intuitively, given two 2-dimensional shared subspaces, the shared alignment compares the angle between the two planes. Likewise, orthogonal planes would result in a shared alignment of 0 and perfectly aligned planes would result in a shared alignment of 1. If the shared subspace consolidates with learning, as has been shown in direct subpopulations [26], we would expect the shared subspace to rotate from the initial subspace over learning. If the shared subspace remains fixed over learning, we would predict that the alignment between the first epoch and later epochs would remain high. We found that the shared alignment decreased from the first epoch for each subpopulation, but the magnitude of decrease was significantly greater for the direct subpopulation (Figure 19 B). This suggests that the direct subpopulation explored and rotated its low-dimensional subspace while the indirect subpopulation remained relatively static and deviated less from its initial subspace (Figure 20 C).

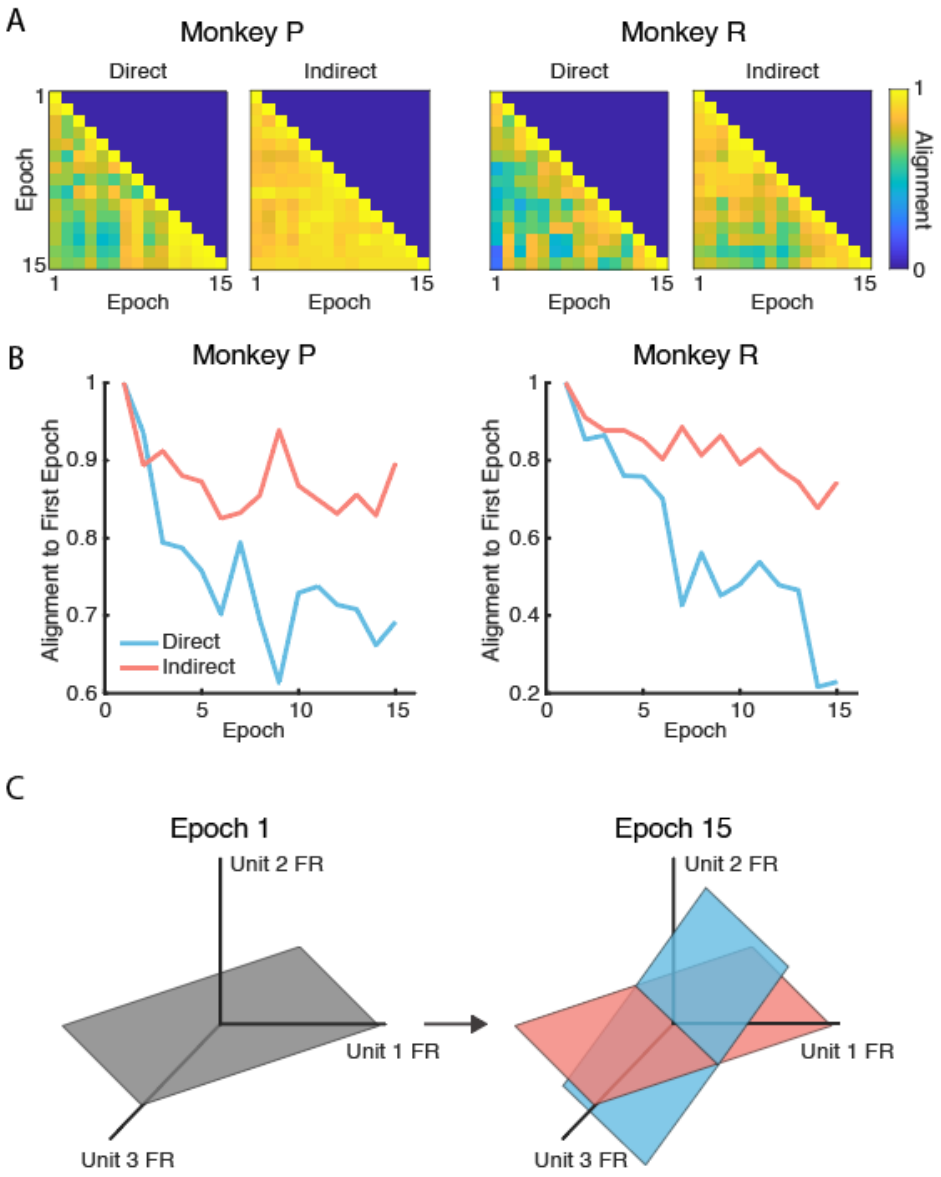


Figure 20 Direct neurons align to a stable low-dimensional solution subspace. (A) Shared alignment was calculated pairwise between epochs for each subpopulation. (B) Alignment diverges from the first epoch in both ensembles, but the magnitude of change is significantly greater for the direct ensemble (Regression on difference; Monkey P: R^2 = 0.367, p = 0.0166; Monkey R: R^2 = 0.708, p = 8.5e-5). (C) Interpretive schematic of changes in low-dimensional subspaces. Neurons initially fire on a low-dimensional subspace (gray plane, left). Over learning, the direct ensemble explores and rotates its low-dimensional subspace (blue plane, right). However, indirect neurons remain relatively static and deviate less from the initial subspace (red plane, right). The changes in coordinate neural dynamics are specific to each subpopulation.

While all of these results show a difference in adaptation between direct and indirect subpopulations, it is still unclear if the relative suppression of exploration of indirect neurons is to limit noise in the network, decrease metabolic costs, or something else. Despite relatively low levels of exploration, we find that the indirect population remains task-relevant throughout learning (Figure 21).

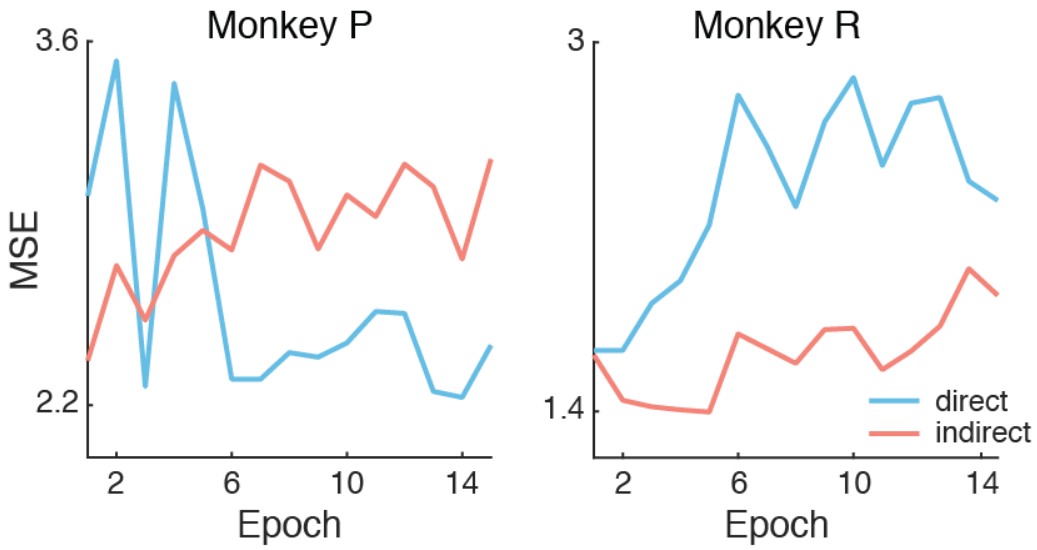


Figure 21 Task-relevant information is present in both the direct and indirect population over the course of learning. Offline linear regressions were used to estimate cursor positions (x and y) based on the binned spike activity for the direct population and indirect populations separately at each epoch. Mean squared error (MSE) was calculated from the residuals, defined as the Euclidean distance between the actual and estimated cursor positions for each time bin.

We performed offline linear regressions on the binned spike activity of the direct subpopulation to predict cursor positions and compared its performance to random samples of the same size selected from the entire recorded population (direct and indirect subpopulations). We calculated the mean squared error (MSE) from the residuals, defined as the Euclidean distance between the actual and estimated cursor positions at each time bin, and found that the MSE for the direct subpopulation was not significantly different from the distribution of MSEs obtained from random samples of neurons (Figure 22). Thus, the smaller change in exploration by the indirect neurons is not a result of activity suppression. Rather, it is indicative of a relative increase in exploration by the direct neurons.

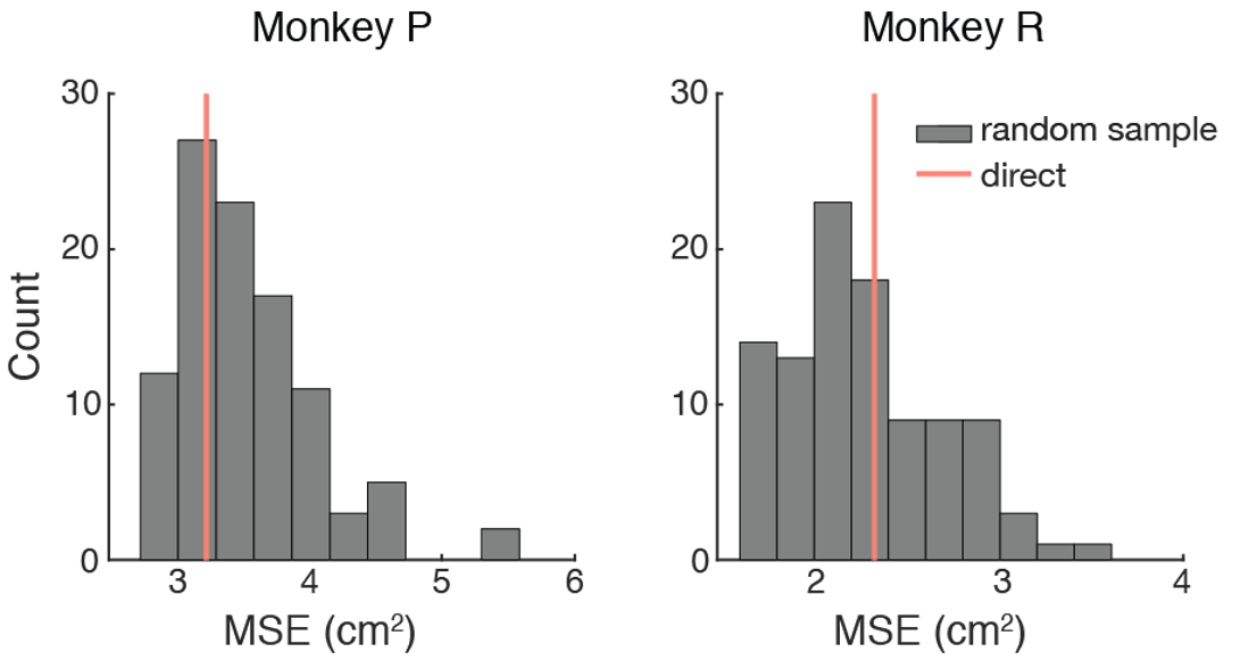


Figure 22 Task-relevant information is present in the entire neural population. Offline linear regressions were used to estimate cursor positions (x and y) based on the binned spike activity for the direct population (Monkey P: N = 15; Monkey R: N = 10). Mean squared error (MSE) was calculated from the residuals, defined as the Euclidean distance between the actual and estimated cursor positions for each time bin. The MSE for the direct subpopulation was within the distribution of MSEs from random sampling of neurons (Permutation test; Monkey P: MSE = 3.23, p = 0.33; Monkey R: MSE = 2.34, p = 0.62).

To examine whether direct and indirect neurons retain these exploratory behaviors with different decoders, we analyzed data from a second experiment in which Monkey P learned to perform the same BMI task with a new decoder following proficient control with the original learned decoder (Figure 23). Eight experimental blocks were performed over the course of four days, alternating between control with the new and previously learned decoder each day.

