Effect of Contextual Task Changes on Motor Cortical Activity and Brain-Computer Interface Performance

by

Angelica Janina Herrera

B.S., Biomedical Engineering, Johns Hopkins University, 2016

Submitted to the Graduate Faculty of the

Swanson School of Engineering in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2022

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Angelica Janina Herrera

It was defended on

June 29, 2022

and approved by

Steven M. Chase, PhD, Professor, Department of Biomedical Engineering, Carnegie Mellon University

Aaron P. Batista, PhD, Professor, Department of Bioengineering

Michael L. Boninger, MD, Professor, Department of Physical Medicine and Rehabilitation

Dissertation Director: Jennifer L. Collinger, PhD, Associate Professor, Department of Physical Medicine and Rehabilitation

Effect of Contextual Task Changes on Motor Cortical Activity and Brain-Computer Interface Performance

Angelica Janina Herrera, PhD

University of Pittsburgh, 2022

The restoration of arm and hand function is a top rehabilitation priority for people with tetraplegia. To address this need, intracortical brain-computer interfaces (BCIs), which decode neural activity to predict desired movement commands, enable people with severe motor impairments to achieve high performance control of assistive devices in laboratory environments. A key challenge limiting the clinical translation of BCIs is their ability to provide reliable control in more naturalistic settings.

Here, we evaluate the impact of contextual changes on neural activity in human motor cortex to provide a framework for the development of more robust decoding schemes. We began this work by first determining if task-related neural activity is stable over time. We found that there exists a stable representation of latent neural activity patterns that can be aligned over time to provide a stable input to BCI decoders for computer access. Our findings demonstrate that users can obtain reliable, accurate BCI control of their assistive devices for months at a time without the need for technician intervention.

Focusing on the restoration of arm and hand function, we next studied how contextual changes during object interaction affect high degree of freedom control of a robotic arm. When participants were instructed to attempt identical reach and grasp movements with and without an object in the workspace, we found that object presence significantly modulates population neural activity from planning to movement completion. We also found that grasp intention coincided with neural activity modulation in response to object presence. Small modulation of neural activity in

primary somatosensory cortex was identified in response to both the presence of an object and grasp intention. Here, we found evidence for movement-related and broader task-context variables in both the motor and somatosensory cortex during attempted movements.

Our results aim to provide researchers a better understanding of the impact of contextual changes on motor cortical activity and provide a framework for developing more contextually invariant BCIs. We hope to advance the clinical translation of BCIs for day-to-day use by providing stable and reliable BCI performance in the face of contextual changes.

Table of Contents

List of Figures

1.0 Introduction

The loss of arm and hand function that occurs with tetraplegia severely limits independence and is the top rehabilitation priority for people with cervical spinal cord injury (SCI) (Anderson, 2004; Collinger, Boninger, et al., 2013; Huggins, Moinuddin, Chiodo, & Wren, 2015). Intracortical brain-computer interfaces (BCIs), which record simultaneously from ~200 neurons in the motor cortex, have enabled people with tetraplegia to control reaching and grasping movements of a robotic arm with the speed and skill approaching that of an able-bodied person (Collinger, Wodlinger, et al., 2013; Wodlinger et al., 2015). Portable systems have also begun to be used in BCI research, providing study participants with independent computer access (Dekleva, Weiss, Boninger, & Collinger, 2021; John D Simeral et al., 2021; Weiss, Gaunt, Franklin, Boninger, & Collinger, 2019). Despite these advancements, recording instabilities and variations in contextual factors, such as changes in behavioral state or task-related parameters, lead to a degradation in BCI control and negatively affect performance (Chestek et al., 2007b; Degenhart et al., 2020; Downey, Schwed, Chase, Schwartz, & Collinger, 2018; Jarosiewicz et al., 2015; Kim, Wood, Fellows, Donoghue, & Black, 2006; Perge et al., 2013; Sussillo, Stavisky, Kao, Ryu, & Shenoy, 2016).

For a BCI system to be reliable and meet users' needs, it must provide stable, accurate control over a desired period (Blabe et al., 2015; Collinger, Boninger, et al., 2013; Huggins et al., 2015). However, we know that neural activity fluctuates over time, causing disruptions in BCI control (Colachis IV et al., 2021; Downey et al., 2018; Jarosiewicz et al., 2015; Perge et al., 2013; Simeral, Kim, Black, Donoghue, & Hochberg, 2011). Various recalibration algorithms have been developed, but many either require knowledge of the user's intention, which limits their use to simple, low degree of freedom control, or require a lot of training data and are computationally

expensive, making them difficult to apply in clinically viable systems (Bishop et al., 2014; Homer et al., 2014; Jarosiewicz et al., 2015; Karpowicz et al., 2022; Li, O'doherty, Lebedev, & Nicolelis, 2011; Orsborn et al., 2014; Sussillo et al., 2016; Wu & Hatsopoulos, 2008; Zhang & Chase, 2013). If we want to provide users with robust BCI control, we need to develop a system that is easy to use, does not require a long setup time, and minimizes interruptions in BCI control for recalibration.

Since our lab focuses on restoring upper limb function through high degree of freedom control of a robotic arm and hand, we also consider the impact of contextual changes on motor cortical activity for reaching and grasping movements. We have previously found that when interacting with an object, there is an increase in global firing rates in motor cortex (M1) compared to when no object is present (Downey et al., 2017). This context-dependent modulation, which does not arise during decoder training due to the controlled nature of the training task, hinders our ability to predict the user's intended movement. One way of overcoming this unexpected neural instability is to scale the neural activity to remove the correlated increase in firing rates of the recorded units (Downey et al., 2017). However, in this process, we gain no further insight into why this contextual modulation occurs, and additionally, we may be removing potentially relevant information related to object contact, grasp, and manipulation as well as other object-related factors. To provide reliable BCI control for a variety of reach and grasp-related movements performed in day-to-day situations, we must investigate the neural mechanisms behind this context-dependent modulation.

The goal of this dissertation is to understand how environment- or task-related contextual changes impact motor cortical activity and characterize the neural mechanisms behind these observed changes. We investigated the cortical representation of movement during kinematically

similar, but contextually different tasks as part of an ongoing intracortical BCI clinical trial with a human subject. These experiments allowed us to isolate and study the patterns of neural activity related to various contextual factors. This body of work will help advance the design of improved real-time tools and decoding algorithms for stable, accurate BCI control across a variety of daily tasks in more natural, as opposed to laboratory, environments.

1.1 Structure and Contents of Dissertation

Chapter 2 provides an overview of the current standard in BCI research for neural recording and decoding for the restoration of movement. Neural signal processing techniques for population level analysis are also presented, with a brief explanation of the methods as well as motivations for their application over the study of individual neurons. Additionally, current barriers to clinical adoption, including fluctuations in neural recordings over time as well as due to changes in context, are reviewed.

Chapter 3 describes the general methods relevant to all experiments presented in this work, including electrophysiology details, the general decoder calibration protocol, and dimensionality reduction techniques for population-level analysis of motor cortical activity.

Chapter 4 focuses on overcoming neural recording instabilities to provide reliable BCI control. We find that task-related neural activity is remarkably stable despite a relatively rapid decrease of stable units. We leverage this stable representation of neural activity to overcome neural instabilities by aligning our estimates of the stable underlying neural activity patterns. Using this neural signal stabilization method, we obtained robust BCI control of a computer cursor over a period of 104 days in a person with tetraplegia. Our method not only outperformed a fixed BCI

mapping, which was not updated over the course of the experiments, but also performed comparably to daily recalibration of the BCI, the current standard in BCI research. Thus, we demonstrate that neural signal stabilization can be used to provide independent BCI operation to potential users.

In chapter 5, we characterize the impact of object presence as well as grasp intention on population neural activity. When characterizing the effect of object presence on M1 activity, we found that neural population activity was significantly modulated throughout planning, reaching, and grasping. Motor cortical modulation in response to the presence or absence of an object also tended to vary concurrently with grasp intention (e.g., keeping the hand open or closing the hand). However, we determined that object presence and grasp intention impact distinct patterns of neural activity. We suggest that this context-dependent modulation corresponds to inputs from the cortical grasping circuit, consisting of the anterior intraparietal area, ventral premotor cortex, and M1, encoding the object properties and the potential grips afforded to this object.

Chapter 6 investigates the role of context-dependent signals on M1 and S1 activity during object interaction. Here, we aim to determine the neural mechanisms underlying the effect of object presence. We first characterized the response of somatosensory cortex (S1) to changes in object presence. We found small, yet significant changes in S1 activity in response to object presence from planning to movement execution. These findings point to communication between M1 and S1 for movement generation, even in the absence of sensory feedback. We next repeated the object interaction task but altered the precision requirements placed on the target location. We found small, yet consistent, modulation in neural activity throughout planning and movement, suggesting that M1 activity not only encodes for kinematic parameters but also for cognitive task-related parameters such as perceived task difficulty.

Together, the results from this work suggest that the continued advancement of BCI systems for the restoration of arm and hand function requires a more holistic perspective of motor cortical activity and its relationship to motor output. By considering the role of neural population activity rather than just single unit recordings as well as both control and context-dependent signals, we can work towards developing the next generation of clinical BCI systems.

2.0 Background

Disorders of the brain, spine, or limbs can disrupt the normal flow of commands from the brain to the body, impairing the neural pathways that provide individuals with control of their muscles. Our study focuses on restoring upper limb function in people with spinal cord injuries (SCI) or other neuromuscular disorders such as spinocerebellar degeneration. In the United States alone, more than 200,000 individuals live with a chronic spinal cord injury (SCI) (National Spinal Cord Injury Statistical Center, 2020). SCI can result in paralysis below the site of injury depending on the severity of the injury. In the case of a severe cervical spinal cord injury, individuals completely lose the ability to voluntarily move their limbs. Since the pathways between the brain and the rest of the body are damaged, this impacts the person's ability to perform critical activities of daily living (e.g. feeding oneself, getting ready, communicating, etc.). To restore function and improve these individuals' quality of life, brain computer interfaces (BCIs) have been developed to bypass these impaired neural pathways and instead relay the brain's commands to assistive devices, such as computer cursors or robotic arms, that directly interact with the environment. To provide control of these assistive devices, neural activity must be mapped to the user's intentions. Since arm and hand kinematics are well represented in motor cortical activity, intracortical BCIs have typically targeted the hand and arm region of M1 and translated this neural activity to control signals (Collinger, Wodlinger, et al., 2013; A. P. Georgopoulos, Kalaska, Caminiti, & Massey, 1982; Hochberg et al., 2006; Taylor, Tillery, & Schwartz, 2002; Velliste, Perel, Spalding, Whitford, & Schwartz, 2008). BCI control has enabled people with tetraplegia to achieve up to simultaneous ten-dimensional control of a robotic arm (Wodlinger et al., 2015). Here, we will

provide a general overview of BCI systems, their role in providing control for assistive devices, and new perspectives on neural computations that generate movement.

2.1 Neural Interfaces and Signal Acquisition

Distinct types of neural interfaces have been developed to record neural activity and thus predict motor intention. This includes non-invasive methods such as electroencephalography (EEG) in which electrodes are placed on the surface of the scalp, semi-invasive subdural electrocorticography (ECoG) arrays which are implanted underneath the skull and lie on the cortical surface, and invasive intracortical microelectrodes that penetrate the cortex. Several other types of neural interfaces have been tested for BCIs, such as functional near-infrared spectroscopy, electromyography, and endovascular devices, among others; however, they fall outside the scope of this review (for a brief review of these interfaces, see Pandarinath & Bensmaia, 2022).

Although non-invasive methods like scalp EEG may initially appear more appealing, practical challenges limit their utility as long-term assistive devices. Since EEG systems have much more limited spatial and temporal resolution than more invasive methods, they have typically been used for binary switches or for spelling words by selecting letters when highlighted on a display rather than for simultaneous high degree of freedom control required for arm and hand movement (Pandarinath & Bensmaia, 2022; Slutzky, 2019). EEG systems also require donning and doffing caps and typically applying gel, making their setup inconvenient and time-consuming.

ECoG arrays have higher spatiotemporal resolution than EEG due to their increased proximity to the recorded neurons. Signals from ECoG arrays also tend to provide greater long-term stability than those recorded using intracortical arrays and can be used reliably for BCI control

over multiple months in non-human primates and multiple days in humans without the need for recalibration (Blakely, Miller, Zanos, Rao, & Ojemann, 2009; Chao, Nagasaka, & Fujii, 2010). Despite these advantages, performance comparisons of 2D cursor control and communication BCIs have demonstrated that intracortical arrays outperform ECoG arrays (Pandarinath & Bensmaia, 2022).

Intracortical arrays, which record action potentials from single or multiple neurons, provide finer signal resolution and are better suited for simultaneous and continuous control of many degrees of freedom. Although intracortical microelectrode arrays provide the best signal for predicting movement intent, they have limited recording longevity (Barrese et al., 2013; Slutzky, 2019). Previous studies in both non-human primates and human subjects have found that signal quality degrades over time (Chestek et al., 2011; Downey et al., 2018; Hochberg et al., 2006; Suner, Fellows, Vargas-Irwin, Nakata, & Donoghue, 2005). Despite this, intracortical microelectrode arrays have been found to reliably record signals for 3-5 years and to provide repeatable, accurate BCI control over 1000 days (Bullard, Hutchison, Lee, Chestek, & Patil, 2020; Colachis IV et al., 2021; Downey et al., 2018; Hughes et al., 2021; Simeral, Kim, Black, Donoghue, & Hochberg, 2011; Sponheim et al., 2021).

2.2 Neural Encoding of Movement

2.2.1 Representational Model for the Encoding of Movement

A landmark study found that individual neurons in arm area of monkey motor cortex were strongly modulated by the direction in which the monkeys reached their arms (Apostolos P.

Georgopoulos, Schwartz, & Kettner, 1986). Neurons tended to fire maximally when the monkey would reach in a particular direction, termed the preferred direction, and would decrease its firing when reaching in other directions. This relationship can be described using a cosine tuning function:

$$f = b_o + b_x d_x + b_y d_y \qquad \textbf{Equation 2.1}$$

where f is the neuron's firing rate, d_x and d_y are the components of a unit vector pointing in the direction of movement, and b_o, b_x, and b_y are regression coefficients. This encoding model can be written as the dot product of two vectors, $\vec{D}$ for movement direction, with components d_x and d_y, and $\vec{B}$ for the neuron's preferred direction, with components b_x, and b_y. We can visualize this function as the linear projection of $\vec{D}$ onto $\vec{B}$, where the magnitude of $\vec{B}$ becomes proportional to the cosine of the angular difference between the neuron's preferred direction and reach direction:

$$f = \left|\vec{B}\right| \cos\theta \qquad \textbf{Equation 2.2}$$

An advantage of this encoding model is that it retains its form as we increase the number of control dimensions and consequently the number of movement parameters encoded by the neuron's firing rate (Collinger, Gaunt, & Schwartz, 2018; S. D. Kennedy & Schwartz, 2019). Thus, this model has been widely used in BCI studies to provide endpoint velocity control.

2.2.2 Limitations of the Representational Model

Early single neuron studies focused on establishing representational models that related single neuron activity to movement parameters (Evarts, 1968; A. P. Georgopoulos et al., 1982; Humphrey, Schmidt, & Thompson, 1970; Morrow, Jordan, & Miller, 2007; Thach, 1978). However, not all individual neurons modulate in response to movement variables, spurring alternative hypotheses that have brought these representational models under question (Churchland et al., 2012; Churchland & Shenoy, 2007; Fetz, 1992; Mussa-Ivaldi, 1988; Scott, 2008). Further limitations of this model include assumptions of independence and linearity of movement parameters and a stable relationship between neural firing rates and movement variables. These assumptions are often violated, given that the signals recorded from motor cortex are noisy and exhibit context-dependent modulation (e.g., changes over time, with behavioral states, with differing environmental cues, etc.) (Downey et al., 2017, 2018; Dunlap, Colachis, Meyers, Bockbrader, & Friedenberg, 2020; Hatsopoulos, Xu, & Amit, 2007; Kennedy, 2011; Kim et al., 2006; Perge et al., 2013; Rasmussen, Schwartz, & Chase, 2017; Sachs, Ruiz-Torres, Perreault, & Miller, 2015; Scott & Kalaska, 2017; Sergio & Kalaska, 1998; Simeral et al., 2011; Steinmetz et al., 2000; Suway et al., 2018; Yu et al., 2007). The continuing advancement of large scale neural recordings alongside emerging data analysis techniques have led researchers to instead begin considering the inherent structure of population neural activity (Briggman, Abarbanel, & Kristan, 2005; Cunningham & Yu, 2014; Gao & Ganguli, 2015; Shenoy, Sahani, & Churchland, 2013; Stevenson & Kording, 2011).

2.2.3 Neural Manifold-Based View of Motor Cortical Activity

Previous studies have demonstrated that the brain's underlying neural structure reflects established synaptic connections that are hard to change (Golub et al., 2018; Oby et al., 2019; Okun et al., 2015; Raichle & Raichle, 2001; Sadtler et al., 2014; Tsodyks, Kenet, Grinvald, & Arieli, 1999). As such, it has been proposed that patterns of neural activity are constrained to a low-dimensional neural subspace, which captures a large amount of variance in neural firing rates (Cunningham & Yu, 2014; Gallego, Perich, Miller, & Solla, 2017; Gao & Ganguli, 2015; Sadtler et al., 2014). Rather than considering neural activity completely independently, this model instead focuses on understanding the co-modulation of groups of neurons working together to generate movement. To visualize this concept, let us first consider the activity of three neurons within a high-dimensional space, where each coordinate axis corresponds to the firing rate of an individual neuron (Figure 1). The activity of all three neurons at each time step is represented by a point in this space. We can define a low-dimensional plane within this high-dimensional space that captures a large amount of variance in the neural firing rates. Shown in Figure 1, as the neural activity unfolds over time, we can trace out a trajectory that is mostly constrained to this low-dimensional neural subspace, or manifold, defined by its own latent coordinate axes. Since recorded neurons belong to a common underlying network and are not likely to be independent, it would make intuitive sense for fewer explanatory variables to reflect the coordinated response of a population of neurons (Cunningham & Yu, 2014).

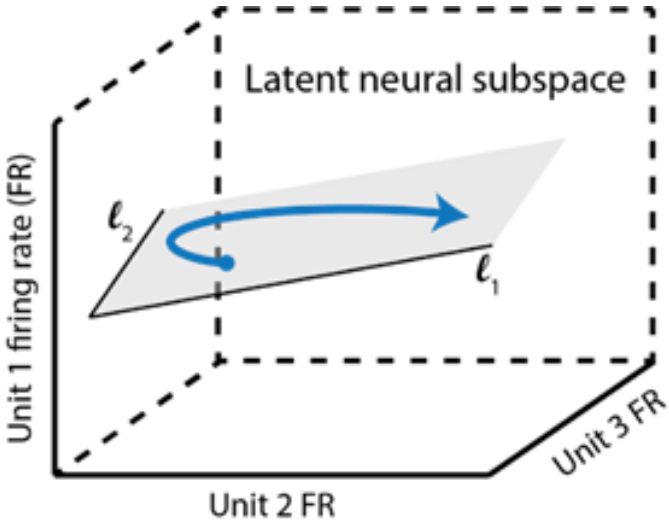

Figure 1 Conceptual illustration of a low-dimensional neural subspace.The blue neural trajectory describes the time-evolution of neural population activity along a low-dimensional plane shown in grey defined by the latent coordinate axes $\boldsymbol{l_1}$ and $\boldsymbol{l_2}$.

2.2.4 Dimensionality Reduction for Population-Level Analyses

Dimensionality reduction techniques have been used to estimate the underlying latent neural subspace and identify its corresponding explanatory variables. Covariance methods such as principal components analysis (PCA) and factor analysis (FA), time series methods such as Gaussian process factor analysis (GPFA) and latent linear dynamical systems (LDS), classification-based methods such as linear discriminant analysis (LDA), and non-linear methods such as Isomap and locally linear embedding (LLE) have been used to analyze the structure of population neural activity across a variety of species and brain areas (Cunningham & Yu, 2014; Gallego et al., 2017). Each method type has its own best use case depending on the questions that are being asked and can yield differing interpretations of the data. For example, PCA attempts to uncover the structure of neural activity in an unsupervised manner by identifying latent axes that

capture the greatest variance in the data. However, if our goal also includes identifying the latent axes that maximize between-class separation (e.g., target direction, stimulus type, decision taken, etc.), we may instead consider applying a supervised method such as LDA. These two methods would produce different latent axes and would yield different interpretations of the neural data.

2.2.5 Dynamical Systems Approach to Encoding of Movement

Numerous studies have taken a dynamical systems approach to understand the neural control of movement (Omrani, Matthew, Kaufman, Hatsopoulos, & Cheney, 2017; Pandarinath, Ames, et al., 2018; Shenoy et al., 2013; Vyas, Golub, Sussillo, & Shenoy, 2020). As opposed to the representational model (Equation 2.1) where each neuron directly relates to a given movement parameter, the dynamical systems model attempts to understand how computations are performed by a population of neurons to generate movement.

A dynamical system is a model that describes how a system changes over time. Let us consider a simple model without noise that operates independently of external factors:

$$\dot{x}(t) = f(x(t)) \qquad \textbf{Equation 2.3}$$

where $x(t)$ describes the state of the system and $f(x(t))$ represents the rules that determine how the system changes over time, specifically in which direction and how quickly the system is changing. By knowing the system's initial state, $x(0)$, and the dynamical rules that control the system's time evolution, $f(x(t))$, we can determine the state of the system at all times. In the case of a cortical system, $x(t)$ describes the latent low-dimensional neural state of a population of

neurons. Since motor cortex also receives and processes information from other brain areas, our model must also consider the addition of external inputs:

$$\dot{x}(t) = f\big(x(t)\big) + u(t) \qquad \textbf{Equation 2.4}$$

As we can see in Equation 2.4, the initial state of a dynamical system broadly drives the system's future state. Thus, in the case of a neural system, we would expect the initial neural state to determine the subsequent patterns of neural activity. Moreover, the same dynamical rules governing the time evolution of the neural state would produce different patterns of neural activity depending on the initial state. Several studies have suggested that preparatory activity in dorsal premotor cortex (PMd) and M1 acts as the initial state of a neural dynamical system for the generation of reaching movements (Churchland et al., 2012; Churchland, Cunningham, Kaufman, Ryu, & Shenoy, 2010; Churchland, Yu, Ryu, Santhanam, & Shenoy, 2006; Sussillo, Churchland, Kaufman, & Shenoy, 2015). This preparatory activity, recorded during a delay period prior to movement onset, could be used to predict the progression of neural activity patterns during movement. Since the dimensionality of neural activity is greater than that of muscle activity, dynamical systems models also predict that not all neural activity patterns need to map to the output of movement (Omrani et al., 2017; Pandarinath & Bensmaia, 2022). Instead, various studies suggest that population activity may also encode internal processes that help provide simpler and smoother dynamics (Churchland et al., 2010; Elsayed, Lara, Kaufman, Churchland, & Cunningham, 2016; Kaufman, Churchland, Ryu, & Shenoy, 2014; Kaufman et al., 2016; Russo et al., 2018).

2.3 Decoding Neural Signals for BCI Control

Intracortical microelectrode arrays record extracellular voltage signals from surrounding neurons and subsequently tend to record activity from more than one unit on a given electrode. Neural recordings have historically been spike sorted, a process that separates action potentials of individual neurons, where the goal has been to characterize the properties of single units. However, to achieve high-performing BCIs, our goal is to obtain as much information as possible from our recordings. Numerous studies have demonstrated the use of multiunit activity (MUA), in which a voltage threshold is used to detect threshold crossings of the voltage waveform, for achieving high-quality BCI control (Fraser, Chase, Whitford, & Schwartz, 2009; Oby et al., 2016; Stark & Abeles, 2007; Trautmann et al., 2019; Ventura, 2008). For the scope of this dissertation, we will focus on the use of MUA for BCI decoding.

2.3.1 Population Vector Algorithm

Based on the representational model described above, each recorded neuron fires proportionally to the angular difference between the unit's preferred direction and the intended movement direction (see section 2.2.1 for details). The broad directional tuning of neurons allows us to develop a framework for predicting intended movement direction, also referred to as decoding neural activity. (Georgopoulos et al., 1986). First, we define a vector containing each neuron's preferred direction based on the neural response to a set of actual or intended movements. We then take the vector sum of the preferred directions, weighted by the neurons' instantaneous normalized firing rates, to produce a population vector pointing in and predicting the movement direction. The population vector algorithm (PVA) has been used to decode monkeys' movement

intentions for 3D cursor control and for 4D control of a robotic arm and gripper for self-feeding (Taylor et al., 2002; Velliste et al., 2008). However, decoding biases arise when the preferred directions of the sample population are not uniformly distributed (Chase, Schwartz, & Kass, 2009a).

2.3.2 Optimal Linear Estimator Decoder

The optimal linear estimator (OLE) decoder has been shown to significantly outperform the population vector algorithm (Chase, Schwartz, & Kass, 2009b; Salinas & Abbott, 1994). Optimal linear estimation can be implemented similarly to PVA, but rather than solely relying on the directional tuning of each unit, it corrects for non-uniformity in the distribution of preferred directions by assigning optimal weights to each unit's population-vector contribution. The highest level of robotic arm and hand control achieved by people with tetraplegia has been obtained using indirect OLE, resulting in significant improvement on tests of upper limb function (Collinger, Wodlinger, et al., 2013; Wang, Chan, Heldman, & Moran, 2007; Wodlinger et al., 2015).

2.3.3 Kalman Filter Approaches

Commonly used in guidance and navigation systems, the Kalman filter is an optimal estimation algorithm used to predict the state of a system from indirect, noisy measurements. The Kalman filter uses two models: (1) the system model, which in the case of a neural system describes the time-evolution of movement, and (2) the observation model, which describes the relationship between neural activity and movement (Kim, Simeral, Hochberg, Donoghue, & Black, 2008; Wu, Gao, Bienenstock, Donoghue, & Black, 2006). By predicting the current movement

based on the movement at the previous time step and updating this prediction using measurements of the neural activity, we obtain an updated estimate of the movement. By taking the dynamics of the decoded movement into account, we obtain smoother control compared to that of an OLE decoder. Kalman filters have been used to achieve the best cursor-based control to date for communication and recreational computer-based applications (Gilja et al., 2015a; Nuyujukian et al., 2018; Pandarinath et al., 2017).

2.3.4 Neural Network-Based BCIs

Increased visibility into the field of deep learning has propagated the use of neural networks. When compared to linear approaches, which are the current standard in BCI research, neural networks have been shown to produce comparable BCI performance (Sussillo et al., 2012, 2016). Neural network-based approaches have also been demonstrated to improve BCI reliability in non-human primates by leveraging large training datasets to overcome various types of recording instabilities (Sussillo et al., 2016). Despite their advantages, neural network-based approaches are computationally expensive and require large sets of training data which make them difficult to apply in a clinical setting.

2.3.5 Neural Dynamical Systems Approaches

Numerous studies have begun investigating whether the use of neural dynamical systems can improve BCI performance (Pandarinath, Ames, et al., 2018). Although several studies have begun incorporating a dynamical systems approach for BCI cursor control, there are no studies to date investigating the use of dynamics for closed-loop BCI control of reach and grasp movements

(Kao et al., 2015; Willett et al., 2016). Offline studies have shown that a dynamical systems-based approach can outperform standard linear approaches used for BCI control of reaching movements (Makin, O'Doherty, Cardoso, & Sabes, 2018; Pandarinath, O'Shea, et al., 2018). However, it remains unclear whether this type of approach would work as effectively for hand and grasp control (Rouse & Schieber, 2018; Suresh et al., 2020).

2.4 Brain-Computer Interfaces for Restoration of Upper Limb Function

In surveys of potential BCI users with SCI, the restoration of arm and hand function has been shown to be of high value, particularly for those with cervical SCI (Anderson, 2004; Blabe et al., 2015; Collinger, Boninger, et al., 2013; Huggins et al., 2015). Various BCI-controlled technologies have been developed to achieve this goal, such as robotic arms, exoskeletons, and functional electrical stimulation (FES) (Aflalo et al., 2015; Ajiboye et al., 2017; Bouton et al., 2016; Carmena et al., 2003; Chadwick et al., 2011; Collinger, Wodlinger, et al., 2013; Ethier, Oby, Bauman, & Miller, 2012; Hochberg et al., 2012, 2006; Moritz, Perlmutter, & Fetz, 2008; Velliste et al., 2008). Early studies demonstrated the use of robot arm control for the restoration of upper limb function in non-human primates (Carmena et al., 2003; Velliste et al., 2008). The field has continued to progress past proof-of-concept demonstrations towards the goal of clinical translation by increasing the dimensionality of simultaneous, continuous BCI control from 4D to 7D to 10D over the past decade (Collinger, Wodlinger, et al., 2013; Hochberg et al., 2012; Wodlinger et al., 2015). These additional degrees of freedom have provided participants with greater wrist and hand control for use in completing functional tasks, such as reaching to, grasping, carrying, and placing an object on a shelf.

Functional electrical stimulation is a technique that applies an electrical stimulus to a muscle or nerve, causing a desired muscle contraction. FES has been used to restore native muscle function in people with SCIs (Ragnarsson, 2007). By restoring muscle movement, studies of both surface and implantable FES systems have demonstrated participants' ability to move their own hand to grasp and release objects and perform daily living tasks such as eating food with a fork (Alon & McBride, 2003; Peckham et al., 2001). Recent advances have led to BCI-FES systems where participants use cortically controlled stimulation to activate their own arm muscles (Ajiboye et al., 2017; Bouton et al., 2016).

2.5 Barriers and Limitations to Clinical Adoption of BCIs

2.5.1 Signal Quality and Stability

As discussed in section 2.1, intracortical microelectrode arrays provide greater signal resolution than less invasive methods, yielding more accurate predictions of movement intention. However, these types of recording arrays require invasive surgeries for implantation, making them difficult to repair or replace if hardware issues arise. Various longitudinal studies have sought to characterize the reliability of these implants in both non-human primates and human subjects (Barrese et al., 2013; Bullard et al., 2020; Chestek et al., 2011; Colachis IV et al., 2021; Downey et al., 2018; Sponheim et al., 2021; Suner et al., 2005). Although one study found that the average lifespan of intracortical microelectrode arrays was 622 days, they also found that human implants tended to last longer than NHP implants (Sponheim et al., 2021). The authors suggested that this difference may arise from greater quality control checks prior to implant, surgical procedures done

by neurosurgeons rather than investigators, and greater care of the connector and surrounding skin throughout the lifetime of the implant in humans. Indeed, long-term studies of BCI control have found that human study participants can continue to effectively use their BCI systems over 1000 days (Colachis IV et al., 2021; Hughes et al., 2021; Simeral et al., 2011; Sponheim et al., 2021).

To provide adequate BCI control across these long time scales, BCIs are typically tested under ideal laboratory conditions, in which a new decoder is calibrated at the beginning of each session for optimal control. However, neural activity recorded by implanted intracortical electrodes can vary over short time spans. Changes in recorded neural activity, or neural recording instabilities, may arise from small movements of the electrode arrays as well as biological reactions to the implanted electrodes. In the former case, we could potentially lose information from neurons that fall in and out of recording range. These types of changes have been found to be prevalent in human subjects, with over half of recorded units becoming unstable from one day to the next (Downey et al., 2018). Since BCI decoders are defined using neural activity recorded during a calibration period, any instabilities impact the decoder's ability to correctly predict the user's intended movement. Failure to account for these neural recording instabilities can lead to degradation of BCI control, often within hours.

Several studies have sought to improve decoder stability and provide more accurate control over long-time scales. Many of these recalibration methods update decoder or tuning parameters while the participant performs a simple BCI task (Bishop et al., 2014; Homer et al., 2014; Jarosiewicz et al., 2015; Li et al., 2011; Orsborn et al., 2014; Wu & Hatsopoulos, 2008; Zhang & Chase, 2013). However, the parameter updates require knowledge of the user's intent. These supervised methods tend to work well for structured tasks with a predetermined set of available movements. However, this does not generalize well outside of controlled laboratory conditions.

Neural network-based decoders have also been implemented, demonstrating stability of provided control in the face of various types of neural recording instabilities by leveraging large amounts of training data (Karpowicz et al., 2022; Sussillo et al., 2016). However, these deep learning methods require new users to repeatedly perform training tasks over many days prior to using their BCI to obtain large training datasets. Additionally, training these types of neural network-based decoders require a lot of computational power and time, taking up to several hours to train a decoder, making this method difficult to implement in clinical practice. Despite the development of several different types of decoders to overcome neural recording instabilities, many of these studies have not tested their methods during online, closed-loop control, where user adaptation can occur, and instead have tested them exclusively offline. As we strive to provide BCI control for increasingly complex tasks, we must keep in mind the effect of neural signal instabilities as the participant uses their BCI system over a period of hours to days and their impact on BCI performance.

