Investigating the structure of sensorimotor cortical computations

Ryan Canfield

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading committee:
Amy Orsborn, Chair
Greg Horwitz
Eric Chudler

Program Authorized to Offer Degree:
Department of Bioengineering

University of Washington

Abstract

Investigating the structure of sensorimotor cortical computations

Ryan Canfield

Chair of the Supervisory Committee:
Professor Amy Orsborn
Electrical & Computer Engineering

Cortical computations underlie how humans interact with their environment through sensory information processing and movement execution, but our understanding of how they are implemented is limited. Computations are mediated by the structured transfer and transformation of information between specific neural populations with precise timing, which is mechanistically facilitated by different cell-types. To understand the structure underlying cortical computations we must understand the spatiotemporal structure, which can be used to make precise hypotheses about the mechanistic role of specific cell-types. This thesis aims to investigate the structure of cortical computations through a spatial, temporal, and transcriptomic lens. I investigated the spatiotemporal structure underlying motor control because the spatiotemporal structure that supports motor computations is not well established. I found that reach information was heterogeneously distributed across frontal motor cortex, yet the neural populations with the most task information also had neural activity patterns that evolved most similarly in time. Next, I assessed the spatiotemporal structure underlying how movement

information transitions between two motor computations: planning and execution. I found spatially distinct, computation-specific neural populations that coexisted with a separate, spatially distributed population involved in both planning and execution. I also investigated the transcriptomic structure of cortical computations in visual cortex, where the spatiotemporal structure is well established. I tested and validated new optogenetic tools to specifically manipulate parvalbumin-expressing (PV+) neurons. Using these tools, I found that PV+ neurons modulated visual sensitivity. Collectively the results from this thesis shed light on cortical computations in areas without a well-understood spatiotemporal structure and begin revealing the role of transcriptomically defined cell types in areas where the spatiotemporal structure is established. Further work will be necessary to integrate novel optogenetic tools that target specific cell types with investigations into cortical computations to reveal how they are mechanistically implemented.

Acknowledgements

It is hard to believe that I decided to pursue a doctorate six and a half years ago. The journey was hard and felt never-ending at times, but looking back, I would not change anything. Every good and bad moment shaped me into the scientist and person I am today. But, the main reason I could persevere and grow through the difficult times was due to the support I received from my advisors, lab colleagues, and family.

Thank you to my co-advisors, Amy Orsborn and Greg Horwitz, for their persistent belief in me and willingness to train me in all aspects necessary to succeed as a scientist. I'm particularly thankful for their patience as they trained me (a mechanical engineer with no biology background), to conduct electrophysiology experiments, perform IHC, and assist with neural surgeries. Then, you both trusted me to help implement a new and complex electrophysiological technique! Greg, I am so thankful for your engagement and feedback as we sat in your office and I ran half-baked ideas by you. You always made time for me and provided a valuable new perspective on my motor research ideas. You have played a significant role in shaping and improving the projects in this dissertation, and how I think about neuroscience questions. Amy, I will forever be thankful for your constant support and encouragement during my time as a PhD student. If I had any questions or concerns, I always felt comfortable bringing them to you. I was confident that you would guide me in the right direction, and you constantly did. You always provided the freedom to investigate and pursue any random scientific idea I found interesting before gently steering back on track. Your ability to communicate complex scientific ideas and concepts is an inspiration, and I so lucky that I have had the opportunity to learn from you. Greg and Amy, you have both significantly improved my life and I am so thankful that I had the opportunity to work, learn, problem solve, and think about brains with you.

I would also like to thank the other members of my committee. Thank you to Eric Chudler and Steve Perlmutter for their guidance and project feedback as I completed my doctoral research.

I am grateful for all of the members of both the Orsborn and Horwitz lab during my time, both past and present. We went through a lot of ups and downs together but learned and grew tremendously during this journey. It was a pleasure planning and conducting experiments with you. You have all been wonderful colleagues. I learned so much from working with you and appreciate the constant support as I finished my research.

Last but not least, I would like to thank my family. I want to thank my mom for providing valuable career guidance throughout my entire life; even before I knew it. You have been a role model my entire life. I saw that you loved your career, and wanted to work hard to succeed. I didn't know it at the time, but this provided the blueprint for how I envisioned my career. I was unsatisfied working in industry because I wanted what you had. Your work ethic has been inspiring. I frequently remember how you juggled career and family whenever I feel overworked. Success as a scientist is difficult, and you encountered more challenges than most. The way you held your values through hard times is also a tremendous inspiration, and something I try strive to achieve. I also want to thank my wife, Demi. Without you I'm not sure I would have had the fortitude to persevere through my doctoral research. You have always ensured that I kept my priorities straight. You constantly reminded me that if I was passionate about my research, that is all that mattered. You then led by example through your passion for your work. I am so thankful to have you, and our pup Bodie, in my life.

Table of Contents

List of figures:

Supplemental figures:

List of tables:

Publications from this thesis

Chapter 2: The spatiotemporal structure of neural activity during reaching

Manuscripts:

R. A. Canfield, T. Ouchi, H. Fang, B. Macagno, L. I. Smith, L. R. Scholl, A. L. Orsborn, “The spatiotemporal structure of neural activity in motor cortex during reaching,” bioRxiv, Oct 2025, doi: 10.1101/2025.10.17.683171 (under review at *Journal of Neuroscience*)

H. Fang, **R. A. Canfield**, T. Ouchi, B. Macagno, E. Shlizerman, A. L. Orsborn, “RPNT: Robust pre-trained neural transformer – a pathway for generalized motor decoding,” arXiv, Jan 2026, doi: https://doi.org/10.48550/arXiv.2601.17641

Conference abstracts:

R. A. Canfield, T. Ouchi, L. R. Scholl, P. Rajeswaran, L. I. Smith, A. L. Orsborn, “Single-neuron and population approaches reveal spatially structured neural dynamics across frontal motor cortex”, Poster presented at Neural Control of Movement (NCM), 2025 Apl 28 – May 2; Panama City, Panama

R. A. Canfield, T. Ouchi, L. R. Scholl, P. Rajeswaran, L. I. Smith, A. L. Orsborn, “Spatially Distributed Sampling of Motor Cortex During Reaching”, Poster presented at Society for Neuroscience (SFN), 2024 Oct 4-9; Chicago, Illinois

R. A. Canfield, L. R. Scholl, P. Rajeswaran, L. I. Smith, A. L. Orsborn, “Data-driven sampling of motor cortical networks during reaching”. Poster presented at Research in Encoding and Decoding of Neural Ensembles (AREADNE) 2024 Jun 25-29; Milos, Greece.

Chapter 3: The spatiotemporal structure of neural activity in motor cortex across motor computations

Conference abstracts

R. A. Canfield, T. Ouchi, B. Macagno, A. L. Orsborn, “A spatially distributed neural population facilitates motor computations through stable dynamics”, Poster presented at Society for Neuroscience (SFN), 2025 Nov 15 – 19; San Diego, California

R. A. Canfield, T. Ouchi, B. Macagno, A. L. Orsborn, “A spatially distributed neural population facilitates motor computations through stable dynamics”, Poster presented at Allen Lake Conference, 2025 Sep 28 – Oct 2; Renton, Washington

Chapter 4: The role of parvalbumin-expressing inhibitory neurons in visual sensitivity

Manuscripts

John K Mich, Smrithi Sunil, Nelson Johansen, Refugio A Martinez, Jiatai Liu, Bryan B Gore, Joseph T Mahoney, Mckaila Leytze, Yoav Ben-Simon, Darren Bertagnolli, Ravi Bhowmik, Yemeserach Bishaw, Krissy Brouner, Jazmin Campos, **Ryan Canfield**, Tamara Casper, Nicholas P Donadio, Nadezhda I Dotson, Tom Egdorf, Amanda Gary, Shane Gibson, Jeff Goldy, Erin L Groce, Kenta M Hagihara, Daniel Hirschstein, Han Hou, Will D Laird, Elizabeth Liang, Luke V Loftus, Nicholas Lusk, Jocelin Malone, Naomi X Martin, Deja Monet, Josh S Nagra, Dakota Newman, Nhan-Kiet Ngo, Paul A Olsen, Victoria Omstead, Ximena Opitz-Araya, Aaron Oster, Christina A Pom, Lydia Potekhina, Melissa Reding, Christine Rimorin, Augustin Ruiz, Adriana E Sedeño-Cortés, Nadiya V Shapovalova, Michael Taormina, Naz Taskin, Michael Tieu, Nasmil J Valera Cuevas, Sharon W Way, Natalie Weed, Vonn Wright, Zizhen Yao, Thomas Zhou, Delissa A McMillen, Michael Kunst, Medea McGraw, Bargavi Thyagarajan, Jack Waters, Trygve E Bakken, Nick Dee, Shenqin Yao, Kimberly A Smith, Karel Svoboda, Kaspar Podgorski, Yoshiko Kojima, Greg D Horwitz, Hongkui Zeng, Tanya L Daigle, Ed S Lein, Bosiljka Tasic, Jonathan T Ting, and Boaz P Levi., “Enhancer-AAVs allow genetic access to oligodendrocytes and diverse populations of astrocytes across species,” bioRxiv, Sep. 2023, doi: 10.1101/2023.09.20.558718

Conference abstracts

R. A. Canfield, S. D. Gibson, and G. D. Horwitz, “Selective modulation of parvalbumin-expressing neurons in macaque visual cortex,” Talk presented at the Biomedical Engineering Society Conference 2023.

Chapter 5:

Reviews

M. C. Dadarlat, **R. A. Canfield**, and A. L. Orsborn, “Neural Plasticity in Sensorimotor Brain–Machine Interfaces,” Annu. Rev. Biomed. Eng., vol. 25, no. 1, pp. 51–76, Jun. 2023, doi: 10.1146/annurev-bioeng-110220-110833.

Others:

Journal manuscripts

T. Ouchi, L. R. Scholl, P. Rajeswaran, **R. A. Canfield**, L. I. Smith, and A. L. Orsborn, “Mapping eye, arm, and reward information in frontal motor cortices using electrocorticography in non-human primates,” J. Neurosci., p. e1536242025, Jan. 2025, doi: 10.1523/JNEUROSCI.1536-24.2025.

N. Roth, G. Chapius, O. Winter, International Brain Laboratory, R. A. Ressmeyer, L. M. Bun, **R. A. Canfield**, G. D. Horwitz, N. A. Steinmetz, “A flexible quality metric for electrophysiological recordings across brain regions and species,” bioRxiv, Mar. 2026, doi: 10.64898/2026.03.06.710130

Perspectives

R. A. Canfield, A. L. Orsborn, and G. D. Horwitz, “Windows and periscopes into primate behavior,” Cell Rep., vol. 36, no. 3, p. 109435, Jul. 2021, doi: 10.1016/j.celrep.2021.109435.

Conference abstracts

R. A. Canfield, L. R. Scholl, P. Rajeswaran, L. I. Smith, A. L. Orsborn, Neural Population Dynamics in Large-Scale Electrocorticography (ECoG) During Reaching. Poster session presented at Neural Control of Movement (NCM) 2023 Apl 17-21; Victoria, BC.

1. Introduction

As humans, we interact with the world to achieve our goals. We react to elements in our environment and act upon them via our sensory and motor systems. To pick up your coffee cup, for example, you integrate visual information about the cup's location, shape, and size to develop a motor plan. Then, your brain produces the neural activity patterns necessary to activate your muscles in a precise spatiotemporal pattern to successfully execute the reach.

Cortical computations play an important role in these actions. Visual information is processed in visual cortex (Hubel and Wiesel, 1959; Tong, 2003), and neural activity that is modulated during reach planning and execution is found in pre-motor and primary motor cortical areas (Georgopoulos et al., 1982; Cisek and Kalaska, 2005; Churchland et al., 2006). However, not all neurons within each area contribute to these computations equally; there is structure (Penfield and Boldrey, 1937; Gattass et al., 1981; Chehade and Gharbawie, 2023; Ouchi et al., 2025). To understand how cortical computations are implemented we must understand this structure by investigating where information about computations is found (spatial), how it evolves over time (temporal), and how specific cell-types contribute (transcriptomic).

Developing a cohesive picture of cortical structure will likely require a two-pronged approach. We must independently understand the spatiotemporal structure that supports cortical computations and the unique anatomical and electrophysiological characteristics of specific cell-types (Tremblay et al., 2016; Bakken et al., 2021; Lee et al., 2023). Then, justifiable hypotheses about the roles of specific cell types in the observed spatiotemporal structure can be made. Experiments to address these hypotheses will help reveal a mechanistic understanding of how cortical computations are implemented.

A comprehensive understanding of the structure underlying cortical computations will be especially important for improving sensorimotor neuroprosthetics such as brain-computer interfaces (BCI). BCIs map neural activity into movement commands for an external device such as computer cursor or prosthetic limb. Implants targeted based on spatial structure can improve performance (Kunigk et al., 2024) and the temporal structure of recorded neural activity can impact performance (Sadtler et al., 2014; Oby et al., 2019, 2025). Further explorations into cortical structure will likely be integral to improve BCI performance and translatability.

This thesis aims to advance our understanding of structure in cortical computations by first focusing on the spatiotemporal structure of motor computations in frontal motor cortex, where it is not fully understood. Next, this thesis investigates transcriptomic structure in visual cortex where the spatiotemporal structure is well established. Last, this thesis examines how neuroprosthetic devices can be improved by understanding cortical structure.

1.1. Experimental tools to study cortical computation structure

Investigating cortical structure has remained difficult because of limited tools to record and manipulate neural activity in primates, whose diverse behavioral repertoire make them a valuable animal model to study cortical computations. Recent developments in primate neural implant technology allow the activity of hundreds of neurons to be simultaneously recorded with high spatial resolution, and for the activity of specific cell types to be precisely manipulated (El-Shamayleh et al., 2017; Jun et al., 2017; El-Shamayleh and Horwitz, 2019; Musk and Neuralink,

2019; Yang et al., 2019). These advances set the stage for new investigations into the structure underlying cortical computation.

1.1.1. Recording neural activity

An initial understanding of cortical structure was revealed even with a limited ability to sample neural activity. Historically, the neural activity of only a single-neuron was recorded while subjects performed simple tasks (Georgopoulos et al., 1982, 1986; Schieber and Hibbard, 1993) or were shown simple stimuli (Hubel and Wiesel, 1959). This approach yielded few neurons but provided flexibility to record neurons from across the cortex and introduced a valuable initial understanding of cortical spatiotemporal structure. However, parallels in structure across cortical areas were unclear from these early studies. Primary visual cortex exhibited a well-defined retinotopic organization (Gattass et al., 1981) but motor cortex had a complex spatial organization where neurons modulated by different movements were intermixed (Evarts, 1968; Schieber and Hibbard, 1993; Fetz, 1994; Kakei et al., 1999; Amirikian and Georgopoulos, 2003; Sergio et al., 2005). In visual cortex, neural responses were time-locked to stimulus onset (Talbot and Marshall, 1941; Gattass et al., 1981), but neurons in motor cortex displayed complex and varied temporal patterns (Churchland and Shenoy, 2007). Single-neuron recordings revealed clear spatiotemporal organization in visual cortex, but not in motor cortex.