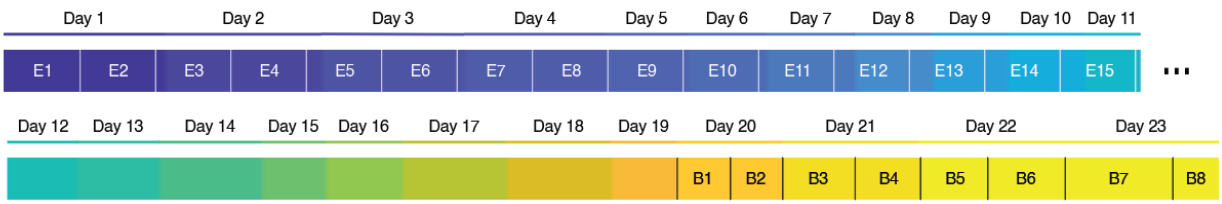


Figure 23 Experiment timeline for Monkey P. Colors indicate different recording days. The lengths of the colored rectangles are proportional to the number of trials performed each day. Vertical black lines show the divisions of training epochs and training blocks. Epochs, labeled E1-E15, each consist of 150 self-initiated trials. Blocks, labeled B1-B8, consist of the last 8 days of the experiments and alternate between the new decoder (odd blocks) and old decoder (even blocks).

Across the eight blocks, the direct subpopulation increased its neural variance over blocks relative to the neural variance in the same neurons during the first training epoch from the first experiment. The indirect subpopulation did not significantly increase its relative neural variance, suggesting a relative inhibition of exploration following initial learning (Figure 24 B). However, exploration did increase in the indirect subpopulation whenever the new decoder was used. Together, these results suggest that the relative inhibition of exploratory behaviors was decoder-specific. New perturbations may evoke new exploratory behaviors in both direct and indirect subpopulations. Furthermore, shared subspace alignment was calculated pairwise between training blocks for both the direct and indirect subpopulations separately. Both direct and indirect subpopulations retained a stable, low-dimensional subspace across all eight blocks despite repeated decoder swaps (Figure 24 C,D).

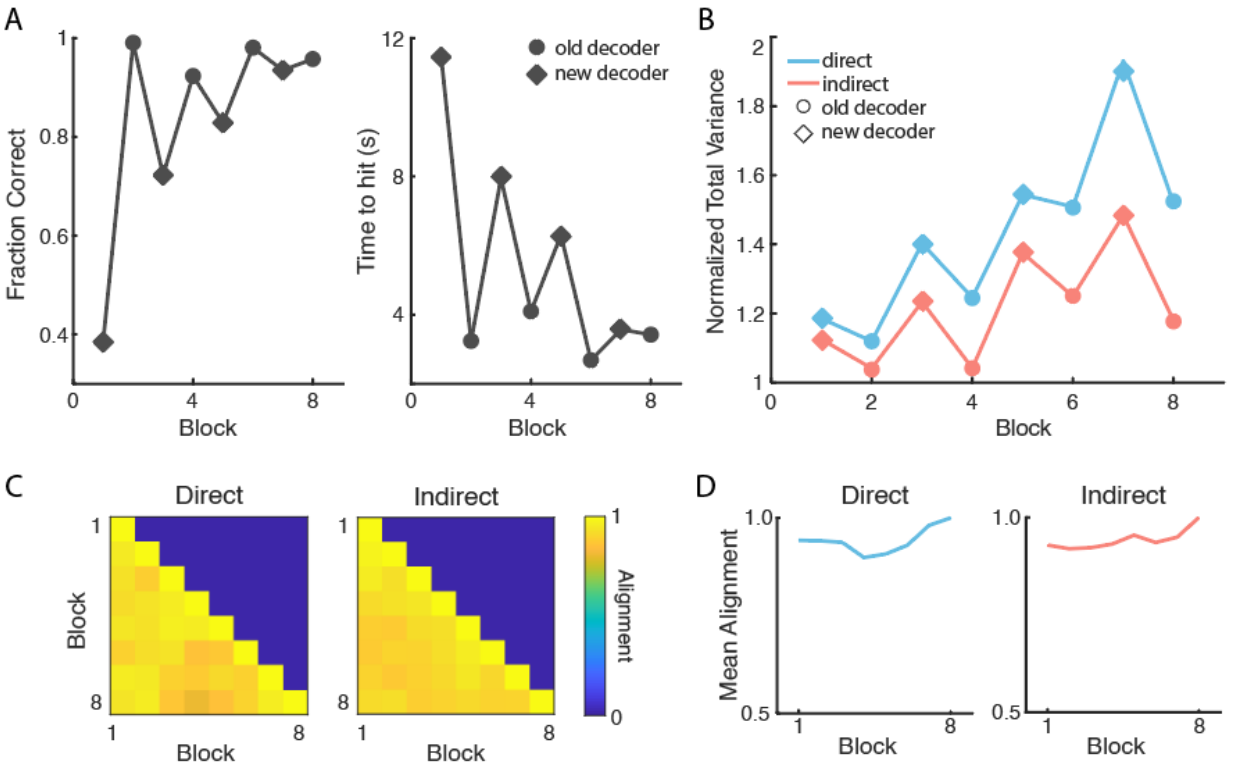


Figure 24 Subpopulations retain roles when learning a new decoder. (A) Monkey P learns to perform BMI with a new decoder following proficient control with the old decoder. 8 experimental blocks were performed, alternating between a new decoder (diamond) and the previously learned decoder (circle). Fraction of initiated trials that were successful increased over training blocks (left) and the time to reach a target decreased over training blocks (right). (B) Direct cells increase their neural variance over blocks relative to Epoch 1 with the old decoder. Indirect cells do not significantly increase their neural variance, suggesting a relative inhibition of exploration following BMI learning with the old decoder. However, neural variance increased while learning to control the new decoder for both subpopulations (square dots). (C) Shared alignment was calculated pairwise between training blocks for both sub-populations. Both direct and indirect ensembles retain a stable low-dimensional subspace. (D) The mean shared alignment within each subpopulation is high and stable across blocks.

4.4 Discussion

Our results demonstrate that subpopulations of neurons, defined by whether or not their activity is a direct input into a BMI, exhibited differential changes in coordination of their neural activity. Although both subpopulations demonstrated some changes in coordination patterns, the subpopulation of neurons whose activity was delegated to control the effector showed larger changes (Figures 19, 20). Importantly, all neurons exhibited task-relevant

activity over learning (Figure 22). Together, our results suggest that neural adaptation to BMIs is driven by network solutions that aim to improve performance through exploring new patterns rather than suppressing pre-existing ones.

We found that as both animals improved control of the BMI, more of the neural variance was captured by shared spaces, indicating a consolidation of neural activity over learning (Figure 19B). However, this increase in consolidation was primarily carried by the direct subpopulation and driven by neural exploration (Figures 19 A,C). Past work has shown that neurons explicitly inducing network interference are suppressed over learning [16]. It has been posited that one potential explanation of the relative decrease in modulation depth seen in the indirect subpopulation is a mechanism to reduce interference and to boost the signal-to-noise ratio [3]. Here, we find evidence that motor cortex forms a network solution that improves control by increasing activity of direct neurons, thereby boosting the signal, rather than suppressing activity of indirect neurons.

Moreover, our results show the effects of neural variance on how neural activity consolidates with learning. Changes in shared alignment and the SOT ratio imply a rotation and consolidation of neural activity that was more pronounced in the direct subpopulation, which had a larger increase in neural variance relative to the indirect subpopulation. These results, along with previous work demonstrating that changes in modulation depth are closely linked with changes in neural variance, suggest that the relative increase in modulation depth of the direct subpopulation [29] is likely a correlate of increased neural exploration of the direct subpopulation. Furthermore, we found similar performance of offline decoding of cursor positions using random samples of neurons from the local network, demonstrating that there is valuable signal present in both direct and indirect neurons. Thus, neural adaptation during BMI learning may be driven by differences between the neural space and control space of the decoder.

Past work has shown that the low-dimensional neural space of direct neurons tends to align with the control space of the decoder [26]. Similarly, other work has demonstrated larger rotations of the neural space when decoders were trained to be explicitly out-of-plane with the neural space [27]. Thus, if we consider indirect neurons as those with decoder weights set to zero, the indirect neurons would always be aligned to the control space defined by zeros and therefore not need to rotate their low-dimensional space. By extension, further work is needed to determine if the amount of neural exploration is driven by misalignment to the control space, magnitude of decoder weights, or both. Our results suggest that the presence of decoder weights have an impact on the degree of neural adaptation. Thus, careful considerations of decoder architecture may be required to maximize the effect BMIs have on neurorehabilitation.

Our results investigated the role of indirect neurons within the scope of motor cortex. However, abstract learning relies on a wide array of innate learning mechanisms, such as planning and error-correction, that engage distributed neural networks, but it remains unknown how these networks are engaged. Previous work in rodents has demonstrated that the striatum is necessary for learning to control a BMI [10], [11]. Furthermore, a human ECoG study showed that a distributed network of cortical areas, including posterior parietal cortex and prefrontal cortex, was activated when subjects learned a BMI task [44]. A logical and necessary extension of this work is to investigate the role of indirect neurons in other cortical and subcortical regions, such as those involved in natural motor adaptation and learning, including the cerebellum and association cortices. Further research is necessary to understand how the indirect neurons in these regions behave and to fully characterize the large-scale networks engaged in BMI learning.

4.5 Methods

4.5.1 Animal Subjects

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 at Berkeley Institutional Animal Care and Use Committee.

Two adult male rhesus monkeys (Macaca mulatta) were chronically implanted in the brain with arrays of 64 microelectrodes (Innovative Neurophysiology, Durham NC) [6]. Monkey P was implanted in the left hemisphere in the arm area of both primary motor cortex (M1) and dorsal premotor cortex (PMd), and in the right hemisphere in the arm area of M1, with a total of 192 microwires across three implants. Monkey R was implanted bilaterally in the arm area of M1 and PMd (256 microwires across four implants). Only activity from M1 was included in the direct ensembles (Monkey P: right M1; Monkey R: left M1) and only activity from the same hemisphere was included in the indirect ensemble.

4.5.2 Electrophysiology

Neural activity was recorded using the MAP system (Plexon). Stable units, to be part of the direct ensemble, were selected based on waveform shape, amplitude, relationship to other units on the same channel, interspike interval distribution, and the presence of an absolute refractory period. Only units from primary motor cortex were used which had a clearly identified waveform with signal-to-noise ratio of at least 4:1. Activity was sorted prior to recording sessions using an online spike-sorting application (Sort Client; Plexon).

Stability of waveforms was confirmed by analyzing the stability of PCA projections over days (Wavetracker; Plexon).

Direct units are defined as the units being used to control the BMI. Indirect units consisted of the remaining recorded units. Only stable units from the same hemisphere were used in our analyses (Monkey P: 15 direct, 14 indirect; Monkey R: 10 direct, 14 indirect; Monkey P, 2-decoder experiments: 14 direct, 10 indirect). Stability in the indirect ensemble was assessed using the methods described in Fraser & Schwartz, 2012 [57].

4.5.3 Experimental Setup and Behavioral Training

Before starting the BMI learning experiments, subjects were over-trained on the task performed with arm movements using a KINARM (BKIN Technologies) exoskeleton which restricted shoulder and elbow to move in the horizontal plane.

4.5.4 BMI Tasks

Data from Ganguly & Carmena, 2009, in which subjects performed a self-initiated, eight-target, center-out reaching task, was analyzed. In these experiments, a cursor on a screen was continuously controlled by neural activity. Subjects self-initiated trials by moving the cursor to a center target. One of the eight peripheral targets was randomly selected each trial. Self-initiated trials consisted of those in which the animal moved the cursor to the center target and held for 250-300 ms. Successful trials required the animal to move the cursor to the peripheral target within 15s of initiating the trial and hold the cursor at the target for 250-300 ms. Successful trials resulted in a juice reward; failed trials were repeated. During BMI control, both arms were removed lightly restrained. Neither animal moved their upper limbs during BMI control.