2.5.2 Context-Dependent Effects on BCI Performance

Throughout this chapter, we have been focusing on motor cortical activity for decoding. Historically, intracortical BCIs have targeted the hand and arm region of M1 since M1 is the primary cortical output for movement generation (E. E. Fetz & Cheney, 1980; Gallego, Makin, & McDougle, 2022; Strick, Dum, & Rathelot, 2021). Many neurons in this region project to spinal motoneurons that then project to muscles which drive movement (Gallego et al., 2022; Strick et al., 2021). Motor cortical activity also tends to correlate with movement parameters (e.g. hand velocity), as described in section 2.2.1 and 2.2.2. However, despite our advances in BCI technology, we have found that the relationship between neural activity and behavior can be context and task-dependent (e.g. whether an object is present in a workspace or the identity of the

object). Previous work from our lab has found that when decoders cannot generalize to these types of contextual changes, BCI performance is negatively impacted (Downey et al., 2017). However, M1 does not exist in isolation, and it is only one piece of a larger network involved in movement generation. There are a number of other brain regions that process movement-related information and are involved in motor output.

Various cortical motor regions outside M1 are thought to contain higher-order information about movement. To generate a sequence of movements, for example, such as when you're playing an instrument or typing, the appropriate commands must be provided not only to perform the individual movements such as hitting an individual key but also hitting the keys in the correct order. In cortex, these types of higher-level sequence representations are encoded in premotor and parietal areas while individual movements are represented in M1 (Gallego et al., 2022; Yokoi & Diedrichsen, 2019). Both of these cortical areas have been used in BCI studies to decode higher-order movement information. Neural activity in premotor cortex has been used to decode planned motor sequences rather than single movements (Shanechi et al., 2012). Although the posterior parietal cortex (PPC) is not considered a motor cortical area, it lies between S1 and M1 and has been found to encode both sensory-related and motor-related properties (Andersen, Aflalo, & Kellis, 2019). This region has been found to encode movement goals and their respective trajectories and has been used for goal-decoding BCIs (Aflalo et al., 2015).

This hierarchical organization of movement-related information points to dedicated motor networks for efficient movement production. For example, effective grasping requires processing an object's location and properties and then transforming them to the appropriate motor commands that guide and shape the hand to grasp of the object (Davare, Kraskov, Rothwell, & Lemon, 2011). The underlying cortical grasping circuit is thought to processes real-time visual information in the

anterior parietal area (AIP) to better inform hand shaping for differentially shaped objects (Murata, Gallese, Luppino, Kaseda, & Sakata, 2000). Ventral premotor cortex (PMv) has also been found to generate sets of possible grasping actions for object manipulation (Giacomo Rizzolatti & Luppino, 2001). Since AIP and PMv are reciprocally connected and receive inputs from ventral visual streams, these areas have access to object identity information that could be used to fine-tune grasping so that it is more appropriate for a given object (Davare et al., 2011). PMv also has reciprocal connections with M1, allowing PMv to contribute to the control of grasping movements (Davare et al., 2011; Davare, Lemon, & Olivier, 2008). Here, the AIP-PMv-M1 circuit demonstrates how incorporating additional motor cortical areas outside of M1 is crucial for fluent motor control.

There are also a number of nonmotor-related signals that impact how we plan and execute our movements, such as sensory feedback. In order to smoothly transition from one phase of movement to another (e.g. contact with an object marks the end of the reach phase and the beginning of the next phase), the brain has to accurately predict the required motor output and expected sensory events (Johansson & Flanagan, 2009a). If a perturbation arises or the sensory predictions are erroneous, corrections can be made based on incoming sensory signals. However, BCIs do not always provide sensory feedback during a task, leaving users to solely rely on visual feedback. Despite the lack of incoming sensory information, S1 may play an additional, broader role in movement production. A recent study found that S1 activity encoded imagined reaches during a motor imagery task (Jafari et al., 2020b). The authors suggested that in the absence of sensory feedback, S1 signals could be used to decode movement intention, indicating another region containing additional motor-related information outside of M1.

Apart from sensory feedback, we should also consider how changes in internal state impact motor control and behavior. For example, M1 has been found to encode signals related to motivation, intrinsic value, and reward expectation (Galaro, Celnik, & Chib, 2019; Gallego et al., 2022; Hennig et al., 2021; Levy et al., 2020; Ramkumar, Dekleva, Cooler, Miller, & Kording, 2016). These inputs potentially arise from various brain areas such as the from the limbic system or from subcortical regions such as the basal ganglia or cerebellum. Based on the latency of reward signals, for example, the basal ganglia-thalamo-cortical loop is thought to project reward-related information to M1 (Ramkumar et al., 2016). Another study cites the role of the ventral striatum as a limbic-motor interface that could impact motor approach and attention in task performance (Galaro et al., 2019). These nonmotor signals that are projected to M1 may correlate with motor-related signals in a way that could add a layer of difficulty to motor decoding.

Here, we have shown that numerous different brain areas and cortical loops are involved in movement generation. Going back to our idea of calibrating robust, generalizable decoders that can output the appropriate kinematic commands across a variety of contexts, we should consider the role of these different brain areas, their projections to M1, and how we can disentangle context-related information in M1 to decode movement signals. This leads us to the crux of this work: how do context-related signals impact M1 activity? Here, we will explore context and task-dependent effects on motor cortical activity as well as BCI performance.

3.0 General Methods

Two participants enrolled in this study as part of a clinical trial conducted under an Investigational Device Exemption granted by the United States Food and Drug Administration (NCT01894802) and was approved by the Institutional Review Boards at the University of Pittsburgh and the Space and Naval Warfare System Center Pacific. Informed consent was obtained from the participants prior to study participation.

3.1 Electrophysiology Methods

3.1.1 Array Implantation

Two participants with tetraplegia enrolled in this study. Participant P2, a 28-year-old (at time of implant) male, had a C5 motor/C6 sensory ASIA B spinal cord injury. Participant P3 was 28 years old at time of implant and had a C6 ASIA B spinal cord injury. Two platinum-coated silicon intracortical microelectrode arrays (Blackrock Microsystems, Salt Lake City, UT), with electrodes arranged in a 10x10 grid, were implanted in the hand and arm region of the participants' left motor cortex. The arrays were 4x4 mm in size, with electrode shank lengths of 1.5 mm. The participants also had two 2.4x4mm, 60-channel intracortical microelectrode arrays, arranged in a 6x10 grid, implanted in area 1 of somatosensory cortex (S1). Electrode tips were coated with sputtered iridium oxide film. For participant P2, 88 of the 100 electrodes on each M1 array and 32

of the 60 electrodes on each S1 array were wired to two percutaneous connectors attached to the skull. Participant P3 had 96 and 32 wired electrodes in each set of M1 and S1 arrays, respectively.

3.1.2 Signal Processing

Neural data were recorded using either patient cables or Cereplex-E digital headstages along with the Neuroport Neural Signal Processor (Blackrock Microsystems, Salt Lake City, UT). Extracellular voltage recordings were sampled at 30 kHz and filtered using a 1st order 750 Hz high-pass filter (Weiss, Flesher, Franklin, Collinger, & Gaunt, 2018). Each time the voltage signal crossed a pre-defined spike threshold, it was counted as a spike for the given channel. Thresholds were set at -4.5 times the root-mean square voltage on each channel at the beginning of each session. Instantaneous firing rates were estimated by convolving spike counts binned at 50 Hz with a 440 ms decaying exponential filter.

3.2 Decoder Calibration

Two types of decoders were used for the experiments presented in this dissertation: an optimal linear estimator (OLE) decoder used for reach and grasp control and a Kalman filter-factor analysis (KF-FA) decoder used for cursor control (details provided in sections 5.2.2 and 4.2.2, respectively). Decoders were trained at the beginning of each test session using a 2-step calibration procedure (Velliste et al., 2008). First, the participants were instructed to observe an automated version of a training task, where either a computer cursor or a virtual arm reached toward a target, achieved the goal at the target location, and proceeded to the next target. The participants were

asked to imagine performing the task while the computer controlled the kinematics of the end effector. This open-loop task was used to estimate the relationship between neural firing rates and ideal intended kinematics. A decoder was then trained and used for the next step in calibration, where the participants repeated the training task. The decoded neural activity controlled the virtual end effector while the computer attenuated command signals orthogonal to the target direction ("orthogonal assistance") (Collinger, Wodlinger, et al., 2013; Velliste et al., 2008). In other words, participants had BCI control and were required to exert velocities to move the end effector but could only move either directly towards or away from the target. A final decoder was calibrated from the recorded neural activity and assisted velocity commands from this second step. This final decoder was used to provide full BCI control to the participants for the subsequent experiments.

3.3 Dimensionality Reduction for Neural Data Analysis

Dimensionality reduction is a statistical method that projects high-dimensional data, such as neural activity recorded from multiple channels, onto a lower-dimensional space. This is useful when one expects that the factors that cause covariation in the measured variables can be explained with a smaller subset of variables, or latent variables. In the case of recorded neurons, we can find the salient features that explain population activity while disregarding noise. In other words, we are taking noisy measurements of individual units and calculating the latent variables representing the neural activity patterns that make up the population response. Here, we present the processing steps and specific dimensionality reduction methods used throughout the studies discussed in this work.

3.3.1 Data Preprocessing

Minimal data preprocessing was done prior to estimating the latent neural subspace and was mainly conducted to provide suitable inputs for dimensionality reduction. Since units with very low (i.e., near zero) variance can cause instabilities in some methods (Cunningham & Yu, 2014), we excluded units with low firing rates (< 1 Hz). Threshold crossings were recorded at a sampling rate of 50 Hz; however, we found that using 20 ms bins yields low spike counts. To overcome this, we increased the bin size to 40 ms. Finally, we smoothed the data using a Gaussian kernel with a 200 ms standard deviation bandwidth. This helped produce smoother neural trajectories within the neural subspace. We found that normalization of the neural activity was not necessary since the methods used throughout this dissertation are scale-invariant. This general preprocessing scheme was used for offline analyses of the neural data, such as those presented in chapters 5.0 and 6.0. Data preprocessing for online cursor control is described in section 4.2.2.

3.3.2 Determining the Dimensionality of a Dataset

For simple movement tasks, studies tend to estimate an ~10D neural subspace, based on previously published dimensionality estimates in non-human primates (Degenhart et al., 2020; Oby et al., 2019; Sadtler et al., 2014). Although our task used a similar number of control dimensions, we opted to estimate the dimensionality of the neural activity for each of our datasets to better capture the patterns of co-modulation among the recorded units. A simple approach to estimating the optimal dimensionality of a dataset involves choosing a cutoff value for the variance explained by the latent dimensions and selecting the top D dimensions (Cunningham & Yu, 2014). However, this cutoff value is typically subjective and varies according to the experimenter.

DataHigh, a graphical user interface for dimensionality reduction and visualization of population neural activity, uses cross-validated metrics to determine the optimal number of dimensions (Cowley et al., 2013). In this work, we used factor analysis for initial dimensionality reduction. Since this is a probabilistic method, we used DataHigh to the identify the dimensionality that maximized the cross-validated data likelihood (Cunningham & Yu, 2014).

3.3.3 Factor Analysis for Dimensionality Reduction

Factor analysis is a covariance-based method that seeks to extract the dimensions that capture variance shared across neurons while ignoring the independent variance of individual neurons (e.g., spiking variability) (Cunningham & Yu, 2014). Unlike principal component analysis, which assumes no observation noise and thus falls subject to spiking variability, FA is better able to capture firing rate variability while ignoring spiking variability using a non-isotropic noise model (Cowley et al., 2013; Cunningham & Yu, 2014). For data analysis purposes, we used DataHigh's implementation of FA for dimensionality reduction to obtain an orthonormalized projection of the neural data.

3.3.4 Linear Discriminant Analysis for Maximal Separation Between Conditions

Linear discriminant analysis is a supervised approach to dimensionality reduction that maintains condition-discriminatory information. LDA not only cares about identifying the latent dimensions that maximize the variance of the data, like in the case of PCA, but also those that maximize the separation between groups (Cunningham & Yu, 2014). To obtain this projection, we want to find the dimensions along which between-class variance is maximized while within-class

variance is minimized. We can solve this optimization problem using singular value decomposition to obtain a set of eigenvectors that span the lower-dimensional subspace. LDA can be used either as a classifier or as a dimensionality reduction technique, reducing the number of dimensions to k-1, where k is the number of classes.

4.0 Stability of Task-Related Neural Activity Over Time

This chapter is based on a first author manuscript which is currently in preparation.

Intracortical brain-computer interfaces (BCIs) can enable control of assistive devices to restore function for people with motor impairments. However, neural recording instabilities can change the relationship between neural activity and movement intention defined by the BCI, thus degrading performance. Here, we overcome these instabilities by stabilizing multielectrode neural recordings such that a stable representation of the low dimensional latent space defined for the neural population is recovered. We demonstrate that neural signal stabilization can provide consistent BCI control of a computer cursor over a period of 104 days in a person with tetraplegia. Our method outperforms a fixed decoder, which is not updated over the course of the experiments, and it performs comparably to a same-day trained decoder. Neural signal stabilization addresses user priorities for independent BCI operation by enabling stable BCI control without supervised recalibration.

4.1 Introduction

Intracortical brain-computer interfaces (BCIs) have enabled people with severe motor impairments to control assistive devices by transforming recorded neural activity into desired movement commands (Ajiboye et al., 2017; Bouton et al., 2016; Collinger, Wodlinger, et al., 2013; Gilja et al., 2015b; Hochberg et al., 2012; Wodlinger et al., 2015). High performance communication has been demonstrated (Pandarinath et al., 2017; Willett, Avansino, Hochberg,

Henderson, & Shenoy, 2021), and intracortical BCIs have begun to transition into home environments enabling independent computer access (Dekleva et al., 2021; Weiss et al., 2019). Intracortical BCIs act by mapping neural activity recorded from multielectrode arrays, specifically firing rates, onto control parameters, such as the 2D velocity of a cursor. BCIs thus operate on the assumption that the relationship between neural activity on a given channel and movement is constant. The presence of neural recording instabilities, or variations in the recorded signals over time, violates this assumption, causing performance to decline (Downey et al., 2018; Jarosiewicz et al., 2013, 2015; Nuyujukian et al., 2014; Perge et al., 2013; Wessberg & Nicolelis, 2004).

Several studies have found that neural recording instabilities pose a problem for continuous decoding in humans over the span of hours or days (Colachis IV et al., 2021; Downey et al., 2018; Jarosiewicz et al., 2015; Perge et al., 2013; Simeral et al., 2011), as opposed to the long-term signal stability seen in non-human primates (Chestek et al., 2007a; Dickey, Suminski, Amit, & Hatsopoulos, 2009; Flint, Wright, Scheid, & Slutzky, 2013). Previous work from our lab has demonstrated that neural recording instabilities occur within a single session, with units becoming unstable at a rate of approximately 11% per hour (Downey et al., 2018). After one day, roughly 60% of units are unstable (Downey et al., 2018). These neural recording instabilities may arise from device factors such as micromovements of the electrode arrays and encapsulation of the electrodes or environmental factors such as noise sources (Degenhart et al., 2020; Perge et al., 2013; Sussillo et al., 2016). Since BCI decoders are defined using neural activity recorded during a calibration period, changes in baseline firing rate or tuning properties will impact the decoder's ability to correctly predict the user’s intended movement (Kim et al., 2006; Santhanam et al., 2007; Wessberg & Nicolelis, 2004). This places a burden on the user by requiring them to stop using the BCI as performance degrades to perform recalibration, a procedure that can take many minutes.

Potential BCI users desire systems that perform with a high degree of accuracy and can be operated independently while minimizing setup time and intervention from a caregiver or technician (Blabe et al., 2015; Collinger, Boninger, et al., 2013; Huggins et al., 2015; Huggins, Wren, & Gruis, 2011). To address users' needs and mitigate this problem, BCI decoders must be capable of providing stable control over time despite changes in neural recordings.

Overt movements are generated through the coordinated activation of neural populations (Churchland et al., 2012; Cunningham & Yu, 2014; Gallego, Perich, Chowdhury, Solla, & Miller, 2020; Gallego et al., 2017; Sadtler et al., 2014). One study demonstrated that for well-practiced movements, the population-level neural representation of movement-related activity is remarkably stable once the impact of single neuron recording instabilities is removed (Gallego et al., 2020). Thus, rather than addressing neural recording instabilities on a unit-by-unit basis, as has been proposed by former re-calibration methods (Bishop et al., 2014; Homer et al., 2014; Jarosiewicz et al., 2015; Li et al., 2011; Orsborn et al., 2014; Wu & Hatsopoulos, 2008; Zhang & Chase, 2013), we have developed a method that recovers a stable representation of the neural population (Degenhart et al., 2020). Our method uses an adaptive stabilizer which provides a stable input to a fixed BCI decoder (Figure 2A). To obtain a stable representation of neural population activity across days, the adaptive stabilizer (orange) uses binned spike counts from all channels to obtain an estimate of the neural subspace for a given time window. In our example, which uses three units for visualization purposes, there is a change in the recorded neural population from Day 1 to Day n. This causes a change in the estimated neural subspace on Day n. Although individual units that we are able to record from may vary over time, the low-dimensional neural subspace is thought to reflect constraints imposed by neural connectivity and thus should contain a stable representation of movement intention (Cunningham & Yu, 2014; Gallego et al., 2017; Golub et al., 2018; Oby et

al., 2019; Sadtler et al., 2014). To align the estimated subspaces over time, the stabilizer identifies electrodes with stable recordings (e.g. Units 1 and 2) and defines the coordinate axes ($\boldsymbol{l_1}$ and $\boldsymbol{l_2}$) of the estimated subspaces with respect to the stable electrodes. This stabilized latent neural activity is then passed to a fixed decoder to produce the user's desired movement. We utilize this stable mapping between latent neural activity and intended movements to provide long-term, stable BCI control.

In previous work, we demonstrated that neural signal stabilization enabled stable BCI control in non-human primates where simulated instabilities were applied to neural recordings (Degenhart et al., 2020). Neural signal stabilization was robust to a variety of instability types including baseline shifts, unit drop-out, changes in tuning properties, and combinations of these three (Degenhart et al., 2020). Although our method was demonstrated in monkeys, it was not tested using real recording instabilities over a period of consecutive weeks. Here, we implement the neural signal stabilization method in a clinical trial involving a human participant with implanted intracortical microelectrode arrays (Blackrock Microsystems, Inc.) to overcome natural recording instabilities over a period of 18 test sessions, spanning 104 days. We demonstrate that neural signal stabilization provides consistent BCI control of a computer cursor and surpasses performance provided by a fixed decoder, trained on the first day of these experiments, across all sessions. Movement direction-specific structure in the stabilized latent representation of population activity is also consistent over time, confirming that neural signal stabilization is accurately aligning daily recordings to a common and consistent neural population space. When compared to a decoder that is calibrated each day, neural signal stabilization enables comparable performance both quantitatively in terms of performance metrics and qualitatively, based on participant feedback. We show that by extracting a stable representation of population-level neural

activity from unstable neural recordings, human BCI control can be sustained for multiple months without the need for clinician intervention.

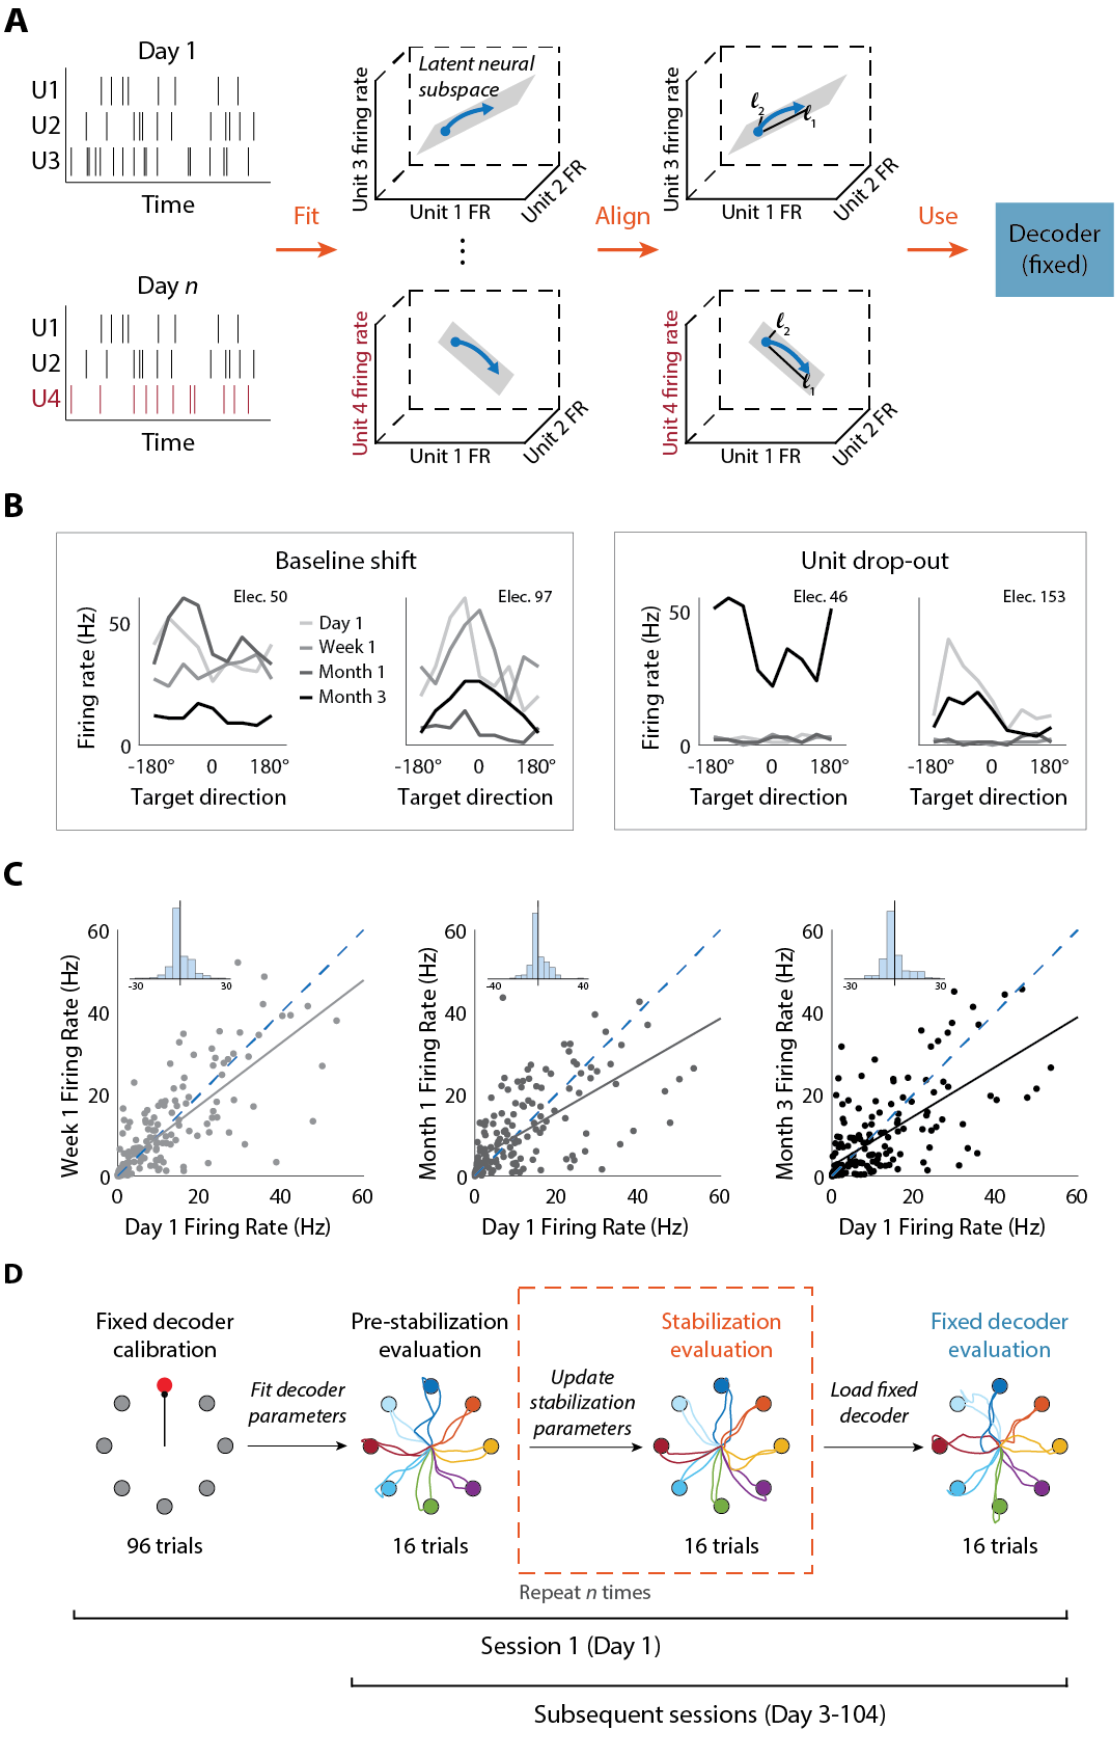
A
Day 1
U1
U2
U3
Time
Day n
U1
U2
U4
Time
Fit
Latent neural subspace
Unit 3 firing rate
Unit 4 firing rate
Unit 1 FR
Unit 2 FR
Align
Use
Decoder (fixed)
B
Baseline shift
Elec. 50
Elec. 97
Day 1
Week 1
Month 1
Month 3
Firing rate (Hz)
Target direction
Unit drop-out
Elec. 46
Elec. 153
C
Week 1 Firing Rate (Hz)
Month 1 Firing Rate (Hz)
Month 3 Firing Rate (Hz)
Day 1 Firing Rate (Hz)
D
Fixed decoder calibration
Fit decoder parameters
Pre-stabilization evaluation
Update stabilization parameters
Stabilization evaluation
Load fixed decoder
Fixed decoder evaluation
96 trials
16 trials
16 trials
16 trials
Repeat n times
Session 1 (Day 1)
Subsequent sessions (Day 3-104)

Figure 2 Neural signal stabilization method and neural recording instabilities. A. Overview of the neural signal stabilization method. Our method provides a stable input to a fixed BCI decoder (blue) by finding a stable representation of neural activity patterns. B. Examples of natural recording instabilities throughout the course of the experiment. Tuning curves are plotted for example units at four different reference time points during the experiment: day 1, day 7 (week 1), day 28 (month 1), and day 83 (month 3). Firing rates were estimated over a 500 ms window starting at reach onset for neural activity recorded during the fourth stabilization update's evaluation block (16 trials per day distributed equally across target directions). C. Comparison of firing rates from each electrode between Day 1 and the other 3 reference time points. Firing rates were estimated using the same data as that in B. The identity line is shown in blue, and the least squares regression line is shown in grey. Inset histograms show the residuals of the corresponding linear regression model. D. Experimental design using neural signal stabilization across multiple test sessions. On Day 1 we calibrated the experiment's initial BCI decoder and evaluated BCI performance during the pre-stabilization evaluation block. This evaluation block was then used to update the stabilizer parameters followed by a stabilization evaluation block. As the session progressed, trials from the stabilization evaluation blocks were continuously added to a sliding buffer. The stabilizer parameters were updated every 16 trials while the decoder parameters stayed constant. We evaluated performance after each update. After the final stabilization update, the fixed decoder parameters were retested at the end the session. Subsequent sessions began using the fifth stabilization update from the previous session for pre-stabilization evaluation.

4.2 Methods

To overcome natural recording instabilities and provide consistent and proficient BCI control over time, we adapted the neural signal stabilization method presented in Degenhart et al. for our human BCI study and tested its efficacy using a 2D cursor center-out task over 18 test sessions, spanning 104 days (Degenhart et al., 2020). Data was collected from participant P2 approximately 3.5 years after implant (1241-1344 days). A threshold was set at -4.5 times the root-

mean square voltage on each recording channel at the beginning of the first test session and was fixed for the remainder of the 104-day experiment.

4.2.1 Characterizing Neural Recording Instabilities

Tuning curves were calculated using data from the fourth stabilization evaluation block across four reference time points: Day 1, Week 1 (day 7), Month 1 (day 28), and Month 3 (day 83). Firing rates were calculated for the first 500 ms after reach onset. Although instances of tuning changes were observed, in which there was a change in the relationship between a unit's activity and the user's intended movement direction, we did not present them in this work because the tuning curves were calculated based on one block of data per day (effectively 2 trials per target direction).

4.2.2 Neural Signal Stabilization and BCI Decoder Calibration

Here we apply our method which we previously demonstrated in non-human primates to enable proficient BCI performance in the face of simulated recording instabilities (Degenhart et al., 2020). To obtain stable BCI control, our method obtains a stable representation of population neural activity and uses this to provide a stable input to a fixed decoder. For online control, we used a Kalman filter-factor analysis (KF-FA) decoder consisting of: (1) an adaptive stabilizer and (2) a fixed decoder. The stabilizer related neural activity, $x_t \in \mathbb{R}^q$, defined as binned spike counts from q electrodes at time step t, to a latent state, $z_t \in \mathbb{R}^{10}$, using factor analysis:

$$z_t \sim \mathcal{N}(0, I)$$ **Equation 4.1**

$$x_t | z_t \sim \mathcal{N}(\Lambda z_t + \mu, \Psi)$$ **Equation 4.2**

where $\Lambda \in \mathbb{R}^{q \times 10}$ is the loading matrix, which relates each electrode to each latent factor, $\mu \in \mathbb{R}^{q}$ is a vector of mean spike counts, and $\Psi \in \mathbb{R}^{q \times q}$ is a diagonal non-isotropic noise matrix such that each electrode can exhibit different levels of observation noise (Degenhart et al., 2020). We used 10 dimensions to define the neural subspace as this value has been previously shown to capture the majority of the shared variance of neural activity recorded during a 2D center-out task (Degenhart et al., 2020; Sadtler et al., 2014). We estimated the parameters Λ, μ, and Ψ from the neural activity and then obtained estimates of latent state $\hat{z}_t$ using the following projection:

$$\hat{z}_t = \beta(x_t - \mu)$$ **Equation 4.3**

where $\beta = \Lambda^T(\Lambda\Lambda^T + \Psi)^{-1}$. The fixed decoder predicted endpoint velocities from the estimated latent neural activity $\hat{z}_t$ using a Kalman filter (Figure 2A).

The KF-FA structure lets us align latent neural activity in the neural subspace rather than change any Kalman filter parameters to maintain stable BCI control. The alignment procedure identified electrodes that maintained a stable mapping with the latent neural subspace by comparing the rows of the loading matrix obtained during initial calibration to the rows of the loading matrix from subsequent blocks of trials (for additional details, see Degenhart et al., 2020). Once the loading matrix was aligned, stabilized estimates of latent space $\hat{z}_t$ were obtained using Equation 4.3.

To adapt the alignment method for our human BCI study, we heuristically determined the proper parameters, including number of alignment channels (130 channels), bin size (60 ms), and isotropic state noise covariance used for the Kalman filter (0.001), needed to provide the user with consistent BCI control. These values differed from those used in (Degenhart et al., 2020) due to the number of channels used for training the decoder, lower firing rates, faster system communication rate, and the unit of measurement used in our system.

A KF-FA decoder was trained on Day 1 to transform neural firing rates to two-dimensional endpoint velocity commands using a two-step calibration method. During the observation block (n = 80 trials), the subject was instructed to observe an automated version of a 2D BCI cursor center-out task and attempt to perform it while the computer controlled the kinematics, after which, the KF-FA decoder was fit. Using this decoder, the subject was then instructed to perform the same task he had previously observed. The decoded neural activity controlled the cursor while the computer attenuated command signals orthogonal to the target direction. After this data was collected (n = 80 trials), a final KF-FA decoder was fit and was used to provide full BCI cursor control to the participant.

4.2.3 Task Design

We used neural signal stabilization to collect a dataset consisting of 18 sessions over 104 days while the participant performed a 2D BCI cursor center-out task. A circular cursor was displayed at the center of the workspace at the beginning of each trial. The participant was then pseudo-randomly presented with one of eight possible targets equally distributed around the center cursor for a period of 1.5 seconds, during which the participant could not move the cursor. After

the presentation period, the participant was given BCI control of the cursor and had 5 seconds to acquire the target.