Technological advances that allowed the coordinated activity of neural populations to be studied suggested new perspectives on the temporal structure underlying motor control. Typically, approximately 100 neurons were recorded simultaneously from microelectrode arrays (MEA) that were chronically implanted into frontal motor cortical areas. Analyzing these neural populations revealed temporal structure in their activity patterns but largely ignored spatial structure due to the physical constraints of the array (Churchland et al., 2012). MEAs record neurons with a grid of electrodes that have rigidly defined spacing and pre-defined lengths, which limits the spatial resolution of neuron sampling.

New developments of high-density laminar MEAs (e.g. Neuropixels probes) have allowed large populations of neurons to be flexibly recorded with high spatial resolution in monkeys (Jun et al., 2017; Musk and Neuralink, 2019; Yang et al., 2019). Leveraging these probes will be necessary to reveal the structure of cortical computations with high spatial resolution across the cortical surface and in depth.

1.1.2. Manipulating neural activity

Manipulating neural activity directly tests the cortical circuits that facilitate computations. This is commonly done by suppressing neural activity with pharmacological injections, or activating neurons with electrical stimulation (Passingham et al., 1983; Matsumura et al., 1991; Tong, 2003). Both suppression and stimulation-based approaches have identified cortical structure by revealing the locations where changes in neural activity impact behavior, suggesting a role in the corresponding computation (Penfield and Boldrey, 1937; Passingham et al., 1983; Tong, 2003). These manipulations, however, are limited in spatial and temporal precision making it difficult to fully understand structure at the single-neuron level.

Neural manipulation tools that allow experimenters to target specific neurons and neural populations with high temporal precision are quickly developing and will be necessary to characterize the structure of cortical computations. Advances in optogenetic tools that rely on transgenic animals have proven useful to study cortical computations in rodents. These tools

allow experimenters to selectively control neural activity with light and record from specific cell types with high temporal precision (Sohal et al., 2009; Chen et al., 2015; Estebanez et al., 2017; Mitani et al., 2018; Vendrell-Llopis et al., 2022). Optogenetically mediated holographic stimulation has further extended experimenter control by providing the capability to manipulate specific neurons and assess their impact on cortical circuits (Finkelstein et al., 2026). However, these tools have not fully translated to primates because transgenic monkeys are impractical.

Primate optogenetic tools are rapidly progressing and will soon become invaluable to reveal the transcriptomic structure underlying cortical computations. Recent developments promise the ability to target specific excitatory and inhibitory cell-types (Vormstein-Schneider et al., 2020; Mich et al., 2021; Yazdan-Shahmorad et al., 2024). This level of targeting will likely prove useful because different cell-types have unique electrophysiological and connectivity properties (Tremblay et al., 2016; Lee et al., 2023). Cell-type specific manipulation is accomplished in primates with viruses that drive opsin expression in a targeted cell-type. Expression of the opsin then allows experimenters to control the activity of those specific cell-types via light without transgenic animals (El-Shamayleh and Horwitz, 2019). Tools to selectively manipulate specific cell types will be critical to mechanistically understand how cortical computations are implemented.

1.2. Analytical tools to study cortical computation structure

Approaches to analyze the collective activity of neural populations have revealed temporal structure in neural activity patterns related to cortical computations (Churchland et al., 2012; Semedo et al., 2019). This is typically accomplished with dimensionality reduction tools such as principal component analysis (PCA) or factor analysis (Cunningham and Yu, 2014). These tools highlight coordinated activity patterns across neurons putatively reflecting the connectivity patterns between neurons (Sussillo et al., 2015). Many implementations pool all recorded neurons and largely overlook the contributions of individual neurons, different cortical areas, or how activity patterns may vary across computations. Neural population analyses, however, can be particularly useful in these scenarios to reveal computational similarities or differences by comparing activity patterns across distinct neural populations or cortical computations. Comparing the low-dimensional structure of distinct neural populations has revealed similarities and differences in temporal structural across cortical areas (Seely et al., 2016; Lara et al., 2018a; Semedo et al., 2019) and similarities across subjects (Safaie et al., 2023). Additionally, comparing the low-dimensional structure across computations, or days, has revealed similarities not captured by traditional single-unit analyses (Elsayed et al., 2016; Gallego et al., 2020). Coupling neural population analyses with the ability to sample neurons with high spatial resolution will be critical to understand how distinct neural populations contribute to the spatiotemporal structure of cortical computations.

1.3. Dissertation overview

This dissertation investigates the structure of cortical computations from different perspectives by leveraging high-density neural recordings and cell-type specific optogenetic tools. Chapter 2 and 3 investigate the spatiotemporal structure underlying motor computations in frontal motor cortex and Chapter 4 investigates transcriptomic structure in primary visual cortex where the spatiotemporal structure is established. Finally, this thesis discusses the translational

implications and importance of understanding the structure of cortical computations for novel therapeutic approaches, focused on brain-computer interfaces (BCI).

Chapter 2: The spatiotemporal structure of neural activity in motor cortex during reaching
This chapter investigates how movement information is spatially distributed across primate frontal motor cortices and how these spatially distributed neural populations coordinate during reaching. I used high-density neural recordings to highlight spatiotemporal complexity of motor representations across frontal motor cortex at the level of neurons and neural populations, where well-learned movements consistently recruit a spatially distributed subset of neurons.

Chapter 3: The spatiotemporal structure of neural activity in motor cortex across motor computations
This chapter extends the findings from chapter 2 across motor computations by comparing how task information is transferred and transformed across spatially distributed neural populations between movement planning and execution. I used high-density recordings to probe the balance of computation-specific and -shared representations. I find evidence of neural populations specific to planning and execution that coexist with a distinct computation-shared population that had similar neural activity patterns across planning and execution, suggesting it played a unique role in the transition.

Chapter 4: The role of parvlbumin-expressing inhibitory neurons in visual sensitivity.
This chapter presents the results of experiments to test and validate optogenetic tools that target parvalbumin-expressing (PV+) neurons in awake, behaving monkeys. I then use the tools to modulate PV+ neurons during a visual contrast detection task and find they play a role in visual sensitivity.

Chapter 5:
This chapter discusses the relevance and importance of understanding the neural circuits underlying cortical computations for improving neuroprosthetic-based rehabilitation therapies, focused on brain-computer interfaces (BCI). I discuss how different components of a BCI could be improved based on an understanding of cortical computation structure.

2. The spatiotemporal structure of neural activity in motor cortex during reaching

2.1. Introduction

Intracortical brain-computer interfaces (BCI) translate neural activity into movement and provide promising therapies for motor rehabilitation (Hochberg et al., 2006; Collinger et al., 2013; Silversmith et al., 2021). Current BCIs leverage existing knowledge about the structure of movement information in the brain to target implants (Dadarlat et al., 2023), often targeting areas where electrical stimulation evokes movement or using functional MRI mapping. These mapping approaches can guide targeting on the scale of several millimeters (Park et al., 2001; Kunigk et al., 2024), which approximately matches the spatial resolution of common implant technologies like multi-electrode arrays with fixed, millimeter-scale geometries (e.g., Utah Arrays are often 4 x 4 mm). New technologies now allow flexible probe targeting to specific neurons and neuronal populations (Jun et al., 2017; Musk and Neuralink, 2019; Yang et al., 2019). The spatiotemporal structure of single neuron and population activity in frontal motor cortices, however, has not been well-characterized, limiting our ability to leverage these technologies for BCIs.

Task information representations in frontal motor cortices appear to be both spatially and temporally complex. Many studies found that the cortical areas where electrical stimulation evoked movements did not have clearly defined spatial boundaries (Asanuma and Rosén, 1972; Park et al., 2001; Chehade and Gharbawie, 2023). For instance, areas that evoked finger movements were intermixed with areas that evoked wrist movement, consistent with anatomical projection tracing that suggests small areas in motor cortex have clustered intracortical connections (Huntley and Jones, 1991). Studies that map relationships between neural activity and behavior at multi-neuron resolution further show that task information is unevenly distributed across frontal motor cortical locations and that neural activity patterns from neighboring locations evolve differently in time (Chehade and Gharbawie, 2023; Ouchi et al., 2025). Single-unit activity during movement also shows a wide range of temporal relationships to reaching (Churchland and Shenoy, 2007), and temporal relationships between individual neurons and behavior suggest they may coordinate via functionally defined networks (Peters et al., 2014; Moore et al., 2024). However, most studies of motor representations measured neuronal populations across limited regions of cortex (Churchland and Shenoy, 2007; Peters et al., 2014; Moore et al., 2024) or mapped individual neuron properties across large cortical areas (Schieber and Hibbard, 1993; Georgopoulos et al., 2007), leaving open the question of how task-relevant information is temporally structured across neural populations in frontal motor cortex.

Understanding the spatiotemporal structure of task representations across motor cortices will likely benefit from neural population analyses, which recently revealed distinct temporal structure in neural activity patterns during reaching (Churchland et al., 2012; Gallego et al., 2018, 2020; Lara et al., 2018a). Existing population analyses, however, pooled neural activity across large areas within and sometimes across multiple motor cortical regions (Churchland et al., 2012; Gallego et al., 2020). Therefore, the spatial structure of population coordination during movement production remains unclear.

Here we examine target-direction information, one aspect of task information, is spatiotemporally organized across the motor cortical surface and depth during reaching. We used a high-density, laminar, micro-electrode array to make spatially precise single-unit recordings from frontal motor cortical neural populations as monkeys performed a well-learned reaching task. These populations span the cortical surface from dorsal pre-motor cortex (PMd) to the putative arm region of primary motor cortex (M1), and cortical depths of nearly 4mm.

Population decoding revealed that task information was unevenly distributed in depth and across the cortical surface. Analyzing population dynamics also showed that the temporal structure of neural populations varied with minimal relationship to spatial location. While spatial position was only marginally predictive of which neural populations had similar temporal structure, the task information represented within the population was strongly predictive. We used task information to define neural populations and found that similar dynamics were primarily driven by a subset of units with high task information rather than spatial location. Linking neural population dynamics to task information, further suggested that these coordinated dynamics contained similar task representations. Our results reveal new details of the spatiotemporal structure in frontal motor cortical representations consistent with emerging ideas of task-specific, spatially distributed functional networks. Our results suggest, in turn, that improving BCIs may require pairing advanced probe technologies with improved functional mapping.

2.2. Methods

2.2.1. Surgical procedures

All procedures were conducted in compliance with the NIH Guide for the Care and Use of Laboratory Animals and were approved by the Institutional Animal Care and Use Committee at the University of Washington. Two male rhesus macaques (*Macaca mulatta*) aged 10 years old (Monkey 1) and 11 years old (Monkey 2) were implanted with a multi-modal chamber stereotaxically targeted over left frontal motor cortex (Fig. 2.1C), using surgical procedures fully described in (Ouchi et al., 2025). The chamber leveraged a semi-chronic artificial dura (AD) that allowed visual and electrophysiological access to a 19 mm diameter window that included premotor cortex and primary motor cortex. A semi-chronic grid with 1 mm diameter holes was rigidly attached to the chamber and implanted over the AD to stabilize cortical movements and allow spatially precise intracortical recordings to repeatable locations.

2.2.2. Behavioral recordings and task

Monkeys performed a self-paced delayed center-out reaching task with targets arranged in a plane (Fig. 2.1A). Movements of the monkey’s head and left arm were restricted during reaches. A glove with retroreflective markers was placed on the monkey’s right hand to allow its 3-dimensional position to be tracked using motion tracking cameras (NaturalPoint, Inc., Corvallis OR) and mapped to the position of a cursor on a screen (Ouchi et al., 2025). Upward hand movements moved the cursor towards the top of the screen and rightward hand movements moved the cursor towards the right of the screen. Trials began with the appearance of a center target. After a brief center hold, a peripheral target appeared 6.5 cm from the center target initiating the delay period. Subjects were required to keep the cursor in the center target during a delay period, the length of which was randomly drawn from a uniform distribution (200 – 600 ms range). Disappearance of the center target (the "go cue") signaled animals to begin reaching to the peripheral target. A successful trial required a brief hold at the peripheral target, after

which a juice or water reward was delivered. Peripheral targets were presented in the pseudo-randomized order to ensure an equal number of successful trials to each target. Task parameters used for each monkey can be found in Table 2.1.

Table 2.1. Center-out task parameters

Parameter	Monkey 1	Monkey 2
Target radius (cm)	1.7	1.7-1.75
Center target hold duration (ms)	200	200
Delay period range (ms)	200 – 600	200 – 600
Peripheral target hold duration (ms)	175	100
Reward duration (ms)	100-200	200-250

2.2.3. Behavioral data analysis

Both monkeys were experts at this task and performed highly-stereotyped movements during all recording sessions. Unrewarded trials and trials exceeding a maximum time threshold were removed from analysis (mean ± SD: 36.9 ± 19.2 (3.7 ± 1.9 %) removed trials for Monkey 1, 18.3 ± 18.6 (2.6 ± 1.9 %) removed trials for Monkey 2). The maximum trial length threshold was calculated to make the trial length distribution approximately symmetric with the following equation:

$$T_{max} = T_{median} + (T_{median} - T_{min}) \quad (1)$$

where T is the trial length, in seconds, between when the cursor enters the center target to initiate the trial to after the trial is rewarded. Median, minimum, and maximum trial length statistics were computed by pooling rewarded trials across all sessions.

We quantified stereotypy of performance across days by correlating the cursor trajectory from 100ms before to 400ms after the cursor left the center target. Data were aligned to movement onset time, approximated by the time when the cursor left the center target. Trials were separated by target direction, then we computed the Pearson correlation coefficient between each trial and all other trials to the same target, across all days. Pearson correlation was computed separately for the X and Y trajectory components, then averaged, weighted by the component variance:

$$r_{traj} = \frac{\sum_{i=1}^{d} \sigma_i^2 r_i}{\sum_{i=1}^{d} \sigma_i^2} \quad (2)$$

where d is the number of trajectory components, σ_i^2 is the trajectory variance, and r_i is the Pearson correlation coefficient between two trajectories in the i^{th} component.

2.2.4. Electrophysiology

Neural recordings

We recorded neural activity using Neuropixels high-density laminar microelectrode arrays (IMEC, Leuven Belgium)(Jun et al., 2017). The Neuropixels probe was held by an oil hydraulic micromanipulator (Narishige International USA, Inc., Amityville, NY). A custom adaptor was designed in-house to securely mount the micromanipulator to the chamber and ensure the Neuropixels probe remained perpendicular to the implanted grid surface, which was approximately perpendicular to the cortical surface at the center of the chamber. We used the micromanipulator to manually translate the Neuropixels probe across the cortical surface until it was visually aligned with a selected grid hole. The probe was then slowly lowered with the micromanipulator to insert it into the brain.