After the initial 19 days of performing the BMI task with a fixed-decoder, Monkey P learned a second decoder over the course of four days. Within each of these four days, Monkey P performed one training block of the new decoder, followed by one training block of the old decoder. Both of these decoders were fixed and used the same direct ensemble as input.

4.5.5 Preprocessing Pipeline

For all analyses, neural data was binned into 100 ms bins to match the decoder timescale. Additionally, learning was analyzed over "training epochs," where each epoch consisted of 150 self-initiated trials. Learning took place over the first 15 training epochs (2,250 self-initiated trials). We chose to analyze the data across training epochs, rather than days,

to eliminate the effect of variable numbers of trials each day. Only the first 15 training epochs (2,250 self-initiated trials) were analyzed; we defined early and late learning as the first and last seven of these training epochs, respectively. Monkey P initiated a total of 3,589 trials. Monkey R initiated a total of 2,357 trials.

4.5.6 Factor Analysis

4.5.6.1 Shared-Over-Total Variance Ratio

Factor analysis (FA) was conducted on the neural population for each epoch to observe underlying correlated neural activity. FA decomposes population signals into correlated and uncorrelated components. For a given neuron i, correlated activity is represented by the shared variance Σ_{ii}^{shared}, and the degree to which the activity was correlated over learning was represented by the ratio of shared-over-total variance of the neural population (SOT). We calculated the SOT ratio according to the methods described in Althaye et al., 2017 [26].

$$SOT = \frac{\mathrm{trace}\left(\Sigma^{shared}\right)}{\mathrm{trace}\left(\Sigma^{total}\right)}$$

4.5.6.2 Total Variance

We also considered how the total variance changed between early and late learning. This was the sum of the private and shared variances ($\Sigma_{ii}^{total} = \Sigma_{ii}^{shared} + \Sigma_{ii}^{private}$).

4.5.6.3 Partial Shared-Over-Total Variance Ratio

We quantified the respective contributions of subpopulations to the SOT ratio using the partial shared-over-total variance (pSOT) ratio. Here, we compared the sum of the shared variance for each subpopulation over the total variance for the entire population. A relative measure was used to account for the fact that the direct and indirect ensembles were different sizes. That is,

$$pSOT_{direct} = \frac{\mathrm{trace}\left(\Sigma_{direct}^{private}\right)}{\mathrm{trace}\left(\Sigma_{all}^{private}\right) + \mathrm{trace}\left(\Sigma_{all}^{shared}\right)}$$

$$pSOT_{indirect} = \frac{\mathrm{trace}\left(\Sigma_{indirect}^{private}\right)}{\mathrm{trace}\left(\Sigma_{all}^{private}\right) + \mathrm{trace}\left(\Sigma_{all}^{shared}\right)},$$

where Σ_i denotes the covariance matrix for the (sub)population i.

4.5.6.4 Shared Space Alignment

We used the ‘‘shared space alignment’’ to measure the similarity between the shared variance of two different training epochs. The shared space alignment is the fraction of shared variance from one epoch captured in the shared space of a second epoch and thus ranges from 0 to 1. We calculated shared alignment according to the methods described in [26]. Given two epochs, A and B, we first compute the projection matrix into Epoch B's shared space, $\mathrm{col}(U^B)$. We then project $\Sigma^{A,\mathrm{shared}}$ onto B's shared space, $P_{U^B}\Sigma^{A,\mathrm{shared}}P_{U^B}^T$. Finally, the alignment is calculated,

$$\text{Shared alignment} = \frac{\mathrm{trace}\left(P_{U^B}\Sigma^{A,\mathrm{shared}}P_{U^B}^T\right)}{\mathrm{trace}(\Sigma^{A,\mathrm{shared}})}$$

4.5.6.5 Offline Decoding of Neural Activity

Linear transforms were used to estimate cursor positions (x and y) based on the binned spike activity for each subpopulation of neurons over all trials and epochs. That is, for a given population of neurons y, we can estimate the cursor positions as follows:

$$y \times \hat{a} = \left[x_{pos}, y_{pos}\right],$$

where a ̂ is the transform matrix.

Residuals were defined as the Euclidean distance between the actual position of the cursor and the estimated position for each time point. To determine the goodness-of-fit, the mean squared error of residuals was calculated. Statistical significance was validated by permutation testing.

4.5.6.6 Quantification & Statistics Analyses

All analyses except for the offline decoding were performed on each training epoch separately. Early and late learning epochs were compared using a two-sample t-test. To determine the significance of decoding of the direct subpopulation, the mean squared error was calculated for random samples of the same size as the number of direct neurons were taken without replacement from the entire recorded population for each animal 100 times.

Chapter 5

Neural Correlates of Control of a Kinematically Redundant BMI

5.1 Abstract

Brain-machine interfaces (BMI) allow users to interact with their external environments by decoding neural signals to control external effectors ranging from computer cursors to robotic arms. Previous work has shown that BMIs can be "two-learner" systems, in which both neurons and decoders co-adapt over the course of learning. However, it is unclear if high levels of decoder adaptation that result in accurate task performance from the onset of training discourage neural adaptation from occurring. In a BMI task where two rhesus macaques were trained to control a kinematically redundant effector, we find that neurons adapt preferred tuning directions over training despite good performance from training onset. Furthermore, population-level firing activity becomes increasingly covaried despite daily re-adaptation of the decoder. Changes in decoder weights did not alter these consolidated patterns but instead increased exploratory neural patterns indicating simultaneous exploration and exploitation of neural activity during BMI learning. Thus, a well-decoded neuroprosthetic device may have potential to be used as both an assistive and rehabilitative device, allowing users to have improved quality of life and facilitating neural adaptation over use.

5.2 Introduction

A major hallmark of natural reaching is the ability to move our arms in complex geometries to achieve relatively simple goals. While our hands move in 3D space, our joints span many more degrees of freedom, thereby making our arms kinematically redundant. This flexibility in joint space allows us to conduct simple tasks with great generalizability and gives us the ability to adapt to situations that may require differing strategies, such as in the face of physical obstacles [58], [59]. Brain-machine interfaces (BMI) [24], [36], [60]–[64] with kinematic redundancy may restore these motor functions, but it is currently unclear how the brain may learn to control such high DOF actuators.

Some groups have used kinematically redundant robotic arms for BMI studies in the past. However, the majority of these studies have either simplified the control of the robot by doing step-wise control, or by forcing subjects to control only the endpoint of the arm, with the inverse kinematics calculated by a computer [8], [64]. Previous work in the lab has asked how kinematic redundancy might be leveraged for BMI control, particularly for subjects who were naïve to the task, but had prior experience with BMI [17]. Specifically, they explored how the extra degrees of freedom (DOF) in a virtual arm were used over learning. While that study gave insight on the behavioral changes that occurred over learning of the BMI, it was left unclear how the brain may have adapted to contribute to these changes. We examined the neural correlates of control in this study. Unraveling the neural bases of this learning would provide insight on how the brain might learn to control high degrees of redundancy.

5.3 Results

5.3.1 Endpoint Trajectories Become More Stereotyped Over Learning

Both monkeys were able to decrease reach times to targets over sessions (Figure 26 B). Learning was tracked by using reach times as a proxy since subjects already achieved a mean accuracy of ~85% across all sessions, including high accuracy in early sessions.

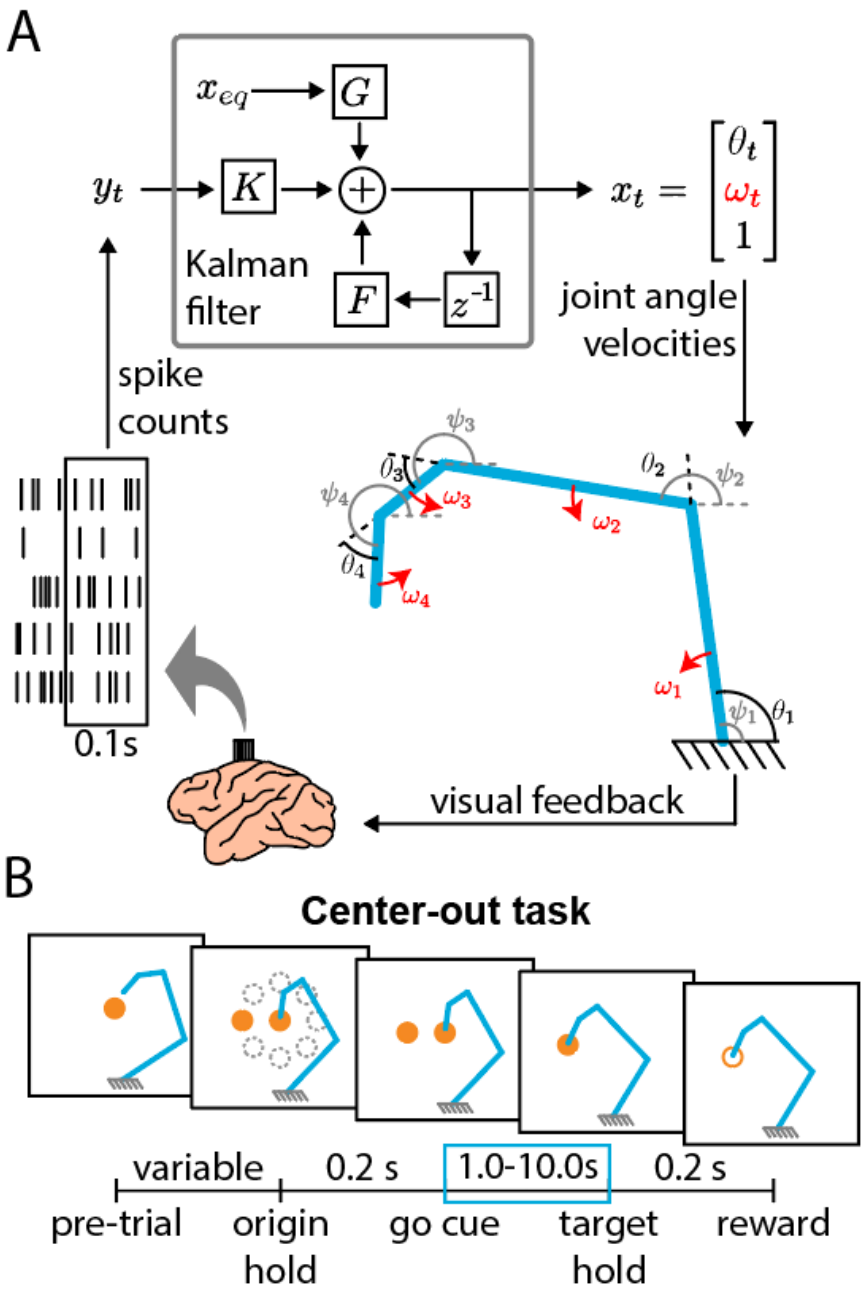


Figure 25 Task design. (A) Closed-loop BMI driven by intracortically-recorded neurons. Neural spikes are decoded with a Kalman filter into a joint angle velocity state-space that controls a virtual 4 DOF arm. (B) Monkeys complete a self-initiated center-out task to eight targets. Only data during movements of correct trials were analyzed (blue rectangle). *Figure adapted from Moorman and Gowda et al., 2017* [17].

We notice that in early learning, while the target acquisition accuracy was high, the movements were less stereotyped than those in late learning, potentially corresponding to the decrease in reach times over learning (Figure 26 A, B). This suggests that arm

configurations may have become more stereotyped over learning as well. To quantify this, we overlay the final arm configurations (i.e. position of the joints at the end of each trial) on top of each other for each session. By collapsing this overlay, we generate a density or heatmap of the final arm configurations (Figure 25C). The final arm configurations in later sessions are much more stereotyped for both animals, corresponding with the more stereotyped traces in Figure 25A.