At the beginning of the first test session, we calibrated the initial decoder. We then ran the 2D cursor center-out task and updated the stabilizer parameters every 16 trials using a 128-trial sliding buffer. This buffer was reset at the beginning of each test session. The participant rested between each block while the stabilizer parameters were updated. During each subsequent session, we began data collection using the fifth stabilization update's parameters, calculated in the previous session, and continued to update the parameters while the subject performed the task (Figure 2D). At the end of each test session, we collected an additional block of 16 trials using the fixed decoder that was calibrated at the beginning of the first test session.

On eleven of the eighteen test sessions, a new KF-FA decoder was trained using a calibration block collected after the participant completed the stabilization evaluation blocks. One block of 16 center-out trials was used to evaluate performance of the same-day trained decoder. Parameters from this decoder were not used to update the stabilizer at any point in time and were completely independent from that part of the experiment.

4.2.4 Evaluation of Stabilizer Performance

We assessed the performance for the pre-stabilization, stabilization, and fixed decoder evaluation blocks for all 18 test sessions. For each of these blocks, we calculated the percent success rate by dividing the number of targets acquired by the total number of targets presented. Target acquisition rate was obtained by dividing the total number of targets presented in each block by the total amount of time the cursor was under BCI control across those trials. We calculated the angular error between the ideal and actual cursor trajectories at each time point and averaged these

values for each trial. The outcome metrics for all the blocks were tested for normality using the Shapiro-Wilk test(Ghasemi & Zahediasl, 2012). Outcome metrics that were found to come from a normal distribution were compared across all block conditions using a one-way ANOVA. If significant differences were found between block conditions, at a 95% confidence level, we performed a Tukey HSD post-hoc test to determine which conditions differed from one another. For outcome metrics that did not come from a normal distribution, we instead ran a Kruskal Wallis test followed by a Dunn-Sidak correction if significant differences were found.

To determine the relationship between time and an evaluation block's outcome metric, we calculated the correlation coefficients and significance of correlation. We calculated the Spearman correlation coefficient for blocks in which the data was not normally distributed, or the assumption of linearity was violated. Otherwise, we calculated the Pearson correlation coefficient. To determine whether the correlation between time and a given outcome metric was significantly different from 0, we compared the p-value to a significance level of 0.05.

4.2.5 Calculating Alignment of Neural Manifolds Over Time

To ascertain the efficacy of alignment over multiple sessions, we visualized the latent signals for both stabilized and unstabilized neural activity on the decoder control plane, the plane read out by the fixed decoder. The stabilized latent neural activity was obtained using each update's stabilization parameters and recorded neural data from each session. For the unstabilized latent neural activity, we related the neural activity recorded when control was provided by each stabilization update to the parameters obtained during initial calibration. This provided us with the latent state that we would have obtained if we had never updated the stabilizer parameters. For each session, we determined how the stabilized and unstabilized latent neural activity drove cursor

velocity at each time point and averaged these values for each trial. The relationship between decoded velocity $\hat{x}_t$ and latent state $\hat{z}_t$, using a Kalman filter decoder, was defined as:

$$\hat{x}_t = K(\hat{z}_t - d) + (I - KC)A\hat{x}_{t-1} \quad \textbf{Equation 4.4}$$

where K is the steady-state Kalman gain(Degenhart et al., 2020). The remaining parameters d, C, and A were fit during initial calibration. The first term, $K(\hat{z}_t - d)$, determines how the latent state at time t drove cursor velocity at time t and the second term, $(I - KC)A\hat{x}_{t-1}$, is a smoothing term. We used the first term to project each set of latent variables to the decoder control plane using Equation 4.3.

To determine whether stabilization updates were optimizing the spread of latent neural activity in the decoder control plane, we found the centroid of stabilized and unstabilized latent neural activity for each block (e.g., pre-stabilization, update 1, update 2, etc). We measured the spatial distribution of latent neural activity by calculating the average Euclidean distance from each block's trial-averaged latent point to its centroid. We ran a one-way ANOVA followed by a Tukey HSD post-hoc test to determine whether latent target-related neural activity changed significantly across stabilization updates. To determine whether control biases were reduced following a stabilization update, we calculated the Euclidean distance between each block's centroid and the origin of the decoder control plane. Since this data was not normally distributed, we ran a Kruskal Wallis test followed by a Dunn-Sidak correction to test for significant differences.

4.3 Results

4.3.1 Characterizing Neural Instabilities

We first characterized the types of neural instabilities that were observed throughout the length of the experiment. In Figure 2B, we show tuning curves for example recordings from four different electrodes at multiple time points during the experiment. We found examples of baseline shift, in which the baseline firing rate of a unit changed, as well as unit drop-out, in which activity from a neuron was no longer recorded, similar to the simulated instabilities in (Degenhart et al., 2020). As shown in Figure 2C, for each subplot, points further from the identity line (indicated by the dashed blue line) corresponded to greater change in firing rate for a particular electrode from Day 1 to the reference time point. The mean firing rate recorded on a given electrode was variable from session to session when compared to Day 1, highlighting the fact that single channel recordings are unstable. We also found that across each time point comparison, there was a large amount of variability for units with higher firing rates, as demonstrated by the linear regression model residuals that were both skewed and exhibited heteroscedasticity. Across all channels, the absolute difference in firing rate was generally small, within 5 Hz for approximately 60% of the channels at Week 1, Month 1, and Month 3 (Figure 3 Changes in firing rate over three months), but was statistically significant ($p<0.0001$, Wilcoxon signed rank test). These types of neural recording instabilities lead to BCI performance loss over time.

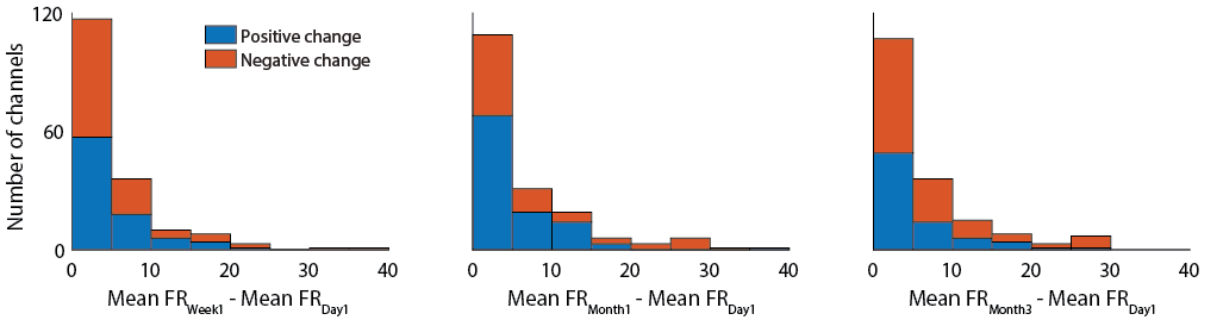

Figure 3 Changes in firing rate over three months. We calculated the absolute difference between mean firing rate on Day 1 from that of Week 1, Month 1, and Month 3 (as defined in Figure 1C). Channels in which firing rate increased from Day 1 to the subsequent time point are plotted in blue (positive change), while those that showed a decrease in firing rate are plotted in orange (negative change). All three distributions were non-normal ($p<0.05$, Shapiro-Wilk test) and had medians that were significantly different from 0 ($p<0.0001$, Wilcoxon signed rank test). The median change in average firing rate also increased across these three reference time points (change from Day 1 to Week 1: median=2.09 Hz, IQR=0.51-6.51 Hz, change from Day 1 to Month 1: median=3.23 Hz, IQR=0.84-8.09 Hz, change from Day 1 to Month 3: median=3.47 Hz, IQR=0.50-7.81 Hz).

4.3.2 Stabilization Maintains BCI Performance Across Days

Over the course of 104 days, we tested the ability of neural signal stabilization to provide consistent BCI control by periodically updating the parameters of the stabilizer while those of the decoder remained constant (Figure 2D). We evaluated BCI cursor control after each stabilizer update. Most test sessions included four updates to the stabilizer parameters, totaling 64 trials. On five of the eighteen test sessions, BCI performance using the fourth stabilizer update was suboptimal, with success rates below 94%, and additional stabilization updates were performed for four of those five sessions (sessions with success rates <85%). The last update for each session was labeled the final update. After the final stabilization update, we retested the fixed decoder,

which was calibrated at the beginning of the first session, and we compared its performance to that of the neural signal stabilization method.

As seen in Figure 4, at the start of each session before updating the stabilizer, the cursor trajectories were irregular and unsteady, particularly on days 28 and 83, and not all targets were reached successfully, which is atypical for a well-practiced task. As the stabilizer was updated, trajectories became smoother, directional biases were reduced, and all targets were successfully acquired. When the fixed decoder was tested at the end of each session, we saw poor BCI cursor performance. Qualitatively, these sample days illustrate general trends that were observed throughout the experiment. Taken together, the results from the pre-stabilization evaluation block each day as well as the fixed decoder comparisons, highlight the decline in BCI performance over time and the need for neural signal stabilization.

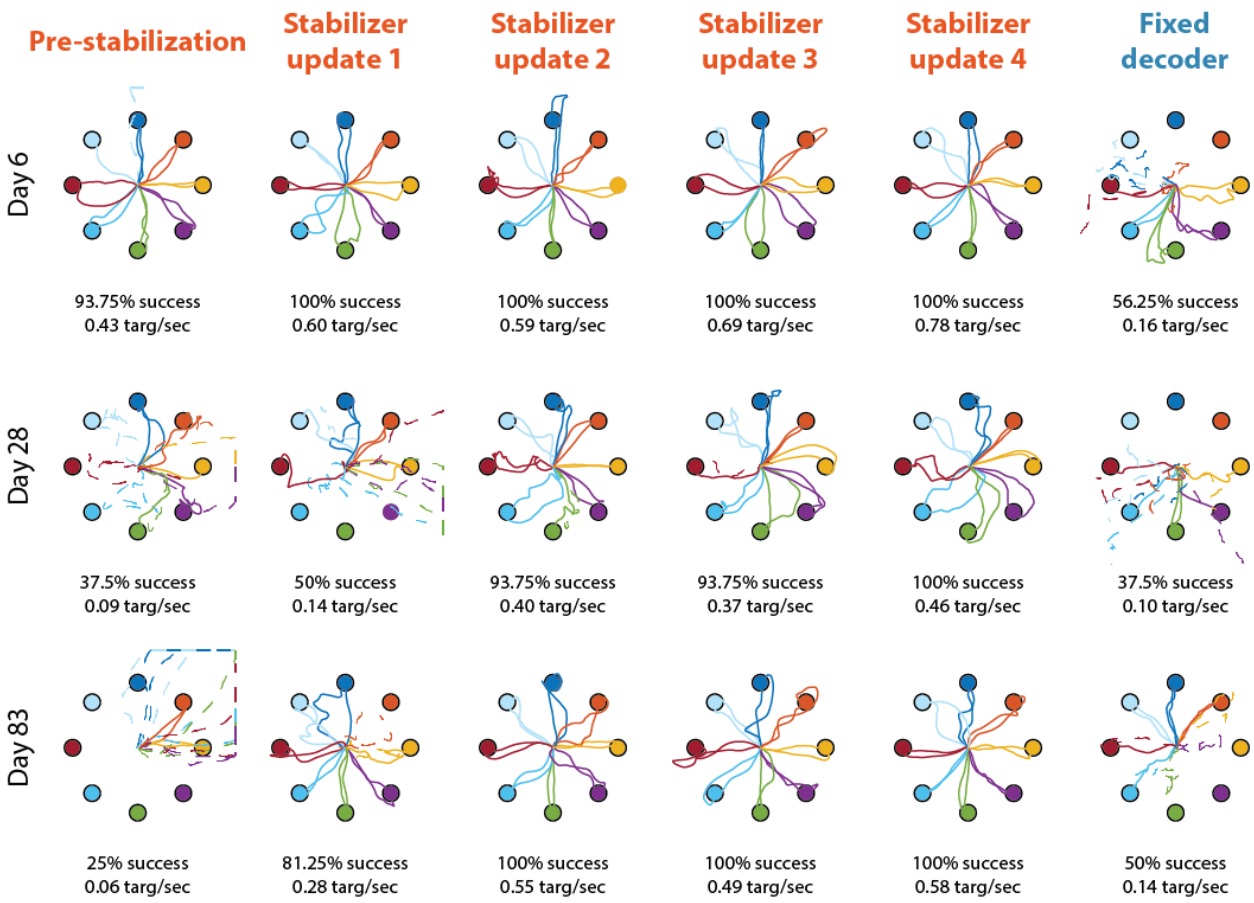

Figure 4 Online cursor trajectories exhibit increased accuracy with stabilization updates. Here we show trajectories from days 6 (session 3), 28 (session 9), and 83 (session 17) of the neural signal stabilization experiments. The 2D BCI cursor task required the participant to reach to 8 targets arranged in a circular pattern. The targets lie on the dashed grey circle in each subfigure. Cursor trajectories for each trial are colored by target location. Dotted lines represent failed trials. Success rate (defined as the percentage of acquired targets divided by the total number of targets presented) and target acquisition rate (defined as the number of targets successfully obtained divided by the total amount of time the cursor was under BCI control) are shown below each block's cursor trajectories.

We quantified changes in BCI performance after each neural signal stabilization update to determine the time course of improvements within each session relative to using the fixed decoder trained on Day 1. To find the optimal number of stabilization updates needed to significantly outperform the fixed decoder, we assessed three different outcome metrics: success rate, target

acquisition rate, and angular error. As seen in Figure 5A, compared to the fixed decoder, success rates were significantly higher after four stabilization updates. The fourth stabilizer update produced a perfect success rate for thirteen out of the eighteen sessions (median=100%, IQR=93.75%-100%), while performance with the fixed decoder had much lower success rates (median=65.6%, IQR=50%-81.25%). For the sessions where four stabilization updates did not provide a perfect success rate, between two and five additional updates were needed to achieve the desired performance, as seen in the "Final update" block. As few as two stabilization updates were required for significantly higher target acquisition rates and lower angular error values compared to the fixed decoder (Figure 5B and 5C). Across all eighteen test sessions, we found that performance began to exceed that of the fixed decoder by the second update with continued improvements through the final stabilization update. This performance improvement can be visualized in the cursor trajectories in Figure 4. As shown in Figure 5, compared to using the most recent decoder from the previous session ("Pre-stabilization" block), performance began to improve after one stabilization update, requiring only 16 trials. The participant was typically able to achieve a 100% success rate with four updates, requiring 64 total trials. These results show that neural signal stabilization can quickly recover stable BCI control without the need to pause for a technician to run a recalibration procedure.

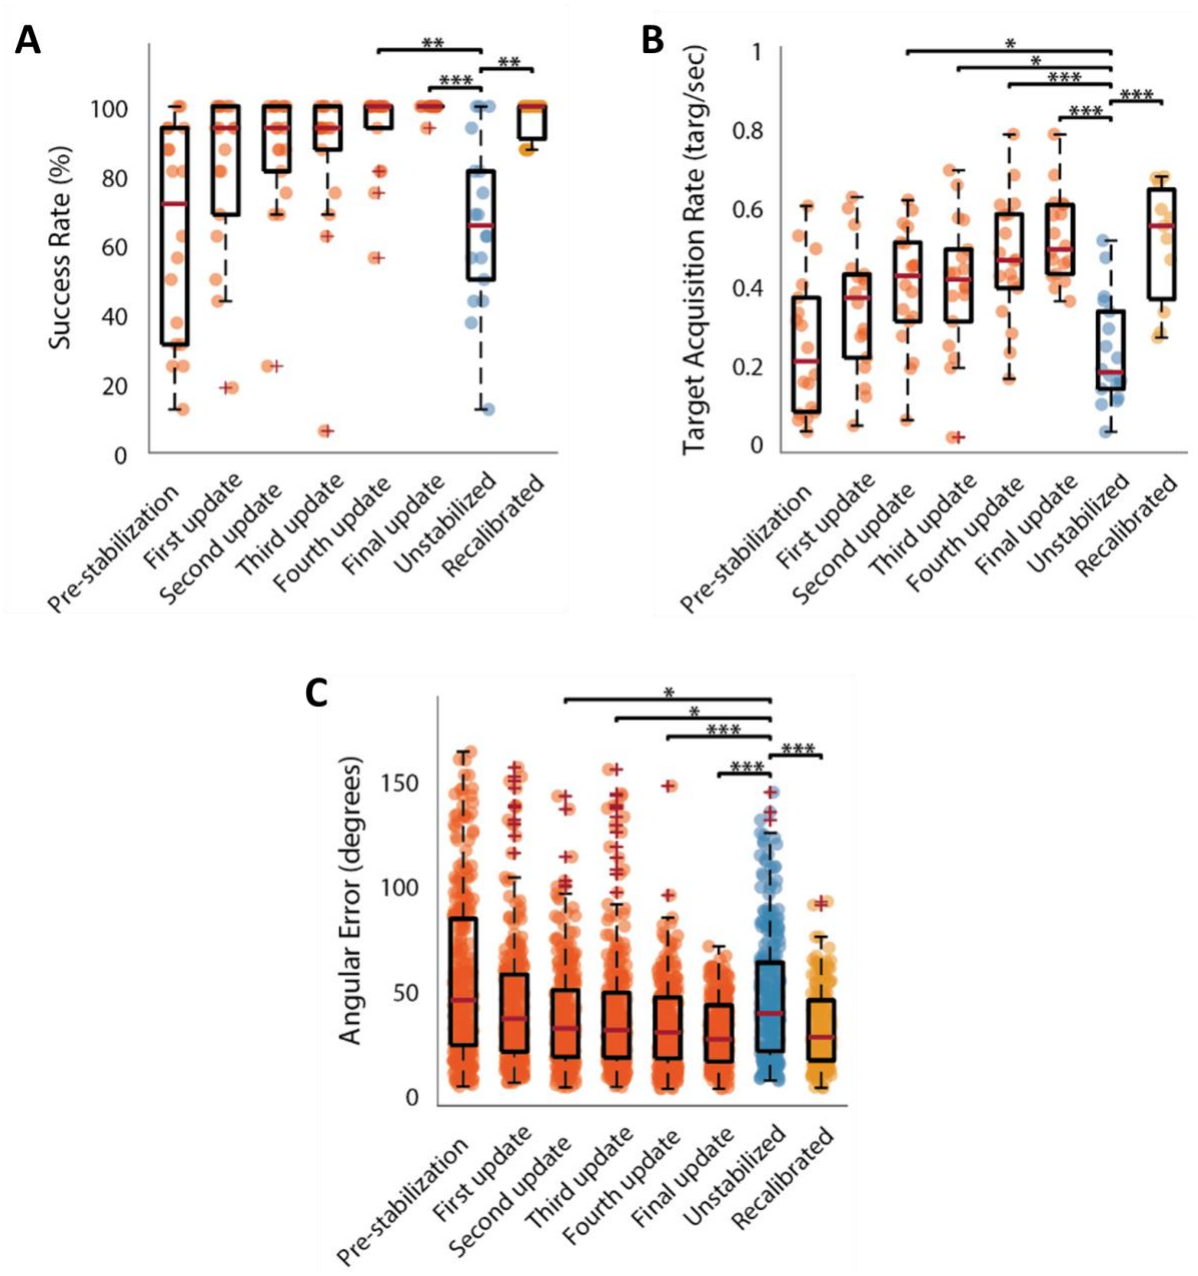

Figure 5 Stabilization updates outperform a fixed decoder and perform comparably to a recalibrated decoder. BCI control using stabilization (orange) begins to exceed the performance of the fixed decoder (blue) after two updates, with continued improvements through the final update. Four outcome metrics were calculated for the pre-stabilization, stabilization, and fixed decoder evaluation blocks throughout the experiment (n = 18 sessions). Outliers are indicated with a '+' sign alongside their corresponding data point. For most evaluation blocks, success rate and angular error were found to come from non-normal distributions ($p<0.05$, Shapiro-Wilk test) unlike target acquisition rate which did not show a significant departure from normality ($p>0.05$, Shapiro-Wilk test). Success rate

(A) was defined as the percentage of acquired targets divided by the total number of targets presented; **p<0.01 and ***p<0.001 (Kruskal Wallis test with post-hoc Dunn-Sidak correction). Target acquisition rate (B) was defined as the number of targets successfully obtained divided by the total amount of time the cursor was under BCI control; *p<0.05 and ***p<0.001 (one-way ANOVA with post-hoc Tukey HSD test). The angular error between the actual and ideal cursor trajectories (C) was calculated at each time point and averaged for each trial; *p<0.05, **p<0.01, and ***p<0.001 (Kruskal Wallis test with post-hoc Dunn-Sidak correction). Since angular error was not averaged across trials, each point in panel C corresponds to each trial's angular error for the 18 sessions (n= 288 trials).

4.3.3 Stabilization Maintains Cursor Trajectory Accuracy and Cursor Velocity

Once we determined that neural signal stabilization could overcome natural recording instabilities from one session to the next, we then tested the consistency of control stabilization provided over the eighteen test sessions. We analyzed trends in the outcome metrics over time to determine whether the quality of BCI control was maintained throughout the course of the experiment. Data from the final update was analyzed from each test session. To test whether neural recording instabilities negatively impacted decoder performance when stabilization was not provided, we also quantified BCI performance using a fixed decoder over time. We found that cursor trajectories obtained after the final update remained accurate for all18 sessions, with high success rates, high target acquisition rates, and consistently low mean angular error, as seen in Figure 6. In comparison, when using a fixed decoder, success rates ($p<0.05$, Pearson correlation) and target acquisition rates ($p<0.01$, Pearson correlation) decreased significantly over time as shown in Figure 6A and 6B. Mean angular error also tended to increase, though this correlation was not significant (Figure 6C). Our results show that neural signal stabilization provides

consistent BCI control that was significantly better than using a fixed decoder across multiple months in the presence of natural neural recording instabilities.

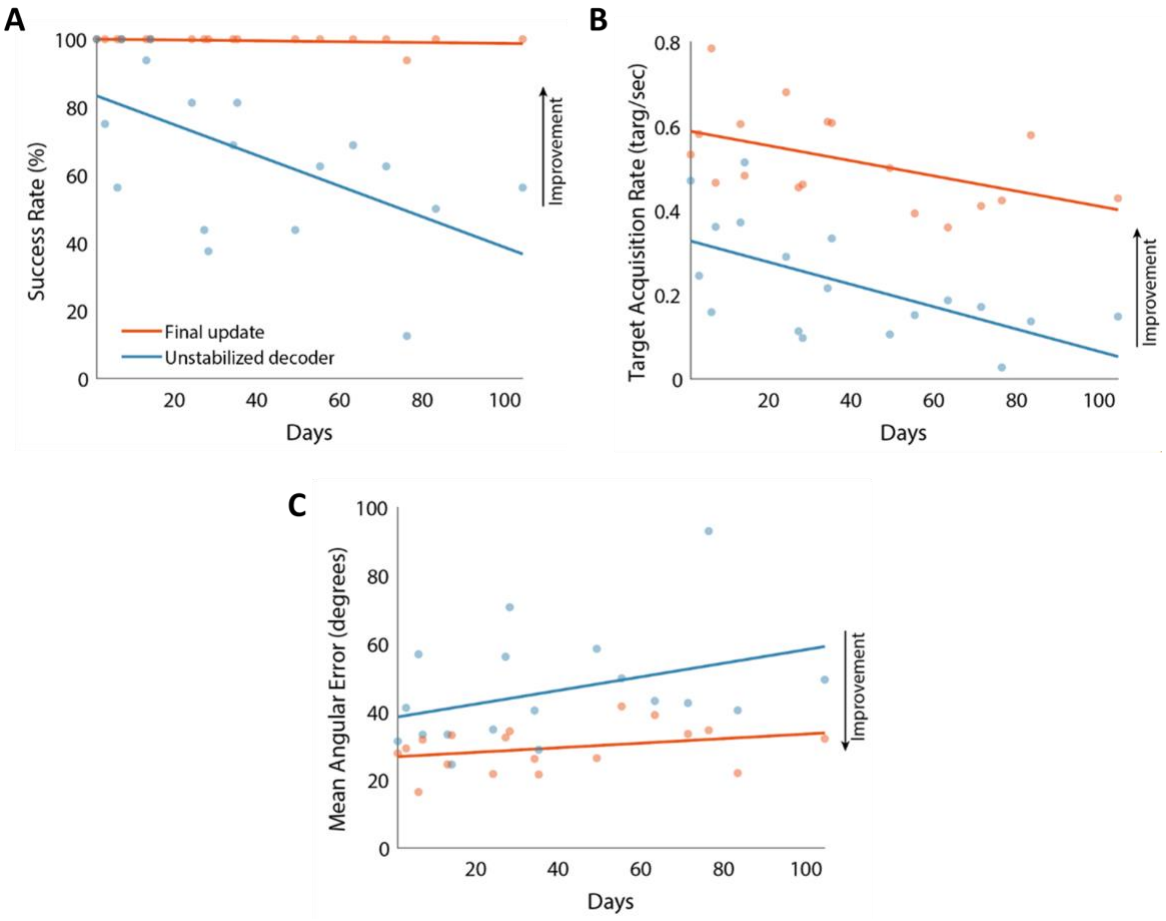

Figure 6 Stabilization maintains successful BCI control over time with minimal loss of movement quality. Success rate (A), target acquisition rate (B), and mean angular error (C) were calculated for the final update's and fixed decoder's evaluation blocks on each of the 18 test sessions. Outcome metrics were calculated as described in Fig. 5. Least-squares lines are shown for the final update (orange) and fixed decoder (blue) conditions. Success rate, target acquisition rate, and mean angular error were not found to significantly change over time when using the final update (orange, $p>0.05$, Pearson and Spearman correlations). Using the fixed decoder (blue), success rate ($p<0.05$, Pearson correlation) and target acquisition rate ($p<0.01$, Pearson correlation) significantly decreased over time. Mean angular error did not exhibit significant correlations over time for the fixed decoder.

4.3.4 Stabilization Preserves Target-Related Neural Activity Patterns

Thus far, we have determined that neural signal stabilization can provide stable, consistent BCI control across 18 sessions spanning three months. The consistency of the BCI control was obtained by aligning estimates of the neural subspace at given time points, thereby reducing the effect of neural recording instabilities and providing a stable input to the fixed decoder (Figure 2A). Here, we examine the extent to which neural signal stabilization appropriately and accurately aligned the low-dimensional neural subspaces associated with movement direction intention. To test this, we analyzed the latent neural activity within the aligned neural subspace, as defined by the stabilization parameters from each update. We also determined what this latent neural activity would have been if the stabilizer had never been updated, which we will refer to as the unstabilized latent neural activity. We visualized and compared trial-averaged stabilized and unstabilized latent neural activity from all updates on the 2-dimensional decoder control plane, i.e., the plane defined by the Kalman filter decoder, representing how the latent neural activity drove cursor velocity. As seen in Figure 7A, the stabilized latent neural activity was clearly separated by target direction, clustering within the same area of the decoder control plane across sessions. On the other hand, the separation based on target direction for the unstabilized latent neural activity did not increase over the course of a session. To evaluate the extent of target separation in the decoder control plane, we classified target direction using a Naïve Bayes classifier using data pooled from the first four stabilization updates across the 18 sessions, and we evaluated classifier performance using 5-fold cross-validation. We found that the stabilized latent neural activity was correctly classified 83.8% of the time (chance level = 12.5%) compared to the unstabilized latent neural activity which was only correctly classified 61.8% of the time. To quantify trends within each session, we measured the dispersion of stabilized and unstabilized latent neural activity from each update

evaluation block. We calculated the distance from each trial-averaged latent point on the decoder control plane to the center of all the trial-averaged latent points for a given update. We found that there was significantly less target separation pre-stabilization (i.e. lighter colored data points remain closer to one another as seen in Figure 7B) compared to later updates, in which the dispersion of stabilized latent neural activity increased within the neural subspace ($p<0.01$, one-way ANOVA, Figure 8A).

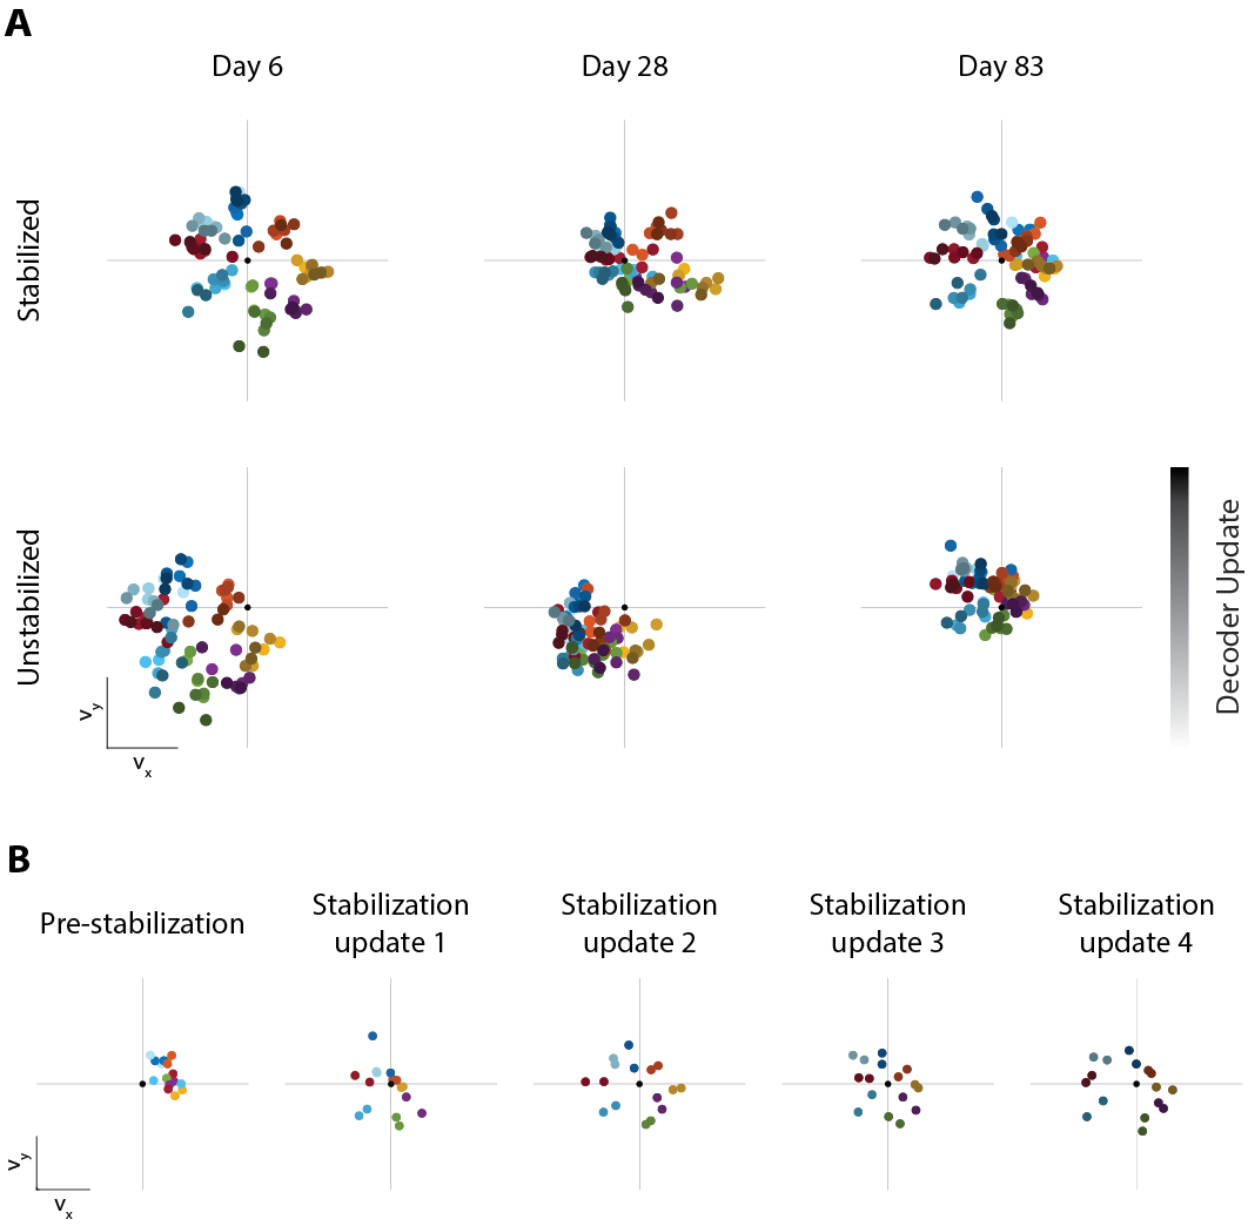

Figure 7 Stabilization of latent neural activity preserves target-related structures for movement decoding. A. The stabilized latent neural activity used online during stabilization (top) and the latent neural activity that would have been used if the base stabilizer parameters had not been updated (bottom), averaged for each trial, are plotted on the decoder control plane, the 2D projection read out by the fixed decoder, for three example days. Target direction is represented as one of eight colors, with lighter colors corresponding to pre-stabilization and earlier updates and darker colors corresponding to updates from later in the session. The coordinate axes and origin of the decoder control plane are represented by the intersecting grey lines and the black dot at their center, respectively. Positive values along $\boldsymbol{v_x}$ correspond to movement to the right and positive values along $\boldsymbol{v_y}$ correspond to upwards

movement. B. Example day (day 83 from panel A) showing the effect of stabilization on the latent neural activity after each update. Prior to running stabilization, we see very little target separation and a large deviation of the centroid from the origin in the pre-stabilization block. After just one stabilization update, target clusters are easier to identify, and the centroid of the stabilized latent neural activity is closer to the origin. This trend continues with subsequent stabilization updates.