We simultaneously recorded from 384 electrodes (of the 960 total electrode contacts on the Neuropixels probe) spanning 3.84 mm in depth during behavioral recordings. Our setup allowed us to insert approximately 4-5 mm (of the 10mm long Neuropixels probe) into the brain. We inserted probes and selected channels following procedures that aimed to align the top of the recorded electrode contacts with the cortical surface and to maximize the number of channels that recorded neural activity. We accomplished this in different ways for each animal. For Monkey 1, we recorded from the electrode contacts spanning the bottom 3.84 mm of the probe. We halted probe insertion once we recorded neurons within the top 0.2 mm of the recorded channels, even though the probe could have been physically advanced further. For Monkey 2, we inserted the probe as far as physically possible into the brain (approximately 4-5 mm), then took advantage of the ability to select the recorded electrodes online (in software). We based our channel selection on neural recordings from an initial baseline session where the 384 channels were configured to span 7.68 mm in depth. We used the rolling root-mean-square of the raw high-frequency (300-10,000 Hz; see below) signal computed by the data acquisition software (OpenEphys, Atlanta, Ga) across all electrodes, to estimate the location of the cortical surface using the empirical observation that electrodes not within the tissue would have clearly distinct properties. Without moving the probe, we reconfigured the channel selection to select 384 channels spanning 3.84 mm of depth such that the most superficial neurons recorded during the baseline session were located within the top 0.2 mm of the recorded channels, thus aiming to match the configuration of Monkey 1. Baseline recordings were used offline to validate the online surface position estimation.

Electrical activity from each Neuropixels channel was hardware filtered into a high-frequency (AP) band (300-10,000 Hz) and sampled at 30 kHz (Jun et al., 2017). Electrical activity was then streamed to a dedicated computer via an external PXIe system (National Instruments, Austin, TX). Data acquisition was controlled on the dedicated neural acquisition computer which saved recorded activity locally (OpenEphys, Atlanta, GA).

We recorded from 14 and 21 unique grid locations, or recording sites, on the cortical surface for Monkey 1 and 2 respectively (Figure 2.1D, E). We conducted these recordings on 16 and 30 non-consecutive days with multiple penetrations made at select grid locations. A majority of the data were collected using a single probe on each day except on four days with Monkey 2, where two probes were used to simultaneously record activity from different grid locations. Recording sessions were included in analyses using the following criteria: 1) they sampled neural activity from at least 2 mm along the Neuropixels probe, and 2) they contained at least 30 trials to each target that passed the trial length threshold. All analysis was conducted using 30 rewarded trials to each target.

Unit quality metrics

The raw AP band signals were post-processed to sort action potentials (spike sorting) using Kilosort 4.0 with default parameters (Pachitariu et al., 2024). Each putative neuron (unit) was then automatically assessed using quality metrics derived from the International Brain Laboratory (Table 2.2) (International Brain Laboratory, 2024). Units were only assessed during the trials selected for inclusion in the analysis. When there were more recorded trials than included in this analysis, we chose a start and end trial to qualitatively maximize stability. For most recordings, the chosen start trial was at the beginning of the behavioral session (median + SD: 0 + 0 trials for Monkey 1, 0 + 68 trials for Monkey 2). Units that passed all quality metrics were then manually inspected to ensure that the waveform did not clearly represent noise and that the firing rate did not contain sharp discontinuities across trials. Firing rate discontinuities were assumed to represent spike sorting errors and units exhibiting them were thus excluded. Only units that passed all quality metrics and manual inspection were used in further analysis.

We computed the drift range of each recording as a post-hoc proxy for recording quality. The drift range is the range of drift computed by Kilosort 4 during spike sorting (Pachitariu et al., 2024). Although Kilosort 4 corrects for drift, the drift range represents an estimation of Neuropixels probe movement relative to the brain and thus the recording instability.

We computed the signal-to-noise ratio (SNR) and true spike percentage for each unit as post-hoc proxies for sort quality. The SNR was defined as the ratio of the signal (average spike amplitude) over the “noise”. The recording noise was estimated by calculating the average median absolute deviation of the hardware filtered AP band data for channels at the same depth as the channel with the highest spike amplitude. The true spike percentage was defined as the percentage of spikes that were not refractory period violations (see Table 2.2).

Table 2.2. Unit quality metric description and acceptance thresholds

Quality Metric Name	Description	Acceptance Threshold
Presence Ratio	The proportion of trials the neuron is active on.	$= 1$ (All trials included in analysis)
Refractory Period Violations	The number of spikes that violate the refractory period of the previous spike.	$\leq 1\%$ of spikes under 1ms
Waveform Amplitude	Peak-to-peak amplitude of the waveform on the channel that recorded the maximum amplitude.	$\geq$ 50 μV
Template Amplitude	The distribution of the waveform template amplitudes used by Kilosort.	Low bin threshold: $\leq 10\%$ Noise cut-off value: ≥ 5

Firing rate estimation

Firing rates of neurons were estimated for all units that passed all quality metrics by binning spikes into 10 ms bins. Firing rates for each unit were then z-scored across all time and trials and smoothed using a Gaussian kernel with σ=50 ms and a total kernel width of 3σ. This process used trial segments from 500 ms before to 1 s after the movement onset time.

Direction tuning analysis

We estimated how much each unit's activity modulated with target direction across the trial by computing the modulation depth (MD). The MD was computed on the smoothed and normalized firing rate with a 300 ms moving window in 10 ms steps and defined as the difference between the maximum and minimum firing rate across target directions.

2.2.5. Decoding analysis

Linear decoding analysis

We quantified the task information represented in neural activity using target direction decoding analyses. Firing rates from 50 ms before to 250 ms after movement onset were averaged across time for each unit and each trial. This produced a matrix of neural activity with size (n_{Tr} ,n_{units}), where n_{Tr} and n_{units} represent the number of trials and units, respectively. Linear decoding was performed on the time-averaged firing rates with linear discriminant analysis (LDA) implemented using the Python library Scikit-learn (Pedregosa et al., 2011) with the eigen solver and automatic shrinkage. Reported decoding accuracies were estimated by averaging the prediction accuracy from 4-fold cross validation. Each fold used 180 randomly selected trials to train the model and the remaining 60 trials to test the model. Chance estimations of decoding accuracy were computed by randomly shuffling the target labels for each trial.

Nonlinear decoding analysis

Our primary results used linear decoding analyses. We supplemented this analysis with additional analyses using nonlinear decoding to quantify the task information as a control for the possibility that encoding nonlinearities may contribute to the observed spatial patterns of task information. Of the many possible non-linear decoding algorithms, we implemented the well-established latent factor analysis via dynamical system (LFADS) algorithm to classify target direction (Pandarinath et al., 2018). The LFADS model was implemented in TensorFlow and trained using an NVIDIA A40 GPU. We followed the same neural data preparation process as the linear decoding analysis (see *2.2.4 Electrophysiology*). To avoid the potential overfitting of nonlinear neural network models we implemented a 0.75/0.25 train/test split with dropout and L2 regularization. We repeated model training and testing for 50 random initializations that randomly selected the units for analysis and randomly assigned trials to the train or test group. The reported decoding accuracy was estimated by averaging the accuracy of 50 random repeats. The total nonlinear analysis took approximately two days (one day for the dataset from each monkey) and the hyperparameter details are listed in Table 2.3.

Table 2.3. Nonlinear decoding model parameters

Parameter	Value	Description
batch_size	16	Number of samples processed before model parameters update
num_layers	2	Number of hidden layers in the network
num_classes	8	Number of output targets for classification
num_epochs	300	Maximum number of training epochs
classifier_dims	64, 32	Dimensions of classifier layers
learning_rate	1e-3	Initial learning rate for optimization
hidden_dim	64	Dimension of hidden representations
n_factors	16	Number of factors
dropout_rate	0.4	Dropout probability for regularization

l2_reg	1e-3	L2 regularization coefficient
clip_value	20.0	Gradient clipping threshold
warmup_epochs	10	Number of warmup epochs for learning rate
min_warmup_lr	1e-7	Minimum learning rate during warmup
min_lr	1e-7	Minimum learning rate for scheduler
lr_patience	50	Learning rate scheduler epochs
lr_factor	0.95	Learning rate reduction factor

Pseudo-population analyses

Our wide sampling across known locations on the cortical surface, paired with laminar recordings from the Neuropixels probes, provided 3-dimensional sampling across a large volume of the brain with precise known locations. We started with unit-level analyses by quantifying decoding accuracy for each unit separately. We then leveraged the animals' consistent behavior across recording sessions to pool recorded units as pseudo-populations to quantify the spatial structure of task information and neural representations using decoding and population dynamics analyses (see below). To achieve our goals, we picked pseudo-population sizes that allowed us to assess the variation across neural populations. When feasible we looked at how trends were impacted by population size and saw that the relative trends persisted across a range of reasonable neural population sizes (Figure S2.5E).

We analyzed the structure of task information across the cortical surface by quantifying target decoding accuracy of neural populations grouped by their cortical surface location. Cortical surface location (i.e., recording site) was defined using the grid location (spatial resolution 1 mm). To control for ensemble size, all populations included 25 units (Table 2.4). When more than 25 units were recorded at a given recording site, 25 units were randomly sampled and used for analysis to ensure decoding accuracy was estimated using the same number of units for each group; if fewer than 25 units were recorded at a recording site it was excluded from further analysis. Random unit sampling and analysis were repeated 1000 times, which was used to estimate the 95% confidence intervals.

We controlled for firing rate differences across recording sites by sub-selecting the data to create similar firing rate distributions across recording sites. The distribution of the unit firing rates from each recording site was discretized with 10 Hz bins. The number of units selected from a bin for decoding accuracy calculations was set to the minimum number of units in each bin (n_{bin}) across all sites. If a recording site had more than n_{bin} units for a given firing rate bin, n_{bin} units were randomly selected. Units were randomly selected and used to estimate decoding accuracy 1000 times to estimate a confidence interval for each recording site. This procedure led to decoding accuracy estimates that used neural ensembles of 16 units across all sites (both Monkey 1 and Monkey 2) with empirically-matched firing rates (R^2: 0.96 for Monkey 1, 0.98 for Monkey 2; Fig. S2.2D).

We analyzed the structure of task information across cortical depths by quantifying target decoding accuracy of neural populations grouped by their location within the cortex relative to the estimated cortical surface. Groups of 25 units were also used when neural populations were organized by cortical depth. Overlapping depth bins were constructed to span 1 mm with a 0.25 mm step size. Units were selected using a similar strategy as described above to maintain a fixed ensemble size of 25 units (Table 2.4). Within each depth bin, units were randomly selected

from across the cortical surface, creating a pseudo-population which was used to compute decoding accuracy. Unit selection and decoding analysis was repeated 1000 times, which was used to estimate the 95% confidence interval.

We also analyzed the structure of task information across both cortical surface location and depth. For these analyses, we repeated the analyses to compute decoding accuracy as a function of depth as described above but adjusted the unit ensemble size to 10 units given the reduced dataset size (Table 2.4). If fewer than 10 units were contained at any given depth, the 3D location was excluded from analyses.

We analyzed the variability in task information from repeated Neuropixels probe penetrations by quantifying target decoding accuracy with units grouped by the penetration. In each monkey, multiple penetrations were made at select recording sites. To control for ensemble size, all populations included 15 units (Table 2.4). When more than 15 units were recorded at a given recording site, 15 units were randomly sampled and used for analysis. When fewer than 15 units were recorded the recording was omitted from analysis.

2.2.6. Neural population analysis

We applied canonical correlation analysis (CCA) to estimate the similarity of temporal dynamics between two neural populations, denoted as populations A and B. Neural populations A and B were created in a variety of ways (see below) and represent pseudo-populations in many cases. CCA was applied to Z-scored and smoothed firing rates (See *2.2.4 Electrophysiology* details) following procedures based on previous work (Gallego et al., 2018). Neural activity from 50 ms before to 250 ms after movement onset of individual trials were concatenated to create an array of size ($nt * nTr$,$nunit$). nt, nTr, and $nunit$ represent the numbers of time points within each trial, trials, and units in the neural population of interest, respectively. Principal Component Analysis (PCA) was applied to project the neural population firing rates onto the top 10 principal components (PC) and estimate the latent dynamics (L) for each neural population A and B, denoted as L_A and L_B, respectively. The number of PCs was chosen based on prior literature (Gallego et al., 2020; Safaie et al., 2023), but we found that this hyper-parameter choice did not notably change our findings (Figure S2.5A, B; Figure S2.7A, B). We also controlled for explained variance by selecting the number of PCs for each neural population that accounted for 60% of the variance, which did not lead to notable changes to our findings (Figure S2.5).

CCA was then applied to "align" the latent dynamics of each neuronal population by finding the new directions that maximize the pairwise correlation between latent trajectories, referred to as the "aligned" axes (Gallego et al., 2018). CCA produces a correlation value for each mode that represents the similarity of the latent trajectories between each neuronal population. We report the CCA correlation using the average computed across the top 4 modes to minimize the possibility of spurious correlations (Sussillo et al., 2015; Gallego et al., 2020). We also used CCA to project the latent trajectories of each neuronal population pair (L_A, L_B) onto the "aligned" axes. After L_A and L_B were each projected onto the "aligned" axis, we then used LDA decoders as described previously (see *2.2.5 Decoding analysis*) to estimate the "aligned decoding accuracy". We trained the decoder using all data from L_A and tested it on all data from L_B to assess decoder generalization to unseen data from a different neural population. The aligned decoding accuracy therefore quantifies the accuracy of a decoder trained on aligned neural activity from population A when tested on the aligned neural activity from population B for prediction.

Shared neural population dynamics between recording sites
We analyzed the spatial structure of population dynamics by performing CCA alignment analyses on different pseudo-populations following a similar approach as the linear decoding analyses (See *2.2.5 Decoding analysis*). Pseudo-populations with a fixed ensemble size were created by randomly selecting 25 units from each recording site (Table 2.4). We used the randomly selected units to perform CCA between the neural populations recorded at each pair of recording sites. Each reported CCA correlation quantifies the similarity of the temporal dynamics across the neural populations. We repeated this analysis 1000 times and randomly selected different units on each iteration to estimate the 95% confidence interval.