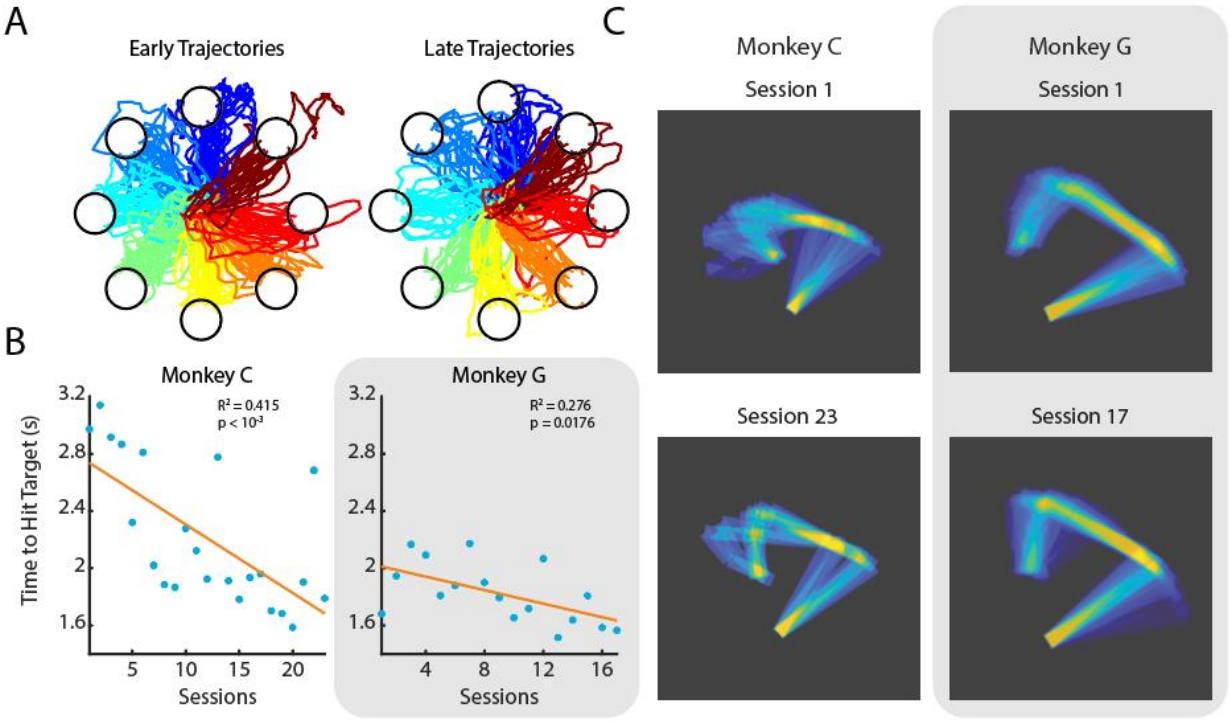


Figure 26 Task performance. Both monkeys exhibit good performance on Session 1. Percent accuracy was ~80% for both animals and remained constant over learning. (a) Example trajectories for Monkey G. (b) Time to Hit (i.e. reach times) decreases for both animals. (c) Heatmaps of final configurations to an example target for the first and last sessions for each animal. Warmer colors indicate more commonly traversed spaces.

To determine how reach times could be decreased over sessions, we then examined if the endpoint kinematics became more stereotyped over learning. To do this, we divided trials into early portions (first 300 ms), corresponding to when the endpoint leaves the center hold target, and late portions (last 300 ms), corresponding to when the endpoint enters the peripheral goal target. The directional variance is calculated, corresponding to the variance in angle the cursor moves in for each trial. We find that the endpoint trajectories become more stereotyped in early portions of the trial for Monkey C (leaving the center hold target), but not significantly different when entering the peripheral target (Figure 27). Conversely, Monkey G improved his stereotypy when entering the peripheral target.

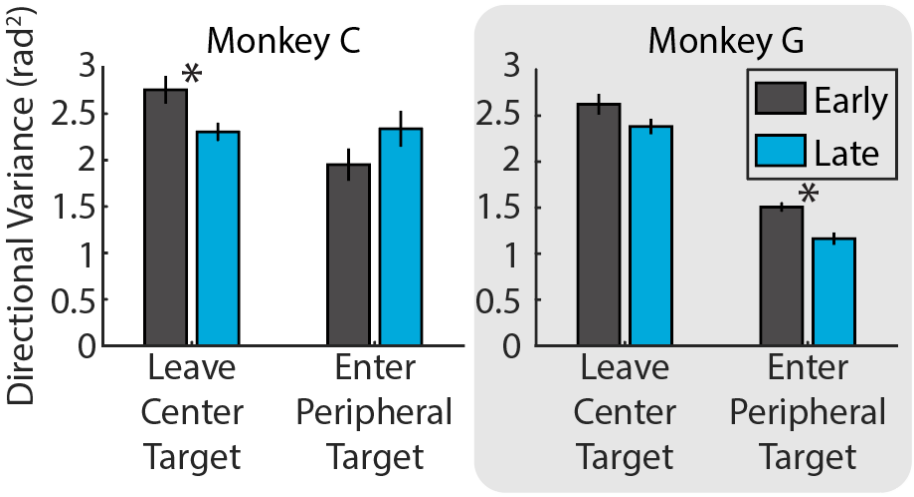


Figure 27 Decrease in directional variance. Endpoint vectors had more stereotyped directions in late learning for both animals, though in different time points of the trials. Monkey C exhibited more stereotyped exits out of the center target while Monkey G exhibited more stereotyped entry into the peripheral targets.

5.3.2 State-Space Linearization

To understand how the animals were able to generate stereotyped movements over learning, we wanted to know in detail how the decoder was transforming the neural activity into state-space actions. As an initial approach, we conducted principal component analysis (PCA) on the observation model of the decoder (i.e. C matrix of the Kalman filter). We found that across all sessions for both animals, the top 2 PCs accounted for roughly 95% of the variance in decoder weights. Since the decoder was trained on visual feedback data, this implied that the animals paid very little attention to the 4 DOFs and potentially cared more about the movements of the endpoint, which only had 2 DOFs.

Consistent with this hypothesis, we found with offline simulations that it was likely animals were attempting to explicitly control the endpoint of the arm rather than the individual joint angles. We trained a decoder offline, relating neural activity with a new state-space model consisting of endpoint velocities in Cartesian space and found that the simulated trajectories were reasonably matched to the targets (Figure 28 A).

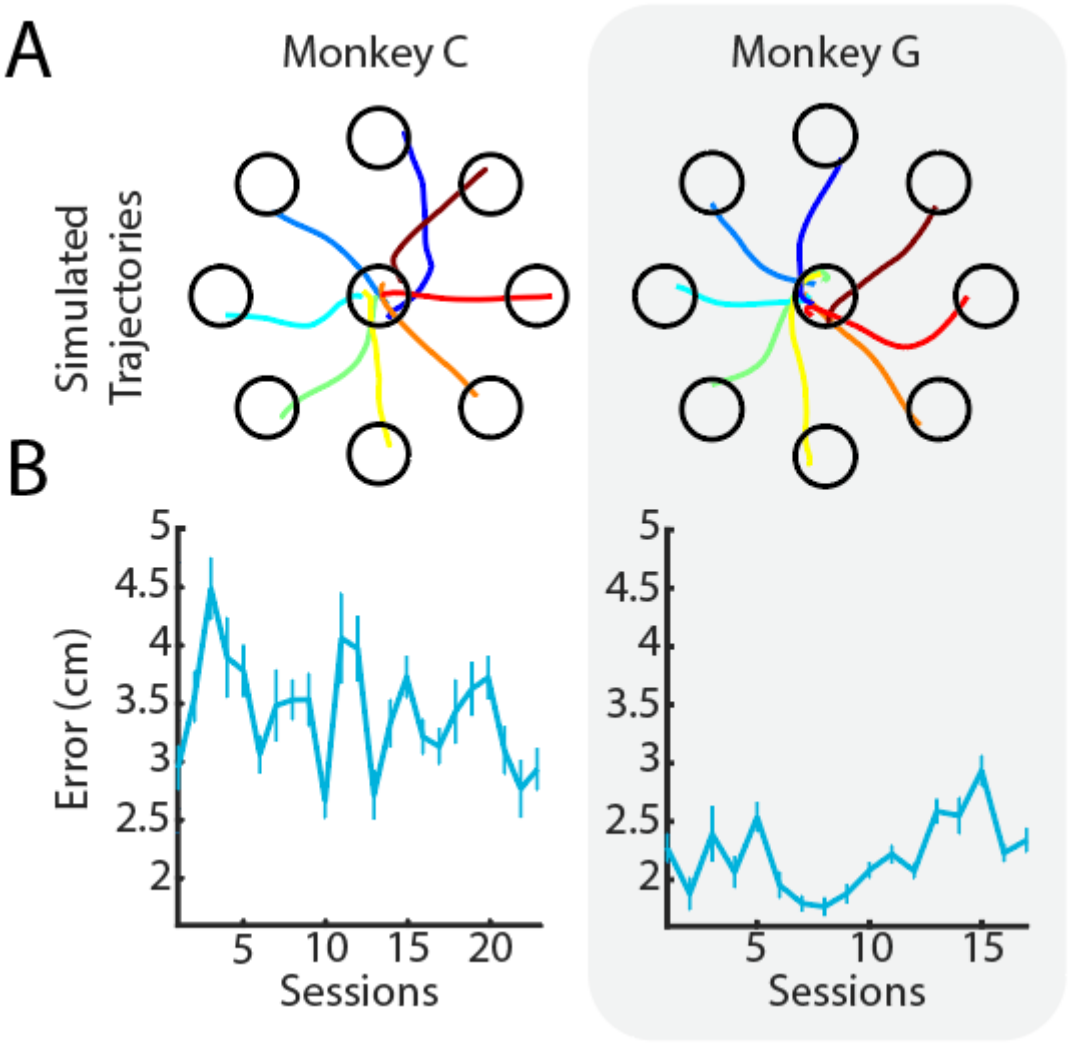


Figure 28 Control is low-dimensional. (a) Neural activity was simulated in an offline decoder to control endpoint velocities. Sample trajectories to each target are shown. (b) Mean distances to targets after each simulated trial. All sessions performed significantly better than chance level ($p < 0.001$, random permutation test).

Next, to more closely examine how the decoder was able to match to the low intended DOFs of the animals, we calculated the tuning directions and modulation depths for both the decoders and neural activity, modifying methods from Orsborn et al., 2014 [16]. In particular, for this particular effector, the state-space is defined as,

$$\omega = [\omega_1, \omega_2, \omega_3, \omega_4]$$

where ω_i corresponds to the i^{th} joint angle velocity ω. Given spiking activity of neuron j, we can then construct an observation model C such that,

$$y_j = C_{j,1}\omega_1 + C_{j,2}\omega_2 + C_{j,3}\omega_3 + C_{j,4}\omega_4.$$

To approximate the observation model in 2 degrees of freedom, we conduct PCA on the C matrix, and can rewrite the output as:

$$y_j \approx C'_{j,1}\omega'_1 + C'_{j,2}\omega'_2$$

where C' and ω' indicate the new values after linear transformations from PCA. Following Orsborn et al, 2014, we can then calculate the decoder tuning directions as,

$$PD_j^d = \arctan\frac{C_{j,2}}{C_{j,1}}.$$

Similarly, the modulation depths of the decoder weights can be approximated as,

$$MD_j^d = \sqrt{C_{j,1}^2 + C_{j,2}^2}.$$

We find that the correlations in decoder preferred tuning directions was incredibly consistent across sessions for both animals, indicating a consistent low-dimensional strategy from both monkeys (Figure 29 A). Furthermore, the correlations between neural MD and decoder MD were consistent across all sessions, with the correlations between neural PD and decoder PD increasing over learning (Figure 29 B, C).