When a control bias was present (i.e. an offset in the movement of the cursor, as seen in Figure 4 on day 28 where cursor trajectories skew right indicating a right bias), we found that stabilizer updates not only increased separation of the latent neural activity across target directions but also centered the spread of latent activity within the neural subspace. By shifting the latent neural activity to the center of the subspace, the neural signal stabilization method attempted to reduce biases in the provided BCI control. We found that the distance from the center of the neural subspace to the centroid of the stabilized neural activity decreased after just one stabilization update; however, this change was not statistically significant (Figure 8B). The centroids of the stabilized compared to the unstabilized latent neural activity were significantly closer to the center of the neural subspace across all updates ($p<0.01$, Kruskal-Wallis), signifying a significant reduction in the amount of control bias present when neural signal stabilization was used. These combined results demonstrate that, when we account for recording instabilities, the task-related activity within the neural subspace is stable over time. Left unaccounted for, the instability of activity within the neural subspace leads to the incorrect interpretation of movement intent by the fixed BCI decoder, leading to poor BCI control. The clear target separation and shifting of the stabilized latent neural activity shows that neural signal stabilization accurately aligns the lower-dimensional representation of recorded neural activity to produce a stable input to the fixed decoder, yielding consistent BCI cursor velocity predictions.

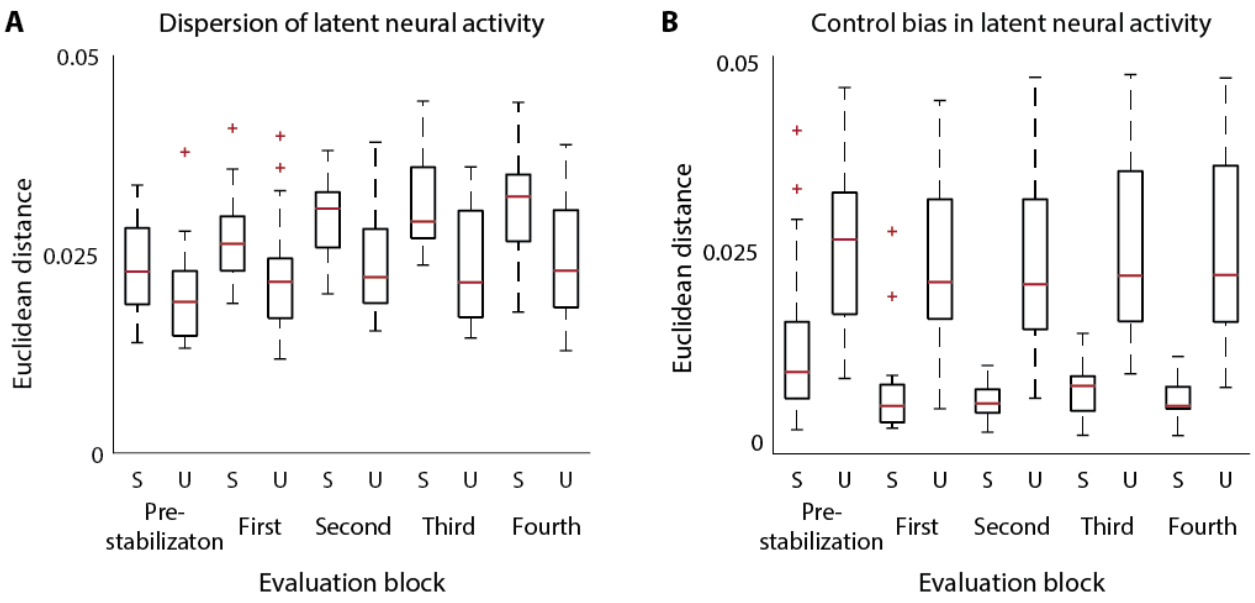

Figure 8 Stabilization updates increases target reach discriminability and minimizes decoder bias. A. Measure of dispersion of stabilized (S) and unstabilized (U) latent neural activity along the decoder control plane. We calculated the distance from each latent point (where each point corresponds to a single trial) to the mean latent activity for a given evaluation block. Increased distance represents greater target separation within the decoder control plane. Stabilized latent neural activity from the third and fourth stabilization blocks had significantly greater target separation than the pre-stabilization block ($p<0.05$, one-way ANOVA), demonstrating increased target separation as updates progressed. Stabilized latent neural activity from the second stabilization block had significantly greater target separation than the unstabilized latent neural activity from the pre-stabilization and first stabilization blocks ($p<0.05$, one-way ANOVA). Finally, stabilized latent neural activity from the third and fourth stabilization evaluation blocks had significantly greater target separation than unstabilized latent neural activity from all blocks ($p<0.05$, one-way ANOVA), signifying that neural signal stabilization provides better representation of movement-based structure within the decoder control plane compared to if we had never stabilized the latent neural activity. B. Measure of reduction in control bias of stabilized (S) and unstabilized (U) latent neural activity along the decoder control plane. We calculated the distance from the origin of the decoder control plane to the mean latent activity for a given evaluation block. Decreased distance corresponds to a decrease in control bias. We found that the centroids of the stabilized latent neural activity were significantly closer to the origin than unstabilized latent neural activity for all evaluation blocks ($p<0.01$, Kruskal Wallis).

4.3.5 Neural Signal Stabilization Provides Comparable Performance to that of a Recalibrated Decoder

Thus far, we have determined that neural signal stabilization outperformed a fixed decoder, reducing the need to perform recalibration. However, the goal of this method was to not only minimize technician intervention but also to provide control comparable to that of the current standard in BCI research, which is to calibrate a new decoder for each session. To test this, we next evaluated how stabilized BCI control compared to performance with a same-day trained decoder. After running the neural signal stabilization paradigm, we calibrated a new decoder at the end of 11 test sessions and collected an additional block of trials (n = 16) for evaluation. We expected that both the fourth stabilization update and the newly trained same-day decoder would yield high levels of BCI control. Across all BCI performance metrics (success rate, target acquisition rate, and angular error), control provided by the final stabilization update was similar to that of the same-day recalibrated decoder ($p>0.05$, Wilcoxon rank sum, Figure 5).

4.3.6 Neural Signal Stabilization Improves Subjective Experience

We further compared neural signal stabilization and daily recalibration by assessing the participant's feedback at the end of each session. The subject was asked to provide verbal feedback regarding the BCI control provided by each stabilization update. On several test sessions, he also provided feedback regarding control obtained from the fixed decoder and the recalibrated decoder. We compared the participant's qualitative assessments along with the quantitative performance

metrics to determine whether neural signal stabilization improved the participant's subjective experience completing the cursor task. For the test sessions in which a new same-day decoder was calibrated and the participant provided feedback on that new decoder (8 sessions), the participant generally reported that the control provided by a stabilization update was similar to that of the new decoder, and he either had no preference or only a slight preference for one over the other (equivalent on 2 days, slight preference for the stabilized decoder on 3 days, slight preference for the new decoder on 3 days). In each of these cases, we found that the participant's qualitative assessments mapped well to the quantitative comparisons (i.e., higher success rate, higher target acquisition rate, and lower angular error were achieved for the decoder he preferred). Overall, neural signal stabilization provided a level of BCI control that was both qualitatively and quantitatively similar to that of a newly calibrated decoder.

4.4 Discussion

This study aimed to overcome neural recording instabilities in humans to provide reliable BCI control without the need for recalibration by implementing neural signal stabilization. Our method is predicated on the idea that neural activity patterns are constrained to a stable low-dimensional neural subspace (Cunningham & Yu, 2014; Degenhart et al., 2020; Golub et al., 2018; Oby et al., 2019; Sadtler et al., 2014). Although our estimates of this latent subspace are made noisy by neural recording instabilities, we can align these estimates over time to obtain stable latent representations of neural activity which are mapped to the user's intended movement. Our previous work with non-human primates used neural signal stabilization to overcome artificial perturbations applied to recorded neural activity over single and five-day periods. Here, we corroborate our

previous findings in NHPs, which showed that stabilization effectively accounts for changes in neural activity and expand upon this work by accounting for natural recording instabilities both within single sessions and over the span of 104 days in a human participant. We found that our method provided better BCI control and reduced performance degradation compared to a fixed decoder. By accurately aligning estimates of the neural subspace over time, our method produced a stable decoder input that preserved target-related structure and reduced the impact of control biases. Degenhart and Bishop et al. showed in an offline comparison that stabilization recovered BCI performance better than any supervised recalibration approaches (Degenhart et al., 2020). Our study took this a step further by comparing online performance between the stabilizer and a same-day trained decoder, the current standard in BCI research. Based on both objective metrics of BCI performance and participant feedback, neural signal stabilization was comparable to daily recalibration of the BCI decoder. Our study shows that neural signal stabilization provides reliable BCI control over months without the need for technician intervention, taking us a step closer to clinically implementing a more independently operated BCI system.

Many alternative approaches have been proposed to maintain the quality of decoding over time while minimizing interruptions in BCI control for recalibration. Adaptive decoding approaches address neural recording instabilities by updating decoder parameters as the user performs a BCI task (Bishop et al., 2014; Homer et al., 2014; Jarosiewicz et al., 2015; Li et al., 2011; Orsborn et al., 2014; Wu & Hatsopoulos, 2008; Zhang & Chase, 2013). These types of methods require some level of supervision or knowledge about the task to infer the user's movement intent at each time step, making them difficult to implement in more naturalistic settings. Deep learning approaches have also been used to either leverage large datasets to train fixed decoders that are robust to neural variability or to align latent neural dynamics over time

(Karpowicz et al., 2022; Sussillo et al., 2016). However, these approaches could require months of training data collection from the user and incur large computational costs, making them difficult to implement for clinical use. In the current study, we were able to maintain BCI performance by performing neural stabilizer updates on small numbers of trials without needing to infer the user's intention at each time step. Future work will need to test whether neural data would provide similar levels of performance for tasks with less evenly distributed kinematics than a 2D-center out task. Finally, other studies have explored using different signal modalities for BCI decoding to provide long-term control (Flint et al., 2013; Milekovic et al., 2018; Vansteensel et al., 2016). Although ECoG- and LFP-based decoders may provide more stable control across months, users obtain much lower performance from these systems. In a study performed in two individuals with neurological disorders, fixed LFP-based BCI decoders produced spelling rates of 3.07 and 6.88 correct characters per minute over 76 and 138 days, respectively, significantly below the state of the art performance levels and considerably lower than potential BCI users' desired speed of 20 letters per minute (Huggins et al., 2015; Milekovic et al., 2018; Pandarinath et al., 2017).

Our previous study demonstrated the performance advantages of using neural signal stabilization over supervised recalibration and alternative adaptive decoding approaches (Degenhart et al., 2020). Here, we expanded upon our previous study's results to determine the effectiveness of neural signal stabilization in minimizing participant burden compared to recalibration. We tested the number of stabilization updates required to overcome the impact of significant instabilities on performance after long periods of disuse. We found that performance began to improve after only one stabilization update (n=16 trials) with continued improvements following subsequent updates. Two stabilization updates (n=32 trials) yielded significantly better performance compared to the fixed decoder across most metrics and four stabilization updates

(n=64 trials) provided BCI control that was on par with that of a same-day trained decoder. Compared to recalibration, for which data collection alone takes approximately 8 minutes, the data needed for one stabilization update takes only about 1.5 minutes to collect. Since each stabilization evaluation block consisted of only 16 trials, there was not much time to allow for any neural adaptation on the part of the participant. If we had provided the participant with more time to explore and adapt to each update, it is possible that performance could have been optimized with even fewer stabilization updates. In the future, we could investigate optimizing the timing of the stabilization updates, in other words testing the number of trials or amount of elapsed time between each stabilization update. An additional advantage of the neural signal stabilization method is that the user would also have access to the BCI even if performance is non-optimal, as opposed to recalibration, which typically requires the user to pause control and perform a separate training task. By reducing not only the number of interruptions required for recalibration but also the amount of time needed to stabilize online control, our method could contribute to the development of a more independently controlled and minimally burdensome BCI.

In our study, neural signal stabilization was used to recover BCI control after long periods of disuse. Since the stabilizer was only updated in the lab during BCI control, when the participant was engaged in the cursor task, we were not able to stabilize population neural activity in between test sessions. We speculate whether neural signal stabilization could provide continuous stable BCI control if the stabilizer was left on to perpetually counteract neural recording instabilities. It remains an open question as to whether data collected during any other brain state could be used for control (Degenhart et al., 2020; Gallego et al., 2018; Kaufman et al., 2016). Results from our previous study in NHPs indicate that neural signal stabilization may be able to restore BCI control even during moments of low user engagement, since the neural manifold is thought to represent

underlying neural connectivity rather than strictly task-related information (Degenhart et al., 2020). Though we did not test this here, we could attempt to reduce the amount of task-dependent neural data needed to achieve optimal control. One option would be to update the stabilizer using "rest" data collected during setup at the beginning of each test session. This could provide better initial BCI control during the task, thereby requiring fewer updates to achieve stable BCI performance.

Our study tested the neural signal stabilization method using a 2D cursor control task, a task typically used for decoder calibration. This task design was chosen to provide a more direct comparison with other automatic recalibration methods, which have predominantly been tested for use in 2D cursor control tasks. Nevertheless, one of our study limitations includes the direct use of data collected during a training task for the stabilization updates. Similar to the concept of using "rest" data described above, we should consider the use of open-ended tasks to provide stabilizer updates. This could include BCI-controlled typing or web browsing, tasks that would more closely resemble the intended use case of BCIs. Additionally, the cursor control task was a lower dimensional task than some of the higher degree of freedom robot arm and hand control we have achieved in our lab previously. Very few studies have tested their methods on decoding arm and hand movements, and those that have, have only tested the efficacy of their methods offline using previously collected data (Gallego et al., 2020; Karpowicz et al., 2022). These methods also have not been tested for use in high degree of freedom tasks that include many more degrees of complexity than highly stereotyped reaches. We previously investigated the effectiveness of neural signal stabilization as the complexity of a task increased (e.g., increasing number of control dimensions) in simulations, and found that our method would continue to work as long as sufficient data was used for neural signal stabilization (Degenhart et al., 2020). Future work could include

increasing task complexity to 2D cursor and click control as well 3 degree of freedom translational arm control with one degree of hand grasp control to further test the capabilities of our method and its stabilization accuracy over time for increasingly complex tasks.

This work has demonstrated the use of a BCI decoder over a period of three and a half months in providing a user with reliable control of a cursor despite the presence of naturally occurring neural recording instabilities. By taking advantage of the stability of latent neural subspaces, we can move closer to the goal of independent operation of BCI technology while minimizing setup time and the need for intervention, as desired by the population of future potential users (Blabe et al., 2015; Collinger, Boninger, et al., 2013; Huggins et al., 2015, 2011).

5.0 Effect of Object Presence on M1 Neural Activity

This chapter is based on a first author manuscript which is currently in preparation.

Our lab previously found that during a BCI-controlled reach-to-grasp task, most recorded neurons showed an increase in firing rate prior to grasp initiation relative to when the same grasping movement was performed in the absence of an object. However, the cause for this global increase in firing rate, which negatively impacts our ability to predict the user's intended movement, was unclear. Here, we explore the impact of object presence and grasp intention (hand open vs. closed) on population neural activity, from planning to movement completion. Our results demonstrate that object presence significantly modulates neural population activity throughout planning, reaching, and grasping. This effect varied concurrently with that of grasp intention; however, object presence impacted neural activity patterns that were distinct to those corresponding to grasp intention. Subsequently, when we changed the task difficulty and precision requirements, we found a small, yet consistent, change in neural activity throughout planning and movement associated with this increased user engagement.

5.1 Introduction

Reaching and grasping movements are vital for achieving activities of daily living such as feeding oneself, brushing one's hair, or communicating with others by picking up the phone or writing a letter. Advancements of arm and hand function restoration have led to the use of BCIs that have enabled people with tetraplegia to control reaching and grasping movements of a robotic

arm with the speed and skill approaching that of an able-bodied person (Collinger, Wodlinger, et al., 2013). The highest level of BCI control achieved to date has been 10-dimensional control of a robot arm and hand, which includes 3 degrees of translation of the arm's endpoint, 3 for orientation of the wrist, and 4 for hand shaping (Wodlinger et al., 2015). Although the participant was able to successfully perform posture matching tasks, she experienced greater difficulty performing functional tasks that required interaction with an object.

Additional work from our lab has shown that the presence of an object causes a global increase in neural firing rates in cases where kinematics are consistent (Downey et al., 2017). This shift in baseline firing rate negatively impacts our ability to predict the user's intended movement, impairing BCI performance. To correct for this, our lab has implemented global normalization of firing rates to remove this baseline shift when approaching an object (Downey et al., 2017). However, the cause for this global increase in firing rate has remained unclear.

Effective grasping requires processing an object's properties, such as size and shape, and transforming these properties to generate the corresponding hand configurations. Numerous studies have focused on identifying and characterizing a cortical network capable of transforming visual signals to motor commands for grasping (Begliomini, Wall, Smith, & Castiello, 2007; Castiello, 2005; Grol et al., 2007; Jeannerod, Arbib, Rizzolatti, & Sakata, 1995; Rizzolatti & Luppino, 2001). This cortical grasping circuit is thought to involve the anterior intraparietal area (AIP), ventral premotor cortex (PMv) and motor cortex. AIP has been found to play an important role in processing visual information to better inform hand shaping for differentially shaped objects (Castiello, 2005; Jeannerod et al., 1995; Murata et al., 2000), whereas PMv has been demonstrated to be more preferentially related to grasping movements (Rizzolatti et al., 1988; Schaffelhofer & Scherberger, 2016). However, the majority of these studies have been conducted in monkeys and

it is still unclear how well the AIP-PMv-M1 circuit directly relates to a neural circuit for grasping in humans (Castiello, 2005).

These studies prompt the question of whether it is simply the presence of the object, the interaction with it, or both that are driving an increase in neural firing rates. They also suggest that movement is not necessary for an object to impact M1 modulation, prompting us to ask how the movement signal interacts with the object signal. To further uncover the basis for this effect, we explored the impact of object presence and grasp intention on population neural activity from planning to movement completion.

5.2 Methods

5.2.1 Object Interaction Task

The object interaction task required the participants to reach to and grasp a target, across a variety of conditions (object present vs. absent, hand open vs. closed). The task was broken up into 4 trial epochs: (1) presentation, where an audio cue instructed the participants to reach to one of two target locations; (2) reach, where the participants moved to the target location; (3) grasp, where a second audio cue instructed the participants to either grasp or keep the hand open while an object was either present or absent at the target location; and (4) hold, where the participants were required to maintain the arm and hand position at the target location for 2 seconds. The participants had a maximum of 20 seconds to reach the target location and an additional 10 seconds to achieve the desired grasp. If the participants were not able to reach the target location or were not able to

successfully maintain the target reach and grasp positions for the entirety of the hold time within the given time constraints, the trial was not considered successful.

Trials were blocked by object and grasp conditions: grasping an object, grasping without an object present, keeping the hand open next to an object, and keeping the hand open without an object present. Two sets of randomized blocks were collected each session, except on session 1 for participant P3 where only one set of blocks was collected. Between 20 and 40 trials of each condition were performed by the participants. Prior to the start of each block of trials, the participants were instructed whether they would have to grasp or keep the hand open and whether an object would be present at the target location. The audio cue at the start of the presentation phase only provided information about the target location. For participant P3, only one target location was used ("left" target, shown in Figure 9). In that case, the audio cue was only used to mark the beginning of the trial. Data from subject P2 was collected over 5 test sessions 1603-1624 days (~4 years) after implant, while data from subject P3 was collected over 4 sessions 281-303 days (~9 months) after implant.

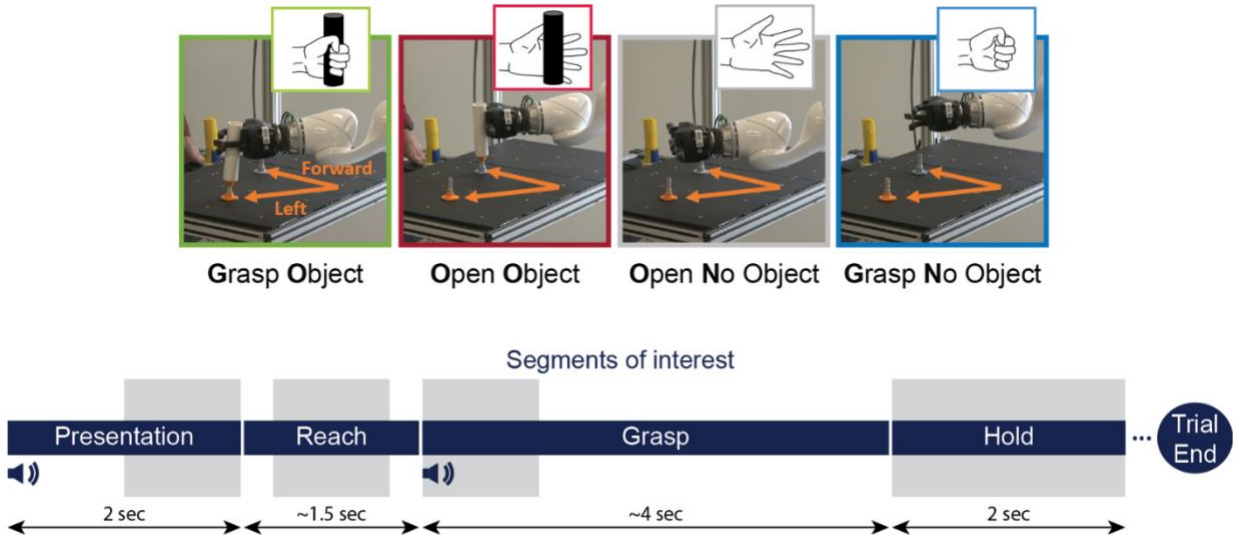

Figure 9 Object interaction task and timeline of trial. Top: The robot arm and hand were used to move to one of two target locations: forward or left, as shown in the right-most sub-figure. For trials where the object was present in the workspace, objects were placed on a spring-based system (shown in orange and gray) by one of the experimenters during the presentation phase. The workspace limits were constrained to the black table.

All data was collected using the KUKA LBR iiwa robot (KUKA Robotics Corporation) and the Right Hand Robotics ReFlex Takktile Gripper (RightHand Robotics, Inc.), which was attached to the KUKA's media flange. Prior to this experiment, the KUKA's joint angles were configured to resemble a right arm, with the "shoulder" mounted on a custom-built stand (not pictured) placed on the right-hand side of the participants. The "elbow" joint was flexed so that the palm of the robotic hand faced the left-hand side of the workspace. At the start of each block, the robot arm was initialized to a "home" position, near the right-hand corner of the workspace closest to the participant. At the end of each trial, the robot returned to this "home" position.

Given the observed movement jitter of the robot arm during the grasp and hold phases, we found that rigid systems that hold the object strictly in place would not be appropriate for this task since this type of system would be more likely to drive the robot into "handguiding" mode.

"Handguiding" mode refers to one of the KUKA robot's safety modes that is turned on when it detects large or sudden forces applied on the arm's sensors. To minimize the occurrence of "handguiding" mode and the frequency of disruptions in BCI control, we developed a flexible spring-based system that would allow for slight displacement of the object. This system consisted of two 3D-printed objects: a spring base and an object base. The workspace table, shown in Figure 9 (top), was drilled to bear a grid of holes with built-in threaded fasteners. The spring base could thus be temporarily attached at a given location on the workspace table. Objects were made using PVC pipes 1 ¼ inch in diameter and were cut down to be approximately 6 inches long. A 3D-printed base was attached to the bottom of the object. This base allowed the object to be easily placed and removed from the spring base.

5.2.2 OLE Decoder Calibration

To map neural activity to movement velocity of the robotic hand, the following model was used:

$$f = b_0 + b_x v_x + b_y v_y + b_z v_z + b_g v_g \qquad \textbf{Equation 5.1}$$

where f is the square-root transformed firing rate of a unit, v_x, v_y, v_z, and v_g are the velocities of each control dimension, and b_0, b_x, b_y, b_z, and b_g are regression coefficients. We can rewrite Equation 5.1 in matrix form as:

$$F = VB \qquad \textbf{Equation 5.2}$$

Here, V contains the velocity components v_x, v_y, v_z, and v_g and is a $t \times d$ array, where t is the number of time points and d is the number of control dimensions. B contains the regression coefficients and is a $d \times n$ array, where n is the number of units.

OLE decoders are implemented similarly to the population vector algorithm (described in section 2.3.1); however, they differentially weight the contribution of each neuron for more optimal decoding. These optimal linear decoding weights, W, directly map neural activity to velocity commands:

$$V = FW \qquad \textbf{Equation 5.3}$$

Here, we used indirect optimal linear estimation (Wang et al., 2007) with ridge regression (Marquardt, 1970) to obtain the W matrix from the coefficient matrix B:

$$B = \frac{(V'V + \lambda_1 I)}{V'F} \qquad \textbf{Equation 5.4}$$

$$W = (B\Sigma B' + \lambda_2 I)^{+} B\Sigma \qquad \textbf{Equation 5.5}$$

where λ is the optimization parameter for ridge regression, I is the identity matrix, X^{+}represents the pseudoinverse of matrix X, X' represents matrix transposition, and Σ is a $n \times n$ diagonal matrix containing the inverse variance of the residuals (Collinger, Wodlinger, et al., 2013). By solving for W in this way, we accounted for the non-uniformity of preferred directions and the differences in variance of each unit (Chase et al., 2009a). In other words, units with more movement-related information were weighted greater and thus contributed more to decoding velocity than units carrying less or more variable information.

Decoders were trained at the beginning of each test session using a 2-step calibration procedure (see details in section 3.2). We used a 4D grasp-and-carry training task (3 degrees of arm translation and 1 degree of hand grasp), where a virtual arm moved toward a transparent ("pass-through") object at one of nine possible target locations. The pass-through object was grasped and held at the target location for 2 seconds and was then carried to the next target location and released. After finishing the calibration procedure, the participants were provided BCI control to move the robotic arm and hand. At this point, experimenters provided any manual gain or bias changes to the decoder as needed. To stabilize kinematics of the arm and minimize variability across trials and conditions, we applied 50% "orthogonal assistance" (see section 3.2 for reference) along translation dimensions and 75% along grasp dimensions of participant P2's BCI control to minimize jitter in endpoint translation. For participant P3, we applied either 50% or 75% assistance along translation dimensions (depending on the accuracy of control on a given session) and 100% assistance along the grasp dimension.

5.2.3 Quantifying Kinematic Behavior and BCI Performance

To quantify behavioral performance, we first excluded unsuccessful and outlier trials for each session. Outlier trials were defined as trials with reach completion times greater than the 85th percentile or lower than the 15th percentile of all reach completion times for a given condition. After removing these trials, we tested whether reach completion times were significantly different across conditions for each session using either a one-way ANOVA or Kruskal-Wallis test, depending on whether the distribution of reach completion times for each condition passed the Shapiro-Wilks normality test. This metric was used as a proxy for comparison of the consistency

of reach trajectories across conditions. Task performance was quantified by success rate, defined as the percentage of completed trials among the total number of trials attempted.

5.2.4 Quantifying Condition-Dependent Modulation of Population Activity

To determine when object presence modulated neural population activity, we identified a latent neural subspace using factor analysis (see section 3.3.3 for additional details) across all trial epochs. Prior to dimensionality reduction, unsuccessful and outlier trials were excluded for each session, as explained in the previous section. On average, we found that the optimal dimensionality needed to explain the motor cortical data was ~15D for participant P2 and ~30D for participant P3. Only data recorded from M1 arrays are discussed here (see section 6.3.1 for analysis of S1 arrays). Additional details regarding data pre-processing for dimensionality reduction is discussed in section 3.3.1.

We obtained separate FA projections of the neural data for each session to avoid confounding dimensions capturing day-to-day variability of neural recordings. We then used the FA latent neural activity as inputs for linear discriminant analysis to obtain the latent dimensions that maximally discriminate based on task condition (e.g., object present vs. absent, grasping vs. keeping hand open, etc.). We decided to first transform the dataset via FA and later calculate the LDA projections to avoid overfitting. By reducing the dimensionality of the neural data prior to using LDA, we attempted to remove some of the spiking variability and noise from our high-dimensional neural dataset. Additionally, to overcome the differences in completion times across trials and conditions, we only used neural data from portions of the trial. These segments of interest were obtained from all task phases: the last second of presentation prior to movement onset, one

second centered around the mid-point of the reach phase, the first second after the cue to grasp was given, and the entirety of the two second hold time (Figure 9 (bottom)).

We visualized the trial-averaged neural trajectories along the top three liner discriminant dimensions and found that the trajectories resided in separate areas of the subspace. To quantify the degree of separation between the condition-dependent, trial-averaged neural trajectories, we calculated the Euclidean distance between paired trajectories (e.g., comparing the trial-averaged neural trajectory for the object present and grasping condition with that of the object absent and grasping condition) at each time point across the entire subspace:

$$d = \sqrt{\sum_{i=1}^{n} (q_i - p_i)^2} \qquad \textbf{Equation 5.6}$$

Here, q and p represent two points in a n-dimensional space. The dimensionality of the subspace, n, was equal to $k - 1$, where k refers to the total number of task conditions (8 for P2 and 4 for P3). We then tested whether the separation we found between paired neural trajectories was significantly different from random noise. To do this, we simulated the expected Euclidean distance values of a null distribution. For each pair of neural trajectories, we scrambled the condition labels to create two random neural data distributions. Trajectories were phase-averaged prior to scrambling the labels so that only one latent neural data point was obtained per phase. We calculated the Euclidean distance between the trial-averaged scrambled-label distributions (d_s). This process was repeated 1000 times. We then determined whether the separation between the trial-averaged, condition-dependent neural trajectories, d_c, was significant by comparing it to the variability found in the scrambled distribution. The p-value was calculated by finding the number

of instances in which the correctly labeled Euclidean distance measurement was greater than that of the scrambled distribution:

$$p = \frac{N - \sum_{i=1}^{N}[d_c > d_s]}{N} \qquad \textbf{Equation 5.7}$$

where N is the number of repeated iterations and the [ ] operator returns a value of 1 if the condition holds, otherwise it returns a value of 0. We further quantified the discriminability between task conditions by training a 10-fold cross-validated LDA classifier using the FA-projected neural data from each phase. Neural data was phase-averaged prior to classification.

5.2.5 Modulation of Latent Neural Activity Along Primary Condition-Specific Axis

To determine whether object presence and grasp intention engage similar patterns of neural activity, we classified object presence along the primary grasp intention axis, defined as the axis that maximally separates hand closed and hand open conditions. We identified the primary grasp intention axis by finding the LDA projection that maximizes separation between grasp intention conditions while ignoring target location and object presence. This LDA projection is defined by a set of eigenvectors (see section 3.3.4 for details), ranked from most to least informative. Thus, the primary grasp intention axis corresponds to the top-ranked eigenvector of the LDA subspace. We then trained a 10-fold cross-validated LDA classifier to discriminate between the presence or absence of an object using neural data that was projected onto the primary grasp intention axis.