We performed three comparisons to assess patterns in CCA correlation between neural populations from different recording sites. First, we compared the absolute distance between recording site grid locations to the CCA correlation of each pair of neural populations (Fig. 2.3C, D). Second, we organized neural populations by their rostral-caudal axis location (Fig 2.3.E, F). To assess how the rostral-caudal location of an individual neural population is related to CCA correlation, we compared the CCA correlation between that neural population and its neighbors: the next most rostral neural population, and the next most caudal neural population. We averaged the CCA correlation between each pair of these three neural populations and report that as the Neighboring CCA correlation based on rostral-caudal position (Fig. 2.3G). Third, we organized neural populations by their decoding accuracy (See *2.2.5 Decoding analysis*; Fig. 2.3H, I). To assess how the task information of an individual neural population is related to CCA correlation, we compared the CCA correlation between that neural population and its neighbors: the neural population with the next lowest decoding accuracy and the neural population with the next highest decoding accuracy (See Fig. 2.2G, H). We averaged the CCA correlation between each pair of these three neural populations and report that as the Neighboring CCA correlation based on decoding accuracy (Fig. 2.3J).

As an additional control, we generalized the rostral-caudal CCA analyses above to look for structure along any axis on the cortical surface. We first defined 1,000 random axes on the cortical surface, then we projected the spatial position of each neural population onto each random axis to get their relative location on the new, random axis. We then regressed the neighboring CCA correlations with the distance along the (randomly chosen) spatial axis and computed the coefficient of determination ($R^2_{spatial_i}$).

We compared whether decoding or spatial organization better predicts the neural populations with similar dynamics by computing the Task Index for each spatial axis:

$$Task\ Index_i = R^2_{decoding} - R^2_{spatial_i} \quad (3)$$

where $R^2_{decoding}$ is the coefficient of determination of CCA with difference in task decoding. Positive Task Index values suggest that similar dynamics are best predicted by decoding accuracy while negative values suggest that similar dynamics are best predicted by spatial organization.

Shared neural population dynamics between task-tuned populations
We assessed the temporal dynamics of task-tuned units by creating neural pseudo-populations comprised of units with either low or high task information. We used the single-unit decoding accuracy (See *2.2.5 Decoding analysis*) to define the amount of task information encoded by

each recorded unit. The 100 units with the highest single-unit decoding accuracy were defined as the *unit-level high task information population*. The 100 units with the lowest single-unit decoding accuracy were correspondingly defined as the *unit-level low task information population*. These units were then arbitrarily split into two groups with low task information (50 units each) and two groups with high task information (50 units each) for further analysis (Table 2.4). We computed the CCA correlation (Fig. 2.4B,C) and aligned decoding performance between each combination of these unit groups (i.e., low-low, high-high, low-high; Figure 2.5B, C). Random unit splitting and analysis was repeated 1000 times to estimate a 95% confidence interval. We assessed the statistical significance of dynamic correlation estimates using CCA lower bounds (see *2.2.7 Statistical analysis*).

To ensure our results were not driven by how units were assigned to populations (i.e. high task information or low task information populations), we applied CCA to neural populations composed of random unit groups, defined as the random unit population. Two distinct neural populations were constructed by randomly selecting recorded units without replacement. We repeated random unit selection and CCA analysis 1000 times to estimate the 95% confidence interval (Fig. 2.4B, C).

We estimated the site-level dynamics similarity by performing CCA on neural populations composed of units recorded from the same site. We constructed two distinct populations of units recorded at the same site by randomly selecting units without replacement. We repeated random unit selection and CCA analysis 1000 times to estimate the 95% confidence interval (Fig. 2.4B, C).

We also replicated our CCA analyses restricting our data to simultaneously recorded neural populations to assess whether pseudo-population analyses underestimated or missed shared dynamics (e.g., due to non-task related neural activity that varies more from session to session). We restricted our analysis to simultaneously recorded units, then split units into high and low task information groups using the procedure described above. We used populations of 20 units to control for population size and recordings with fewer than 80 simultaneously recorded units were excluded from analysis. Random unit selection and CAA analysis was repeated 1000 times to estimate a 95% confidence interval.

2.2.7. Statistical analysis

CCA lower bound estimation

We sub-selected data from timepoints outside the task to assess the specificity of correlated dynamics between two neural populations to the task (Safaie et al., 2023). We used random timepoints selected from inter-trial interval data, which was defined as the time interval between the end of the previous trial and when the center target appeared to initiate the next trial session (mean ± SD: 0.41 ± 0.88 s for Monkey 1, 0.39 ± 0.72 s for Monkey 2). The neural activity of randomly selected units during these task intervals was concatenated and analyzed using the same procedure as if aligned to behavior (See *2.2.6 Neural population analysis*). Units were randomly chosen 1000 times to estimate the 95% confidence interval of CCA correlation.

Statistical significance estimation

We used bootstrapping to estimate the underlying statistical distributions of decoding, CCA correlation, and aligned decoding metrics. We considered two distributions significantly different if their 95% confidence intervals had zero overlap. Analyses results were considered not

statistically different from chance levels if the 95% confidence interval included chance performance. Chance performance for a classifier was assumed to be equivalent to random target selection (12.5% decoding accuracy for an 8-target task with approximately uniform target presentation).

We quantified the difference in decoding accuracy between depths at each recording site using the discriminability index (d') (Williams et al., 2013). The d' metric quantifies the separation between two statistical distributions (groups 1 and 2) with means μ_1, μ_2 and standard deviations σ_1, σ_2 via:

$$d' = \frac{\mu_1 - \mu_2}{\sqrt{\frac{1}{2}(\sigma_1^2 + \sigma_2^2)}}$$

The d' value represents how many standard deviations apart the means of the two distributions are, weighted for unequal variances. We used bootstrap resampling to estimate the underlying statistical distribution of decoding accuracy for two groups at each recording site: 1. units in the depth range with the highest decoding accuracy (See *Pseudo-population analysis*) and 2. all remaining units. We then computed the d' of these two distributions (Figure 2.2I). We also used bootstrap resampling to estimate a chance distribution of d' from each recording site. Units were randomly assigned to two groups (A, B), mirroring the 'high decoding depth' and 'remaining units' groups described above. Random unit assignment and the chance d' calculation were repeated 1000 times for each recording site. We tested if the empirical d' was significantly greater than chance by computing the one-sided p-value (alpha = 0.05) separately for each site. A significant p-value would therefore suggest that units in the 'high decoding depth' had significantly higher decoding accuracy than the remaining units. Recording sites with less than 10 units outside the chosen depth were excluded from this analysis.

Statistical analysis across multiple penetrations

We assessed the potential variability in our results that could be attributed to the random sampling of neurons at a given recording site inherent in electrophysiology by analyzing data collected from the same site across different sessions (i.e., repeated insertions of the Neuropixels probe into the same approximate location on different days). We compared the decoding accuracy difference between multiple penetrations within the same recording site to the decoding accuracy difference between penetrations at other recording sites (Figure S2.2E, F). We estimated the decoding accuracy difference for each pair of repeated penetrations at the same recording site with

$$Within\ site\ \Delta\mu_{i_1,i_2} = \left|\mu_{i_1} - \mu_{i_2}\right| \pm \sqrt{\sigma_{i_1}^2 + \sigma_{i_2}^2} \quad (4)$$

where μ_{i_1}, μ_{i_2} and σ_{i_1} σ_{i_2}represent the mean decoding accuracy and standard deviation, respectively, across unit resampling for two separate penetrations at recording site i_j.

We estimated the decoding accuracy difference across penetrations from different recording sites for penetration i_1 by computing the decoding accuracy difference between penetration i_1 and all other penetrations performed at different recording sites.

$$\overrightarrow{\Delta\mu_{i_1}} = \mu_{i_1} - \overrightarrow{\mu_{s \neq i}} \quad (5)$$

$$Across\ site\ \Delta\mu_{i_1} = mean(\overrightarrow{\Delta\mu_{i_1}} - \min(\overrightarrow{\Delta\mu_{i_1}})) \pm SD(\overrightarrow{\Delta\mu_{i_1}} - \min(\overrightarrow{\Delta\mu_{i_1}})) \quad (6)$$

In equation 5, μ_{i_1} represents the mean decoding accuracy from penetration 1 at recording site i across unit resampling, and $\overrightarrow{\mu_{s\neq i}}$ represents the vector of mean decoding accuracy across unit resampling from all penetrations at all recording sites, s, except when conducted at the same recording site i. We used equation 5 to compute the reported across-site decoding accuracy difference and standard deviation through equation 6, where $mean$ and SD indicate that the mean and the standard deviation were computed across all recording sites, s.

Task-responsive unit estimation

We estimated if individual units were modulated by the task (i.e. task responsive) by comparing the activity around movement onset with baseline activity on a single trial basis using a signal detection theory approach (Banerjee et al., 2010). Briefly, this approach computes the cumulative time-varying log-likelihood ratio (AccLLR) to test whether the single-trial firing rate (before normalization and smoothing) better fits the mean baseline activity or the mean movement onset activity. The activity from 500 ms to 100 ms before movement onset was considered the baseline window. The activity from movement onset to 400 ms after movement onset was considered the response window. We assessed whether a unit had a significant response on a trial by trial basis by comparing the AccLLR during the response window to a detection threshold calculated using receiver operator characteristic (ROC) analysis, which compared the probability of correct to incorrect response detection. Units were considered task-responsive if the area under the curve (AUC) from the ROC analysis was significantly greater than the null hypothesis of AUC = 0.5 using a one-sided Z-test with false discovery rate correction ($p \leq 0.05$).

Table 2.4. Analysis parameter sample size

Parameter	Samples
LDA decoding units analyzed per site	25 Units
LDA decoding units analyzed per depth	25 Units
LDA decoding units analyzed per depth per site	10 Units
LDA decoding units analyzed per penetration	15 Units
CCA units analyzed per recording site	25 Units
Task responsive CCA units analyzed per recording site	20 Units
CCA units analyzed per task information group	50 Units
Simultaneously recorded CCA units analyzed per task information group	20 Units
Iterations per random unit selection	1000 Iterations

2.3. Results

We trained two monkeys to perform a center-out reaching task (Fig. 2.1A). Both monkeys were proficient in the task and produced highly stereotyped movements (reach time mean ± SD: 0.19 ± 0.07 s for Monkey 1, and 0.18 ± 0.07 s for Monkey 2). The average pairwise correlation between the individual trial movement trajectories across all trials and all days was near 1 for both the 2D cursor position (mean ± SD: 0.96 ± 0.04 for Monkey 1, 0.94 ± 0.07 for Monkey 2; Fig. 2.1B, S2.1A) and 3D hand position (mean ± SD: 0.96 ± 0.05 for Monkey 1, 0.94 ± 0.07 for Monkey 2; Fig. S2.1B), confirming consistent reach trajectories across all recording days.

We recorded single-unit extracellular activity from chambers implanted over frontal motor cortices while the monkeys performed center-out reaches (Fig. 2.1C; See *2.2.2. Behavioral recordings and task*). We used high-density laminar microelectrode arrays (Neuropixels) to make spatially precise recordings of neural populations from a large 3D cortical volume (Fig. 2.1D). Neural recordings covered a cortical surface area of 9.9 × 6.7 mm in Monkey 1 and 8.1 × 9.3 mm in Monkey 2 (rostral–caudal × medial–lateral; Figure 2.1D, E). Each probe penetration recorded activity from the cortical surface to a depth ranging from 2.0 - 3.8 mm. Collectively, we recorded 1,022 units from Monkey 1 and 1,441 units from Monkey 2. The behavior stereotypy, spatially precise recordings across a broad cortical region, and large quantity of units allowed us to group neurons into pseudo-populations, then compare the spatial organization of task related information and the temporal structure of neural dynamics across frontal motor cortex.

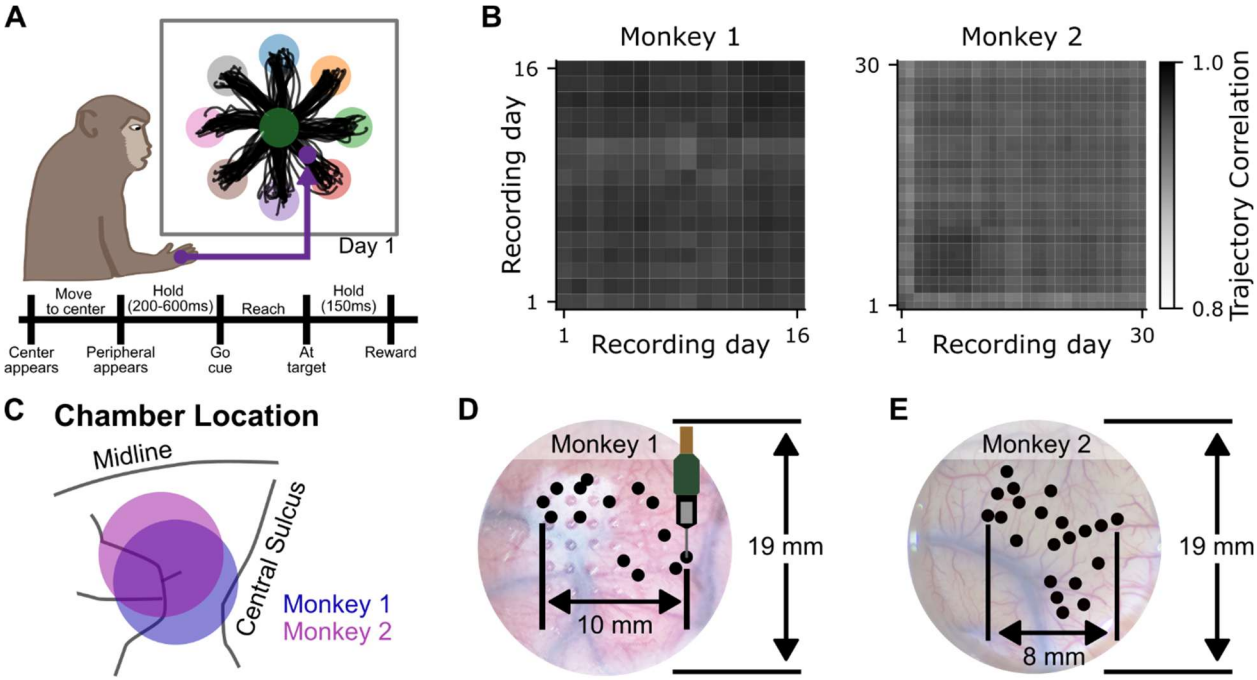

Figure 2.1. Large-scale recordings are made during highly stereotyped reaches. (A) *Schematic of the experimental setup. Hand position was tracked and mapped to the position of the cursor on a visual display. The monkey performed a center-out reaching task with similar trajectories on trials to the same target (Monkey 1). Bottom: trial structure.* ***(B)*** *Average cursor trajectory correlation across days for each monkey.* ***(C)*** *Chamber location for each monkey.* ***(D, E)*** *Neuropixels probe insertion locations (black dots) on the cortical surface (recording sites) for Monkey 1 and Monkey 2.*

2.3.1. Task information is heterogeneously distributed across frontal motor cortices

We first examined how much individual units modulated their activity with reach direction to assess the neural encoding of one key component of task information. Trial-averaged and target-specific peri-stimulus time histograms revealed clear differences in target direction modulation across units: some units showed task modulation, but qualitatively little difference in neural activity between targets (Fig. 2.2A), while other units had qualitatively distinct responses during reaches to different targets (Fig. 2.2B). We computed the modulation depth (MD) for each unit across time within the trial (See *2.2.4 Electrophysiology*) and found that these example units were representative of broader differences in target direction modulation across hundreds of other units (Fig. 2.2C, D). This qualitative evidence suggests that variability in the degree to which individual units modulated their activity with target direction, which we will refer to as task information, was present across all recorded units.