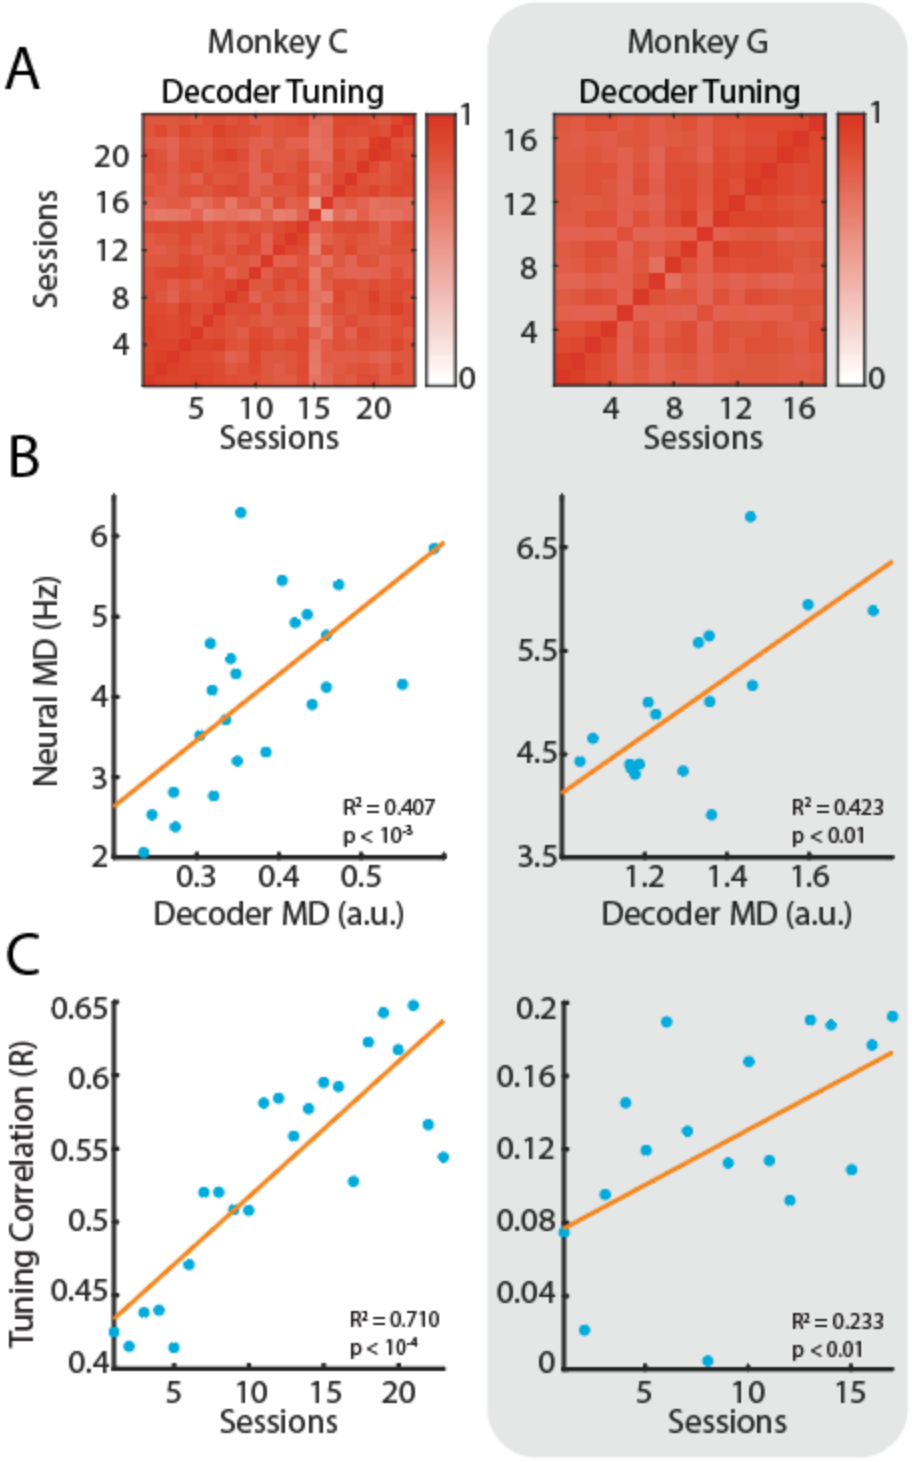


Figure 29 Decoder estimates neural intentions. (A) The observation model is approximated by the top two PCs. Decoder tuning in PC space is stable across all sessions. (B) Decoder modulation depth (MD) is correlated with neural modulation depth. (C) Over learning, neural tuning becomes more correlated with decoder tuning.

These results demonstrate the consistency in which the decoder is decreasing the dimensionality of the task. However, even with these linear transformations, the decoder state-space cannot be converted to a Cartesian coordinate system. To illustrate this point, consider a 1 DOF task where targets are placed around a circle (Figure 30).

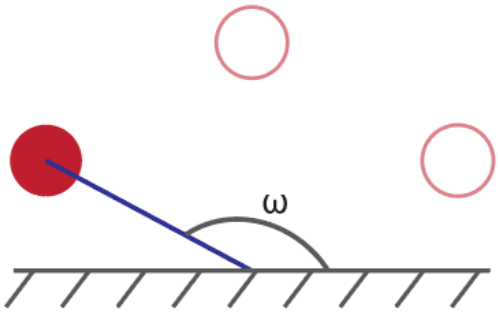


Figure 30 Toy model of linearization. In this example, targets are placed along a circle (red) and the state-space is one-dimensional, [ω]. Converting this state-space to a Cartesian space is possible since each target is linearly separable along a horizontal axis.

A polar coordinate state-space could be easily mapped onto a Cartesian state-space in this situation since the targets are linearly separable in one dimension. That is, the targets are 1-to-1 in Cartesian and polar coordinate spaces. However, if targets were placed around a circle, such as the original experiment, the targets would no longer be separable in one dimension since any given horizontal position could correspond to multiple targets. Thus, while the decoder used in this experiment was able to decrease the dimensionality of the task, the linearization is not mathematically possible for all endpoint locations. Implied in this result is that the brains must have been computing the linear approximations for a given endpoint position (e.g. Jacobian transformation).

5.3.2 Neurons Form a Stable Map to a Converging Decoder

Since the decoder tuning was stable across learning and the correlations between neural tuning and decoder tuning increased over sessions, there must have been some neural consolidation that occurred over learning. Furthermore, while the decoder weights changed each session, the consistency in the PCA models indicates that we could treat the decoder as roughly fixed in 2D space. Given these approximations, we would then expect neural activity to consolidate and form a neuroprosthetic map similar to results shown in previous studies [5], [6], [16], [26], [29].

Indeed, we find that the tuning profiles of the neurons changes in early learning but stabilizes in late learning (Figure 31 A). The correlations between these tuning maps also increases over learning (Figure 31 B).

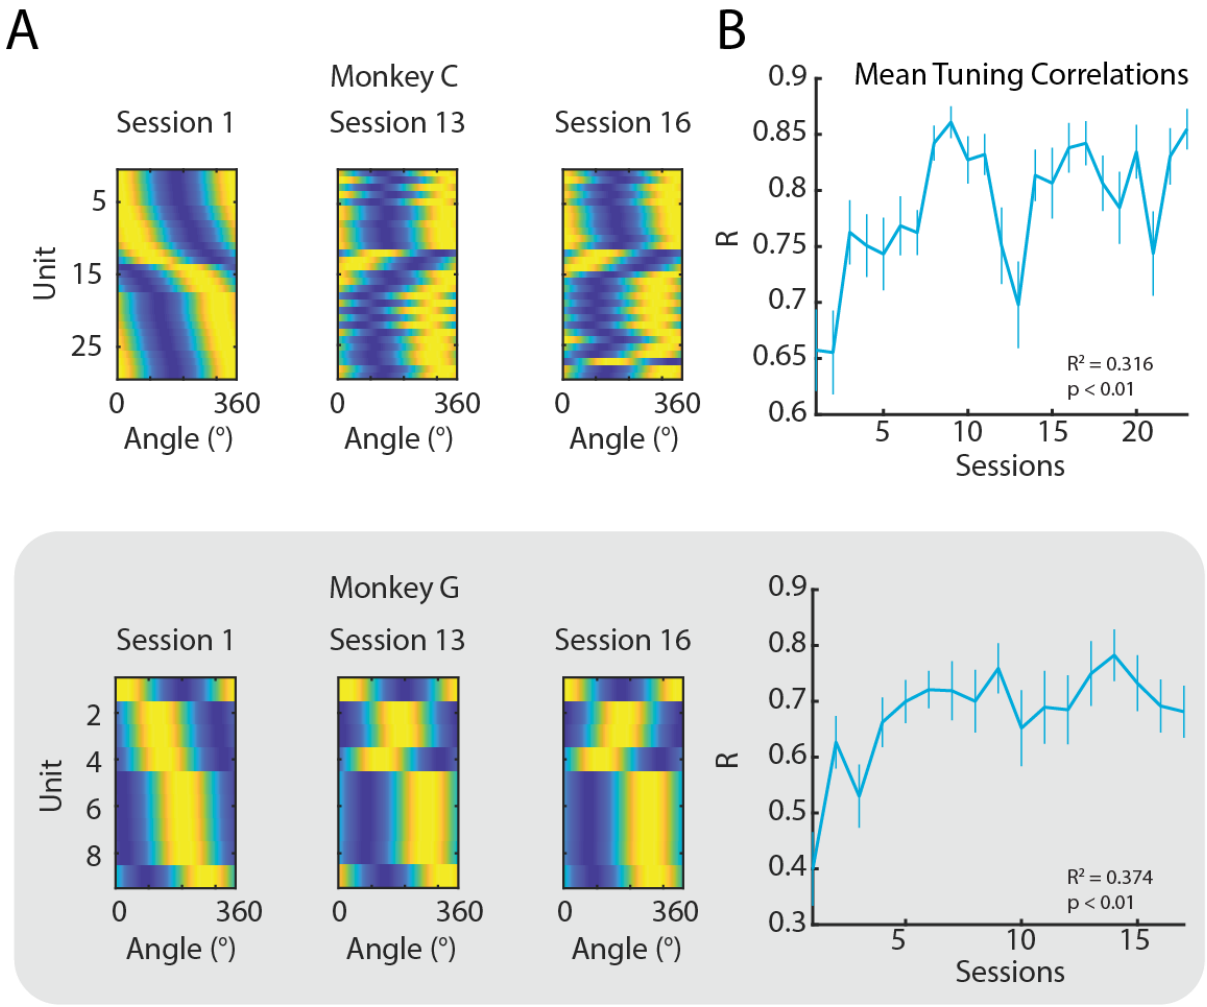


Figure 31 Neuroprosthetic map forms despite daily changes in decoder weights. (A) Tuning curve model comparisons from the first session, a middle session, and a late session for each animal. (B) Stabilization of tuning maps can be characterized by mean correlations of all sessions to each session.

Similarly, to observe how the neural activity was consolidating in late learning, we conducted factor analysis on the neural activity. Factor analysis (FA) separates the variance of the activity into private (non-covarying) and shared (covarying) variance. Intuitively, the ratio of shared to total variance (SOT) represents how correlated the firing activity is [26]. Athalye et al., 2017 showed increased SOT over BMI learning (of a 2D cursor center-out task), linking to faster, more ballistic movements of the BMI and consolidation of a neural manifold [26]. Here, we find the same result despite re-adapting

the decoder each day (Figure 32 A). Thus, despite the BMI decoding neural activity into joint angle velocities, over learning, the high-dimensional neural activity became more correlated, leading to faster endpoint movements. Furthermore, the dimensionality reduced over learning (Figure 32 B). As noted in Chapter 2, the private variance was lower for decoders with weights more similar to that of the first decoder they used. This fluctuation of variance was characterized by flexible changes in modulation depth as well (Figure 32 C, D).