5.3 Results

5.3.1 Characterizing Task Performance

Both participants were able to perform the object interaction task with great success. Participant P2 achieved a 98% average success rate across all conditions. On sessions 1 and 2, his success rate dropped to approximately 86% for trials in which he had to grasp an object. However, for all subsequent sessions, he had no difficulty grasping the object and was able to achieve a 100% success rate. Reach kinematics of the robot arm were similar across all conditions, as seen in Figure 10. Participant P2 displayed highly stereotyped reaches towards both targets, yielding short and consistent reach completion times. We found no significant differences in reach completion times between conditions on any of the five sessions ($p>0.05$, one-way ANOVA and Kruskal Wallis). When pooled together across conditions, we found that session 5 had significantly shorter reach completion times than all other sessions, and session 3 had significantly longer reach completion times than session 1 ($p<0.01$, Kruskal Wallis).

Participant P3 had more variable reach trajectories on average, and generally exhibited more difficulty controlling the robot arm. However, he was able to maintain a 91% average success rate across conditions. On session 3, participant P3 performed exceptionally well, achieving perfect success for all trials. He reported that BCI control felt more intuitive and required less focus, allowing him to quickly and smoothly reach the target location. On sessions 2 and 3, no significant differences in reach completion times were found across conditions. However, on session 4, the participant obtained significantly slower reach completion times when keeping his hand open without an object present (ON) compared to when grasping without an object present

(GN) ($p<0.001$, Kruskal Wallis). These ON trials were also significantly slower than when grasping an object (GO) on both sessions 1 and 4 ($p<0.05$, Kruskal Wallis).

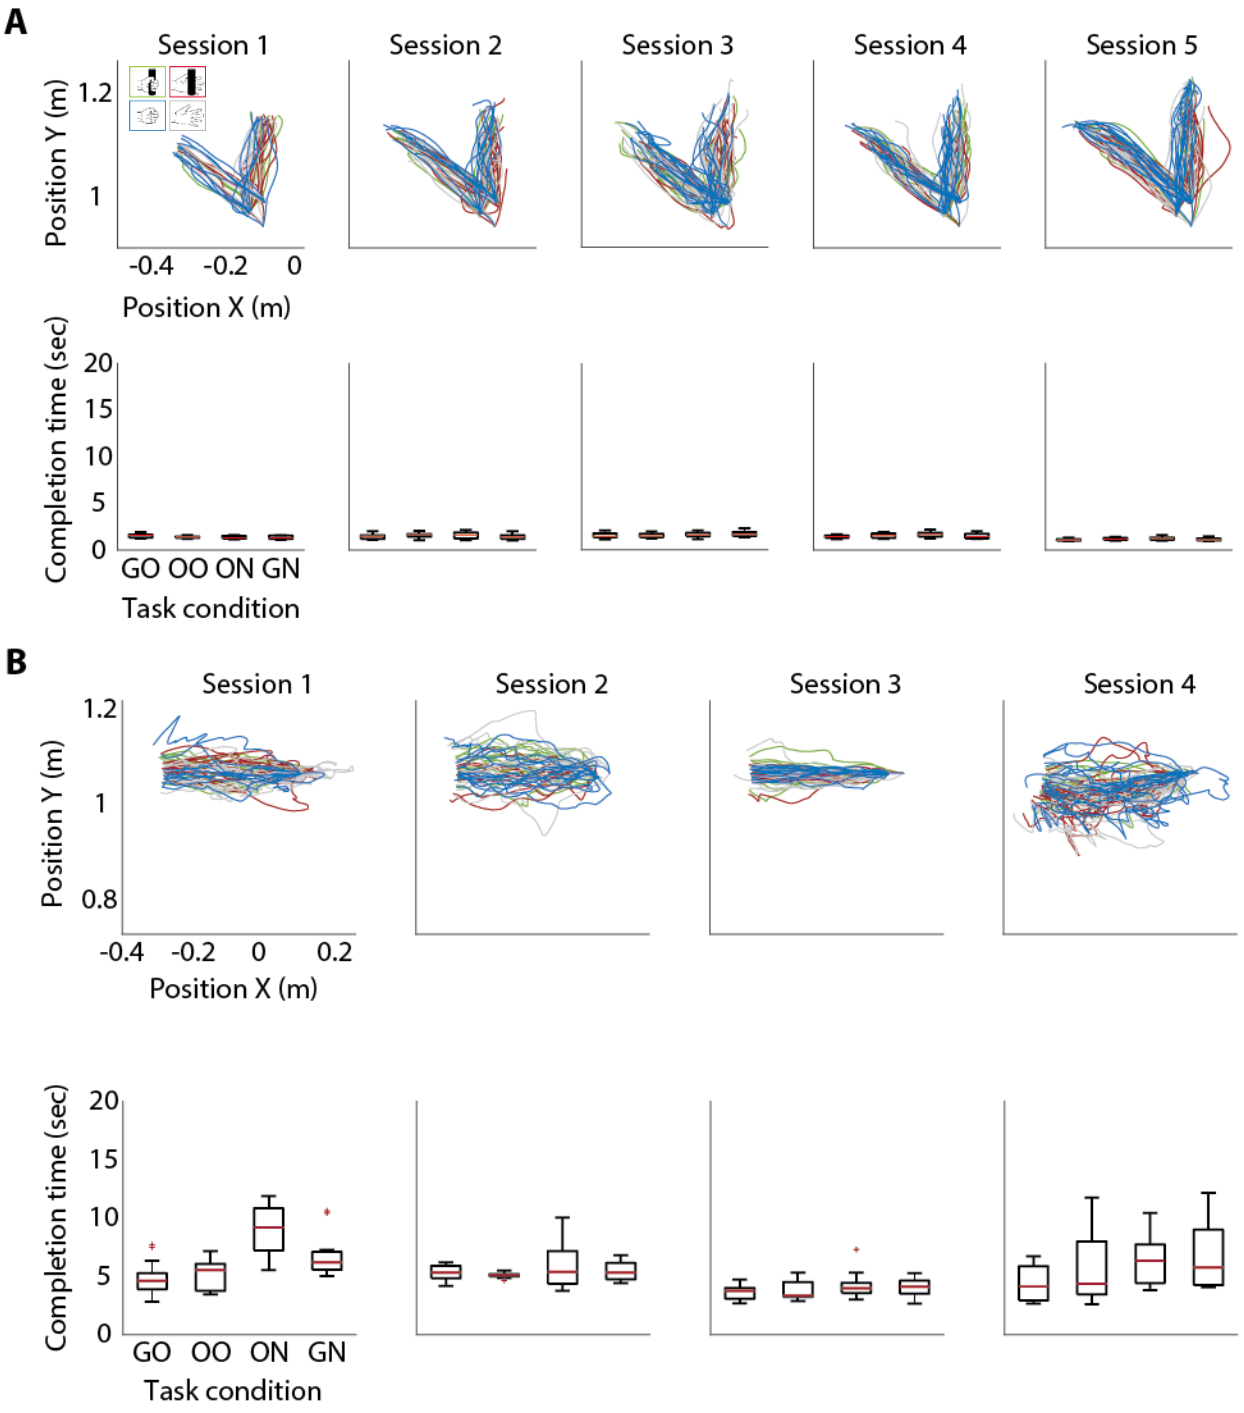

Figure 10 Performance metrics for the object interaction task. Kinematic trajectories, presented from a bird's-eye view (top), and completion times (bottom) are shown for participants P2 (A) and P3 (B). The blue trajectories correspond to grasping without an object, green to grasping an object, red to keeping the hand open with an object, and grey to keeping the hand open without an object. Outliers in reach completion time are denoted by a '+' sign. Reach completion times were combined across target locations for participant P2.

5.3.2 Object Presence Modulates Neural Population Activity in M1 Throughout All Task Phases

To characterize how object presence modulates population motor cortical activity for kinematically similar reach and grasp movements, we first visualized the patterns of neural activity within the estimated neural manifold. In Figure 11, we present the trial-averaged neural trajectories, corresponding to each condition, from a sample session completed by participant P2. Here, we can immediately see that there exists some separation between the condition-dependent neural trajectories across all phases of the trial. Interestingly, conditions that modulate neural trajectories either earlier or later in the trial seem to be less active as the movement changes from reach to grasp. Prior to movement onset, during the presentation phase, we found that neural trajectories were primarily separated by target location and object presence. This marked separation persisted throughout the reach phase. However, as the participant shifted from performing a reaching to a hand-centric movement (e.g., closing the hand or keeping it open) at the start of the grasp phase, we found less discriminability by target location and increased separation in response to both object presence and grasp command. Similar separation was present during the hold phase.

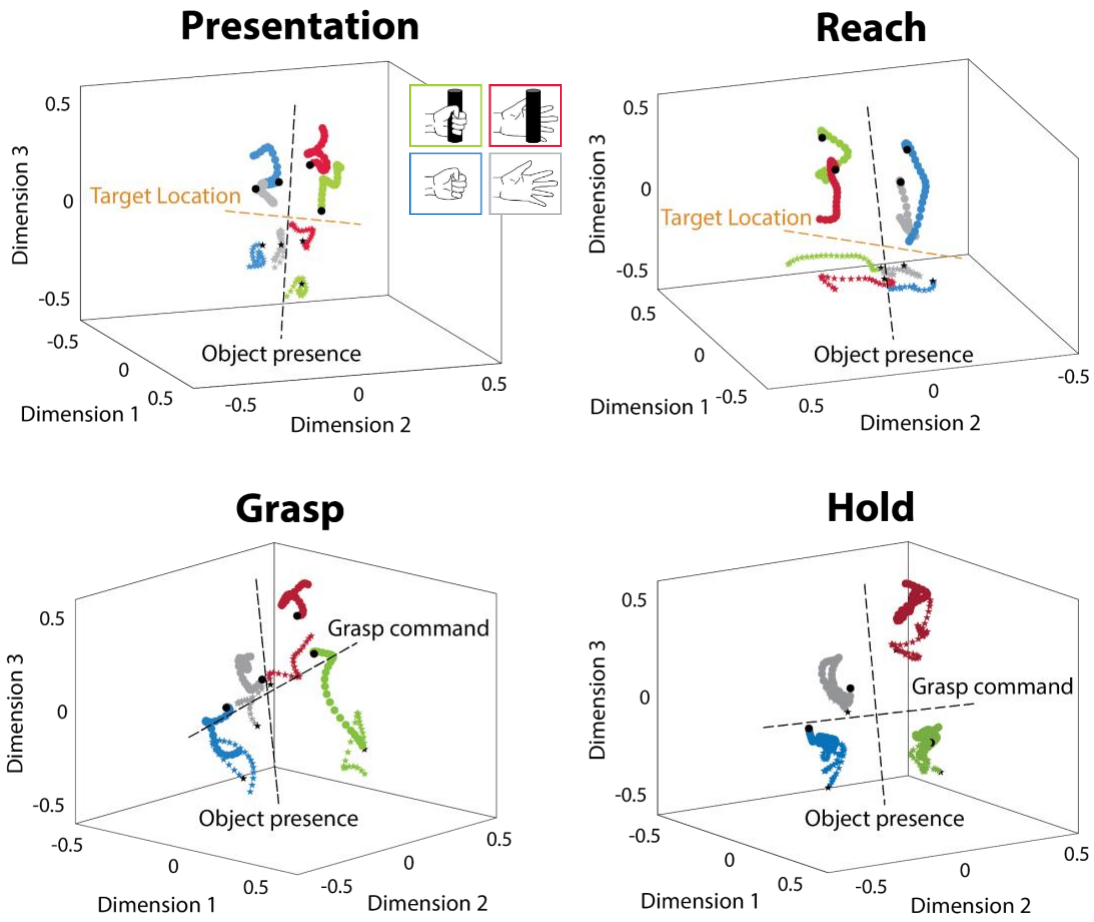

Figure 11 Population neural activity varies in a condition-dependent manner across trial epochs. Example session for participant P2. Trial-averaged neural trajectories are shown for each condition. Colors represent different object and grasp intention conditions, as described in **Figure 10**. Stars represent the "left" target location and circles represent the "forward" target location. The black stars and circles mark the beginning of the neural trajectory for each phase.

5.3.3 Object Presence and Grasp Intention Concurrently Modulate M1 Activity

To quantify the condition-dependent differences within each epoch's subspace, we calculated the Euclidean distance between pairs of trial-averaged neural trajectories at each time point. We determined whether the separation between neural trajectories was significant, in other words, we calculated the p-values presented in this section, by comparing the number of instances

the phase-averaged Euclidean distance measurement was greater than that of the null distribution (represented by the red bars in Figure 12 and Figure 13).

For participant P2, we found significant separation in trial-averaged neural trajectories for "left" and "forward" target conditions earlier in the trial, during the presentation phase and peaking during the reach phase ($p<0.05$ on all sessions). This separation steeply decreased as the trial progressed to the grasp and hold phases. We also found clear separation between object-present and object-absent neural trajectories for both subjects, though they exhibited different levels of baseline modulation (Figure 12). For subject P2, this separation was present at trial onset ($p<0.05$ for n=4 sessions) and increased throughout the remainder of the movement for both hand-open and hand-closed trials, with peak separation occurring during the grasp and hold phases ($p<0.01$, Figure 12A and 13B). Subject P3 showed high levels of sustained separation throughout all trial phases ($p<0.001$, Figure 12C and 13D), with peak separation occurring during the grasp phase.

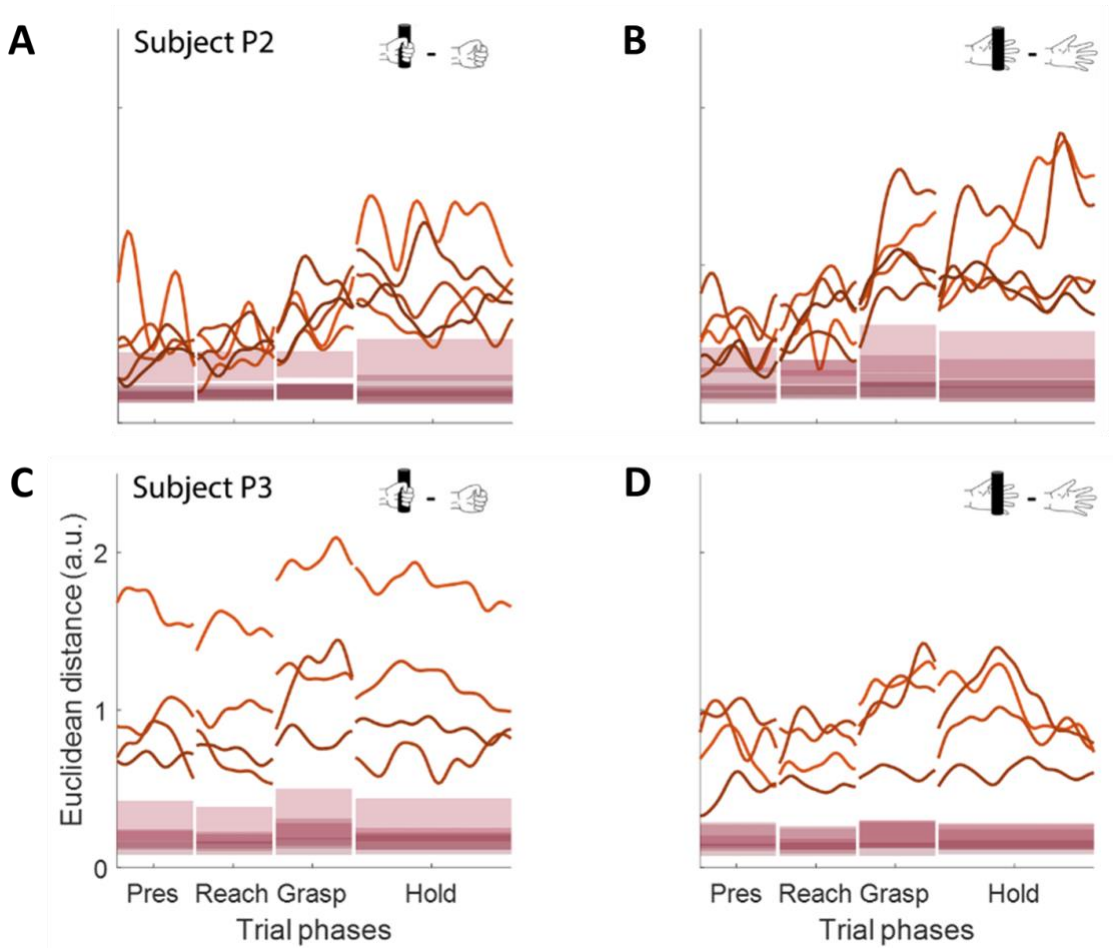

Figure 12 Separation between trial-averaged neural trajectories by object presence. Separation of object-present and object-absent neural trajectories are presented for trials in which the hand was closed (left) or kept open (right) for both participants P2 (A and B) and P3 (C and D). Each orange line represents the Euclidean distance between a pair of trial-averaged neural trajectories. The different shading corresponds to the different sessions data was collected. The red bars on the bottom of the plot correspond to the null distribution of the distance metric, in other words, the amount of variation we would expect to see between two random distributions (see section 5.2.4 for additional details).

Although we found separation of neural trajectories by object presence to occur for both hand open and hand closed conditions, we sought to determine whether grasp intention similarly impacts neural activity. For subject P3, we found that grasp intention significantly modulated

neural activity across all task phases for both object present and absent conditions ($p<0.001$, Figure 13C and 14D). For both subjects, grasp intention modulated neural population activity most prominently late in the trial, during the grasp and hold phases, regardless of object presence ($p<0.01$, Figure 13). We also found greater separation of grasp intention during the presentation phase when an object was present ($p<0.05$ for n = 3 sessions, Figure 13A) compared to absent ($p<0.05$ for n = 1 session, Figure 13B) for participant P2 (Figure 12B), which we did not see for participant P3.

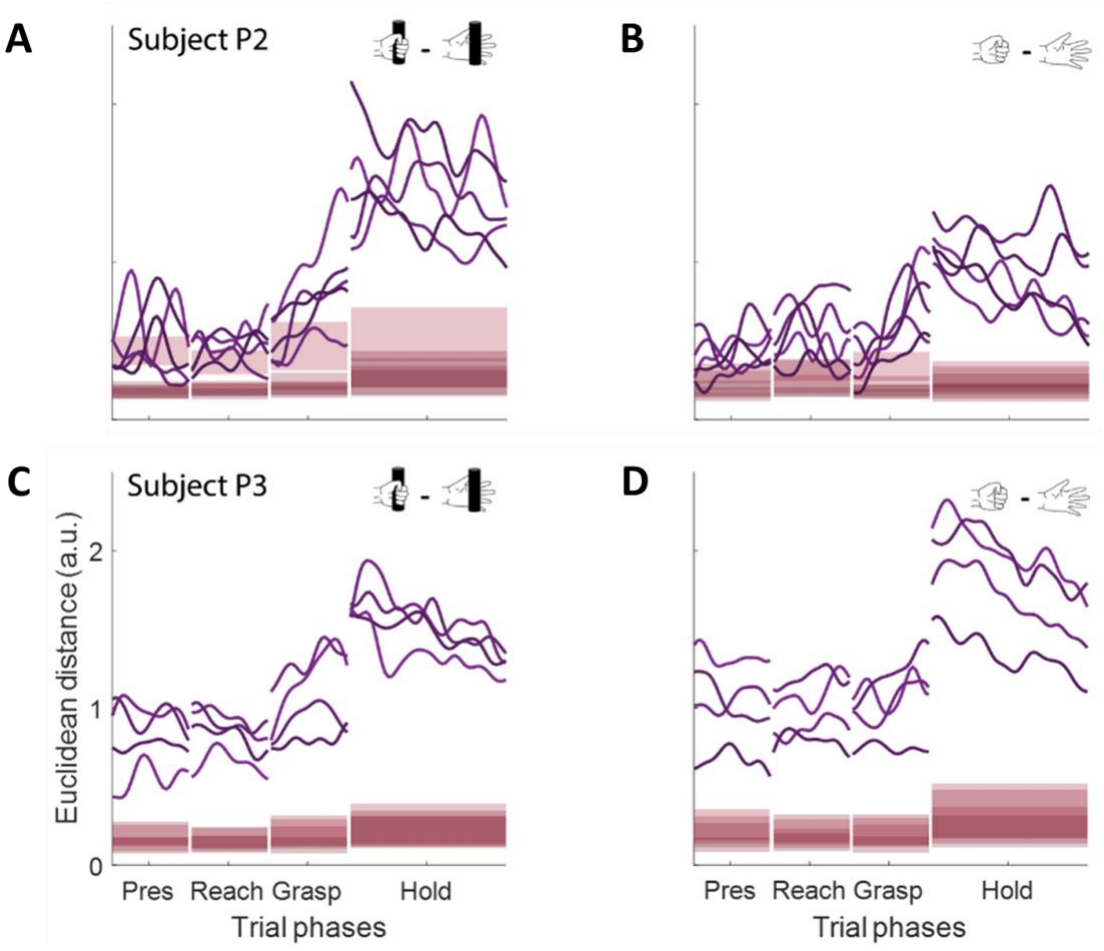

Figure 13 Separation between trial-averaged neural trajectories by grasp intention. Separation by grasp intention for participants P2 (A and B) and P3 (C and D) is shown both when an object was present (left) or absent (right) from the target location. The red bars correspond to the null distribution for each comparison.

We further quantified the significance of the observed separation by classifying the phase-averaged latent neural activity based on the presence or absence of an object as well as whether the hand was closed or kept open. In line with our results, we could classify object presence with accuracies of 75% during planning, 78% during reach, and 91% during the grasp and hold phases, averaged across all five sessions for subject P2. Similarly for subject P3, we could classify the presence or absence of an object with 87% accuracy during planning and reach, 96% during grasp, and 89% during hold phases. These results demonstrate that interaction with an object substantially modulates neural activity not solely at grasp onset but from planning to movement execution.

Classification accuracies for grasp intention (hand open vs. closed) were lower for subject P2, but above chance (50%), during planning and reach (~65%) and increased during the grasp (80%) and hold (96%) phases using data for trials with and without an object present. Subject P3, on the other hand, had larger separation between hand closed and hand open neural trajectories which positively impacted classification with accuracies of 88% during planning, 89% during reach, 94% during grasp, and 98% during hold. Overall, we found that there is a baseline separability between hand closed and hand open neural trajectories that increases to near perfect classification during the hold phase. We expected high classification accuracies late in the trial since the robot kinematics are directly driven by changes in the participants' neural activity.

When all data were combined, the four task conditions could be classified above chance (25%) with increasing accuracy as the hand approached the object ranging from approximately 56% and 87% during planning, 57% and 88% during reach, 77% and 94% during grasp, and 93% and 91% during the hold phase for subjects P2 (Figure 14A) and P3 (Figure 14B), respectively. These data suggest that both object presence and grasp intention are separable effects on motor cortical activity.

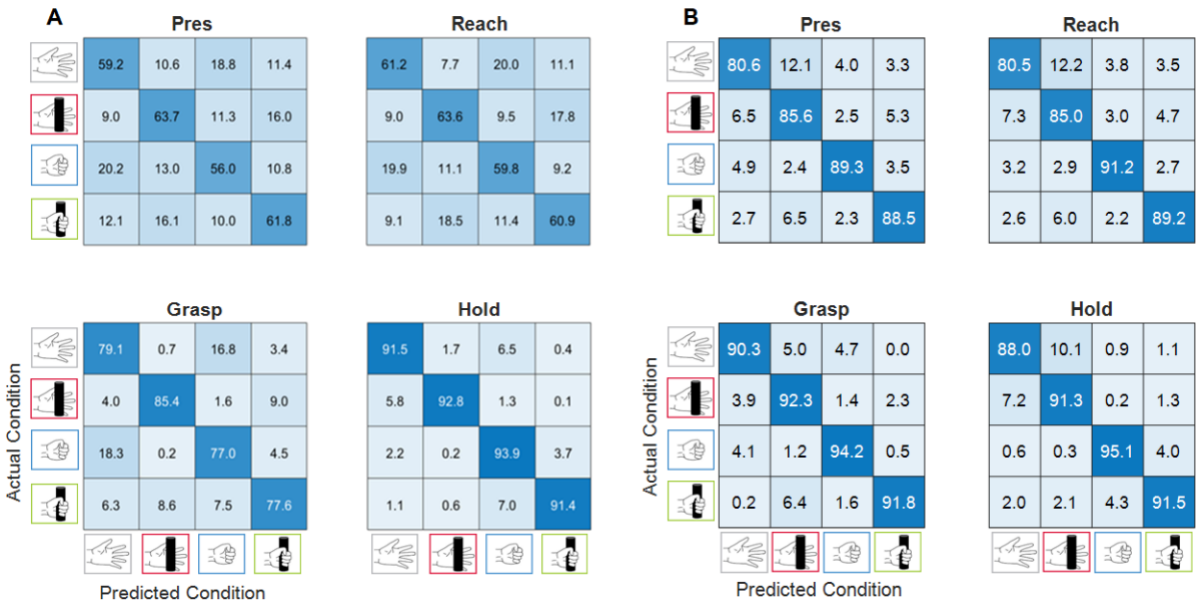

Figure 14 Classification accuracies of latent neural activity by object presence and grasp condition**.** Confusion matrices for each phase, averaged across sessions, are presented for participants P2 (A) and P3 (B).

5.3.4 Object Presence Engages Patterns of Neural Activity Distinct to Grasp Intention

Since both the object presence and grasp intention effects were most pronounced during late trial epochs, we investigated whether these two effects modulated neural activity in a similar manner. We identified a single dimension of the latent neural space that best separated the data according to grasp intention (i.e., a linear discriminant axis). If the object presence effect is driven primarily by changes in hand kinematics rather than the presence or absence of an object, we would expect high classification accuracies for both object presence and grasp intention along the primary grasp intention axis. However, along this dimension, object presence classification was near chance level (50%) across all trial epochs for both subjects, demonstrating that object presence is represented on a separate latent dimension from grasp intention (Figure 15).

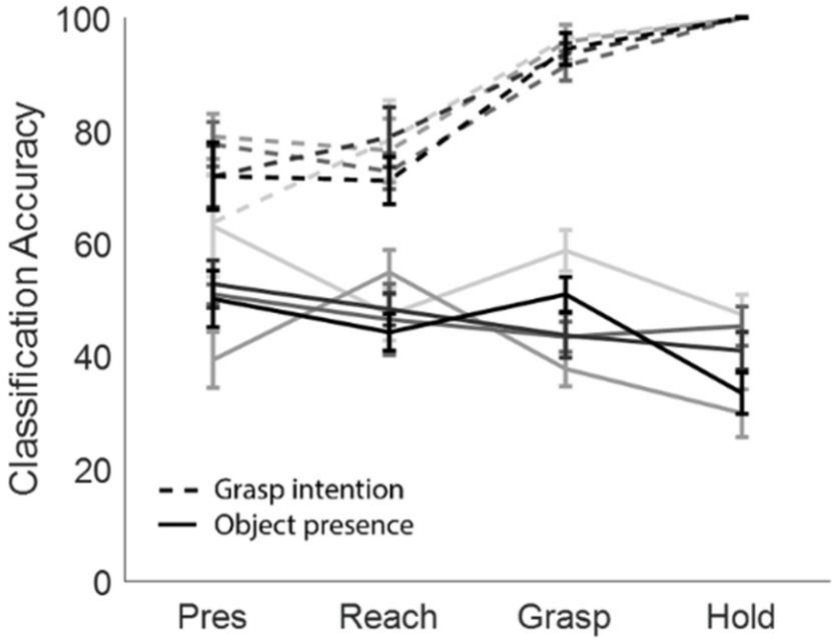

Figure 15 Classification along primary linear discriminant axes. Results are shown for participant P2. An LDA classifier was trained on grasp intention using neural data projected onto the primary grasp intention axis. Accuracies are presented for classification of object presence (solid lines) or grasp intention (dashed lines). Standard error bars are shown for each data point, whose shading represents session order.

Given that object presence is primarily modulated along a separate axis, we sought to determine whether the primary object presence and primary grasp intention axes were orthogonal to one another. If these axes are orthogonal to one another, this would signify that neural activity patterns modulated by the presence or absence of an object are separate to those modulated by hand kinematics. In other words, this would be a non-kinematic driven effect within M1. However, we ultimately concluded that the latent dimensions modulated by object presence and grasp intention were not entirely independent. This can be visualized in Figure 16A where neural trajectories corresponding to all 8 location, object presence, and grasp intention conditions for subject P2 are plotted on the top two linear discriminant axes during the hold phase. Here, we can

see that the separation between neural trajectories corresponding to hand closed and hand open trials is greater when an object was present (green and red trajectories) than when an object was absent (blue and gray trajectories). To determine whether this was generally the case within the entire latent subspace, we compared the difference in the Euclidean distances between hand closed and hand open trials when an object was either present or absent for both subjects (Figure 16B and 17C). We found greater modulation along the grasp intention dimension for object-present trials than for object-absent trials for subject P2. In Figure 16B, the difference in grasp intention separation between object present and object absent conditions was significantly different to a zero median during the presentation, reach (n=4 sessions), grasp, and hold phases ($p<0.05$, Wilcoxon signed rank). For subject P3, the difference in grasp intention separation was also significantly different during the presentation, reach (n=2 sessions), grasp, and hold phases ($p<0.05$, Wilcoxon signed rank, Figure 16C); however, this difference tended to be greater when an object was absent. These findings suggest that the primary object presence and grasp intention axes are not orthogonal to one another and that changes in hand kinematics impacts neural co-activation patterns associated with the presence of an object.

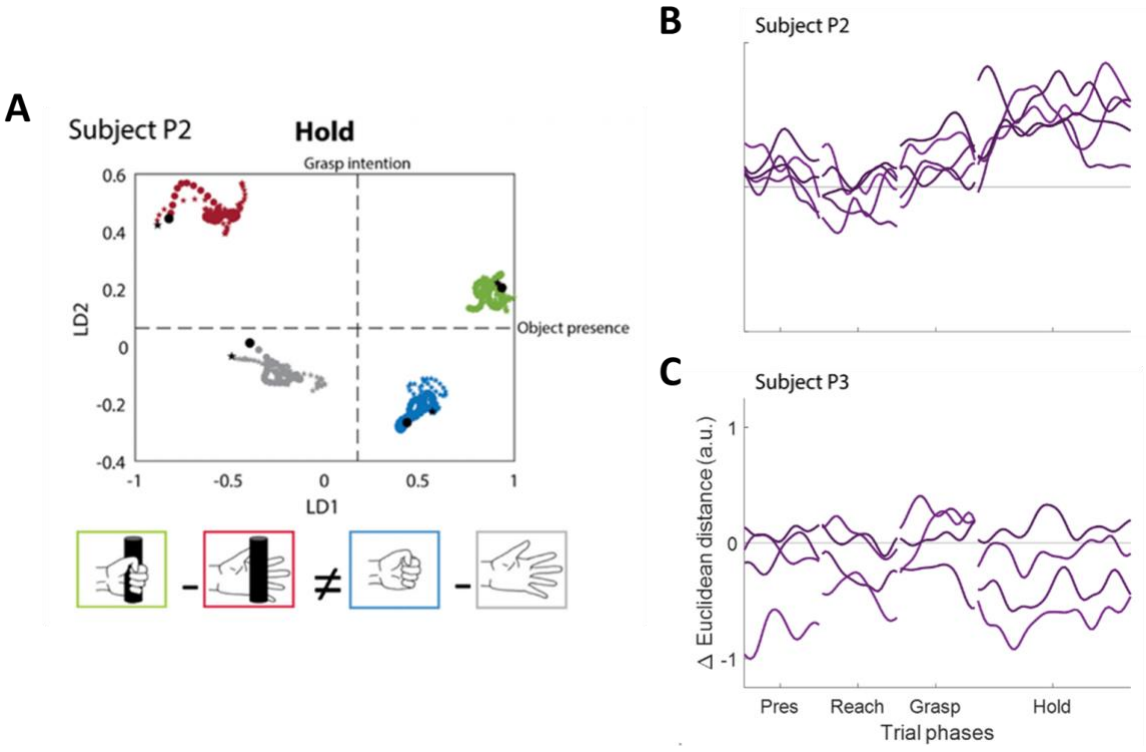

Figure 16 Effect of object presence varies with that of grasp intention. A. Trial-averaged neural trajectories projected onto the top two dimension of the neural subspace estimated for the hold phase of a sample session performed by P2. Dashed vertical lines indicate condition-dependent separation along that axis. B. Comparison of the effect of object presence on the separation of neural trajectories in response to grasp intention for participants P2 and C. P3.