We quantified the distribution of task information across all recorded units by computing how well each individual unit predicted target direction (decoding accuracy; see *2.2.5* D*ecoding analysis*). The distribution of single-unit decoding accuracy was non-normal (Kolmogorov-Smirnov test: $p < 0.0001$ for Monkey 1, $p < 0.0001$ for Monkey 2) and showed a statistically distinct, rightward skew compared to the chance distribution (Skewness ± S.E.: 1.17 ± 0.08 vs. chance: 0.19 ± 0.08 for Monkey 1; 1.19 ± 0.06 vs. chance: 0.25 ± 0.06 for Monkey 2, Fig. 2.2E). Single-unit decoding accuracy was not strongly correlated with the average unit firing rate (R^2: 0.07 for Monkey 1, and 0.07 for Monkey 2; Fig. S2.2A) or proxies for spike-sorting quality like signal-to-noise ratio (SNR; Fig. S2.3A) or the percentage of true spikes based on refractory period violations (Fig. S2.3B). These results align with previous findings that many units in motor cortex are untuned (Georgopoulos et al., 1982; Best et al., 2016). We next leveraged our high density, spatially-precise recordings to investigate the spatial organization of task information across neural populations while controlling for population size.

We grouped units into neural populations based on the depth they were recorded from, irrespective of cortical surface location (i.e., recording site). Our dense sampling found that no single depth consistently encoded more task information than others (decoding accuracy confidence intervals (CI) overlapped across depths; see *2.2.7 Statistical analysis*; Fig. 2.2F), in contrast to previous reports suggesting depth-dependent patterns among more sparsely sampled neurons in motor cortices (Mollazadeh et al., 2011). We then regrouped units into neural populations defined by the recording site, independent of depth. Although all recording sites showed decoding accuracy above the 12.5% chance level, some locations had significantly higher decoding performance than others (no CI overlap across locations; see *Statistical analysis*; Fig. 2.2G, H). To rule out the possibility that these differences were due to our methods, which assume linear encoding, we repeated this analysis using a nonlinear neural network (see *2.2.5 Decoding analysis*; Fig. S2.4). We found a similar range of population decoding accuracy across sites with this alternate decoding approach. Additional controls confirmed that task information decoding was not strongly correlated with the total number of units (R^2: 0.20 for Monkey 1, and 0.01 for Monkey 2; Fig. S2.2B), the mean firing rate (R^2: 0.37 for Monkey 1, and 0.01 for Monkey 2; Fig. S2.2C), or proxies for recording quality like the recording drift range estimated by Kilosort 4 (R^2: 0.05 for Monkey 1, and 0.14 for Monkey 2; Fig. S2.3C), suggesting the decoding accuracy differences were primarily driven by recording site. Moreover, the decoding accuracy variability across recording sites was larger than that for repeated penetrations at the same recording site, for each of 5 recording sites with multiple penetrations (Fig S2.3D, E).

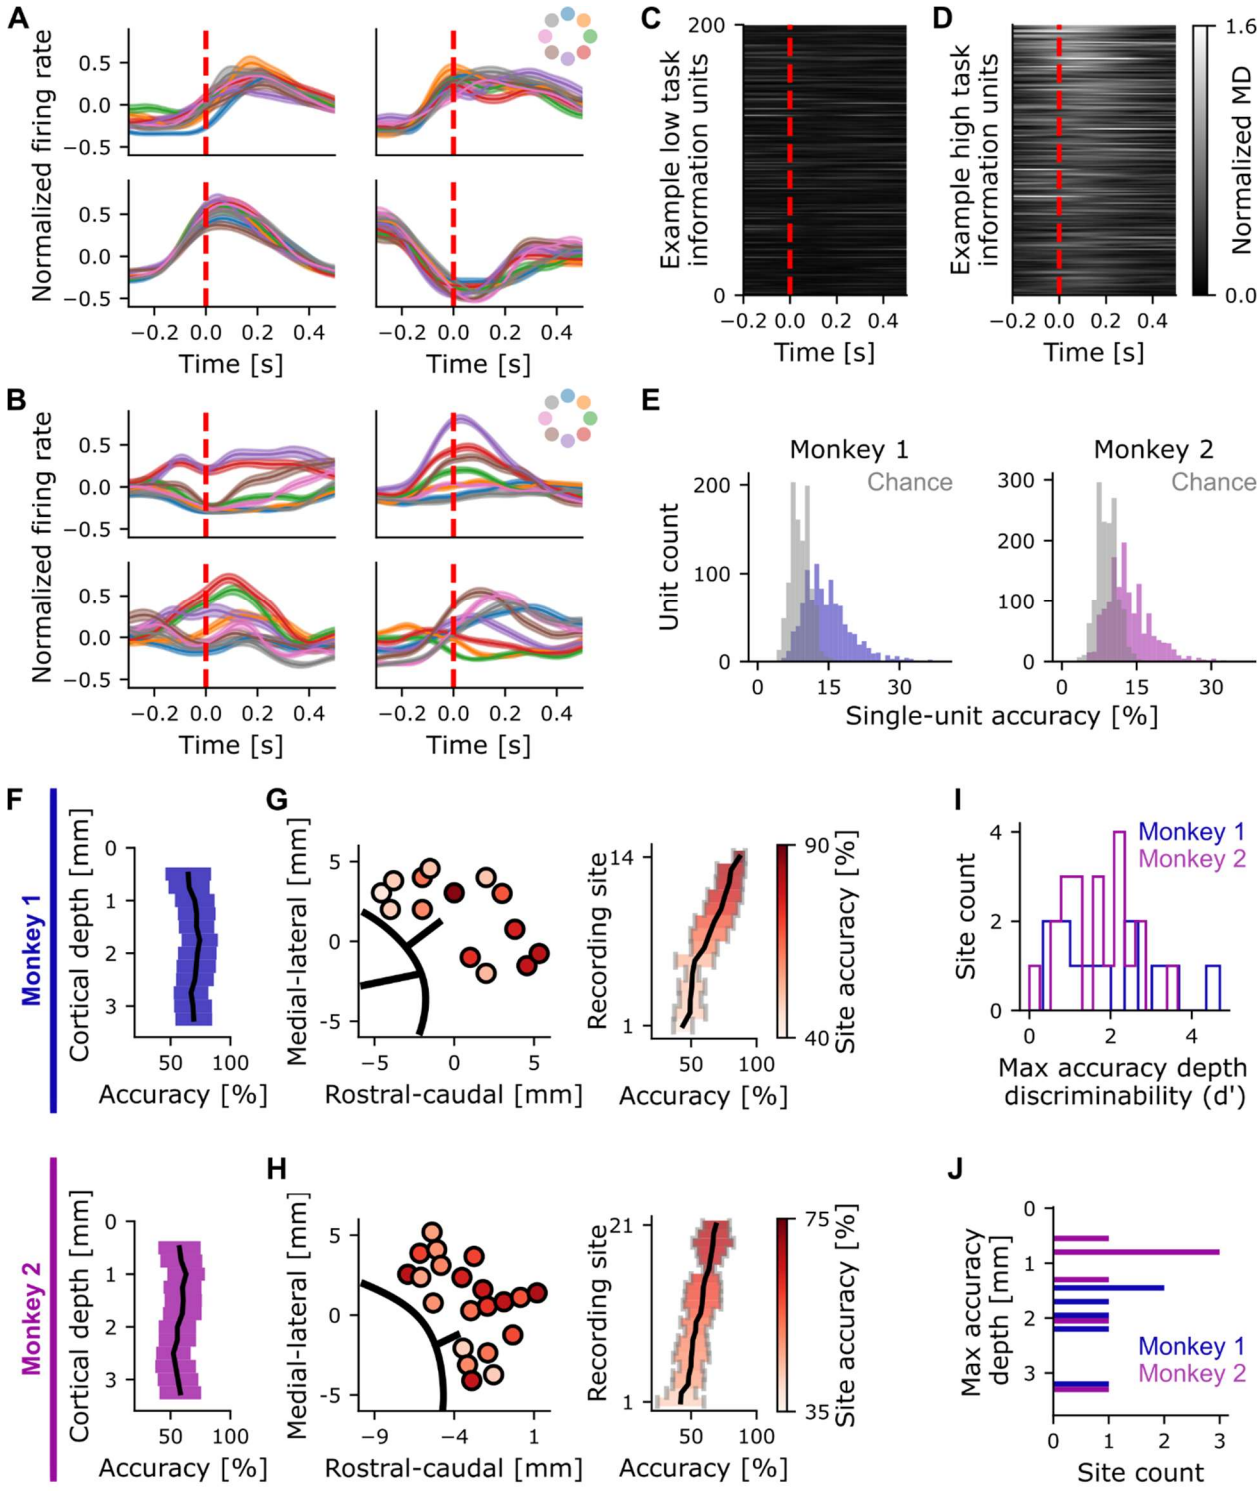

Figure 2.2. Task information is heterogeneously distributed across frontal motor cortices. (A) *The average response of four example units at movement onset of reaches towards each target with qualitatively little separation in the response across different targets (low task information). Each color represents the trial-averaged neural activity to different targets. Red dashes indicate the time of movement onset. Inset in the upper right represents the target that each color is assigned to. Solid lines indicate the mean with shading indicating the standard error of the mean across trials.* ***(B)*** *Same as (A) but for different example units with qualitatively large separation in response to reaches towards different targets (high task information, see 2.2.6 Neural population analysis).* ***(C)*** *Modulation depth (MD)*

(calculated using normalized firing rates) as a function of time for 200 units whose activity was the least predictive of the reach target direction (See 2.2.5 Decoding analysis) during movement. Firing rates were time-aligned to movement onset, indicated by the red dashed line. ***(D)*** *Same as (C) but for a different population of 200 example units whose activity was the most predictive of the reach target direction.* ***(E)*** *Distribution of single-unit decoding accuracy across all recorded units for each monkey (monkey 1, left, blue; monkey 2, right, pink) overlayed with the chance decoding distribution (grey).* ***(F)*** *Distribution of population decoding accuracy across depths. Units at the same depth (+/- 500 µm) across recordings were compiled into distinct neural populations from which decoding accuracy was estimated. Black lines indicate the mean; the shaded region represents the 95% confidence interval (CI).* ***(G)*** *Left: Distribution of population decoding accuracy across the cortical surface (site accuracy) for monkey 1. Units from the same recording site were compiled into distinct neural populations from which decoding accuracy was estimated. The black lines represent the arcuate sulcus. Right: Population decoding accuracy across recording sites, sorted by the mean decoding accuracy. Black lines indicate the mean; the shaded region represents the 95% CI. Colors indicate the site decoding accuracy (color bar) for both right and left panels.* ***(H)*** *Same as (G) for monkey 2.* ***(I)*** *Distribution of discriminability index (d') between the depth range with the maximum decoding accuracy and the remaining units (See 2.2.7 Statistical analysis) across recording sites (N=12 for monkey 1; N=21 for monkey 2).* ***(J)*** *Histogram of the depth (+/- 500 µm) where the highest decoding accuracy was observed for recording sites where decoding depth significantly varied across depths (N=6 for monkey 1; N=7 for monkey 2).*

Based on the variability in task information across recording sites, we revisited the influence of cortical depth, now analyzing depth-dependence individually for each recording site. We identified the 1 mm depth range with the highest decoding accuracy for each site, which we compared to the decoding accuracy obtained with all other units using the discriminability index (d') (See *Statistical significance estimation*). This metric captures the separation between these two distributions (in terms of standard deviations, weighted for unequal variance). We observed a range of discriminability across recording sites, suggesting that many sites had a depth with higher decoding than the remaining units (Fig. 2.2I). We tested whether this depth-dependence was statistically significant by comparing the observed d' to a chance distribution (See *2.2.7. Statistical analysis*). We found that some sites had minimal depth-dependence (N = 6 of 12 for Monkey 1, and N = 14 of 21 for Monkey 2). For sites with depth-dependence (N = 6 of 12 for Monkey 1, and N = 7 of 21 for Monkey 2), the highest decoding accuracy spanned multiple millimeters for each Monkey (Fig. 2.2J). Our observations reveal that the depth-decoding accuracy relationship varied across recording sites. Collectively, our results suggest that task information was heterogeneously distributed across both the motor cortical surface and in depth.

2.3.2. Task information predicts similar dynamics better than spatial position

The prior analyses suggest that task information was unevenly distributed at the level of single neurons. Temporal patterns of motor cortex activity during movement production similarly appeared unstructured at the level of single neurons (Churchland and Shenoy, 2007), yet have been shown to underly more coordinated dynamics at the level of neural populations (Churchland et al., 2012; Kaufman et al., 2014; Elsayed et al., 2016; Lara et al., 2018a). We therefore asked whether coordination at the level of population dynamics may shed light on the spatial heterogeneity of task information in individual neurons and neural populations.

We first compared population temporal dynamics across different neural populations, grouping neurons based on their recording site, which partially captured the spatial variation in task information at the single-unit level (see Fig. 2.2). We assessed the similarity of population dynamics using canonical correlation analysis (CCA) (see *Neural population analysis*; Fig. 2.3A). We estimated each neural population's latent dynamics and then compared the latent trajectories across each pair of neural populations with CCA. We found that the similarity of population dynamics notably varied across population pairs grouped by recording site (mean ± SD: 0.44 ± 0.05 for Monkey 1, 0.41 ± 0.03 for Monkey 2; Fig. 2.3B). Estimated CIs on

population similarity showed no overlap between the most (95% CI: [0.51, 0.61] for Monkey 1, [0.43, 0.54] for Monkey 2) and least (95% CI: [0.30, 0.38] for Monkey 1, [0.28, 0.38] for Monkey 2) correlated neural populations across the datasets, suggesting statistically meaningful differences. We note that our reported CCA values are lower than common reports in motor areas (Gallego et al., 2020; Safaie et al., 2023) due to our controls for neural population size, which influences CCA correlations (Fig. S2.5E).

We investigated three hypotheses to identify structure that might explain these seemingly heterogeneous temporal dynamics. First, we hypothesized that dynamics similarity was related to spatial proximity of neural populations (Fig. 2.3C), motivated by past observations that nearby neurons are likely to be anatomically connected (Gatter and Powell, 1978) and tend to have similar task information encoding properties (e.g., preferred directions (Amirikian and Georgopoulos, 2003)). We compared the pairwise CCA correlation of two neural populations to the absolute distance across the cortical surface separating those populations, which revealed no clear trends (R^2: 0.006 for Monkey 1, 0.014 for Monkey 2; Fig. 2.3D). Both near and distant neural populations had correlated dynamics, suggesting spatial proximity alone does not explain correlated dynamics.