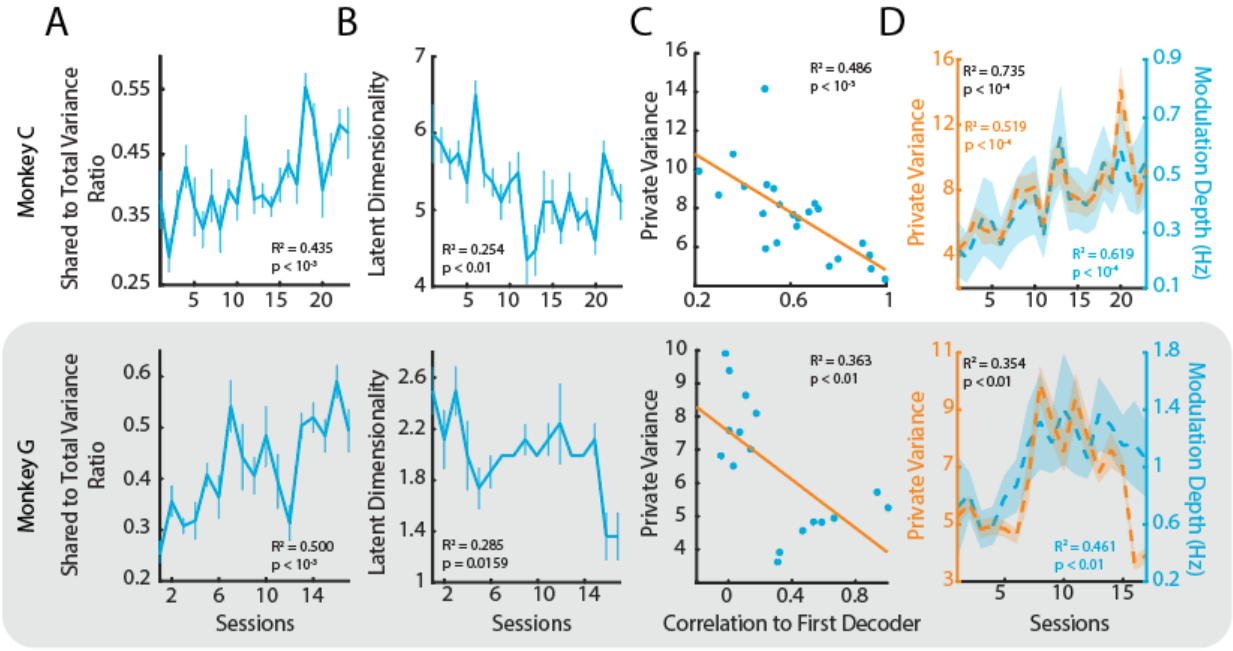


Figure 32 Neural exploration and exploitation occur concomitantly. (A) Ratio of shared over total variance increases over sessions indicating increased correlated activity over skill refinement. (B) Neural dimensionality decreases. (C) Private variance is correlated with changes in the decoder. Larger differences in decoded observation models led to increased neural exploration. (D) The increase in private variance (and total variance) is potentially due to flexible changes in modulation depths of neurons.

5.4 Discussion

Kinematic redundancy in natural motor control is critical for individuals to adapt reaches in different situations, such as in the face of physical obstacles. Thus, if one of the goals of BMIs is to restore upper limb motor function, it is crucial to understand how kinematic redundancy is controlled in the context of BMIs. While many groups have shown high DOF BMI robots being controlled, these groups have largely simplified the tasks. Many only control the endpoint of the effector, leaving the inverse kinematics to be solved by

computers [8], [64]. In this study, individual joints of the arm are controlled by the BMI, allowing us to explore how the neural activity changed through improved control.

From the behavioral results, we found that the endpoint kinematics became increasingly stereotyped over learning with increased endpoint velocities. This result suggests that it was possible the endpoint was an important parameter for the subject to control. However, the movements of this endpoint were the result of individual joint angle velocities being controlled. This suggests that either the monkey only attempted to control the endpoint and the decoder was able to fit to these intentions or the monkey was able to consolidate the control of the joint angles to be represented in smoother endpoint kinematics. Moorman and Gowda et al., 2017 previously showed that all degrees of freedom were utilized for faster reaches, and furthermore, reach times increased when the arm was hidden leaving only the endpoint. Thus, the monkey was able to leverage the redundant DOFs for better task performance [17]. It seems likely then that even if the decoder assisted in simplifying the task for the monkey, the monkey still learned to control the joint angle velocities in such a way to improve endpoint kinematics and likely required linearization of the polar coordinates.

Additionally, we found that the shared variance (over total variance, SOT) increased over learning. In Athalye et al. 2017, it was shown that the increase in SOT for a BMI task related to BMI learning [26]. Specifically, the shared variance contributed to fast, ballistic movements of the BMI. We echoed those results here, with reach times improving as SOT increased. The neural activity, which controlled the joint angle velocities, consolidated over learning, leading to more stereotyped reach trajectories. This indicates that neural activity coordinated over learning to stereotype the behavior of the endpoint, potentially through a high-level representation of the virtual arm, not dissimilar from natural motor reaching.

Embedded in this result is the idea of a 2-learner system where there is a simultaneous convergence of decoder weights and a stabilization of a neuroprosthetic map [6], [16]. In this experiment, a new decoder was constructed each day, allowing the monkey to achieve high accuracy each time. It is plausible that this level of decoding would inhibit the decoded neurons from needing to adapt and conversely, encouraged the decoder to be stable across days. However, we found both to be untrue. While the subject improved his reach times to targets, the Kalman gain also increased in magnitude, implying a more confident fit to the neural firing patterns of the subject. Additionally, the mean firing rate, modulation depth, and SOT all increased, indicating a co-adaptation between decoder and neural activity.

This preliminary work suggests that given a kinematically redundant BMI reaching task, endpoint-control seemed to be important even though the direct control of the limb

stemmed from the joint angles. There was a neural consolidation of the joint angle movements over learning that produced this behavior, though that may not be represented in motor cortex. Furthermore, increase in shared variance despite daily changes to the decoder indicates a convergence between decoder adaptation and neuroprosthetic learning. The ability of the subject to control this virtual arm (which has no proprioceptive feedback) gives promising implications for BMIs to be used as rehabilitative tools in complex reaching tasks. However, further work is needed to determine if control of a kinematically redundant BMI is possible due to good decoding algorithms alone, or if the brain is truly able to consolidate control of the joint angles for endpoint representation. The ability to control the joint angles in such a flexible way would allow high-DOF BMIs to be used for more generalized tasks, which would be required for reaching in the real world where obstacles may be present.

5.5 Methods

5.5.1 Experimental Setup

Two adult male rhesus macaques (Macaca mulatta) was implanted in primary sensory (S1), primary motor (M1), and dorsal premotor (PMd) cortices with chronic microwire electrode arrays (128 channels in each hemisphere, 256 total; Innovative Neurophysiology, Durham, NC). Neural activity was recorded using an OmniPlex system (Plexon Inc., Dallas, TX) and single- and multi-unit activity was sorted online. Both single- and multi-units were used for decoding in the BMI.

The monkey was first instructed to reach to targets on a screen using hand movements for a juice reward. A Kalman filter (KF) with closed-loop decoder adaptation (CLDA) was then used to transform neural activity to movements of a cursor in 2D space. The monkey was instructed to move this cursor to targets in a center-out task. After the subject was able to successfully control this BMI, he was introduced to a 4-link kinematic chain. The KF decoded neural activity into four joint angles' velocities (ω_i in Figure 24 A). Important to note is that the joint angle velocities were relative to the preceding joint. That is, the absolute joint angle $\psi(\theta) = \Sigma_{i=1}^{i}\theta_i$, where $\theta = [\theta_1, \theta_2, \theta_3, \theta_4]^{\mathrm{T}}$ defined the relative joint angles (Figure 24A).

The subject reached to eight targets with radii 1.5 cm, spaced uniformly around a circle of radius 8 cm (Figure 24B). The subject was required to move the kinematic chain such that the endpoint entered the peripheral target. Trials were considered successful if the endpoint was held within the peripheral target for 200 ms. The authors encourage the reader to reference Moorman and Gowda et al., 2017 for more detailed methods [17].

5.5.2 Analysis Pipeline

Only rewarded trials were analyzed for this study. Furthermore, to understand how the BMI was controlled during movement, only segments of trials during which the endpoint was still moving were considered (i.e. from leaving the center hold target to entering the peripheral goal target). Only sessions 1-23 for Monkey C were used for analyses since session 24 contained too few trials for certain targets, a requirement for some of the analyses. All analyses were conducted for each target, then averaged across targets to arrive at a mean value for that session, with the exception of the calculation of the modulation depth. The modulation depth was calculated by finding the amplitude of tuning curves for neurons across pooled targets.

5.5.3 Factor Analysis

Factor analysis (FA) was conducted on the neural population each day to observe underlying correlated neural activity. FA decomposes population signals into correlated and uncorrelated components. This correlated activity is represented by the shared variance, and the degree to which the activity was correlated over learning was represented by the ratio of shared variance to total variance of the neural population (SOT).

For neural activity,

$$x = \mu + x^{private} + x^{shared},$$

where x denotes the neural activity, μ is the average firing rate, and $x^{private}$ and x^{shared} are the private and shared components, respectively.

If we denote the covariance matrices by Σ, then neuron i's shared variance is Σ_{ii}^{shared} and private variance is $\Sigma_{ii}^{private}$. The SOT is then calculated as

$$SOT = \frac{\text{trace}\left(\Sigma_{ii}^{shared}\right)}{\text{trace}\left(\Sigma_{ii}^{private}{}_{\Sigma_{ii}^{shared}}\right)}.$$

More detailed methods of FA can be referenced from Athalye et al. 2017 [26].

Chapter 6

Discussion

In these studies, we have explored how neurons adapt in brain-machine interface task paradigms we refer to as 2-learner systems. While studies have shown neurons to be able to adapt to fixed decoders [5], [6], [26], [32], current decoding modalities allow us to fit decoder weights to the neural activity in BMI tasks with greater accuracy in short amounts of time [14], [35]. Previous work has demonstrated the ability for animals to improve BMI performance alongside changes in decoder readaptation [16]. These improvements were correlated with changes in preferred tuning directions and modulation depth. Additional work examining neural activity in low-dimensional spaces has shown neurons to consolidate onto these subspaces over learning and rotations in the subspaces occur when the control space is poorly aligned with the readout space of the neurons [26], [28], [32]. Here, we show a more complete picture, contextualizing these analyses in 2-learner BMI experiments.

6.1 Neural adaptation in 2-learner systems

We find that neurons consolidate onto low-dimensional spaces even when the decoder weights are readapted over learning. Previous work has shown neurons to be constrained on low-dimensional spaces, suggesting that the control space for a BMI should be closely aligned with that of the intrinsic manifold of neural space [27]. Moreover, decoder changes that are within the manifold (i.e. within plane) are easier for neurons to adapt to compared with decoder changes that are explicitly out of manifold (i.e. out of plane). In our experiments, we are agnostic to these distinctions. However, since SmoothBatch and

CLDA act to fit to the population of neurons via linear models, one can assume that our decoder changes are generally within-plane.

In all of our studies, we find that consolidation of neural activity occurs on multi-day timescales alongside changes in decoder weights. Surprisingly, we find that there are also changes that occur on shorter timescales. In Chapter 2, we discovered fluctuations in private variance that occurred on adjacent days after decoders were refit. Importantly, these fluctuations were not at the cost of decreases in shared over total variance (SOT). Instead, both consolidation and exploration worked together in concert over learning as animals improved performance.