5.4 Discussion

Here, we sought to determine the basis for the modulation of neural activity in response to the presence or absence of an object. Our results demonstrate that object presence significantly modulates neural population activity in motor cortex throughout planning, reaching, and grasping. We find that even when arm and hand kinematics are the same (e.g. reaching and grasping an object vs. reaching and closing the hand), this effect is present. We further corroborated that this

effect was not entirely driven by arm and hand kinematics by analyzing patterns of activity in response to changes in grasp intention. We found that although modulation due to both object presence and grasp intention covaried, these two effects engaged neural activity along different latent axes. These axes, however, were not orthogonal to one another, indicating that these two effects are not entirely independent. Based on these results, we infer that modulation of neural activity due to object presence may be encoding object properties necessary for object interaction.

We found that the effect of object presence and grasp intention modulate motor cortical activity concurrently, either as the participant switches from reach to grasp, as was the case for participant P2, or throughout all phases of the trial, as seen with participant P3. Of interest, for participant P2, we found significantly greater separation of grasp-related neural activity patterns when an object was present in the workspace compared to absent (Figure 12B). Prior studies have demonstrated that merely the presence of a distractor object can impact grasp kinematics (Castiello, 1999; Jackson, Jackson, & Rosicky, 1995). Thus, it may be possible that this context-dependent modulation corresponds to extraction of the object properties and ensuing preparation of a grasping movement, even in cases without an intent to interact with the object, as was the case in the hand open trials. Based on the proposed cortical grasping circuit, this would point to interactions between M1 and other brain regions such as anterior intraparietal area or ventral premotor cortex (Castiello, 2005; Schaffelhofer & Scherberger, 2016).

Here, we found key differences in the participants' baseline performance levels. Participant P3 generally exhibited larger task-context differences which persisted throughout the trial. On the other hand, for participant P2, these differences were smaller at trial onset and increased as the trial progressed from reach to grasp and hold. There are several different possibilities for these differences. Our lab has found that the array placements for participant P3 yield significantly more

units tuned to hand and wrist movements than to elbow and shoulder movements. Participant P2, on the other hand, exhibits significantly more modulation in response to shoulder and elbow movements. These tuning differences impact our ability to decode translation of the arm for participant P3 and was responsible for the large kinematic variation of the robot throughout this experiment. Both the differences in kinematics and perceived difficulty of the task, as well as the differentially tuned units analyzed, may lead to the task-context differences we found between the two participants.

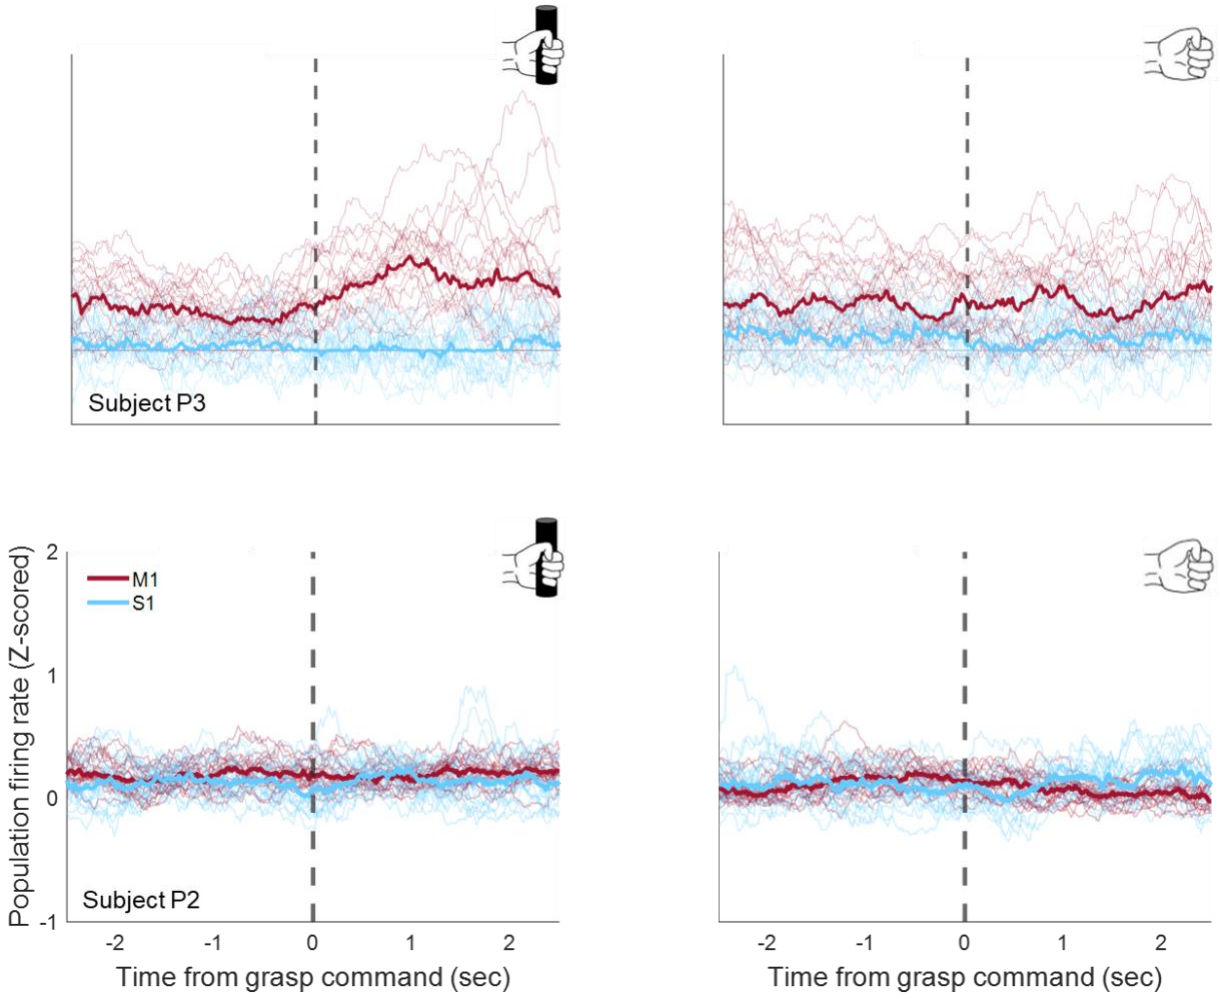

Figure 17 Neural activity in response to object interaction. Z-scored firing rates of both M1 (red) and S1 (blue) units for participants P2 (bottom) and P3 (top). Firing rates from individual units are shown by the less opaque lines and the average response across all units is represented by the thicker lines. The vertical dashed line represents the

time at which the audio cue to grasp was given. Time points to the left of the dashed line correspond to the reaching phase and those to the right of the dashed line correspond to the start of the grasp phase.

Additionally, we speculate that participant P2 may have trained himself out of modulating certain neural activity patterns to achieve better BCI performance during object interaction tasks. P2 has been routinely performing functional tasks that require interaction with objects of various sizes and shapes for the past 6 years. This extensive training may have allowed him to optimize the generated patterns of neural activity to avoid de-stabilizing BCI control during object interaction. Although Downey et al. found a global increase in P2's M1 firing rate during an objection interaction experiment collected in 2017 (Downey et al., 2017), we find little to no object-related response in single units recorded from either M1 or S1 arrays for our dataset (Figure 17 (bottom)). However, when we analyze P3's M1 and S1 activity, we do find increased modulation across a majority of units in M1 in response to the presence of an object (Figure 17 (top)). It is possible that since P3 has had much less practice controlling the robot arm and hand, both for general movements as well as for functional tasks, he continues to generate patterns of neural activity specific to object presence, which negatively impact BCI control.

To better determine whether participant P2 had previously exhibited greater population level modulation due to object presence, similar to that of participant P3, we repeated part of the analysis presented in this chapter using two historical datasets from Downey et al.'s study. We calculated the separation between object-present and object-absent neural trajectories, shown in Figure 18, and found that the magnitude of separation more closely resembled that of P3. This supplementary analysis further bolsters the thought that participant P2's neural activity modulation has changed over time. To investigate whether this change arises from a decline in signal quality, a shift in population neural activity, or learning required to suppress these object-related signals,

in the future, we could conduct a more detailed analysis comparing historical to more recently collected object interaction datasets.

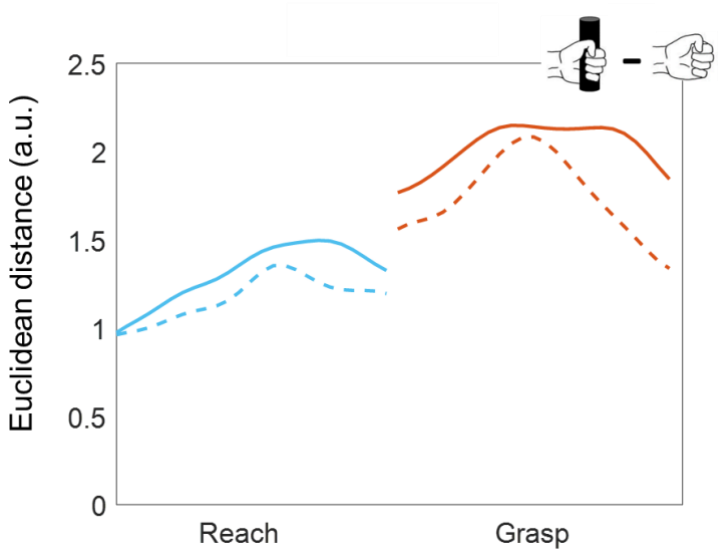

Figure 18 Population activity response to object presence from historical object interaction datasets. We quantified the population level response to object presence during an object interaction task performed by participant P2 in 2017. In this task, the participant was instructed to either grasp an object at a target location or close his hand at the target location with no object present. The same segments of data as those used for Euclidean distance analyses presented in this chapter were used: 1 second centered around the midpoint of reach and 1 second after the cue to grasp. The solid and dashed lines represent two different sessions of data collection.

Since participant P3 was still learning how to achieve stable high degree of freedom control and how to best maneuver the robot arm and hand to smoothly approach the presented target location, there was significant session-to-session variability in his BCI control. This caused changes in the number of target locations experimenters presented, the number of trials collected per condition, and the amount of orthogonal assistance needed to smooth out jittery kinematics on a given session. Depending on the session, P3's varying levels of fatigue, and thus task engagement, may have also impacted decoder calibration and consequently BCI performance.

However, participant P3's novice status also provides us the opportunity to study how robot arm control training from an early stage impacts context-dependent modulation in M1. Future work could involve tracking the degree of object present-dependent modulation, both on a single neuron and population level scale, to analyze how this effect differentially impacts M1 activity over time. This process may also give us insight into learning or adaptation effects that may take place as the participant attempts to overcome control instabilities during object interaction.

6.0 Role of Context-Dependent Signals on M1 and S1 Activity

6.1 Introduction

High-performance brain-computer interfaces have continued to push the boundaries of the gold standard in BCI research to further improve the quality of life for people with tetraplegia (Ajiboye et al., 2017; Flesher et al., 2021; Willett et al., 2021). To achieve this level of control, the vast majority of human BCI studies target motor cortex for decoding movement intention (Collinger, Wodlinger, et al., 2013; Hochberg et al., 2012; Pandarinath et al., 2017; Simeral et al., 2011). However, as we continue to advance BCI capabilities, we are faced with increasingly complex hurdles. In a study of 10-dimensional BCI control of a robotic arm, a participant exhibited difficulty grasping objects if the decoder calibration did not include any objects, although this context-dependency was not as prevalent for 7D control (Collinger, Wodlinger, et al., 2013; Wodlinger et al., 2015). The interaction between these types of context-dependent signals and the user's control signal currently complicates our decoders' ability to reliably predict the user's intended movement. However, context-dependent signals may provide important information that could ultimately improve our estimate of the user's command signal.

Somatosensory feedback, for example, is vital for performing dexterous object manipulation tasks, such as grasping small objects. Despite having visual feedback, without cutaneous or proprioceptive inputs, these types of tasks are very difficult to accomplish because sensory predictions and afferent signals help guide movements (Johansson & Flanagan, 2009b). Researchers have begun focusing on developing bidirectional BCIs to improve motor control through the restoration of sensation. A recent study from our lab demonstrated significant

improvement in a functional object manipulation task when intracortical microstimulation (ICMS)-evoked tactile feedback was provided (Flesher et al., 2021), further demonstrating the importance of sensory feedback for improved motor control capabilities. Performance improvements came from reduced time needed to successfully grasp the object presented. This suggests that the sensory feedback received upon object contact signaled the achievement of the first object interaction subgoal and a transition to the next action phase in the motor plan (Flesher et al., 2021; Johansson & Flanagan, 2009b). S1 has also recently been found to encode imagined reaches during a motor imagery task, in other words, in the absence of sensory feedback. This may suggest that S1 activity reflects an efference copy of movement signals from motor cortices and may play a role in motor production (Jafari et al., 2020a). In the first half of this chapter, we will focus on characterizing the impact of object presence as well as grasp intention on S1 activity. Our goal is to determine whether expectations of sensory feedback or sensory predictions in response to the presence of an object are responsible for the context-dependent response found in M1 activity.

Among the set of nonmotor signals that modulate M1 activity, such as afferent inputs from S1 as described above, here we also consider the impact of cognitive signals corresponding to a attention or arousal (Gallego et al., 2022; Hennig et al., 2021). We propose a variation on the object interaction task described in Chapter 5 to include varying precision restrictions surrounding the target location. This type of task would help us determine whether object-dependent modulation of M1 activity arises from increased attention when reaching towards an object. This attentional increase may arise from the expectation of interaction with an object in the workspace, even in cases where there is no intention to do so, because the hand must be positioned in such a way to support a stable grasp. In this precision-restricted object interaction task, we are interested in

determining if subjects place increased precision requirements when approaching and grasping the object. This would help us determine whether neural trajectories corresponding to trials with lower precision restrictions, in other words easier trials, more closely resemble neural trajectories corresponding to object absent trials from the previous object interaction experiment detailed in Chapter 5.

In the previous chapter we discussed the cortical grasp circuit. Here, we delve deeper to characterize S1 activity in response to task-related parameters, namely object presence, and explore additional nonmotor task variables.

6.2 Methods

6.2.1 Analysis of Object Interaction Dataset Using S1 Recordings

To determine whether there was S1 modulation during the object interaction task, and if so, whether it reflected information about object presence and grasp intention, we repeated the analyses described in sections 5.2.4 and 5.2.5. Please refer to those sections for detailed descriptions of all analyses presented in this chapter. A key difference to note is that the optimal dimensionality needed to explain the S1 data dropped from ~15D in participant P2 to ~7D and from ~30D in participant P3 to ~8D. This is unsurprising as there are significantly fewer S1 (n = 64) than M1 (n = 176-192) electrodes available.

6.2.2 Precision-Restricted Object Interaction Task

The precision-restricted object interaction experiment followed the same structure as that of the original object interaction experiment described in section 5.2.1. However, here we no longer included an open hand condition and instead tested the effect of varying precision requirements on target acquisition. Prior to the start of the experiment, participant P2 was instructed to perform the object interaction experiment while adhering to varying levels of precision restrictions, defined by the area of the error bounds placed around the target location for successful completion of each trial (Figure 19 – top). Trials were blocked by object and precision conditions, e.g., grasping an object with strict error bounds, grasping without an object present with relaxed error bounds, etc. Either one or two sets of randomized blocks were collected each session. Between 10 and 20 trials of each condition were performed by participant P2. Data was collected for participant P2 over 5 sessions 2158 to 2229 days (~6 years) post implant.

As depicted in Figure 19, a relaxed precision restriction, “easy” block, corresponded to relaxed error bounds placed on the target location, allowing the participant to deviate from the target during the 2-second hold time without penalty. The “medium” block implemented the same error bounds used in the previous object interaction experiment. The “hard” block, with the strictest precision requirements, required the participant to stay as close as possible to the target location during the 2-sec grasp hold time with very little forgiveness to movement jitter.

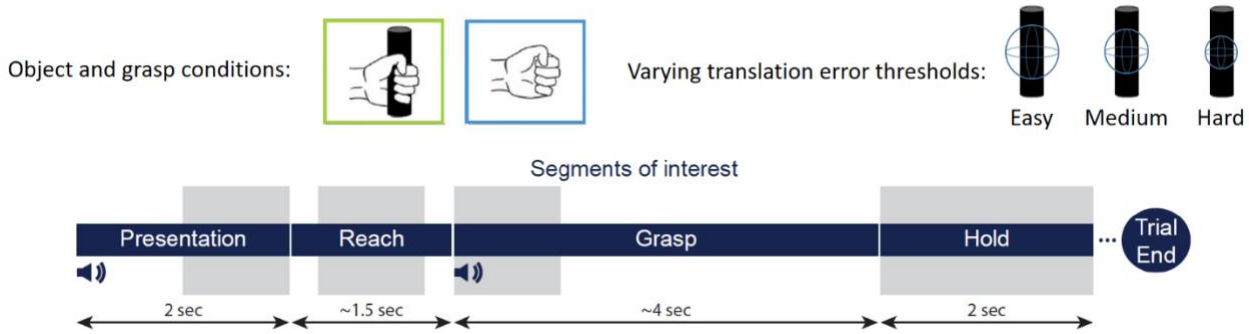

Figure 19 Precision restrictions placed on object interaction task. Only two of the four object and grasp conditions from the original object interaction experiment were tested here, grasping an object and grasping the hand without an object present. The black cylinders on the left represent the object and the circles surrounding the center of each cylindrical object correspond to the error bounds placed for each precision restriction. The same trial structure was used as that described in **Figure 9**.

All behavioral and neural data analyses were conducted in the same manner as that described in sections 5.2.3 and 5.2.4, respectively.

6.3 Results

6.3.1 Object Presence Produces Small Modulation of S1 Activity

To determine whether patterns of S1 neural activity modulate in response to the presence or absence of an object, we quantified the separation between paired trial-averaged neural trajectories for both participants P2 and P3, similar to our analysis in section 5.3.3. For subject P2, although we found some significant separation between object-present and object-absent neural trajectories, it was not consistent across sessions (Figure 20). However, most sessions exhibited significant separation between object-present and object-absent neural trajectories when the participant was instructed to grasp ($p<0.05$, $n=4$ sessions, Figure 20A) rather than keep the hand

open ($p<0.05$, n=1 session, Figure 20B). Separation between hand-closed and hand-open neural trajectories varied for the presentation, reach, and grasp phases; however, significant separation was found during the hold phase for both object-present ($p<0.05$ for n=4 sessions) and object-absent ($p<0.05$ for n=3 sessions) conditions, as can be seen in Figure 21A and 22B. Consistent with these results, we found that an LDA classifier could not easily distinguish between object present and object absent conditions, yielding classification accuracies of approximately 60% across all trial phases (chance level = 50%). Classification by grasp intention yielded marginally better results: 61% during planning, 62% during reach, 63% during grasp and 67% during the 2-second hold. When all four conditions were combined, we obtained ~40% classification accuracy, which while being above chance level (25%), did not indicate consistent significant modulation in S1.

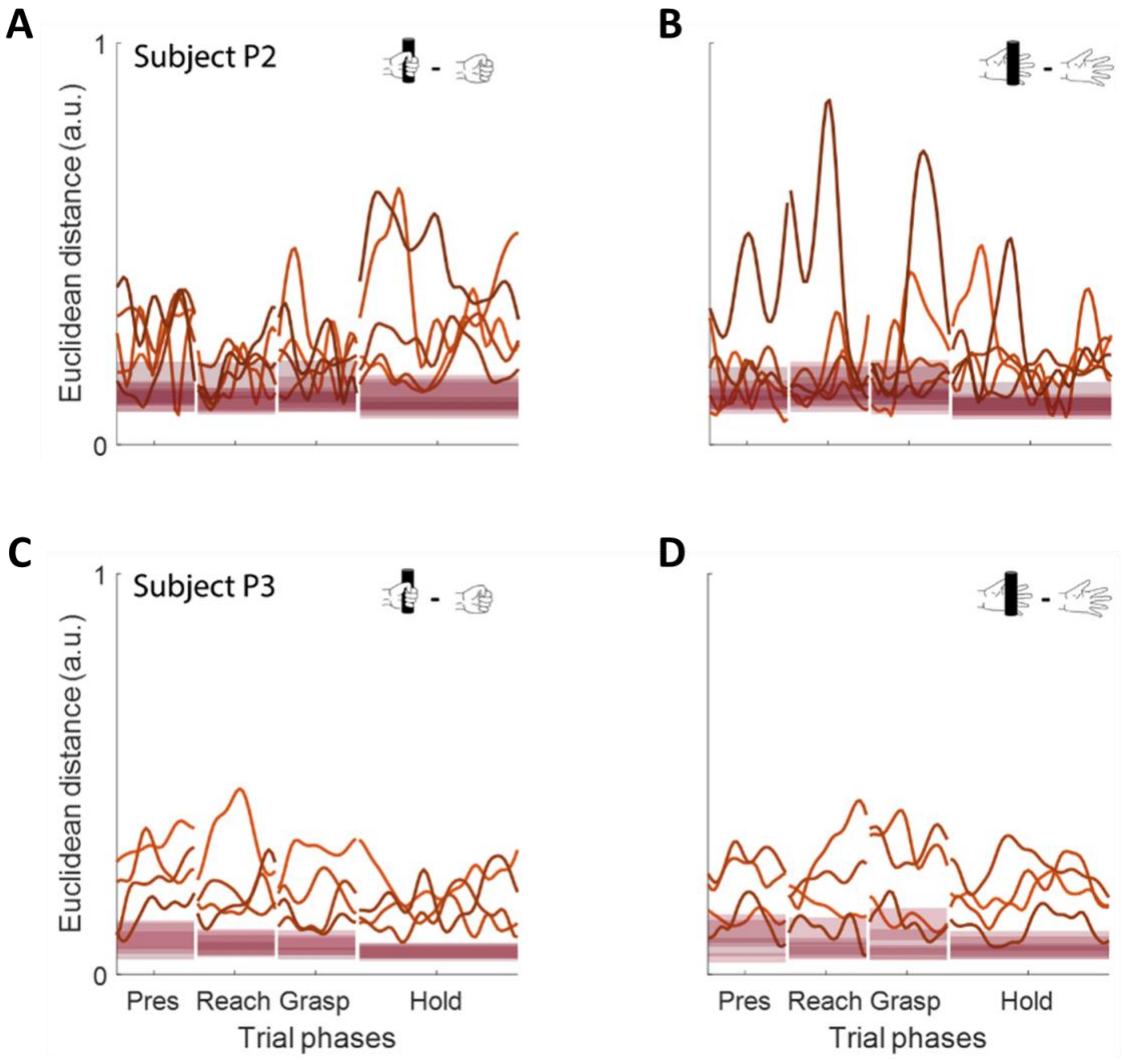

Figure 20 Separation of trial-averaged S1 neural trajectories by object presence. Here we compare separation between trial-averaged neural trajectories when an object was either present or absent for hand closed (A for participant P2 and C for participant P3) and hand open (B for participant P2 and D for participant P3) trials. In each subplot, the red bars represent the theoretical null distribution for each phase (as described in section 5.2.4).

Subject P3 exhibited small but significant changes in S1 modulation in response to both object presence ($p<0.05$, for most phases across sessions, Figure 20C and 21D) and grasp intention ($p<0.05$ for most phases across sessions, Figure 21C and 22D). S1 exhibited less task-dependent modulation on session 4, which had the most inconsistent and unstable kinematic trajectories as

well as the longest completion times out of all sessions collected (Figure 10). We obtained higher classification accuracies for object presence as well as grasp intention discrimination (~68%) for all sessions and trial phases compared to those of subject P2. When the four task conditions were classified, we obtained classification accuracies of 52% during planning, which then marginally increased to 54% for reach and grasp and decreased back down to 52% during hold. These results indicated that although there was some small context-dependent modulation from S1 population activity, it did not tend to vary with task phases, such as the switch from reaching to grasping.

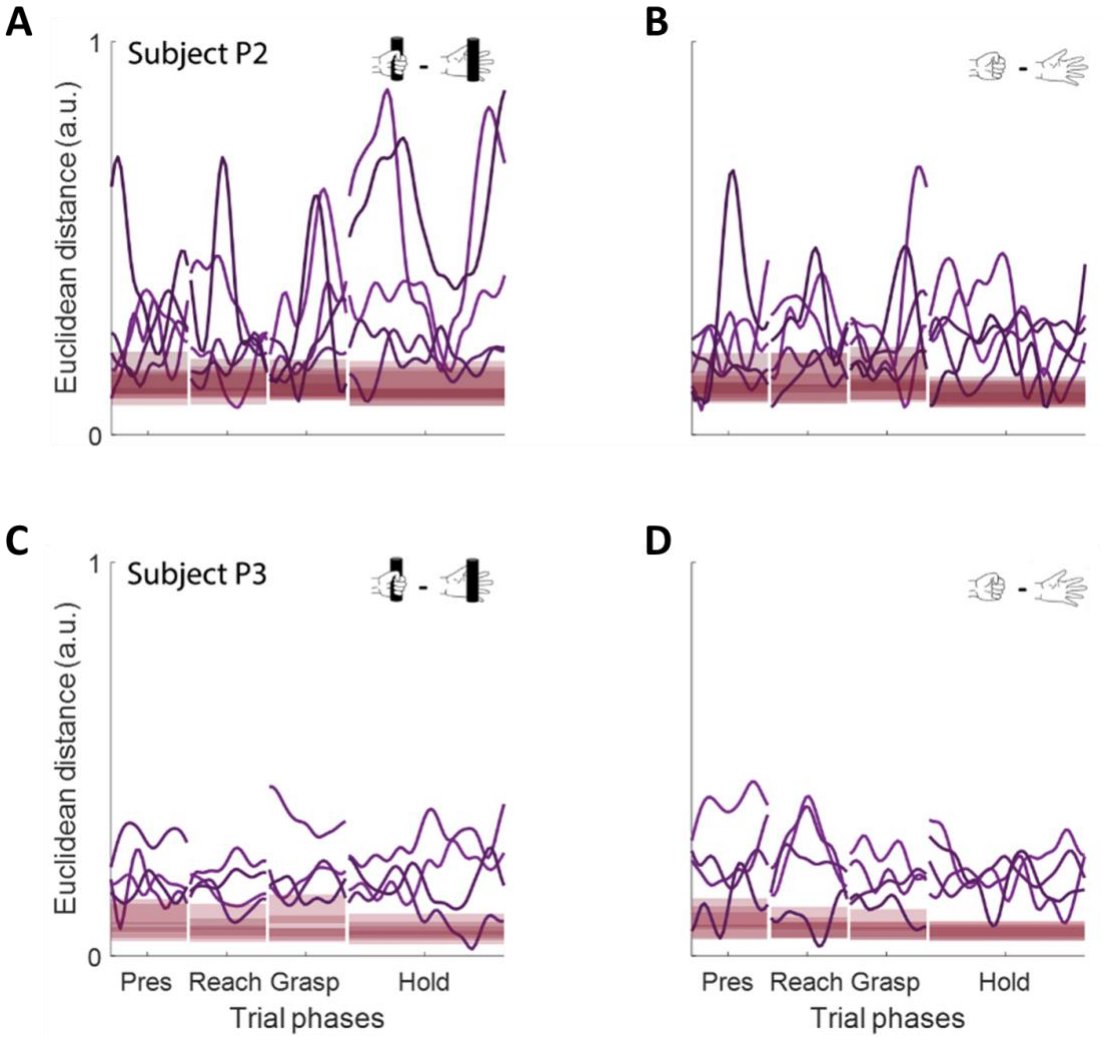

Figure 21 Separation of trial-averaged S1 neural trajectories by grasp intention. Here we show the separation between hand-closed and hand-open neural trajectories both when interacting with an object (A for participant P2 and C for participant P3) or when performing the movement without an object present (B for participant P2 and D for participant P3). The red bars correspond to a theoretical null distribution.

6.3.2 Effect of Nonmotor Task Variables on M1 Activity During Object Interaction

6.3.2.1 Characterizing Task Performance of Precision-Restricted Object Interaction Task

Data for this experiment was only collected with participant P2. As described in Chapter 5, P2 generally demonstrated smooth control of the robotic arm, as showcased by the kinematic

trajectories in Figure 22. Here, we notice slightly different kinematic trajectories to the target location on session 1 compared to all other sessions. We speculate that this was due to an update to the robot's "home" position, the position the robot starts each trial. This new position was used for the entirety of the experiment and did not negatively affect the participant's ability to successfully complete the task. Additionally, we found that the average grasp position for all conditions changed minimally and generally stayed in the same place. This suggests that the participant did not use the larger error threshold bounds to explore the target boundary space. Instead, he consistently reached to the same reference target location.

Generally, the participant reported that the hard difficulty condition was much more difficult than the easy and medium ones, particularly in earlier sessions. Difficulty arose from having to keep the hand closed and within the target location for 2 seconds. Movement jitter and rapid changes in the instantaneous velocity control of the arm made it hard to keep this posture for the required hold time. On average, the participant took significantly more time to reach to the target location for hard trials and significantly shorter for the easy trials across most sessions ($p<0.05$, Kruskal Wallis).

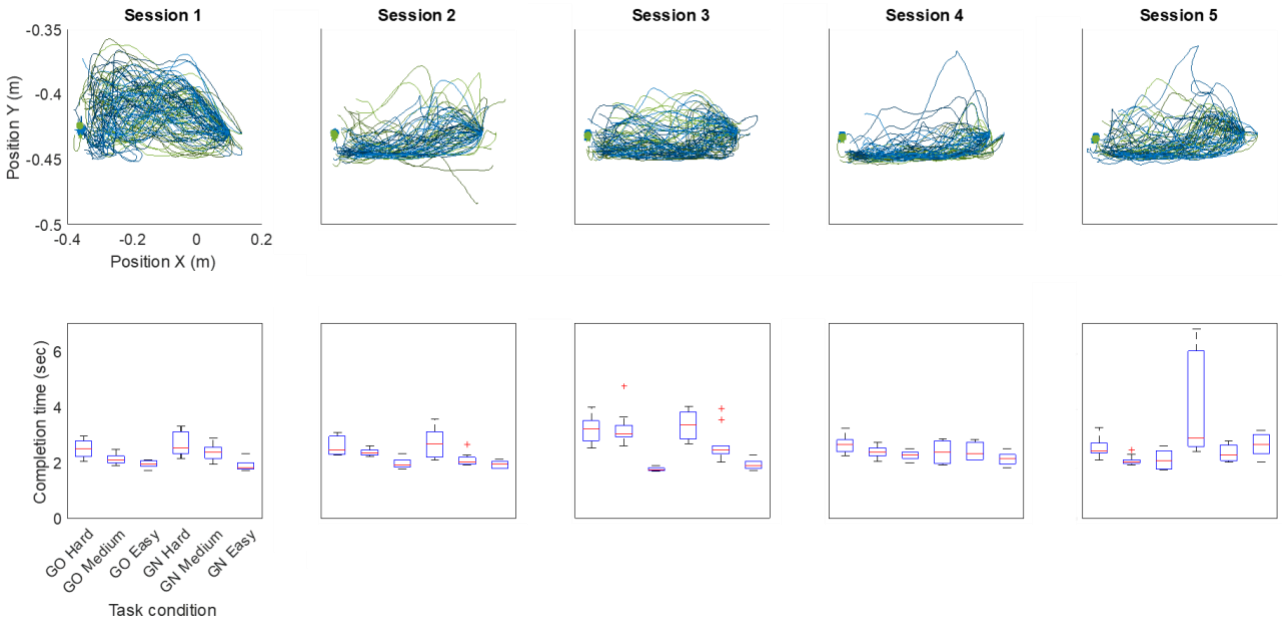

Figure 22 Performance metrics for the precision-restricted object interaction task. Kinematic trajectories, presented from the participant's view (top), and completion times (bottom) are shown for participant P2. The blue trajectories correspond to grasping without an object (GN) and green to grasping an object (GO). Variations in shading correspond to different precision requirements, with the darkest shading representing the strictest error bounds. Outliers in reach completion time are denoted by a '+' sign.

6.3.2.2 Changes in Task Precision Requirements Impact M1 Activity

We began by characterizing how object presence modulates population motor cortical activity, regardless of the precision restriction applied. As seen in Figure 23A, we continued to find significant modulation by object presence ($p<0.001$) for all task phases. We also found that the severity of the precision restriction placed on the target location significantly modulated M1 activity, regardless of object presence, as seen in Figure 23B and 24C. A further comparison of neural trajectory separation in response to precision restriction determined that there was slightly greater separability across object absent (Figure 23B) compared to object present trials (Figure

23C). This was most prominent when comparing the neural trajectory separation between “easy” and “hard” conditions.