Second, we hypothesized that dynamics similarity was related to the location of the neural populations along an anatomically-defined axis across the cortical surface (Fig. 2.3E). We focused on the rostral-caudal axis based on the anatomical and functional organization of brain regions in frontal motor cortices (Gatter and Powell, 1978; Riehle and Requin, 1989; Dum and Strick, 2002; Cisek and Kalaska, 2005). We found that neural populations organized by their location along the rostral-caudal axis showed some relationship to dynamics similarity, but the relationships were not consistent across monkeys. Caudal neural populations had qualitatively more similar dynamics than rostral neural populations, consistent with our hypothesis, in Monkey 1. However, this pattern was not observed in Monkey 2. (Fig. 2.3F). We quantified this pattern by computing the Neighboring CCA correlation, which computes how correlated the dynamics of a given neural population are to neural populations with nearby rostral-caudal positions (see *2.2.6 Neural population analysis*; Fig. 2.3G). This revealed a clear correlation between rostral-caudal location and population dynamics in Monkey 1 ($R^2 = 0.64$), but no trend in Monkey 2 ($R^2 = 0.01$), potentially related to slightly different chamber locations (see Fig. 2.1C). These results suggest that the rostral-caudal location of a neural population likely impacts how similar the dynamics of two neural populations are, yet the inconsistent results across monkeys suggest other factors are involved.

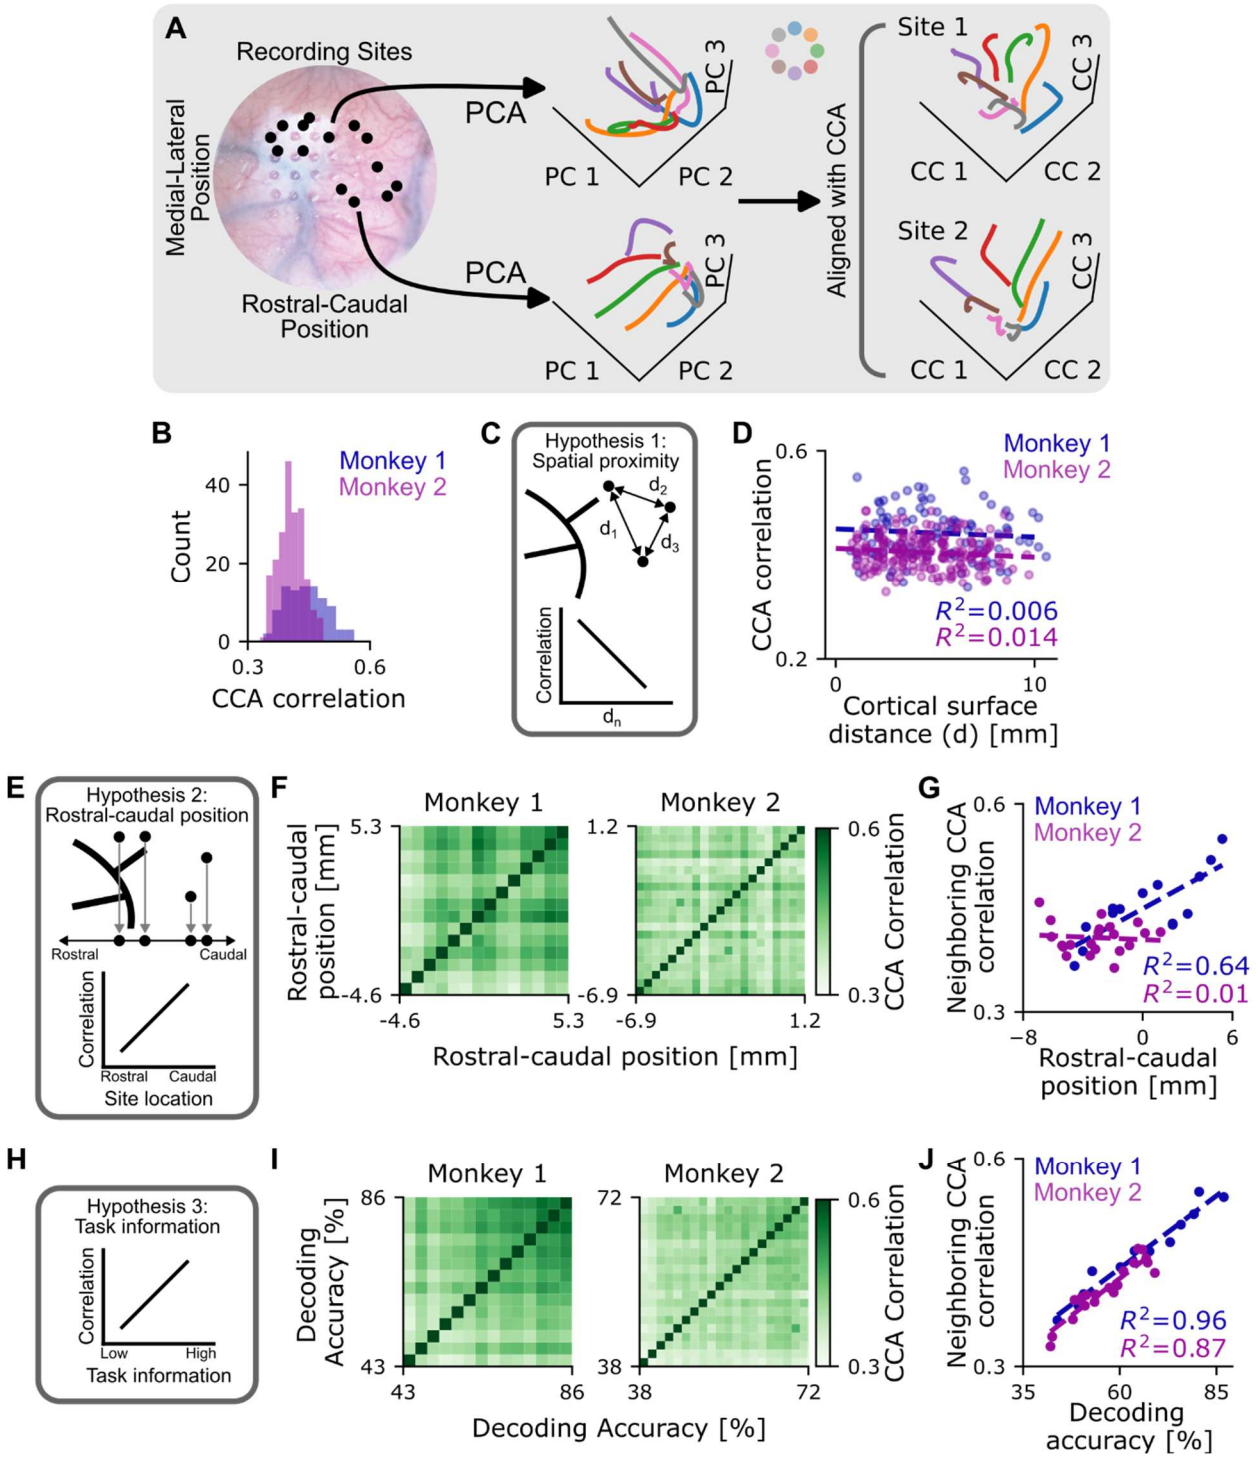

Figure 2.3. Task information predicts similar dynamics better than spatial position. (A) *Diagram of canonical correlation analysis (CCA) of latent dynamics to quantify their similarity. Recording site locations for Monkey 1 (left), the latent dynamics of distinct neural populations each composed of units recorded at different recording sites (middle), and the latent dynamics of each neural population projected onto the shared axes found with CCA (right). Each color represents the trial-averaged dynamics to a different target (inset).* ***(B)*** *The distribution of CCA correlation between pairs of neural populations from different recording sites.* ***(C)*** *Diagram depicting hypothesis 1. Each dot (top) represents a recording site and d_n represents the absolute distance between a pair of recording sites. This hypothesis predicts that the CCA correlation decreases as two neural populations become further apart (bottom).* ***(D)*** *Relationship between CCA correlation and the spatial proximity of two distinct neural populations for monkey 1 (blue)*

and monkey 2 (pink). Each dot represents a pair of neural populations (N=91 for monkey 1; N=210 for monkey 2); dashed lines represent the linear regression. ***(E)*** *Diagram depicting hypothesis 2, formatted as in (C). This hypothesis predicts that CCA correlation increases as location moves caudally along the rostral-caudal axis.* ***(F)*** *Heatmap of CCA correlation between neural populations organized by rostral-caudal position for Monkey 1 (left) and Monkey 2 (right). Each pixel represents a pair of neural populations (N=91 for monkey 1; N=210 for monkey 2).* ***(G)*** *Neighboring CCA correlation based on rostral-caudal position for each distinct neural population (see Neural population dynamics between recording sites) for monkey 1 (blue) and monkey 2 (pink). Each dot represents a recording site (N=14 for monkey 1; N=21 for monkey 2).* ***(H)*** *Diagram depicting hypothesis 3, formatted as in (C). This hypothesis predicts that CCA correlation increases as the amount of information contained within a population increases.* ***(I)*** *Heatmap of CCA correlation between distinct neural populations organized by the amount of task information for Monkey 1 (left) and Monkey 2 (right). Format as in (F).* ***(J)*** *Neighboring CCA correlation based on task information for each distinct neural population, formatted as in (G).*

Third, we hypothesized that dynamics similarity was related to task information encoding within the neural population (Fig. 2.3H). This hypothesis is consistent with populations of neurons coordinating to perform functional computations, potentially independent of spatial location. Such an architecture is supported by past observations of mixtures of spatially local and distributed task information in multi-neuron resolution functional mapping studies (Chehade and Gharbawie, 2023; Ouchi et al., 2025) and anatomical connections spanning multiple millimeters that could support spatially-distributed coordination (Gatter and Powell, 1978; Huntley and Jones, 1991). We organized neural populations by their task information and observed that neural populations with high task information had correlated dynamics while neural populations with low task information tended to have distinct dynamics (Fig. 2.3I). We quantified this pattern by computing the Neighboring CCA correlation with neural populations organized by task information, rather than spatial position, and saw a strong trend in both animals (R^2 = 0.96 for Monkey 1, R^2 = 0.87 for Monkey 2).This analysis highlights that spatially distributed neural populations had correlated dynamics if they were task modulated, suggesting that the temporal dynamics between neural populations is best predicted by the task information of the neural populations rather than spatial location. To validate that these results were not influenced by units that were unresponsive to the task, we performed controls to replicate these analyses with neural populations composed of only task-responsive units (See *2.2.7 Statistical analysis*). Sub-selecting task-responsive units did not qualitatively alter our findings, suggesting that task-responsiveness does not explain the patterns of dynamics similarity we observed (Fig. S2.6B-E). To ensure these results were not specific to the particular coordinate system we chose (i.e., anatomical axis), we repeated the spatial analysis using different axes across the cortical surface. We found that task information was always a better predictor of dynamics similarity than any axis across the cortical surface (Fig. S2.7C, D).

Collectively, our results revealed that the dynamics of neural populations varied, but similar dynamics were best predicted by the underlying variation in task information that they contained. Distinct neural populations had correlated dynamics if they had similarly high amounts of task information, even if the neural populations were spatially separated. The relationship between dynamics similarity and task information we observed does not appear to reflect a trivial correlation between these potentially-related metrics. Conceptually, correlated dynamics and task information can be independent (Fig. S2.8A-C). Empirically, we observed only weak correlations between task information and dynamics similarity when generating neural populations via random selection (R^2: 0.14 for Monkey 1, and 0.15 for Monkey 2; Fig. S2.8D, E).

2.3.3. Similar dynamics are driven by a subset of units with high task information

Our results suggest that task information played a role in producing correlated dynamics in distinct neural populations. However, there are multiple possible explanations for the correlation we observed between the site-level task information and population dynamics. For instance, this link could be at the unit level, where a select few highly task modulated units drove the coordinated dynamics, or at the site level, where all units at a recording site collectively contributed to coordinated dynamics. To distinguish these alternatives, we pooled all recorded units and then separated them only by their task information, which was estimated using single-unit decoding accuracy (see *2.2.6 Neural population analysis*; Fig. 2.4A). We created two pseudo-populations with high and low task information by grouping together the 100 units with the highest and lowest single-unit decoding accuracy, respectively. As an internal control, we further separated these unit-level high and low task information neural populations into two distinct populations (random assignment) to create two high and two low neural populations (four total). We analyzed how task information influenced neural population dynamics by comparing each neural population (Fig. 2.4A) while controlling for the number of units in each neural population (Fig. S2.5E).

We found that unit-level high task information populations had correlated dynamics (95% CI: [0.84, 0.87] for Monkey 1, [0.83, 0.85] for Monkey 2; High vs High in Fig. 2.4B). Unit-level low task information populations, in contrast, exhibited more distinct dynamics (95% CI: [0.38, 0.46] for Monkey 1, [0.36, 0.42] for Monkey 2; Low vs Low in Fig. 2.4B). Similarly, the dynamics of unit-level high task information populations were distinct from those for unit-level low task information populations (95% CI: [0.38, 0.47] for Monkey 1, [0.33, 0.42] for Monkey 2; Low vs High in Fig. 2.4B). The dynamics of unit-level high task information populations were also significantly more correlated than that of randomly selected populations (95% CI: [0.62, 0.72] for Monkey 1, [0.55, 0.67] for Monkey 2; Rand vs Rand in Fig. 2.4B). Units with high task information had correlated dynamics across higher CC modes, providing evidence that the correlated dynamics were unlikely to be driven by trivial correlations captured by the first few CC modes (Fig. 2.4C). The link between task information and dynamics is also consistent with the observation that spatial distributions of the units with similar dynamics and population decoding accuracy (see Fig. 2.2G, H) were qualitatively (Fig. S2.9A) and quantitatively (Fig. S2.9B) similar.

If the link between task information and correlated dynamics is indeed at the unit-level, we would expect that neural populations selected based on task information would have more similar dynamics than neural populations selected based on other factors, including recording site. To test this, we randomly selected neural populations composed of units recorded at the same recording site and compared their dynamics. We found that the dynamics similarity of neural populations selected based on recording site was significantly lower than neural populations selected based on task information (95% CI: [0.65, 0.75] for Monkey 1, [0.55, 0.72] for Monkey 2; Within site in Fig. 2.4B; Fig. 2.4B, C). We also controlled for the possibility that pseudo-population analyses could obscure shared dynamics (e.g., by failing to capture correlations related to non-task relevant signals that vary across sessions). We replicated our unit-level group analyses restricted to simultaneously recorded populations (which were also from the same recording site), and found that high task information populations consistently had significantly higher dynamics similarity than other sub-populations (no CI overlap across populations at any recording site; see *2.2.7. Statistical analysis*; Fig. 2.4D).