6.2 Mechanisms of neural adaptation

Adaptations on short timescales have been well-explored in the field of psychophysics where subjects are able to adapt to perturbations such as visual-motor rotations within minutes. In these types of perturbations, adaptation involves modifying control patterns of effectors that the subjects are already experts at controlling (e.g. modifying angles of arm movements to throw a ball). Past work has shown that these types of adaptations are driven by the cerebellum [65]. It is possible that what we observed in these studies follows similar adaptation mechanisms. That is, neural adaptation in 2-learner systems may be driven by changes in cerebellum. Subsequently, changes that we see in motor cortex are themselves correlates of changes in more upstream structures. Some studies have suggested that this may be the case [9], [10].

Although upstream structures may be important for neural adaptation, in Chapter 4, we found specificity in which cortical neurons were adapted onto specific subspaces. Neurons that were used as decoder inputs (direct neurons) exhibited greater increases in firing rate variance compared to all other recorded units in motor cortex (indirect neurons). Similarly, in Chapter 2, increases in private variance were linked with exploratory patterns due to decoder changes. These results suggest that fluctuations in firing rate variance are important for learning. Note that these fluctuations did not decrease the amount of variance that was already consolidated onto shared spaces. Thus, we posit that one possible mechanism for this simultaneous exploration and exploitation is that of neural modulation depth.

Intuitively, the modulation depth of a neuron corresponds to the range of the tuning curve. Thus, a higher modulation depth indicates a higher dynamic range in firing rates. In contrast, an increase in mean firing rate does not necessarily increase the neural variance. In general, increasing the modulation depth of the neurons increases the velocities for a cursor in a BMI task since the state-space update is based on the neural

activity multiplied by a constant decoder weight. This is a viable strategy for exploration then, similar to a multi-armed bandit in machine learning and A/B testing.

Additionally, neurons that have no impact to this multi-armed bandit model would quickly drop activity. Indirect neurons can be thought of as direct neurons with decoder weights set to 0. It makes sense that these neurons have suppressed exploration since their activity does not directly impact the decoder output and continuous firing would be metabolically expensive.

6.3 Open questions

BMIs are often cited as having high potential as assistive devices as they bypass the peripheral nervous system and do not require users to have fully functioning endogenous effectors. BMIs also have the potential to act as rehabilitative devices. Studies (including those presented here) have shown changes in neural firing patterns over learning. However, our results suggest that these adaptations may not be changes in neuroplasticity. Rather, neuroprosthetic maps may be representations of specific changes in control by upstream structures akin to how a puppeteer (cerebellum) controls their puppet (motor cortex).

Furthermore, in Chapter 6.1 we posited the multi-armed bandit model where neurons can quickly explore and exploit neural patterns important for control. These models aim to identify the most important features and exploit them for improved control. While this is a potentially efficient solution for the brain, it may hinder any effects of neural adaptation by preferentially adapting neurons with the highest readout "score." Therefore, additional regularization in the decoder weights may facilitate more distributed adaptations among all neurons.

The approach taken by Sadtler et al. is one way to create a perturbation that requires neural activity out of the manifold [27]. However, what if one wanted to induce neural adaptation within manifold (e.g. increasing the modulation depth of neurons)? One possible regularization model would be to maintain the decoder tuning directions for each neuron, but instead change the decoder modulation depths. Such a scenario would keep the control space aligned with the neural space but require different neurons to fire differently to achieve the same output control. These types of decoders with regularizations may help BMIs strike the balance between assistive and rehabilitative devices.

Prior to this work, it was unclear if fitting decoder weights to as closely match the neural population as possible was detrimental for neural adaptation. Our results, highlighted in all the chapters, but particularly Chapter 3, indicate that neurons engage

in similar adaptation mechanisms and rates regardless of the decoder adaptation level. As a corollary then, it would behoove subjects to have saturated decoders to improve initial performance without affecting neural adaptation.

6.4 Dissertation Contributions and Publications

1. **Simultaneous exploration and exploitation of neural patterns underly neuroprosthetic learning.** In Chapter 2, we use factor analysis to separate neural firing variance into coordinated (shared) and independent (private) components as a proxy for consolidated and exploratory patterns, respectively. Our results indicate that when decoder changes are made during the time-course of BMI learning, neurons increase private variance on within-day timescales, while shared variance increases over longer timescales. These changes that vary on different timescales suggest a neural learning mechanism that relies on modifying learned neural patterns rather than generating new ones *de novo*. Furthermore, these results are similar to those of natural motor adaptation and more work is required to determine if neural adaptation in a 2-learner BMI paradigm engages similar network changes as those in natural motor adaptation.

   Publications:

   - **You, Albert K.**, Bing Liu, Abhimanyu Singhal, Suraj Gowda, Helene Moorman, Amy Orsborn, Karunesh Ganguly, Jose M. Carmena. "Flexible modulation of neural variance facilitates abstract skill learning." *under review.*
   - **You, Albert K.**, Bing Liu, Abhimanyu Singhal, Suraj Gowda, Helene Moorman, Amy Orsborn, Karunesh Ganguly, Jose M. Carmena. "Flexible modulation of neural variance facilitates neuroprosthetic skill learning." *bioRxiv* (2019): 817346.

2. **The rate of neural adaptation is invariant to initial levels of decoder adaptation.** In Chapter 3, we find that across multiple time series where the initial performance of a BMI varied, the rates of neural adaptation remained consistent. In addition, subjects continued to improve at the task in all cases, indicating that neural adaptation mechanisms are consistent during skill refinement, and not just initial skill acquisition. These results highlight the ability of CLDA to improve subjects' initial performance with a BMI without hindering the rate of neural adaptation.

   Publication:

- Liu, Bing*, **Albert K. You***, Amy Orsborn, and Jose M. Carmena. "The rate of neuroprosthetic learning is invariant to initial levels of closed-loop decoder adaptation." *in preparation.*

3. **BMI learning is driven by differential adaptation of decoder output neurons.** In Chapter 4, we explore how neurons in a local network adapt differently depending on if they are the output neurons for a decoder (i.e. direct neurons) or neurons in the supporting network (indirect neurons). We find that while both subpopulations of neurons retain a high decoding ability for the BMI task over learning, changes in coordinated firing patterns and rotations in the neural space are observed more strikingly in direct neurons. The combination of these results suggests a local network solution for BMI that involves boosting the signal of direct neurons by changing their firing patterns rather than suppressing noise that may come from indirect neurons.

   Publications:
   - Zippi, Ellen L.*, **Albert K. You**, and Jose M. Carmena. "Local Network Adaptation Underlying Neuroprosthetic Learning." *in preparation.*
   - **You, Albert K.***, Ellen L. Zippi*, and Jose M. Carmena. "Large-Scale Neural Consolidation in BMI Learning." *2019 9th International IEEE/EMBS Conference on Neural Engineering (NER)*. IEEE, 2019.

4. **Neurons perform a state space linearization in a kinematically redundant BMI.** In Chapter 5, where two monkeys are instructed to control a kinematically redundant BMI effector in 2D space, we find that CLDA is able to approximate the control of the effector in 2 dimensions. However, the dimensionality reduction exists in joint space, and thus still requires the monkeys to perform the task with an angular velocity state space. Interestingly, we find that that neural activity approximates a Cartesian state space in offline simulations. This linearization must therefore be driven by neural activity.

   Publication:
   - **You, Albert**, Abhimanyu Singhal, Helene Moorman, Suraj Gowda, Jose M. Carmena. "Neural Correlates of Control of a Kinematically Redundant Brain-Machine Interface." *2019 9th International IEEE/EMBS Conference on Neural Engineering (NER)*. IEEE, 2019.

Bibliography

[1] "WHO | Neurological disorders affect millions globally: WHO report," *WHO*, 2010.

[2] E. E. Fetz, "Operant conditioning of cortical unit activity," *Science (80-.).*, 1969, doi: 10.1126/science.163.3870.955.

[3] K. V. Shenoy and J. M. Carmena, "Combining decoder design and neural adaptation in brain-machine interfaces," *Neuron*, vol. 84, no. 4. Cell Press, pp. 665–680, 19-Nov-2014, doi: 10.1016/j.neuron.2014.08.038.

[4] V. R. Athalye, J. M. Carmena, and R. M. Costa, "Neural reinforcement: re-entering and refining neural dynamics leading to desirable outcomes," *Current Opinion in Neurobiology*, vol. 60. pp. 145–154, 2020, doi: 10.1016/j.conb.2019.11.023.

[5] J. M. Carmena *et al.*, "Learning to control a brain-machine interface for reaching and grasping by primates," *PLoS Biol.*, vol. 1, no. 2, pp. 193–208, 2003, doi: 10.1371/journal.pbio.0000042.

[6] K. Ganguly and J. M. Carmena, "Emergence of a stable cortical map for neuroprosthetic control," *PLoS Biol.*, vol. 7, no. 7, 2009, doi: 10.1371/journal.pbio.1000153.

[7] D. M. Taylor, S. I. H. Tillery, and A. B. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science (80-.).*, 2002, doi: 10.1126/science.1070291.

[8] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz, "Cortical control of a prosthetic arm for self-feeding," *Nature*, vol. 453, no. 7198, pp. 1098–1101, 2008, doi: 10.1038/nature06996.

[9] K. B. Clancy, A. C. Koralek, R. M. Costa, D. E. Feldman, and J. M. Carmena, "Volitional modulation of optically recorded calcium signals during neuroprosthetic learning," *Nat. Neurosci.*, vol. 17, no. 6, pp. 807–809, 2014, doi: 10.1038/nn.3712.

[10] A. C. Koralek, X. Jin, J. D. Long, R. M. Costa, and J. M. Carmena, "Corticostriatal plasticity is necessary for learning intentional neuroprosthetic

skills," *Nature*, vol. 483, no. 7389, pp. 331–335, 2012, doi: 10.1038/nature10845.

[11] R. M. Neely, A. C. Koralek, V. R. Athalye, R. M. Costa, and J. M. Carmena, "Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia," *Neuron*, vol. 97, no. 6, pp. 1356-1368.e4, Mar. 2018, doi: 10.1016/j.neuron.2018.01.051.

[12] G. Kumar *et al.*, "Designing Closed-Loop Brain-Machine Interfaces Using Model Predictive Control," *Technologies*, vol. 4, no. 2, p. 18, Jun. 2016, doi: 10.3390/technologies4020018.

[13] B. P. Olson, J. Si, J. Hu, and J. He, "Closed-loop cortical control of direction using support vector machines," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 13, no. 1, pp. 72–80, Mar. 2005, doi: 10.1109/TNSRE.2004.843174.

[14] S. Dangi, A. L. Orsborn, H. G. Moorman, and J. M. Carmena, "Design and analysis of closed-loop decoder adaptation algorithms for brain-machine interfaces," *Neural Comput.*, vol. 25, no. 7, pp. 1693–1731, Jul. 2013.

[15] A. L. Orsborn, S. Dangi, H. G. Moorman, and J. M. Carmena, "Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 20, no. 4, pp. 468–477, 2012, doi: 10.1109/TNSRE.2012.2185066.

[16] A. L. Orsborn, H. G. Moorman, S. A. Overduin, M. M. Shanechi, D. F. Dimitrov, and J. M. Carmena, "Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control," *Neuron*, vol. 82, no. 6, pp. 1380–1393, 2014, doi: 10.1016/j.neuron.2014.04.048.

[17] H. G. Moorman, S. Gowda, and J. M. Carmena, "Control of Redundant Kinematic Degrees of Freedom in a Closed-Loop Brain-Machine Interface," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 25, no. 6, pp. 750–760, 2017, doi: 10.1109/TNSRE.2016.2593696.

[18] G. Santhanam, S. I. Ryu, B. M. Yu, A. Afshar, and K. V. Shenoy, "A high-performance brain-computer interface," *Nature*, vol. 442, no. 7099, pp. 195–198, 2006, doi: 10.1038/nature04968.

[19] W. Wu, Y. Gao, E. Bienenstock, J. P. Donoghue, and M. J. Black, "Bayesian population decoding of motor cortical activity using a Kalman filter," *Neural Comput.*, vol. 18, no. 1, pp. 80–118, 2006, doi: 10.1162/089976606774841585.