We further quantified the degree of neural modulation in response to object presence and precision requirement through LDA classification. As expected, classification of object presence yielded high classification accuracies, similar to those found in the original object interaction dataset: 76% during the presentation phase, 86% during reach, 96% during grasp, and 92% during hold. The precision restriction level (easy, medium, and hard) was classified above chance level (33%) and maintained an accuracy of approximately 56% throughout all trial phases. The combined classification of object presence and precision restriction performed much better than chance (16.6%) with lower accuracies during presentation (50%) and reach (56%), which peaked during the grasp phase (62%), before decreasing again during the hold phase (56%).

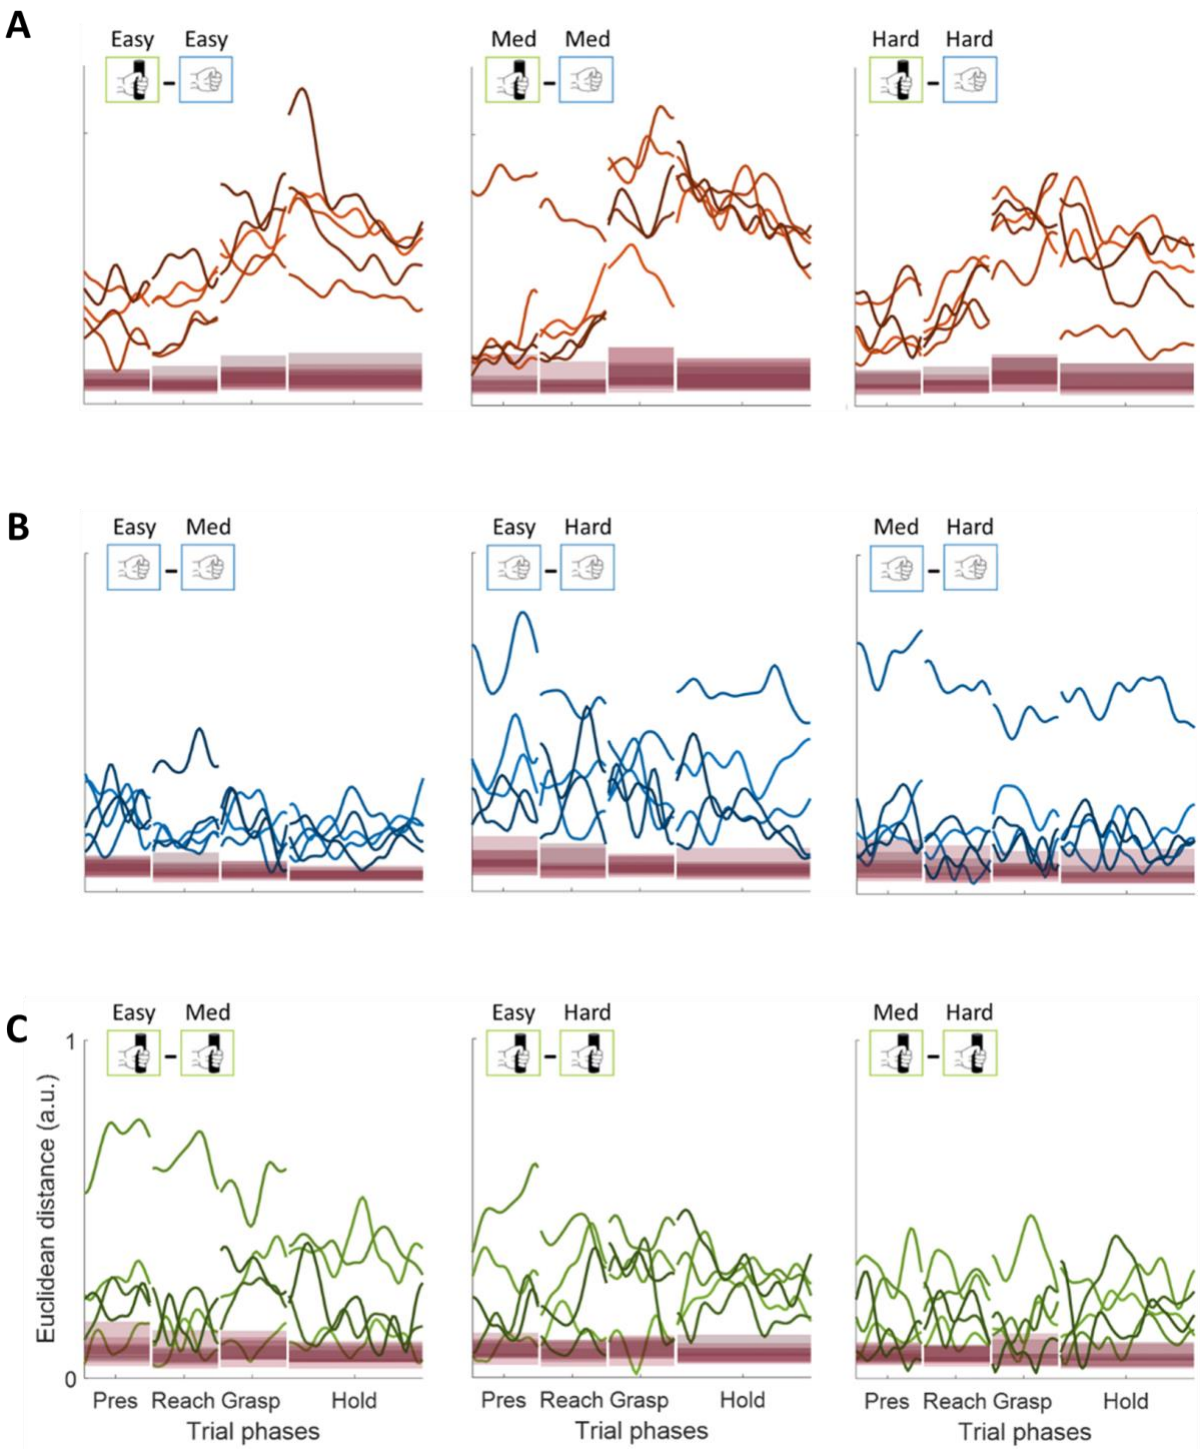

Figure 23 Measurements of M1 modulation in response to object presence and precision restrictions placed on an object interaction task. Top: Euclidean distance measurements of the separation between trial-averaged neural trajectories corresponding to grasping an object and grasping without an object. The left-most column corresponds to trials collected under the strictest precision requirement, "hard", the middle column to a neutral

precision requirement, “medium”, and the right-most column to the least stringent requirement, “easy”. Middle and bottom rows: Euclidean distance measurements of the separation between trial-averaged neural trajectories corresponding to varying levels of precision requirements. The middle row shows this for object absent trials and the bottom row for object present trials. As shown in previous figures, the red bars correspond to theoretical null distributions for comparison purposes.

6.4 Discussion

In this chapter, we have explored both context-dependent modulation in other brain areas (S1) as well as the effect of more cognitive task-related variables on motor cortical activity. In our study of S1 we found, particularly for subject P3, small, yet significant changes in S1 activity in response to both object presence and grasp intention across all task phases. This S1 modulation may correspond to bidirectional communication with M1 for the optimization of movement execution. A recent study found that S1 neurons encoded intended reaches during an imagined reach task in the absence of overt movement or sensory feedback (Jafari et al., 2020a). The authors reflected that this suggests that S1 activity may reflect an efference copy of movement signals originating from M1. Furthermore, it suggests that S1 may play a role in motor production in the absence of sensory feedback, similar to the conditions of the object interaction task in which the participant does not perceive cutaneous sensation upon object contact. Although we did not find a similar level of S1 modulation in participant P2, this may be due to differences in array placement between the two participants, as described in section 5.4. Participant P2’s decrease in motor cortical single-unit response to object presence (Figure 17) may also indicate a lack of object-related information projected from S1.

Our study did not provide sensory feedback through intracortical microstimulation of the S1 arrays upon object contact. However, our lab has previously demonstrated that bidirectional BCIs that can evoke tactile percepts on the fingers can significantly improve functional task performance (Flesher et al., 2021). As the technological development of BCIs continues to advance, the incorporation of sensory feedback will be crucial to provide a clinically viable neuromodulation device for use in daily living activities (C. Hughes, Herrera, Gaunt, & Collinger, 2020). Thus, we must also understand how motor cortical activity responds to cutaneous feedback as well as feedback from unexpected behavioral errors (e.g. incorrectly grasping an object). We would expect these afferent signals to change M1 from a feedforward to a feedback-driven system (Perich et al., 2020). To better understand the impact of afferent inputs to M1 for motor execution, future work would look to characterize the changes in neural modulation as well as the impact on BCI performance in response to a simple, linear encoding scheme applied to one electrode corresponding to the index finger. This study could then be expanded to analyze the inclusion of groups of electrodes as well as more biomimetic encoding schemes to provide more natural percepts.

Next, we tested whether the population-wide increase in firing rates we saw in response to an object may also be driven by changes in precision requirements, or perceived task difficulty. In our precision experiment, we found some modulation of neural activity in response to task difficulty both when an object was present and when it was absent from the workspace. Unlike modulation in response to grasp intention or object presence, this effect stayed constant throughout the entire trial and did not change as the subject approached the target location or closed the hand. This effect was most notable between easy and hard object-absent trials, potentially suggesting increased levels of attention needed to precisely reach to the target location without an object

serving as a reference marker. Since we find that changes in task difficulty impact neural activity, we suspect that the object presence effect may in part be driven by changes in cognitive state such as increased perceived precision requirements for object-present trials.

Although precision requirements were varied throughout the object interaction task, we found that the participant inherently placed the same, strict requirements for all trials, regardless of the true error boundaries around the target. If the participant's goal was to achieve a high success rate, it would make sense for him to consistently reach to the same location, regardless of increased or decreased precision requirements. It would be more efficient and simpler to implement this strategy than to explore the space and determine the limits of each target boundary. However, although he was always attempting to stay within a strict error boundary around the target location, reach completion times demonstrated that the knowledge of the difficulty of the task significantly impacted the speed of his reaching movement. Thus, it's possible that the knowledge of an increased restriction on the task either caused him to reach towards the target more slowly to minimize trajectory error or impacted his BCI decoder in such a way that movement velocities were decreased for these harder trials. To further test the potential impact of task engagement or attention, we could develop an additional experiment that requires the participant to directly engage in the varying difficulty levels of the task. For example, we could instruct the participant to perform a balancing task that would inherently become more difficult if the object to be balanced is placed on increasingly smaller surfaces. We could also incorporate pupil size measurements, which have been found to correlate with internal state, to test whether the pupils were more dilated for more difficult trials (Hennig et al., 2021; McGinley et al., 2015).

Based on the results presented in this and the previous chapter, we can infer that modulation of motor cortical activity due to object presence is not solely driven by changes in kinematics and instead may be driven by a combination of kinematic, sensory, and cognitive aspects.

7.0 Conclusions and Future Work

7.1 Summary

7.1.1 Stability of Task-Related Activity Over Time

We demonstrated that our method of neural signal stabilization could overcome natural recording instabilities to provide stable BCI control of a computer cursor over a period of 104 days in a person with tetraplegia. Our method is predicated on the concept that there exists an underlying stable representation of population neural activity, which is thought to reflect constraints imposed by neural connectivity. Rather than overcoming neural recording instabilities on a unit-by-unit basis, we instead stabilized our neural recordings using the stable representation of latent neural activity. We found that our method not only outperformed a fixed decoder, which was not updated over the course of the experiment, but it also performed comparably to a same-day trained decoder, the current standard in BCI research. Neural signal stabilization provides users with stable, uninterrupted BCI control that does not require manual intervention for performing a recalibration procedure. Thus, our method provides independent BCI operation for potential users.

7.1.2 Effect of Context-Dependent Task Parameters on M1 and S1 Activity

To understand the context-dependent modulation exhibit by M1 during object interaction, we characterized the impact of object presence as well as grasp intention on population neural activity. We found that object presence significantly modulated population neural activity

throughout planning, reaching, and grasping. This effect varied concurrently with that of grasp intention, such that increases in modulation to object presence also corresponded to increases in modulation to grasp intention. Although, the effects of object presence and grasp intention were not completely independent, we found they modulated distinct patterns of neural activity. These results suggested to us that the context-dependent modulation in response to object presence may be linked to higher-order cortical processing of grasp affordance.

To obtain a better understanding of the neural basis for M1's modulation in response to object presence, we next investigated the role of context-dependent signals on both M1 and S1 activity during object interaction. For the object interaction task, described in Chapter 5, we found some modulation in S1 activity corresponding to the presence of an object throughout planning, reaching, and grasping. Although this task did not incorporate sensory feedback during object interaction, we find that our results may still point to communication between M1 and S1 during movement execution. We also tested the effects of a cognitive-based task parameter, the variation of a precision requirement for target acquisition, on motor cortical activity. We found that this context-dependent modulation was most notable between easy and hard trials, potentially suggesting that increased attention or task engagement helps to reach to the target location more precisely. This further suggests that M1 activity not only encodes for kinematic parameters but may contain a combination of kinematic, sensory, and cognitive-related signals.

7.2 Future Work

7.2.1 Generalization of Neural Signal Stabilization for Continuous Use

The current design of our neural signal stabilization method relies heavily on design choices made in preclinical work conducted in animal models. Although this helped provide insight to determine a reasonable initial set of parameters that was ultimately able to recover BCI control over a period of three months, we consider the amount of improvement that could be achieved by further optimizing various method parameters. This could include testing the ideal timing of stabilization updates using either trial-based or open-ended tasks while the stabilization method is used. We could also envision an additional algorithm that estimates the most favorable number of alignment channels, bin size, latent dimensionality, noise variance, etc. either across a broad set of users or personalized for each separate user to provide optimal results. Additionally, within our study design, our method was not fully integrated within our system to run continuously throughout the day, thereby lacking the ability to employ BCI control in unstructured tasks. Through the use of a portable system that can be operated independently by the user at home, a system which we have previously presented (Dekleva et al., 2021; Weiss et al., 2019), we could further generalize our method to be used at the discretion of the user. For example, the user could consider turning on stabilization 5 minutes prior to needing BCI control and then turning it off once they achieve their desired level of control or, conversely, they may prefer to have the stabilization algorithm running continuously in the background and not have to think about turning it on or off. Lastly, although neural signal stabilization was able to provide accurate, functional BCI control over a period of 3 months, we also may expect to see decreases in cursor velocity or trajectory accuracy over longer time spans. Rather than viewing our method as a replacement to

decoder recalibration, we anticipate users to enable neural signal stabilization for performance maintenance, followed by the calibration of a new decoder every few months. We can compare this to placing a computer on sleep mode on a day-to-day basis and then restarting or shutting down the computer once a week to fully optimize performance. We expect these future directions to advance our ability to provide a generalized version of the neural signal stabilization method that can be used with continuous BCI control across a large set of users.

7.2.2 Understanding Context-Dependent Modulation in Real-World Scenarios

The studies presented in this body of work aim to explain the effects of context-dependence on motor cortical activity during a standard reach and grasp motion under ideal laboratory conditions. However, to provide robust BCI control that can be used across a variety of settings, we must look beyond simple use cases. In chapter 5.0, we found evidence that suggests that participants may learn to overcome control instabilities arising from changes in task context. We speculate that participants innately attenuate neural activity that highly differs from that recorded during the decoder calibration period. This would allow the participant to more effectively approach and interact with an object without needing to overcome instantaneous deviations in arm and hand velocity. To further explore this possibility, we suggest tracking and analyzing recorded neural activity as new users begin to use BCI-controlled robot arms for object manipulation tasks. This could provide insight into any learning or adaptation effects that may arise as a participant modulates their neural activity to drive the decoder to produce the desired movement output.

The development of bidirectional BCIs, which incorporate both motor control of assistive devices and sensory feedback from interaction with a user's environment, poses additional opportunities for context-dependent modulation of M1 activity. Given that the integration of

sensory feedback is vital to providing more accurate motor control, we may expect to find that changes in afferent inputs to M1 during intracortical microstimulation of S1 areas will lead to changes in neural activity patterns. As suggested in section 6.4, future work could characterize the changes in neural modulation as well as the impact on BCI performance when incorporating intracortical microstimulation during object interaction.

Although we found some evidence to support the idea that the modulation of motor cortical activity due to task context may be driven in part by cognitive aspects, such as task engagement, we find that the potential impact of task engagement must be tested further. Section 6.4 highlights additional experiments to be conducted that would more directly require the participant to engage in a task that has distinct and obvious levels of difficulty. By incorporating additional metrics, such as pupil size measurements or reaction time, we could use these markers to verify the participant's state of arousal or engagement in the task. While this section does not constitute an exhaustive list of possible behavioral or task contexts, we do consider these additions to span a large set of eventualities, leading us one step closer to providing robust BCI control in real-world applications.

7.2.3 Dynamical Systems Approach to Decoding Context-Dependent Signals

Here we found that motor cortical activity not only contains a motor signal that readily maps to movement output, but also encodes cognitive and somatosensory-related information. We have briefly mentioned that these signals provide rich information about the task and could serve to improve BCI decoding. However, it may be difficult to isolate the control from the condition-dependent signal. The representational encoding model widely used for BCI decoding today directly relates neural activity to movement parameters. Thus, when context-dependent signals arise, these decoders are unable to provide a good prediction of the user's intended movement.

A recent study from Rouse et al. found that neural dimensions corresponding to location and object conditions progressively shifted as the trial progressed from reach to grasp and hold (Rouse & Schieber, 2018). The authors found that in a reach-grasp-manipulation task, the firing rate variance could be split up into general task, location, and object subcomponents. When PCA was fit to the general task component of neural activity, they found rotational dynamics from presentation to reach to grasp and hold, demonstrating that this structure exists not only in reaching but also in grasping movements. However, this projection did not show clear separation between different objects and their target locations. Instead, each trajectory followed similar neural dynamics in the general subspace. When projected onto separate movement and contact subspaces, they found separation between conditions. Thus, rather than having modulation on a fixed dimension, changes due to object type and target location shifted throughout the trial.

From a dynamical systems perspective, it would make sense for neural activity to encode for differences in object context. This would allow the motor cortex to prepare for movement or differentially encode for various task parameters important for context-dependent reaching and grasping without explicitly producing these movements (Elsayed et al., 2016; Kaufman et al., 2014). Unlike the representational model, the dynamical systems approach does not assume that movement-related parameters are the only variables modulating neural activity. Instead, this approach considers the both the state of the system as well as its initialization point. We can consider various context-dependent parameters to correspond to different initial points for our neural system. Depending on the context of the task, e.g., whether we plan to grab a bottle or merely close our hand, a dynamical systems-based decoder would in theory be able to generate the same intended kinematic movement, simply from different starting points within the neural system.

A dynamical-based decoder was recently shown to achieve significantly better performance compared to an OLE decoder, a non-dynamical, standard-use BCI decoder, in a cursor control target acquisition task (Kao et al., 2015). However, no studies have incorporated a dynamical systems-based decoder for decoding of arm and hand movements. As the field moves toward developing bidirectional BCIs that incorporate sensory feedback, we may start to see M1 modulation in response to intracortical stimulation, in which case, using a dynamical systems approach may improve high degree of freedom robot control.

7.2.4 Leveraging Additional Brain Areas for Improved BCI Performance

Although motor cortical activity is well-tuned to a host of movement parameters (Kalaska, 2009), making it an appealing choice for BCI applications, M1 is only one of many brain regions responsible for motor output (Gallego et al., 2022). Beyond the cortical grasping network, described in Chapter 5, other cortical areas play important roles in processing cognitive or contextual information about movement (Fine & Hayden, 2022; Yokoi & Diedrichsen, 2019). By incorporating neural recordings from multiple brain areas, including areas traditionally thought of as representing higher-order variables for motor control such as premotor cortex, we may be able to use this rich set of information to provide contextually generalizable BCI systems more suitable to daily living tasks.

Bibliography

Aflalo, T., Kellis, S., Klaes, C., Lee, B., Shi, Y., Pejsa, K., … Andersen, R. A. (2015). Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. *Science*, *348*(6237), 906–910. https://doi.org/10.1126/SCIENCE.AAA5417/SUPPL_FILE/AFLALO.SM.PDF

Ajiboye, A. B., Willett, F. R., Young, D. R., Memberg, W. D., Murphy, B. A., Miller, J. P., … Kirsch, R. F. (2017). Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: a proof-of-concept demonstration. *The Lancet*, *389*(10081), 1821–1830. https://doi.org/10.1016/S0140-6736(17)30601-3

Alon, G., & McBride, K. (2003). Persons with C5 or C6 tetraplegia achieve selected functional gains using a neuroprosthesis. *Archives of Physical Medicine and Rehabilitation*, *84*(1), 119–124. https://doi.org/10.1053/APMR.2003.50073

Andersen, R. A., Aflalo, T., & Kellis, S. (2019). From thought to action: The brain-machine interface in posterior parietal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(52), 26274–26279. https://doi.org/10.1073/PNAS.1902276116/ASSET/439E2333-1950-498F-9AD5-184CC483187F/ASSETS/GRAPHIC/PNAS.1902276116FIG05.JPEG

Anderson, K. D. (2004). Targeting Recovery: Priorities of the Spinal Cord-Injured Population. *Journal of Neurotrauma*. https://doi.org/10.1089/neu.2004.21.1371

Barrese, J. C., Rao, N., Paroo, K., Triebwasser, C., Vargas-Irwin, C., Franquemont, L., & Donoghue, J. P. (2013). Failure mode analysis of silicon-based intracortical microelectrode arrays in non-human primates. *Journal of Neural Engineering*, *10*(6), 066014. https://doi.org/10.1088/1741-2560/10/6/066014

Begliomini, C., Wall, M. B., Smith, A. T., & Castiello, U. (2007). Differential cortical activity for precision and whole-hand visually guided grasping in humans. *European Journal of Neuroscience*, *25*(4), 1245–1252. https://doi.org/10.1111/J.1460-9568.2007.05365.X

Bishop, W., Chestek, C. C., Gilja, V., Nuyujukian, P., Foster, J. D., Ryu, S. I., … Yu, B. M. (2014). Self-recalibrating classifiers for intracortical brain–computer interfaces. *Journal of Neural Engineering*, *11*(2), 026001. https://doi.org/10.1088/1741-2560/11/2/026001

Blabe, C. H., Gilja, V., Chestek, C. A., Shenoy, K. V, Anderson, K. D., & Henderson, J. M. (2015). Assessment of brain-machine interfaces from the perspective of people with paralysis Related content. *Journal of Neural Engineering*, *12*(4), 043002. https://doi.org/10.1088/1741-2560/12/4/043002

Blakely, T., Miller, K. J., Zanos, S. P., Rao, R. P. N., & Ojemann, J. G. (2009). Robust, long-term control of an electrocorticographic brain-computer interface with fixed parameters. *Neurosurgical Focus*, *27*(1), E13. https://doi.org/10.3171/2009.4.FOCUS0977

Bouton, C. E., Shaikhouni, A., Annetta, N. V., Bockbrader, M. A., Friedenberg, D. A., Nielson, D. M., … Rezai, A. R. (2016). Restoring cortical control of functional movement in a human with quadriplegia. *Nature*, *533*(7602), 247–250. https://doi.org/10.1038/nature17435

Briggman, K. L., Abarbanel, H. D. I., & Kristan, W. B. (2005). Optical imaging of neuronal populations during decision-making. *Science*, *307*(5711), 896–901. https://doi.org/10.1126/SCIENCE.1103736/SUPPL_FILE/BRIGGMAN-SOM.PDF

Bullard, A. J., Hutchison, B. C., Lee, J., Chestek, C. A., & Patil, P. G. (2020). Estimating Risk for Future Intracranial, Fully Implanted, Modular Neuroprosthetic Systems: A Systematic Review of Hardware Complications in Clinical Deep Brain Stimulation and Experimental Human Intracortical Arrays. *Neuromodulation: Technology at the Neural Interface*, *23*(4), 411–426. https://doi.org/10.1111/NER.13069

Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., … Nicolelis, M. A. L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, *1*(2). https://doi.org/10.1371/journal.pbio.0000042

Castiello, U. (1999). Mechanisms of selection for the control of hand action. *Trends in Cognitive Sciences*, *3*(7), 264–271. https://doi.org/10.1016/S1364-6613(99)01346-7

Castiello, U. (2005). The neuroscience of grasping. *Nature Reviews Neuroscience 2005 6:9*, *6*(9), 726–736. https://doi.org/10.1038/nrn1744

Chadwick, E. K., Blana, D., Simeral, J. D., Lambrecht, J., Kim, S. P., Cornwell, A. S., … Kirsch, R. F. (2011). Continuous neuronal ensemble control of simulated arm reaching by a human with tetraplegia*. *Journal of Neural Engineering*, *8*(3), 034003. https://doi.org/10.1088/1741-2560/8/3/034003

Chao, Z. C., Nagasaka, Y., & Fujii, N. (2010). Long-term asynchronous decoding of arm motion using electrocorticographic signals in monkeys. *Frontiers in Neuroengineering*, *3*(MAR), 3. https://doi.org/10.3389/FNENG.2010.00003/

Chase, S. M., Schwartz, A. B., & Kass, R. E. (2009a). Bias, optimal linear estimation, and the differences between open-loop simulation and closed-loop performance of spiking-based brain-computer interface algorithms. *Neural Networks : The Official Journal of the International Neural Network Society*, *22*(9), 1203. https://doi.org/10.1016/J.NEUNET.2009.05.005

Chase, S. M., Schwartz, A. B., & Kass, R. E. (2009b). Bias, optimal linear estimation, and the differences between open-loop simulation and closed-loop performance of spiking-based brain-computer interface algorithms. *Neural Networks*, *22*(9), 1203–1213. https://doi.org/10.1016/j.neunet.2009.05.005

Chestek, C. A., Batista, A. P., Santhanam, G., Yu, B. M., Afshar, A., Cunningham, J. P., … Shenoy, K. V. (2007a). Single-neuron stability during repeated reaching in macaque premotor cortex. *Journal of Neuroscience*, *27*(40), 10742–10750. https://doi.org/10.1523/JNEUROSCI.0959-07.2007

Chestek, C. A., Batista, A. P., Santhanam, G., Yu, B. M., Afshar, A., Cunningham, J. P., … Shenoy, K. V. (2007b). Single-neuron stability during repeated reaching in macaque premotor cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *27*(40), 10742–10750. https://doi.org/10.1523/JNEUROSCI.0959-07.2007

Chestek, C. A., Gilja, V., Nuyujukian, P., Foster, J. D., Fan, J. M., Kaufman, M. T., … Shenoy, K. V. (2011). Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. *Journal of Neural Engineering*, *8*(4), 045005. https://doi.org/10.1088/1741-2560/8/4/045005

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., … Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature 2012 487:7405*, *487*(7405), 51–56. https://doi.org/10.1038/nature11129

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., & Shenoy, K. V. (2010). Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron*, *68*(3), 387–400. https://doi.org/10.1016/J.NEURON.2010.09.015/ATTACHMENT/49F68993-BB39-42F9-9825-A84597BA96F3/MMC1.PDF

Churchland, M. M., & Shenoy, K. V. (2007). Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *Journal of Neurophysiology*, *97*(6), 4235–4257. https://doi.org/10.1152/JN.00095.2007/ASSET/IMAGES/LARGE/Z9K0050782070017.JPEG

Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G., & Shenoy, K. V. (2006). Neural Variability in Premotor Cortex Provides a Signature of Motor Preparation. *Journal of Neuroscience*, *26*(14), 3697–3712. https://doi.org/10.1523/JNEUROSCI.3762-05.2006

Colachis IV, S. C., Dunlap, C. F., Annetta, N. V, Tamrakar, S. M., Bockbrader, M. A., & Friedenberg, D. A. (2021). Long-term intracortical microelectrode array performance in a human: a 5 year retrospective analysis. *Journal of Neural Engineering*, *18*(4), 0460d7. https://doi.org/10.1088/1741-2552/AC1ADD

Collinger, J. L., Boninger, M. L., Bruns, T. M., Curley, K., Wang, W., & Weber, D. J. (2013). Functional priorities, assistive technology, and brain-computer interfaces after spinal cord injury. *Journal of Rehabilitation Research and Development*, *50*(2), 145–160. https://doi.org/10.1682/JRRD.2011.11.0213

Collinger, J. L., Gaunt, R. A., & Schwartz, A. B. (2018). Progress towards restoring upper limb movement and sensation through intracortical brain-computer interfaces. *Current Opinion in Biomedical Engineering*, *8*, 84–92. https://doi.org/10.1016/J.COBME.2018.11.005

Collinger, J. L., Wodlinger, B., Downey, J. E., Wang, W., Tyler-Kabara, E. C., Weber, D. J., … Schwartz, A. B. (2013). High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet*. https://doi.org/10.1016/S0140-6736(12)61816-9

Cowley, B. R., Kaufman, M. T., Butler, Z. S., Churchland, M. M., Ryu, S. I., Shenoy, K. V., & Yu, B. M. (2013). DataHigh: graphical user interface for visualizing and interacting with high-dimensional neural activity. *Journal of Neural Engineering*, *10*(6), 066012. https://doi.org/10.1088/1741-2560/10/6/066012

Cunningham, J. P., & Yu, B. M. (2014, October 28). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, Vol. 17, pp. 1500–1509. https://doi.org/10.1038/nn.3776

Davare, M., Kraskov, A., Rothwell, J. C., & Lemon, R. N. (2011). Interactions between areas of the cortical grasping network. *Current Opinion in Neurobiology*, *21*(4), 565–570. https://doi.org/10.1016/J.CONB.2011.05.021

Davare, M., Lemon, R., & Olivier, E. (2008). Selective modulation of interactions between ventral premotor cortex and primary motor cortex during precision grasping in humans. *The Journal of Physiology*, *586*(11), 2735–2742. https://doi.org/10.1113/JPHYSIOL.2008.152603

Degenhart, A. D., Bishop, W. E., Oby, E. R., Tyler-Kabara, E. C., Chase, S. M., Batista, A. P., & Yu, B. M. (2020). Stabilization of a brain–computer interface via the alignment of low-dimensional spaces of neural activity. *Nature Biomedical Engineering*, 1–14. https://doi.org/10.1038/s41551-020-0542-9

Dekleva, B. M., Weiss, J. M., Boninger, M. L., & Collinger, J. L. (2021). Generalizable cursor click decoding using grasp-related neural transients. *Journal of Neural Engineering*, *18*(4), 0460e9. https://doi.org/10.1088/1741-2552/AC16B2

Dickey, A. S., Suminski, A., Amit, Y., & Hatsopoulos, N. G. (2009). Single-unit stability using chronically implanted multielectrode arrays. *Journal of Neurophysiology*, *102*(2), 1331–1339. https://doi.org/10.1152/JN.90920.2008/ASSET/IMAGES/LARGE/Z9K0080996310011.JPEG

Downey, J. E., Brane, L., Gaunt, R. A., Tyler-Kabara, E. C., Boninger, M. L., & Collinger, J. L. (2017). Motor cortical activity changes during neuroprosthetic-controlled object interaction. *Scientific Reports*. https://doi.org/10.1038/s41598-017-17222-3

Downey, J. E., Schwed, N., Chase, S. M., Schwartz, A. B., & Collinger, J. L. (2018). Intracortical recording stability in human brain-computer interface users. *Journal of Neural Engineering*. https://doi.org/10.1088/1741-2552/aab7a0

Dunlap, C. F., Colachis, S. C., Meyers, E. C., Bockbrader, M. A., & Friedenberg, D. A. (2020). Classifying Intracortical Brain-Machine Interface Signal Disruptions Based on System Performance and Applicable Compensatory Strategies: A Review. *Frontiers in Neurorobotics*, *14*, 76. https://doi.org/10.3389/FNBOT.2020.558987/BIBTEX

Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M., & Cunningham, J. P. (2016). Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications 2016 7:1*, *7*(1), 1–15. https://doi.org/10.1038/ncomms13239

Ethier, C., Oby, E. R., Bauman, M. J., & Miller, L. E. (2012). Restoration of grasp following paralysis through brain-controlled stimulation of muscles. *Nature 2012 485:7398*, *485*(7398), 368–371. https://doi.org/10.1038/nature10987

Evarts, E. V. (1968). Relation of pyramidal tract activity to force exerted during voluntary movement. *Https://Doi.Org/10.1152/Jn.1968.31.1.14*, *31*(1), 14–27. https://doi.org/10.1152/JN.1968.31.1.14

Fetz, E. E., & Cheney, P. D. (1980). Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells. *Https://Doi.Org/10.1152/Jn.1980.44.4.751*, *44*(4), 751–772. https://doi.org/10.1152/JN.1980.44.4.751