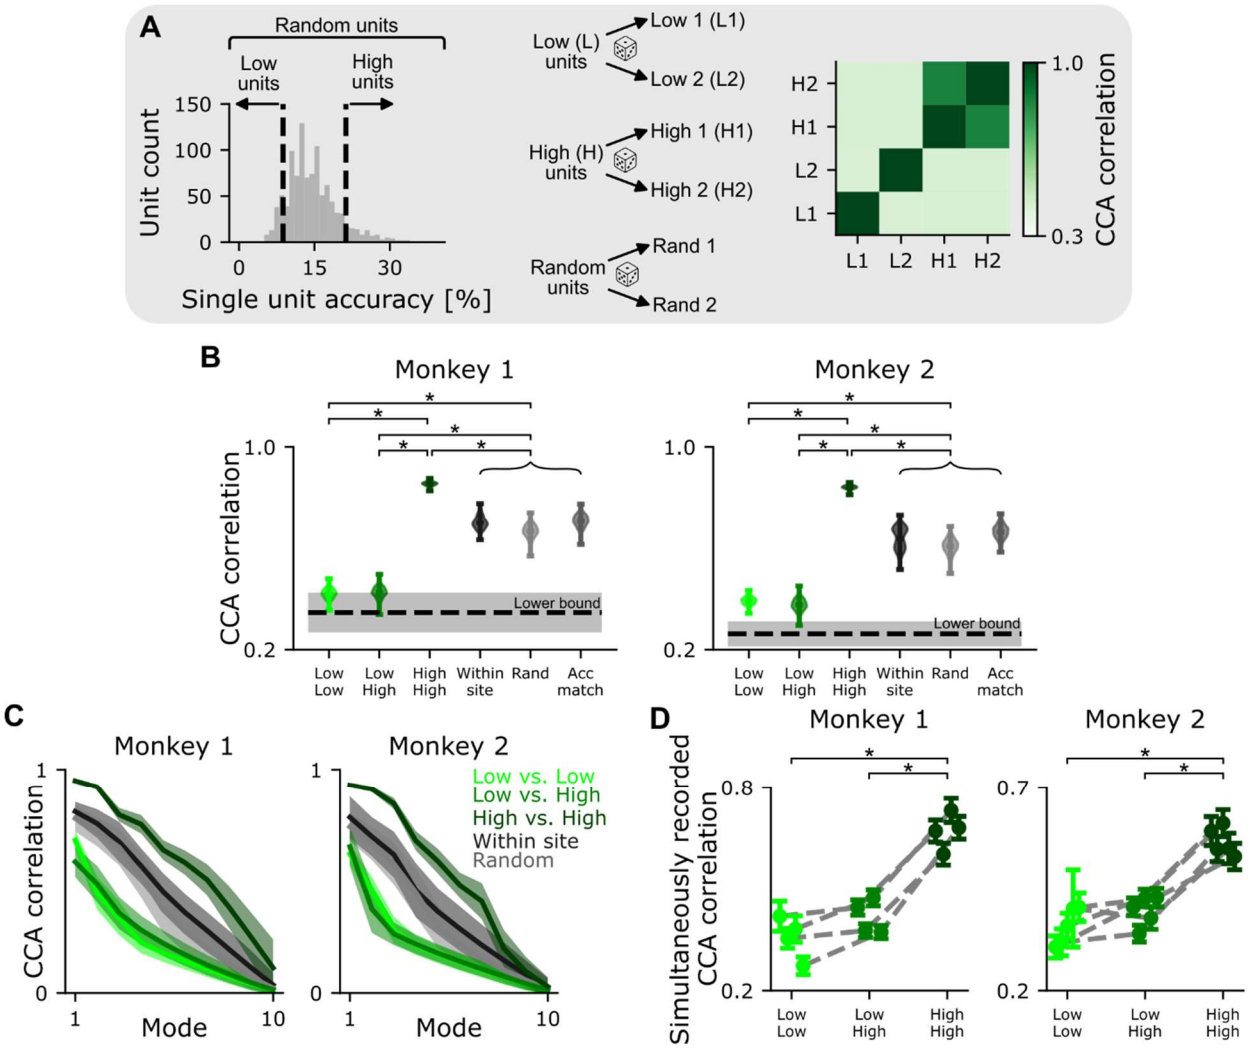

Figure 2.4. Neural populations composed of units with high task information have similar dynamics. (A) *Diagram depicting how distinct neural populations were composed and compared. Units were grouped based on their single-unit decoding accuracy into high (H) and low (L) task information groups (left). Each of these groups were randomly split into distinct neural populations as an internal control (L1, L2, H1, H2). We then compared the dynamics between each combination of these four distinct neural populations using CCA for all population combinations defined by task information (right). Neural populations composed of units drawn from all recorded units were also composed and compared (Rand).* ***(B)*** *CCA correlation between distinct neural populations defined by task information (L and H), recording site (Within site), random selection (Rand), or random selection while matching population decoding accuracy to the high task information populations (Acc match) for monkey 1 (left) and monkey 2 (right). The distribution for each comparison represents the variability from different random unit selections (N=1000). Each distribution depicts the CCA correlation results from all random unit selections and is centered on the mean. The black dashed line represents the lower bound estimated from inter-trial interval activity (see 2.2.7 Statistical analysis) and the gray shading represents the 95% CI across random unit selections.* ***(C)*** *CCA correlation as a function of CC mode for select neural population comparisons. The shaded regions represent the 95% confidence interval for each mode across random unit selections.* ***(D)*** *CCA correlation restricted to populations of simultaneously recorded units. Units were split based on their task information into new high and low task information groups for comparison (See Neural population analysis). The dashed line connects comparisons from the same recording site and error bars represent the 95% CI across random unit selections. Significance labels represent when all distributions, matched for recording site, were significantly different.*

The correlations we observed within high task-information neural populations were significantly larger during the task than during inter-trial intervals, a commonly used a lower-bound estimate on population correlations (see *2.2.7 Statistical analysis*; Fig. 2.4B). Low task-information

populations, in contrast, remained near the lower bounds, consistent with a link between task-related activity and correlated neural dynamics. Correlated dynamics within the unit-level high task-information population, however, were unlikely to be a trivial result of the relationship between task information and neural dynamics. As noted above, empirically, we found that CCA correlation and decoding accuracy were weakly correlated when neural populations were composed of randomly selected units. Yet, the dynamics similarities observed within our high task information sub-populations exceed those expected from this correlation alone (Fig. S2.8D,E). In fact, randomly selected neural populations that were decoding accuracy-matched to the high task information populations all had lower CCA correlations than the high task information sub-populations (95% CI: [0.66, 0.76] for Monkey 1, [0.61, 0.72] for Monkey 2; Acc match in Fig. 2.4B; S2.8F, G). Together, our results indicated that correlated dynamics of two neural populations were primarily driven by highly task modulated units.

2.3.4. Spatially distributed populations of units have similar task representations

Our previous analysis revealed that neural populations were more likely to have similar dynamics when they contain units with high task information. Yet, it remains unclear if those similar dynamics also reflect a shared task representation. If two neural populations have similar task representations, a decoder trained on one will generalize to the other. We therefore analyzed patterns of decoding generalization across neural populations defined by task information. Specifically, we aligned the dynamics of each neural population with CCA, then trained a decoder on one neural population and tested it on the other (Fig. 2.5A) (Gallego et al., 2020; Safaie et al., 2023).

We found that decoders trained on the aligned dynamics of unit-level high task information populations consistently generalized to other unit-level high task information populations (95% CI: [98.7, 100.5] for Monkey 1, [96.7, 100.2] for Monkey 2; High vs High in Fig. 2.5B), even though the units themselves were spatially distributed across frontal motor cortices (see Fig. 2.4D). This was not the case between two unit-level low task information populations, which performed at or near chance level (95% CI: [10.8, 21.5] for Monkey 1, [10.1, 18.8] for Monkey 2; Low vs Low in Fig. 2.5B). We also found that decoders trained on the aligned dynamics from low task information populations generalized above chance to high task information populations, though performance was variable across populations (95% CI: [54.9, 83.6] for Monkey 1, [49.3, 77.6] for Monkey 2; Low vs High in Fig. 2.5B). This is consistent with a distribution of task information (see Fig. 2.2) where even the neural populations with the lowest relative task information had enough information to train decoders that performed above chance level. These shared task representations did not appear to be a trivial result of correlated dynamics, which can be independent conceptually (Fig. S2.10A-C). Similar to our within-population decoding analyses, we found a weak correlation between CCA correlation and aligned decoding accuracy when neural populations were composed of randomly selected units (R2: 0.22 for Monkey 1, 0.27 for Monkey 2; Fig S2.10D). Yet, the dynamics similarities observed within our high task information sub-populations exceed that expected from this correlation alone (Fig. S2.10D,E).

Our results suggest that neural populations in frontal motor cortices contain a mixture of distinct and shared dynamics, and the exact proportion of distinct and shared dynamics vary across distinct populations. High task information populations primarily exhibited shared dynamics driven by a similar task representation. Other neural populations had a higher proportion of

distinct dynamics potentially from participating in other motor related computations and processes.

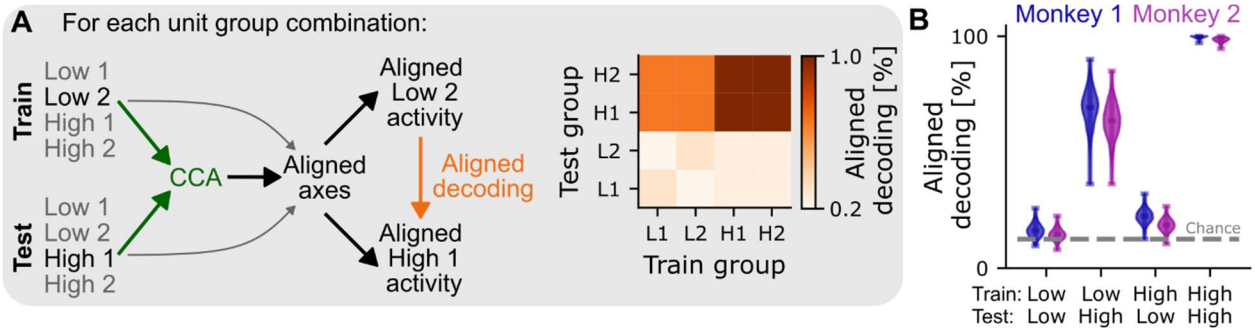

Figure 2.5. Neural populations composed of spatially distributed units with high task information have similar task representations. (A) *Diagram depicting analyses to assess whether neural populations have similar task representations in their latent dynamics (left). We applied CCA to the neural populations composed of units with either high or low task information to find the axes that best align the individual latent dynamics ("aligned axes"). We then projected the latent dynamics from each neural population onto the aligned axes (gray arrow), trained a decoder using aligned activity from one population, then tested the resulting decoder using aligned activity from the other population ("aligned decoding"). Heatmap representation of the resulting aligned decoding between all combinations of neural populations (right).* ***(B)*** *Aligned decoding for each neural population comparison. The distribution represents the variability of aligned decoding from different random unit selection (N=1000). Each distribution depicts the aligned decoding results from all random unit selections and is centered on the mean. The gray dashed line represents chance-level decoding.*

2.4. Discussion

Limits in our ability to sample neural population activity have made it difficult to fully resolve the structure of task information in frontal motor cortices. We used precise targeting of high-density laminar microelectrode arrays to better map spatiotemporal dynamics of individual neurons and neural populations across frontal motor cortices areas. Decoding analyses revealed variability across both the cortical surface and depths, with significantly more task information at some locations than others and no consistent pattern in depth relationships across cortical locations. Population dynamics analysis also revealed variability in the temporal structure of task-related population activity across the cortical surface. Neurons that were most informative about the task were distributed in space, yet also had activity patterns that evolved coherently in time. Our findings highlight the complex spatiotemporal structure of frontal motor cortex and reveal potential links between heterogeneity in task information and population dynamics. Our largescale neuron-resolution maps provide new perspectives on the seeming contradictions across studies mapping motor cortices at different scales, while also highlighting the need to consider spatiotemporal complexity for translational applications like BCIs.

Spatiotemporal heterogeneity in task decoding and dynamics

Our results extend previous findings that single-unit tuning notably varies across neurons (Schieber and Hibbard, 1993; Georgopoulos et al., 2007) to neural populations distributed across frontal motor cortices. We found that task information was spatially distributed (Fig. 2.2G, H) but primarily clustered in units recorded at the same cortical surface location (Fig. S2.9D, E). Task information was similarly distributed across all cortical depths, where the depth with the most task information varied across cortical surface locations (Fig. 2.2I, J). Our depth analysis was limited by the complicated sulcal topography of frontal motor cortex. We inserted probes perpendicular to the implanted grid surface, preventing us from confidently estimating the

location of layer boundaries. Despite these limitations, our findings are consistent with previous work suggesting minimal spatial organization in motor cortical activity at the level of single neurons.

Population-level perspectives that consider how neurons coordinate to control movements, meanwhile, have suggested a more unified picture of motor cortices. Our results add to this literature by revealing significant variability in population dynamics across frontal motor cortices. Caudal areas near the central sulcus are known to be predominantly involved in movement execution compared to rostral areas (Riehle and Requin, 1989; Cisek and Kalaska, 2005; Ouchi et al., 2025). Consistent with this, we found a weak correlation between rostral-caudal location and population dynamics (Fig. 2.3F, G). However, patterns in population dynamics were more strongly linked to the spatial distribution of task information across the cortical surface (Fig. 2.3I, J), potentially facilitated by strong coordination specifically among the highly task modulated units (Fig. 2.4). These results motivate future investigations into co-occurring neural dynamics related to non-task relevant activity, which will be best resolved by widespread simultaneous recordings. Furthermore, our results open questions about how spatially distributed neurons may coordinate to produce movements and highlights the need to integrate neuron-level and population-level approaches to fully characterize the neural representations of movement.