[20] D. Sussillo *et al.*, "A recurrent neural network for closed-loop intracortical brain-machine interface decoders," 2012, doi: 10.1088/1741-2560/9/2/026027.

[21] E. C. Leuthardt, G. Schalk, J. R. Wolpaw, J. G. Ojemann, and D. W. Moran, "A brain-computer interface using electrocorticographic signals in humans," *J. Neural Eng.*, vol. 1, no. 2, pp. 63–71, 2004, doi: 10.1088/1741-2560/1/2/001.

[22] G. Townsend, B. Graimann, and G. Pfurtscheller, "Continuous EEG classification during motor imagery - Simulation of an asynchronous BCI," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 12, no. 2, pp. 258–265, Jun. 2004, doi:

10.1109/TNSRE.2004.827220.

[23] K. Cassady, A. You, A. Doud, and B. He, "The impact of mind-body awareness training on the early learning of a brain-computer interface," *TECHNOLOGY*, vol. 02, no. 03, pp. 254–260, Sep. 2014, doi: 10.1142/s233954781450023x.

[24] K. Ganguly and J. M. Carmena, "Neural correlates of skill acquisition with a cortical brain-machine interface," *J. Mot. Behav.*, vol. 42, no. 6, pp. 355–360, 2010, doi: 10.1080/00222895.2010.526457.

[25] T. Gulati *et al.*, "Robust Neuroprosthetic Control from the Stroke Perilesional Cortex," *J. Neurosci.*, vol. 35, no. 22, pp. 8653–8661, 2015, doi: 10.1523/JNEUROSCI.5007-14.2015.

[26] V. R. Athalye, K. Ganguly, R. M. Costa, and J. M. Carmena, "Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control," *Neuron*, vol. 93, no. 4, pp. 955-970.e5, 2017, doi: 10.1016/j.neuron.2017.01.016.

[27] P. T. Sadtler *et al.*, "Neural constraints on learning," *Nature*, vol. 512, no. 7515, pp. 423–426, 2014, doi: 10.1038/nature13665.

[28] M. D. Golub *et al.*, "Learning by neural reassociation/631/378/1595/631/378/2629 article," *Nat. Neurosci.*, vol. 21, no. 4, pp. 607–616, 2018, doi: 10.1038/s41593-018-0095-3.

[29] K. Ganguly, D. F. Dimitrov, J. D. Wallis, and J. M. Carmena, "Reversible large-scale modification of cortical networks during neuroprosthetic control," *Nat. Neurosci.*, vol. 14, no. 5, pp. 662–669, 2011, doi: 10.1038/nn.2797.

[30] V. R. Athalye, F. J. Santos, J. M. Carmena, and R. M. Costa, "Evidence for a neural law of effect," *Science (80-.).*, vol. 359, no. 6379, pp. 1024–1029, 2018, doi: 10.1126/science.aao6058.

[31] M. M. Churchland *et al.*, "Neural population dynamics during reaching," *Nature*, vol. 487, no. 7405, pp. 51–56, 2012, doi: 10.1038/nature11129.

[32] E. R. Oby *et al.*, "New neural activity patterns emerge with long-term learning," *Proc. Natl. Acad. Sci.*, p. 201820296, Jun. 2019, doi: 10.1073/pnas.1820296116.

[33] K. C. Ames and M. M. Churchland, "Motor cortex signals for each arm are mixed across hemispheres and neurons yet partitioned within the population response," *Elife*, vol. 8, Oct. 2019, doi: 10.7554/eLife.46159.

[34] S. Dangi, A. L. Orsborn, H. G. Moorman, and J. M. Carmena, "Design and analysis of closed-loop decoder adaptation algorithms for brain-machine interfaces," *Neural Comput.*, vol. 25, no. 7, pp. 1693–1731, 2013.

[35] V. Gilja *et al.*, "A high-performance neural prosthesis enabled by control algorithm design," *Nat. Neurosci.*, 2012, doi: 10.1038/nn.3265.

[36] S. Gowda, A. L. Orsborn, S. A. Overduin, H. G. Moorman, J. M. Carmena, and S. Member, "Designing Dynamical Properties of Brain – Machine Interfaces to Optimize Task-Speci fi c Performance," *Ieee Trans. Neural Syst. Rehabil. Eng.*, vol. 22, no. 5, pp. 911–920, 2014.

[37] J. Izawa, T. Rane, O. Donchin, and R. Shadmehr, "Motor adaptation as a process of reoptimization," *J. Neurosci.*, vol. 28, no. 11, pp. 2883–2891, Mar. 2008, doi: 10.1523/JNEUROSCI.5359-07.2008.

[38] K. P. Kording, J. B. Tenenbaum, and R. Shadmehr, "The dynamics of memory as a consequence of optimal adaptation to a changing body," *Nat. Neurosci.*, vol. 10, no. 6, pp. 779–786, Jun. 2007, doi: 10.1038/nn1901.

[39] R. Shadmehr, M. A. Smith, and J. W. Krakauer, "Error correction, sensory prediction, and adaptation in motor control.," *Annu. Rev. Neurosci.*, vol. 33, pp. 89–108, 2010, doi: 10.1146/annurev-neuro-060909-153135.

[40] K. Wei and K. Körding, "Relevance of error: What drives motor adaptation?," *J. Neurophysiol.*, vol. 101, no. 2, pp. 655–664, Feb. 2009, doi: 10.1152/jn.90545.2008.

[41] M. D. Golub, S. M. Chase, A. P. Batista, and B. M. Yu, "Brain-computer interfaces for dissecting cognitive processes underlying sensorimotor control," *Current Opinion in Neurobiology*, vol. 37. Elsevier Ltd, pp. 53–58, 01-Apr-2016, doi: 10.1016/j.conb.2015.12.005.

[42] S. M. Chase, R. E. Kass, and A. B. Schwartz, "Behavioral and neural correlates of visuomotor adaptation observed through a brain-computer interface in primary motor cortex," *J. Neurophysiol.*, vol. 108, no. 2, pp. 624–644, Jul. 2012, doi: 10.1152/jn.00371.2011.

[43] B. Jarosiewicz, S. M. Chase, G. W. Fraser, M. Velliste, R. E. Kass, and A. B. Schwartz, "Functional network reorganization during learning in a brain-computer interface paradigm," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 105, no. 49, pp. 19486–19491, Dec. 2008, doi: 10.1073/pnas.0808113105.

[44] J. D. Wander *et al.*, "Distributed cortical adaptation during learning of a brain-computer interface task," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 110, no. 26, pp. 10818–10823, Jun. 2013, doi: 10.1073/pnas.1221127110.

[45] M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, and J. P. Donoghue, "Instant neural control of a movement signal," *Nature*, vol. 416, no. 6877, pp. 141–142, 2002, doi: 10.1038/416141a.

[46] A. P. Georgopoulos, R. E. Kettner, and A. B. Schwartz, "Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population," *J. Neurosci.*, vol. 8, no. 8, pp. 2928–2937, 1988.

[47] D. A. Nowak, D. Timmann, and J. Hermsdörfer, "Dexterity in cerebellar agenesis," *Neuropsychologia*, vol. 45, no. 4, pp. 696–703, Jan. 2007, doi: 10.1016/j.neuropsychologia.2006.08.011.

[48] R. M. Costa, D. Cohen, and M. A. L. Nicolelis, "Differential corticostriatal plasticity during fast and slow motor skill learning in mice," *Curr. Biol.*, vol. 14, no. 13, pp. 1124–1134, Jul. 2004, doi: 10.1016/j.cub.2004.06.053.

[49] B. P. Ölveczky, T. M. Otchy, J. H. Goldberg, D. Aronov, and M. S. Fee, "Changes in the neural control of a complex motor sequence during learning," *J.*

Neurophysiol., vol. 106, no. 1, pp. 386–397, Jul. 2011, doi: 10.1152/jn.00018.2011.

[50] M. S. Fee and J. H. Goldberg, "A hypothesis for basal ganglia-dependent reinforcement learning in the songbird," *Neuroscience*, vol. 198. pp. 152–170, 2011, doi: 10.1016/j.neuroscience.2011.09.069.

[51] A. Neuringer, C. Deiss, and G. Olson, "Reinforced variability and operant learning," *J. Exp. Psychol. Anim. Behav. Process.*, vol. 26, no. 1, pp. 98–111, 2000, doi: 10.1037/0097-7403.26.1.98.

[52] J. Wessberg and M. A. L. Nicolelis, "Optimizing a linear algorithm for real-time robotic control using chronic cortical ensemble recordings in monkeys," *J. Cogn. Neurosci.*, vol. 16, no. 6, pp. 1022–1035, Jul. 2004, doi: 10.1162/0898929041502652.

[53] C. A. Chestek *et al.*, "Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex.," *J. Neural Eng.*, vol. 8, no. 4, p. 045005, Aug. 2011, doi: 10.1088/1741-2560/8/4/045005.

[54] L. Shpigelman, H. Lalazar, and E. Vaadia, "Kernel-ARMA for Hand Tracking and Brain-Machine Interfacing During 3D Motor Control."

[55] E. E. Fetz, "Volitional control of neural activity: implications for brain-computer interfaces," *J. Physiol.*, vol. 579, no. 3, pp. 571–579, Mar. 2007, doi: 10.1113/jphysiol.2006.127142.

[56] D. M. Taylor, "Direct Cortical Control of 3D Neuroprosthetic Devices," *Science (80-.).*, vol. 296, no. 5574, pp. 1829–1832, 2002, doi: 10.1126/science.1070291.

[57] G. W. Fraser and A. B. Schwartz, "Recording from the same neurons chronically in motor cortex," *J. Neurophysiol.*, vol. 107, no. 7, pp. 1970–1978, 2012, doi: 10.1152/jn.01012.2010.

[58] J. F. Yang, J. P. Scholz, and M. L. Latash, "The role of kinematic redundancy in adaptation of reaching," *Exp. Brain Res.*, vol. 176, no. 1, pp. 54–69, 2007, doi: 10.1007/s00221-006-0602-8.

[59] J. Baillieul, "Avoiding obstacles and resolving kinematic redundancy," *IEEE Int. Conf. Robot. Autom.*, pp. 1698–1704, 1986, doi: 10.1109/ROBOT.1986.1087464.

[60] J. M. Carmena, "Advances in Neuroprosthetic Learning and Control," *PLoS Biol.*, vol. 11, no. 5, pp. 1–4, 2013, doi: 10.1371/journal.pbio.1001561.

[61] S. J. Bensmaia and L. E. Miller, "Restoring sensorimotor function through intracortical interfaces: Progress and looming challenges," *Nat. Rev. Neurosci.*, vol. 15, no. 5, pp. 313–325, 2014, doi: 10.1038/nrn3724.

[62] R. E. Kass, G. W. Fraser, M. Velliste, B. Jarosiewicz, A. B. Schwartz, and S. M. Chase, "Functional network reorganization during learning in a brain-computer interface paradigm," *Proc. Natl. Acad. Sci.*, 2008, doi: 10.1073/pnas.0808113105.

[63] S. Kim, J. D. Simeral, L. R. Hochberg, J. P. Donoghue, and M. J. Black, "NIH Public Access," vol. 5, no. 4, pp. 455–476, 2010, doi: 10.1088/1741-2560/5/4/010.Neural.

[64] J. L. Collinger *et al.*, "High-performance neuroprosthetic control by an individual

with tetraplegia," *Lancet*, vol. 381, no. 9866, pp. 557–564, 2013, doi: 10.1016/S0140-6736(12)61816-9.

[65] O. Donchin *et al.*, "Cerebellar regions involved in adaptation to force field and visuomotor perturbation," *J. Neurophysiol.*, vol. 107, no. 1, pp. 134–147, Jan. 2012, doi: 10.1152/jn.00007.2011.