Fetz, Eberhard E. (1992). Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and Brain Sciences*, *15*(4), 679–690. Retrieved from https://depts-washington-edu.pitt.idm.oclc.org/fetzweb/assets/codingmvmtparam.pdf

Fine, J. M., & Hayden, B. Y. (2022). The whole prefrontal cortex is premotor cortex. *Philosophical Transactions of the Royal Society B*, *377*(1844), 20200524. https://doi.org/10.1098/RSTB.2020.0524

Flesher, S. N., Downey, J. E., Weiss, J. M., Hughes, C. L., Herrera, A. J., Tyler-Kabara, E. C., … Gaunt, R. A. (2021). A brain-computer interface that evokes tactile sensations improves robotic arm control. *Science*, *372*(6544), 831–836. https://doi.org/10.1126/SCIENCE.ABD0380/SUPPL_FILE/ABD0380S4.MP4

Flint, R. D., Wright, Z. A., Scheid, M. R., & Slutzky, M. W. (2013). Long term, stable brain machine interface performance using local field potentials and multiunit spikes. *Journal of Neural Engineering*, *10*(5), 056005. https://doi.org/10.1088/1741-2560/10/5/056005

Fraser, G. W., Chase, S. M., Whitford, A., & Schwartz, A. B. (2009). Control of a brain–computer interface without spike sorting. *Journal of Neural Engineering*, *6*(5), 055004. https://doi.org/10.1088/1741-2560/6/5/055004

Galaro, J. K., Celnik, P., & Chib, V. S. (2019). Motor Cortex Excitability Reflects the Subjective Value of Reward and Mediates Its Effects on Incentive-Motivated Performance. *Journal of Neuroscience*, *39*(7), 1236–1248. https://doi.org/10.1523/JNEUROSCI.1254-18.2018

Gallego, J. A., Makin, T. R., & McDougle, S. D. (2022). Going beyond primary motor cortex to improve brain–computer interfaces. *Trends in Neurosciences*, *45*(3), 176–183. https://doi.org/10.1016/J.TINS.2021.12.006

Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A., & Miller, L. E. (2020). Long-term stability of cortical population dynamics underlying consistent behavior. *Nature Neuroscience*, *23*(2), 260–270. https://doi.org/10.1038/s41593-019-0555-4

Gallego, J. A., Perich, M. G., Miller, L. E., & Solla, S. A. (2017, June 7). Neural Manifolds for the Control of Movement. *Neuron*, Vol. 94, pp. 978–984. https://doi.org/10.1016/j.neuron.2017.05.025

Gallego, J. A., Perich, M. G., Naufel, S. N., Ethier, C., Solla, S. A., & Miller, L. E. (2018). Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications*, *9*(1), 1–13. https://doi.org/10.1038/s41467-018-06560-z

Gao, P., & Ganguli, S. (2015). On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology*, *32*, 148–155. https://doi.org/10.1016/J.CONB.2015.04.003

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, *2*(11), 1527–1537. https://doi.org/10.1523/JNEUROSCI.02-11-01527.1982

Georgopoulos, Apostolos P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*. https://doi.org/10.1126/science.3749885

Ghasemi, A., & Zahediasl, S. (2012). Normality tests for statistical analysis: A guide for non-statisticians. *International Journal of Endocrinology and Metabolism*, *10*(2), 486–489. https://doi.org/10.5812/ijem.3505

Gilja, V., Pandarinath, C., Blabe, C. H., Nuyujukian, P., Simeral, J. D., Sarma, A. A., … Henderson, J. M. (2015a). Clinical translation of a high-performance neural prosthesis. *Nature Medicine*, *21*(10), 1142–1145. https://doi.org/10.1038/nm.3953

Gilja, V., Pandarinath, C., Blabe, C. H., Nuyujukian, P., Simeral, J. D., Sarma, A. A., … Henderson, J. M. (2015b). Clinical translation of a high-performance neural prosthesis. *Nature Medicine*, *21*(10), 1142–1145. https://doi.org/10.1038/nm.3953

Golub, M. D., Sadtler, P. T., Oby, E. R., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., … Yu, B. M. (2018). Learning by neural reassociation. *Nature Neuroscience*, *21*(4), 607–616. https://doi.org/10.1038/s41593-018-0095-3

Grol, M. J., Majdandžić, J., Stephan, K. E., Verhagen, L., Dijkerman, H. C., Bekkering, H., … Toni, I. (2007). Parieto-Frontal Connectivity during Visually Guided Grasping. *Journal of Neuroscience*, *27*(44), 11877–11887. https://doi.org/10.1523/JNEUROSCI.3923-07.2007

Hatsopoulos, N. G., Xu, Q., & Amit, Y. (2007). Encoding of Movement Fragments in the Motor Cortex. *Journal of Neuroscience*, *27*(19), 5105–5114. https://doi.org/10.1523/JNEUROSCI.3570-06.2007

Hennig, J. A., Oby, E. R., Golub, M. D., Bahureksa, L. A., Sadtler, P. T., Quick, K. M., … Yu, B. M. (2021). Learning is shaped by abrupt changes in neural engagement. *Nature Neuroscience 2021 24:5*, *24*(5), 727–736. https://doi.org/10.1038/s41593-021-00822-8

Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., … Donoghue, J. P. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, *485*(7398), 372–375. https://doi.org/10.1038/nature11076

Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., … Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature 2006 442:7099*, *442*(7099), 164–171. https://doi.org/10.1038/nature04970

Homer, M. L., Perge, J. A., Black, M. J., Harrison, M. T., Cash, S. S., & Hochberg, L. R. (2014). Adaptive offset correction for intracortical brain-computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *22*(2), 239–248. https://doi.org/10.1109/TNSRE.2013.2287768

Huggins, J. E., Moinuddin, A. A., Chiodo, A. E., & Wren, P. A. (2015). What Would Brain-Computer Interface Users Want: Opinions and Priorities of Potential Users With Spinal Cord Injury. *Archives of Physical Medicine and Rehabilitation*, *96*(3), S38-S45.e5. https://doi.org/10.1016/J.APMR.2014.05.028

Huggins, J. E., Wren, P. A., & Gruis, K. L. (2011). What would brain-computer interface users want? Opinions and priorities of potential users with amyotrophic lateral sclerosis. *Amyotrophic Lateral Sclerosis*, *12*(5), 318–324. https://doi.org/10.3109/17482968.2011.572978

Hughes, C., Herrera, A., Gaunt, R., & Collinger, J. (2020). Bidirectional brain-computer interfaces. *Handbook of Clinical Neurology*, *168*, 163–181. https://doi.org/10.1016/B978-0-444-63934-9.00013-5

Hughes, C. L., Flesher, S. N., Weiss, J. M., Downey, J. E., Boninger, M., Collinger, J. L., & Gaunt, R. A. (2021). Neural stimulation and recording performance in human sensorimotor cortex over 1500 days. *Journal of Neural Engineering*, *18*(4), 045012. https://doi.org/10.1088/1741-2552/AC18AD

Humphrey, D. R., Schmidt, E. M., & Thompson, W. D. (1970). Predicting Measures of Motor Performance from Multiple Cortical Spike Trains. *Science*, *170*(3959), 758–762. https://doi.org/10.1126/SCIENCE.170.3959.758

Jackson, S. R., Jackson, G. M., & Rosicky, J. (1995). Are non-relevant objects represented in working memory? The effect of non-target objects on reach and grasp kinematics. *Exp Brain Res*, *102*, 519–530.

Jafari, M., Aflalo, T., Chivukula, S., Kellis, S. S., Salas, M. A., Norman, S. L., … Andersen, R. A. (2020a). The human primary somatosensory cortex encodes imagined movement in the absence of sensory information. *Communications Biology 2020 3:1*, *3*(1), 1–7. https://doi.org/10.1038/s42003-020-01484-1

Jafari, M., Aflalo, T. N., Chivukula, S., Kellis, S. S., Salas, M. A., Norman, S. L., … Andersen, R. A. (2020b). Neural correlates of cognitive motor signals in primary somatosensory cortex. *BioRxiv*, 2020.04.20.041269. https://doi.org/10.1101/2020.04.20.041269

Jarosiewicz, B., Masse, N. Y., Bacher, D., Cash, S. S., Eskandar, E., Friehs, G., … Hochberg, L. R. (2013). Advantages of closed-loop calibration in intracortical brain-computer interfaces for people with tetraplegia. *Journal of Neural Engineering*, *10*(4). https://doi.org/10.1088/1741-2560/10/4/046012

Jarosiewicz, B., Sarma, A. A., Bacher, D., Masse, N. Y., Simeral, J. D., Sorice, B., … Hochberg, L. R. (2015). Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface. *Science Translational Medicine*, *7*(313). https://doi.org/10.1126/scitranslmed.aac7328

Jeannerod, M., Arbib, M. A., Rizzolatti, G., & Sakata, H. (1995). Grasping objects: the cortical mechanisms of visuomotor transformation. *Trends in Neurosciences*, *18*(7), 314–320. https://doi.org/10.1016/0166-2236(95)93921-J

Johansson, R. S., & Flanagan, J. R. (2009a). *Coding and use of tactile signals from the fingertips in object manipulation tasks*. https://doi.org/10.1038/nrn2621

Johansson, R. S., & Flanagan, J. R. (2009b). Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience 2009 10:5*, *10*(5), 345–359. https://doi.org/10.1038/nrn2621

Kalaska, J. F. (2009). From intention to action: Motor cortex and the control of reaching movements. *Advances in Experimental Medicine and Biology*, *629*, 139–178. https://doi.org/10.1007/978-0-387-77064-2_8/FIGURES/8_9_978-0-387-77064-2

Kao, J. C., Nuyujukian, P., Ryu, S. I., Churchland, M. M., Cunningham, J. P., & Shenoy, K. V. (2015). Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nature Communications 2015 6:1*, *6*(1), 1–12. https://doi.org/10.1038/ncomms8759

Karpowicz, B. M., Ali, Y. H., Wimalasena, L. N., Sedler, A. R., Keshtkaran, M. R., Bodkin, K., … Pandarinath, C. (2022). Stabilizing brain-computer interfaces through alignment of latent dynamics. *BioRxiv*, 2022.04.06.487388. https://doi.org/10.1101/2022.04.06.487388

Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience 2014 17:3*, *17*(3), 440–448. https://doi.org/10.1038/nn.3643

Kaufman, M. T., Seely, J. S., Sussillo, D., Ryu, S. I., Shenoy, K. V., & Churchland, M. M. (2016). The largest response component in the motor cortex reflects movement timing but not movement type. *ENeuro*, *3*(4), 85–101. https://doi.org/10.1523/ENEURO.0085-16.2016

Kennedy, P. (2011). Changes in emotional state modulate neuronal firing rates of human speech motor cortex: A case study in long-term recording. *Http://Dx.Doi.Org/10.1080/13554794.2010.532137*, *17*(5), 381–393. https://doi.org/10.1080/13554794.2010.532137

Kennedy, S. D., & Schwartz, A. B. (2019). Distributed processing of movement signaling. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(52), 26266–26273. https://doi.org/10.1073/PNAS.1902296116/ASSET/FED3738F-4A7B-45D7-A55C-47D415728680/ASSETS/GRAPHIC/PNAS.1902296116FIG07.JPEG

Kim, S. P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P., & Black, M. J. (2008). Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia*. *Journal of Neural Engineering*, *5*(4), 455. https://doi.org/10.1088/1741-2560/5/4/010

Kim, S. P., Wood, F., Fellows, M., Donoghue, J. P., & Black, M. J. (2006). Statistical analysis of the non-stationarity of neural population codes. *Proceedings of the First IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics, 2006, BioRob 2006*, *2006*, 811–816. https://doi.org/10.1109/BIOROB.2006.1639190

Levy, S., Lavzin, M., Benisty, H., Ghanayim, A., Dubin, U., Achvat, S., … Schiller, J. (2020). Cell-Type-Specific Outcome Representation in the Primary Motor Cortex. *Neuron*, *107*(5), 954-971.e9. https://doi.org/10.1016/J.NEURON.2020.06.006

Li, Z., O'doherty, J. E., Lebedev, M. A., & Nicolelis, M. A. L. (2011, December). Adaptive decoding for brain-machine interfaces through bayesian parameter updates. *Neural Computation*, Vol. 23, pp. 3162–3204. https://doi.org/10.1162/NECO_a_00207

Makin, J. G., O'Doherty, J. E., Cardoso, M. M. B., & Sabes, P. N. (2018). Superior arm-movement decoding from cortex with a new, unsupervised-learning algorithm. *Journal of Neural Engineering*, *15*(2), 026010. https://doi.org/10.1088/1741-2552/AA9E95

Marquardt, D. W. (1970). Generalized Inverses, Ridge Regression, Biased Linear Estimation, and Nonlinear Estimation. *Technometrics*, *12*(3), 591. https://doi.org/10.2307/1267205

McGinley, M. J., Vinck, M., Reimer, J., Batista-Brito, R., Zagha, E., Cadwell, C. R., … McCormick, D. A. (2015). Waking State: Rapid Variations Modulate Neural and Behavioral Responses. *Neuron*, *87*(6), 1143–1161. https://doi.org/10.1016/J.NEURON.2015.09.012

Milekovic, T., Sarma, A. A., Bacher, D., Simeral, J. D., Saab, J., Pandarinath, C., … Hochberg, L. R. (2018). Stable long-term BCI-enabled communication in ALS and locked-in syndrome using LFP signals. *Journal of Neurophysiology*, *120*(1), 343–360. https://doi.org/10.1152/JN.00493.2017/SUPPL_FILE/MOVIE

Moritz, C. T., Perlmutter, S. I., & Fetz, E. E. (2008). Direct control of paralysed muscles by cortical neurons. *Nature 2008 456:7222*, *456*(7222), 639–642. https://doi.org/10.1038/nature07418

Morrow, M. M., Jordan, L. R., & Miller, L. E. (2007). Direct comparison of the task-dependent discharge of M1 in hand space and muscle space. *Journal of Neurophysiology*, *97*(2), 1786–1798. https://doi.org/10.1152/JN.00150.2006/ASSET/IMAGES/LARGE/Z9K0020779640008.JPEG

Murata, A., Gallese, V., Luppino, G., Kaseda, M., & Sakata, H. (2000). Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area AIP. *Journal of Neurophysiology*, *83*(5), 2580–2601. https://doi.org/10.1152/JN.2000.83.5.2580/ASSET/IMAGES/LARGE/9K0400917019.JPEG

Mussa-Ivaldi, F. A. (1988). Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neuroscience Letters*, *91*(1), 106–111. https://doi.org/10.1016/0304-3940(88)90257-1

National Spinal Cord Injury Statistical Center. (2020). *Facts and Figures at a Glance*. Retrieved from www.msktc.org/sci/model-system-centers.

Nuyujukian, P., Albites Sanabria, J., Saab, J., Pandarinath, C., Jarosiewicz, B., Blabe, C. H., … Henderson, J. M. (2018). Cortical control of a tablet computer by people with paralysis. *PLOS ONE*, *13*(11), e0204566. https://doi.org/10.1371/JOURNAL.PONE.0204566

Nuyujukian, P., Kao, J. C., Fan, J. M., Stavisky, S. D., Ryu, S. I., & Shenoy, K. V. (2014). Performance sustaining intracortical neural prostheses. *Journal of Neural Engineering*, *11*(6), 066003. https://doi.org/10.1088/1741-2560/11/6/066003

Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., … Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences*, 201820296. https://doi.org/10.1073/pnas.1820296116

Oby, E. R., Perel, S., Sadtler, P. T., Ruff, D. A., Mischel, J. L., Montez, D. F., … Chase, S. M. (2016). Extracellular voltage threshold settings can be tuned for optimal encoding of movement and stimulus parameters. *Journal of Neural Engineering*, *13*(3), 036009. https://doi.org/10.1088/1741-2560/13/3/036009

Okun, M., Steinmetz, N. A., Cossell, L., Iacaruso, M. F., Ko, H., Barthó, P., … Harris, K. D. (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature 2015 521:7553*, *521*(7553), 511–515. https://doi.org/10.1038/nature14273

Omrani, M., Matthew, X., Kaufman, T., Hatsopoulos, N. G., & Cheney, P. D. (2017). Perspectives on classical controversies about the motor cortex. *J Neurophysiol*, *118*, 1828–1848. https://doi.org/10.1152/jn.00795.2016.-Primary

Orsborn, A. L., Moorman, H. G., Overduin, S. A., Shanechi, M. M., Dimitrov, D. F., & Carmena, J. M. (2014). Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron*, *82*(6), 1380–1393. https://doi.org/10.1016/j.neuron.2014.04.048

Pandarinath, C., Ames, K. C., Russo, A. A., Farshchian, A., Miller, L. E., Dyer, E. L., & Kao, J. C. (2018). Latent Factors and Dynamics in Motor Cortex and Their Application to Brain–Machine Interfaces. *Journal of Neuroscience*, *38*(44), 9390–9401. https://doi.org/10.1523/JNEUROSCI.1669-18.2018

Pandarinath, C., & Bensmaia, S. J. (2022). THE SCIENCE AND ENGINEERING BEHIND SENSITIZED BRAIN-CONTROLLED BIONIC HANDS. *Physiological Reviews*, *102*(2), 551–604. https://doi.org/10.1152/PHYSREV.00034.2020/ASSET/IMAGES/MEDIUM/PRV-00034-2020R01.PNG

Pandarinath, C., Nuyujukian, P., Blabe, C. H., Sorice, B. L., Saab, J., Willett, F. R., … Henderson, J. M. (2017). High performance communication by people with paralysis using an intracortical brain-computer interface. *ELife*, *6*. https://doi.org/10.7554/eLife.18554

Pandarinath, C., O'Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., … Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, *15*(10), 805–815. https://doi.org/10.1038/s41592-018-0109-9

Peckham, P. H., Keith, M. W., Kilgore, K. L., Grill, J. H., Wuolle, K. S., Thrope, G. B., … Wiegner, A. (2001). Efficacy of an implanted neuroprosthesis for restoring hand grasp in tetraplegia: A multicenter study. *Archives of Physical Medicine and Rehabilitation*, *82*(10), 1380–1388. https://doi.org/10.1053/APMR.2001.25910

Perge, J. A., Homer, M. L., Malik, W. Q., Cash, S., Eskandar, E., Friehs, G., … Hochberg, L. R. (2013). Intra-day signal instabilities affect decoding performance in an intracortical neural interface system. *Journal of Neural Engineering*, *10*(3). https://doi.org/10.1088/1741-2560/10/3/036004

Perich, M. G., Conti, S., Badi, M., Bogaard, A., Barra, B., Wurth, S., … Milekovic, T. (2020). Motor cortical dynamics are shaped by multiple distinct subspaces during naturalistic behavior. *BioRxiv*, 2020.07.30.228767. https://doi.org/10.1101/2020.07.30.228767

Ragnarsson, K. T. (2007). Functional electrical stimulation after spinal cord injury: current use, therapeutic effects and future directions. *Spinal Cord 2008 46:4*, *46*(4), 255–274. https://doi.org/10.1038/sj.sc.3102091

Raichle, M. E., & Raichle, M. E. (2001). Searching for a baseline: Functional imaging and the resting human brain. *Nature Reviews Neuroscience*. https://doi.org/10.1038/35094500

Ramkumar, P., Dekleva, B., Cooler, S., Miller, L., & Kording, K. (2016). Premotor and Motor Cortices Encode Reward. *PLOS ONE*, *11*(8), e0160851. https://doi.org/10.1371/JOURNAL.PONE.0160851

Rasmussen, R. G., Schwartz, A., & Chase, S. M. (2017). Dynamic range adaptation in primary motor cortical populations. *ELife*. https://doi.org/10.7554/elife.21409

Rizzolatti, G., Camarda, R., Fogassi, L., Gentilucci, M., Luppino, G., & Matelli, M. (1988). Functional organization of inferior area 6 in the macaque monkey. *Experimental Brain Research 1988 71:3*, *71*(3), 491–507. https://doi.org/10.1007/BF00248742

Rizzolatti, Giacomo, & Luppino, G. (2001). The Cortical Motor System. *Neuron*, *31*(6), 889–901. https://doi.org/10.1016/S0896-6273(01)00423-8

Rouse, A. G., & Schieber, M. H. (2018). Condition-Dependent Neural Dimensions Progressively Shift during Reach to Grasp. *Cell Reports*. https://doi.org/10.1016/j.celrep.2018.11.057

Russo, A. A., Bittner, S. R., Perkins, S. M., Seely, J. S., London, B. M., Lara, A. H., … Churchland, M. M. (2018). Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron*, *97*(4), 953-966.e8. https://doi.org/10.1016/J.NEURON.2018.01.004

Sachs, N. A., Ruiz-Torres, R., Perreault, E. J., & Miller, L. E. (2015). Brain-state classification and a dual-state decoder dramatically improve the control of cursor movement through a brain-machine interface. *Journal of Neural Engineering*, *13*(1), 016009. https://doi.org/10.1088/1741-2560/13/1/016009

Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., … Batista, A. P. (2014). Neural constraints on learning. *Nature*. https://doi.org/10.1038/nature13665

Salinas, E., & Abbott, L. F. (1994). Vector reconstruction from firing rates. *Journal of Computational Neuroscience*, *1*(1–2), 89–107. https://doi.org/10.1007/BF00962720

Santhanam, G., Linderman, M. D., Gilja, V., Afshar, A., Ryu, S. I., Meng, T. H., & Shenoy, K. V. (2007). HermesB: A continuous neural recording system for freely behaving primates. *IEEE Transactions on Biomedical Engineering*, *54*(11), 2037–2050. https://doi.org/10.1109/TBME.2007.895753

Schaffelhofer, S., & Scherberger, H. (2016). Object vision to hand action in macaque parietal, premotor, and motor cortices. *ELife*, *5*(2016JULY). https://doi.org/10.7554/ELIFE.15278

Scott, S. H., & Kalaska, J. F. (2017). Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. *Journal of Neurophysiology*, *73*(6), 2563–2567. https://doi.org/10.1152/jn.1995.73.6.2563

Scott, Stephen H. (2008). Inconvenient Truths about neural processing in primary motor cortex. *The Journal of Physiology*, *586*(5), 1217–1224. https://doi.org/10.1113/JPHYSIOL.2007.146068

Sergio, L. E., & Kalaska, J. F. (1998). Changes in the temporal pattern of primary motor cortex activity in a directional isometric force versus limb movement task. *Journal of Neurophysiology*, *80*(3), 1577–1583. https://doi.org/10.1152/JN.1998.80.3.1577/ASSET/IMAGES/LARGE/JNP.SE52F3.JPEG

Shanechi, M. M., Hu, R. C., Powers, M., Wornell, G. W., Brown, E. N., & Williams, Z. M. (2012). Neural population partitioning and a concurrent brain-machine interface for sequential motor function. *Nature Neuroscience 2012 15:12*, *15*(12), 1715–1722. https://doi.org/10.1038/nn.3250

Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Http://Dx.Doi.Org.Pitt.Idm.Oclc.Org/10.1146/Annurev-Neuro-062111-150509*, *36*, 337–359. https://doi.org/10.1146/ANNUREV-NEURO-062111-150509

Simeral, J. D., Kim, S. P., Black, M. J., Donoghue, J. P., & Hochberg, L. R. (2011). Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array. *Journal of Neural Engineering*, *8*(2), 025027. https://doi.org/10.1088/1741-2560/8/2/025027

Simeral, John D, Hosman, T., Saab, J., Flesher, S. N., Vilela, M., Franco, B., … Reeves, B. (2021). Home Use of a Percutaneous Wireless Intracortical Brain-Computer Interface by Individuals With Tetraplegia. *IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING*, *68*(7). https://doi.org/10.1109/TBME.2021.3069119

Slutzky, M. W. (2019). Brain-Machine Interfaces: Powerful Tools for Clinical Treatment and Neuroscientific Investigations. *Neuroscientist*, *25*(2), 139–154. https://doi.org/10.1177/1073858418775355

Sponheim, C., Papadourakis, V., Collinger, J. L., Downey, J., Weiss, J., Pentousi, L., … Hatsopoulos, N. G. (2021). Longevity and reliability of chronic unit recordings using the Utah, intracortical multi-electrode arrays. *Journal of Neural Engineering*, *18*(6), 066044. https://doi.org/10.1088/1741-2552/AC3EAF

Stark, E., & Abeles, M. (2007). Predicting Movement from Multiunit Activity. *Journal of Neuroscience*, *27*(31), 8387–8394. https://doi.org/10.1523/JNEUROSCI.1321-07.2007

Steinmetz, P. N., Roy, A., Fitzgerald, P. J., Hsiao, S. S., Johnson, K. O., & Niebur, E. (2000). Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature 2000 404:6774*, *404*(6774), 187–190. https://doi.org/10.1038/35004588

Stevenson, I. H., & Kording, K. P. (2011). How advances in neural recording affect data analysis. *Nature Neuroscience 2011 14:2*, *14*(2), 139–142. https://doi.org/10.1038/nn.2731

Strick, P. L., Dum, R. P., & Rathelot, J. A. (2021). The Cortical Motor Areas and the Emergence of Motor Skills: A Neuroanatomical Perspective. *Https://Doi.Org/10.1146/Annurev-Neuro-070918-050216*, *44*, 425–447. https://doi.org/10.1146/ANNUREV-NEURO-070918-050216

Suner, S., Fellows, M. R., Vargas-Irwin, C., Nakata, G. K., & Donoghue, J. P. (2005). Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *13*(4), 524–541. https://doi.org/10.1109/TNSRE.2005.857687

Suresh, A. K., Goodman, J. M., Okorokova, E. V., Kaufman, M., Hatsopoulos, N. G., & Bensmaia, S. J. (2020). Neural population dynamics in motor cortex are different for reach and grasp. *ELife*, *9*, 1–16. https://doi.org/10.7554/ELIFE.58848

Sussillo, D., Churchland, M. M., Kaufman, M. T., & Shenoy, K. V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience 2015 18:7*, *18*(7), 1025–1033. https://doi.org/10.1038/nn.4042

Sussillo, D., Nuyujukian, P., Fan, J. M., Kao, J. C., Stavisky, S. D., Ryu, S., & Shenoy, K. (2012). A recurrent neural network for closed-loop intracortical brain–machine interface decoders. *Journal of Neural Engineering*, *9*(2), 026027. https://doi.org/10.1088/1741-2560/9/2/026027

Sussillo, D., Stavisky, S. D., Kao, J. C., Ryu, S. I., & Shenoy, K. V. (2016). Making brain–machine interfaces robust to future neural variability. *Nature Communications*, *7*(1), 13749. https://doi.org/10.1038/ncomms13749

Suway, S. B., Orellana, J., McMorland, A. J. C., Fraser, G. W., Liu, Z., Velliste, M., … Schwartz, A. B. (2018). Temporally Segmented Directionality in the Motor Cortex. *Cerebral Cortex*, *28*(7), 2326–2339. https://doi.org/10.1093/CERCOR/BHX133

Taylor, D. M., Tillery, S. I. H., & Schwartz, A. B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science*, *296*(5574), 1829–1832. https://doi.org/10.1126/SCIENCE.1070291/SUPPL_FILE/1070291S3.MOV

Thach, W. T. (1978). Correlation of neural discharge with pattern and force of muscular activity, joint position, and direction of intended next movement in motor cortex and cerebellum. *Https://Doi.Org/10.1152/Jn.1978.41.3.654*, *41*(3), 654–676. https://doi.org/10.1152/JN.1978.41.3.654

Trautmann, E. M., Stavisky, S. D., Lahiri, S., Ames, K. C., Kaufman, M. T., O'Shea, D. J., … Shenoy, K. V. (2019). Accurate Estimation of Neural Population Dynamics without Spike Sorting. *Neuron*, *103*(2), 292-308.e4. https://doi.org/10.1016/J.NEURON.2019.05.003/ATTACHMENT/0E1D0D1C-0CAC-47DF-96D8-55EDE84DD09F/MMC1.PDF

Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science*, *286*(5446), 1943–1946. https://doi.org/10.1126/SCIENCE.286.5446.1943/ASSET/8C46CB92-15B7-4940-9A83-E6F3D8C056A7/ASSETS/GRAPHIC/SE4798046003.JPEG

Vansteensel, M. J., Pels, E. G. M., Bleichner, M. G., Branco, M. P., Denison, T., Freudenburg, Z. V., … Ramsey, N. F. (2016). Fully Implanted Brain–Computer Interface in a Locked-In Patient with ALS. *New England Journal of Medicine*, *375*(21), 2060–2066. https://doi.org/10.1056/NEJMOA1608085/SUPPL_FILE/NEJMOA1608085_DISCLOSURES.PDF

Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature 2008 453:7198*, *453*(7198), 1098–1101. https://doi.org/10.1038/nature06996

Ventura, V. (2008). Spike Train Decoding Without Spike Sorting. *Neural Computation*, *20*(4), 923–963. https://doi.org/10.1162/NECO.2008.02-07-478

Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation Through Neural Population Dynamics. *Https://Doi-Org.Pitt.Idm.Oclc.Org/10.1146/Annurev-Neuro-092619-094115*, *43*, 249–275. https://doi.org/10.1146/ANNUREV-NEURO-092619-094115

Wang, W., Chan, S. S., Heldman, D. A., & Moran, D. W. (2007). Motor Cortical Representation of Position and Velocity During Reaching. *J Neuro-Physiol*, *97*, 4258–4270. https://doi.org/10.1152/jn.01180.2006

Weiss, J. M., Flesher, S. N., Franklin, R., Collinger, J. L., & Gaunt, R. A. (2018). Artifact-free recordings in human bidirectional brain–computer interfaces. *Journal of Neural Engineering*, *16*(1), 016002. https://doi.org/10.1088/1741-2552/AAE748

Weiss, J. M., Gaunt, R. A., Franklin, R., Boninger, M. L., & Collinger, J. L. (2019). Demonstration of a portable intracortical brain-computer interface. *Brain-Computer Interfaces*, *6*(4), 106–117. https://doi.org/10.1080/2326263X.2019.1709260

Wessberg, J., & Nicolelis, M. A. L. (2004). Optimizing a linear algorithm for real-time robotic control using chronic cortical ensemble recordings in monkeys. *Journal of Cognitive Neuroscience*, *16*(6), 1022–1035. https://doi.org/10.1162/0898929041502652

Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M., & Shenoy, K. V. (2021). High-performance brain-to-text communication via handwriting. *Nature 2021 593:7858*, *593*(7858), 249–254. https://doi.org/10.1038/s41586-021-03506-2

Willett, F. R., Pandarinath, C., Jarosiewicz, B., Murphy, B. A., Memberg, W. D., Blabe, C. H., … Ajiboye, A. B. (2016). Feedback control policies employed by people using intracortical brain–computer interfaces. *Journal of Neural Engineering*, *14*(1), 016001. https://doi.org/10.1088/1741-2560/14/1/016001

Wodlinger, B., Downey, J. E., Tyler-Kabara, E. C., Schwartz, A. B., Boninger, M. L., & Collinger, J. L. (2015). Ten-dimensional anthropomorphic arm control in a human brain-machine interface: Difficulties, solutions, and limitations. *Journal of Neural Engineering*, *12*(1). https://doi.org/10.1088/1741-2560/12/1/016011

Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., & Black, M. J. (2006). Bayesian Population Decoding of Motor Cortical Activity Using a Kalman Filter. *Neural Computation*, *18*(1), 80–118. https://doi.org/10.1162/089976606774841585

Wu, W., & Hatsopoulos, N. G. (2008). Real-time decoding of nonstationary neural activity in motor cortex. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *16*(3), 213–222. https://doi.org/10.1109/TNSRE.2008.922679

Yokoi, A., & Diedrichsen, J. (2019). Neural Organization of Hierarchical Motor Sequence Representations in the Human Neocortex. *Neuron*, *103*(6), 1178-1190.e7. https://doi.org/10.1016/J.NEURON.2019.06.017

Yu, B. M., Kemere, C., Santhanam, G., Afshar, A., Ryu, S. I., Meng, T. H., … Shenoy, K. V. (2007). Mixture of trajectory models for neural decoding of goal-directed movements. *Journal of Neurophysiology*, *97*(5), 3763–3780. https://doi.org/10.1152/JN.00482.2006/ASSET/IMAGES/LARGE/Z9K0050781670010.JPEG

Zhang, Y., & Chase, S. M. (2013). A stabilized dual Kalman filter for adaptive tracking of brain-computer interface decoding parameters. *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, 7100–7103. https://doi.org/10.1109/EMBC.2013.6611194