Our analyses suggest that subsets of neurons, identifiable from their unit-level properties (e.g., high task information), tightly coordinate during the task consistent with recent mechanisms of population coordination (Perich et al., 2018; Semedo et al., 2019; Veuthey et al., 2020). Yet, our results also highlight that this sub-population captures neural dynamics shared across all neurons in motor cortices, as revealed by patterns of decoding generalization (Fig. 2.5). The presence of shared task-relevant dynamics across all neurons is consistent with the ability to decode movement information through population recordings with non-targeted sampling (e.g., fixed geometry multi-electrode arrays) (Georgopoulos et al., 1986; Churchland et al., 2012). Neurons sharing some task-relevant dynamics are also likely related to observations that stable latent cortical dynamics underlie consistent behavior (Gallego et al., 2020). Past studies focused on the collective dynamics of all neurons (Churchland et al., 2012; Elsayed et al., 2016; Gallego et al., 2018, 2020), but the contributions of individual neurons to dynamics will likely need to be further considered. Our results suggest that the strength of the shared dynamics between two neural populations will vary, predicting that neurons that participate in other computations would coordinate differently. This generates testable hypotheses for future studies given the many forms of information about movement variables observed in motor cortices (e.g. planning vs. execution (Riehle and Requin, 1989; Cisek and Kalaska, 2005); arm vs. eye movement (Cisek and Kalaska, 2002; Pesaran et al., 2010; Ouchi et al., 2025); or movement variables encoded in different coordinate systems (Evarts, 1968; Fetz, 1994; Kakei et al., 1999; Sergio et al., 2005)). While focusing on target information, our results add new insights into the variable temporal structure observed in single-neuron activity (Churchland and Shenoy, 2007) by suggesting that the similarity of dynamics between distinct neural populations is driven by different combinations of a shared task-relevant component.

Spatially-distributed functional networks

Spatially focal recordings suggest that performing well-learned movements recruits consistent subsets of neurons with precise timing (Rubino et al., 2006; Peters et al., 2014; Best et al., 2016). Patterns in spike-timing between neurons also show that the spatiotemporal structure of these recruited neural populations are task-specific, suggesting they form a functionally defined

network (Moore et al., 2024). Consistent with this, we found a subset of spatially distributed units encoded significant amounts of task information and did so with highly correlated dynamics, which were distinct from the dynamics of other neural populations (Fig. 2.4B, C). Our results extend previous findings by suggesting that neurons participating in these functionally defined networks are spatially distributed across cortical areas and depths during well-learned center-out reaching tasks in macaques. Whether and how the functional network structure varies across species and behavioral tasks is an important question for future studies.

The potential relationship between population dynamics and connections between neurons (Sussillo et al., 2015) raises questions about the structural basis of functional networks. Anatomical studies found long-range connections across frontal motor cortex supporting cortico-cortical communication (Gatter and Powell, 1978; Huntley and Jones, 1991; Ninomiya et al., 2019). Functional studies suggest that subcortical structures contribute to functional network coordination (Ghanayim et al., 2025; Ramot et al., 2025). Both communication pathways would provide mechanisms for spatially distributed neurons to form functional networks, and future research investigating links to anatomical connectivity will further improve our understanding of motor representations.

If well-learned motor tasks recruit specialized functional networks, this raises questions about whether and how neurons flexibly contribute to multiple computations and, if so, how multiple networks form. Different movements, or muscle activation patterns, seem to display distinct correlation patterns between neurons (Moore et al., 2024), consistent with the possibility that distinct computations involve distinct networks. Our analyses focused on neural activity during movement execution, but frontal motor cortices also exhibit distinct, but linked, population-level activity during movement planning (Elsayed et al., 2016). Our results make the testable prediction that these computations might be implemented by distinct, spatially distributed functional networks. Recent work also finds that motor learning shapes the spatiotemporal recruitment of individual neurons (Peters et al., 2014). Such plasticity mechanisms may contribute to forming functionally-defined networks, though significant work remains to investigate how multiple functional networks coexist and coordinate during complex motor tasks.

Implications for BCI

Our results also highlight the need to consider the spatial distribution of movement information when targeting implants for BCIs. Current BCIs have primarily targeted mesoscale microelectrode array implants (with fixed geometry) based on functional cortical mapping (Hochberg et al., 2006; Collinger et al., 2013; Willett et al., 2021; Kunigk et al., 2024). Our results (Fig. 2.2) highlight that the functional organization of task information may be at a smaller scale than accessible by functional mapping techniques like functional MRI and intracortical micro-stimulation. These mapping approaches, then, may be limited in their utility for targeting spatially-precise, flexible neural implants (Jun et al., 2017; Musk and Neuralink, 2019; Yang et al., 2019). Optimizing task information decoding for BCIs may require new mesoscale and causal mapping techniques to precisely target implants to functionally-relevant networks, or dense non-targeted recording strategies paired with computational methods to identify functionally-relevant networks within large datasets (e.g., feature selection).

Improved mapping approaches may also be needed to target coordinated neural populations. Our results revealed heterogeneity in the population temporal dynamics across motor cortices, which influence BCI performance. A mismatch between the temporal structure of neural activity

with those expected by a BCI can reduce performance (Sadtler et al., 2014; Oby et al., 2019, 2025), suggesting that BCIs may benefit from targeting recordings to neural populations that share common dynamics. Successfully mapping temporal dynamics will also likely help ensure robust BCI performance over long timescales in the presence of recording instabilities. Recent approaches to maintain performance when neural measurements change specifically leverage the temporal structure of neural population activity (Degenhart et al., 2020; Karpowicz et al., 2025). Our results suggest that functionally defined networks have distinct dynamics, which could obscure the ideal temporal dynamics for maintaining stable BCI task performance. Methods to identify functional networks and shared task-relevant dynamics among neural populations will likely improve longitudinal BCI decoding approaches.

Together, our study reveals the spatiotemporal complexity of neural activity patterns across frontal motor cortical areas at the level of individual neurons and populations. New insights into the spatiotemporal distribution of motor representations in the brain will be critical for leveraging advanced neural probe technologies for BCIs.

2.5. Appendix

2.5.1. Supplemental figures

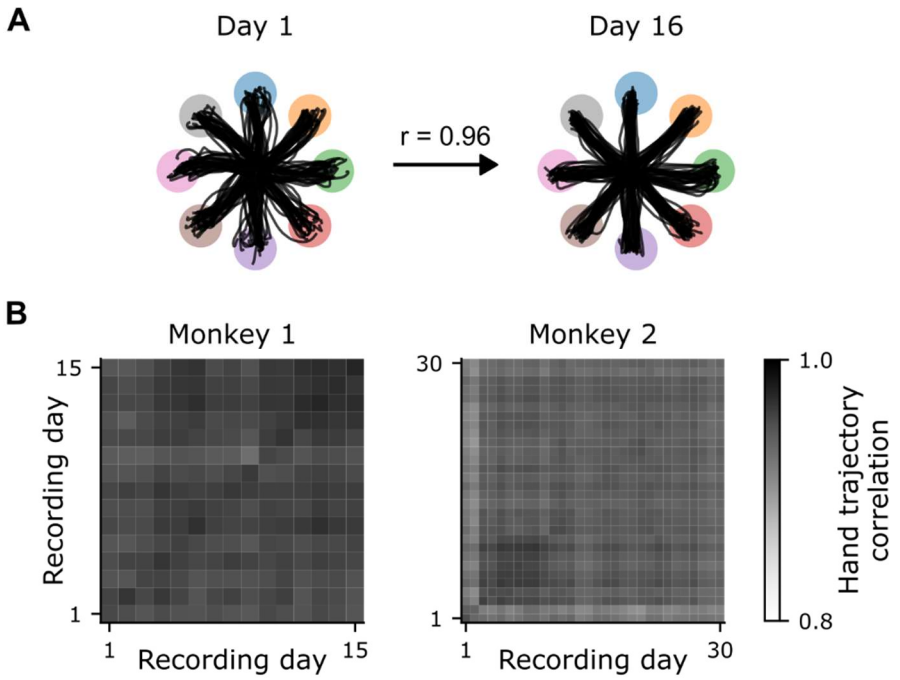

Figure S2.1. Three-dimensional hand trajectories are correlated across days. (A) *Individual trial cursor trajectories for Monkey 1 on two different example days (day 1, left; day 16, right), and the average cursor trajectory correlation between those days (r = 0.96, n = 240).* ***(B)*** *Average hand trajectory correlation across all combinations of day pairs for each monkey. Only 15 days are reported for Monkey 1 because of a technical error where 3D kinematics were not recorded for a single day.*

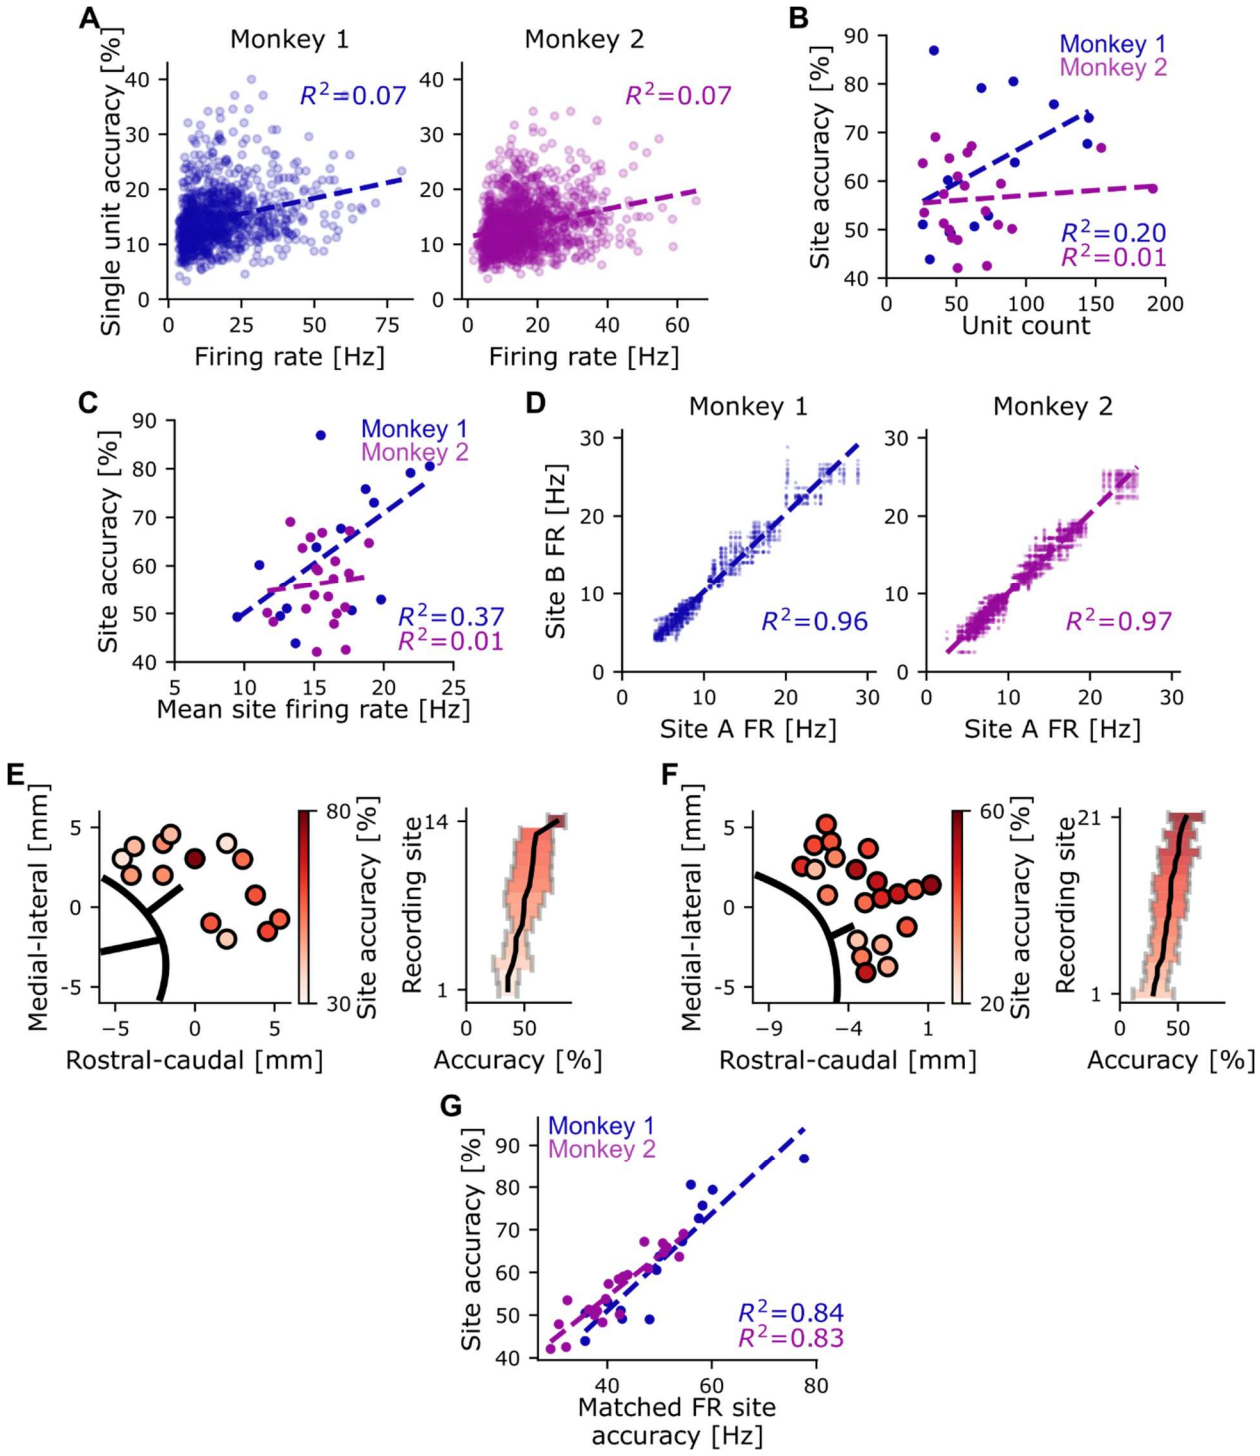

***Figure S2.2. Decoding accuracy variability is not explained by firing rate**. **(A)** Relationship between single-unit decoding accuracy and the average firing rate of each recorded unit (N=1,022 for monkey 1; N=1,441 for monkey 2). **(B)** Relationship between number of recorded units and site accuracy for each recording site (N=14 for monkey 1; N=21 for monkey 2). **(C)** Relationship between the average firing rate of all recorded units at each site and the site accuracy for each recording site (N=14 for monkey 1; N=21 for monkey 2). **(D)** Relationship of unit firing rates for all pairs of firing rate-matched recording site populations (referred to as A, B) (N=1456 for Monkey 1; N=3360 for Monkey 2). **(E)** Left: Distribution of firing rate-matched population decoding accuracy across the cortical surface (site accuracy) for monkey 1. The black lines represent the arcuate sulcus. Right: Population decoding accuracy across recording sites, sorted by the mean decoding accuracy. Black lines indicate the mean; the shaded region represents*

the 95% CI. Colors indicate the site decoding accuracy (color bar) for both right and left panels. ***(F)*** *Same as (E) but for Monkey 2.* ***(G)*** *Relationship of decoding accuracies estimated in the firing rate-matched condition and the non-matched condition (N=14 for monkey 1; N=21 for monkey 2). For all panels, dashed lines represent the linear regression.*

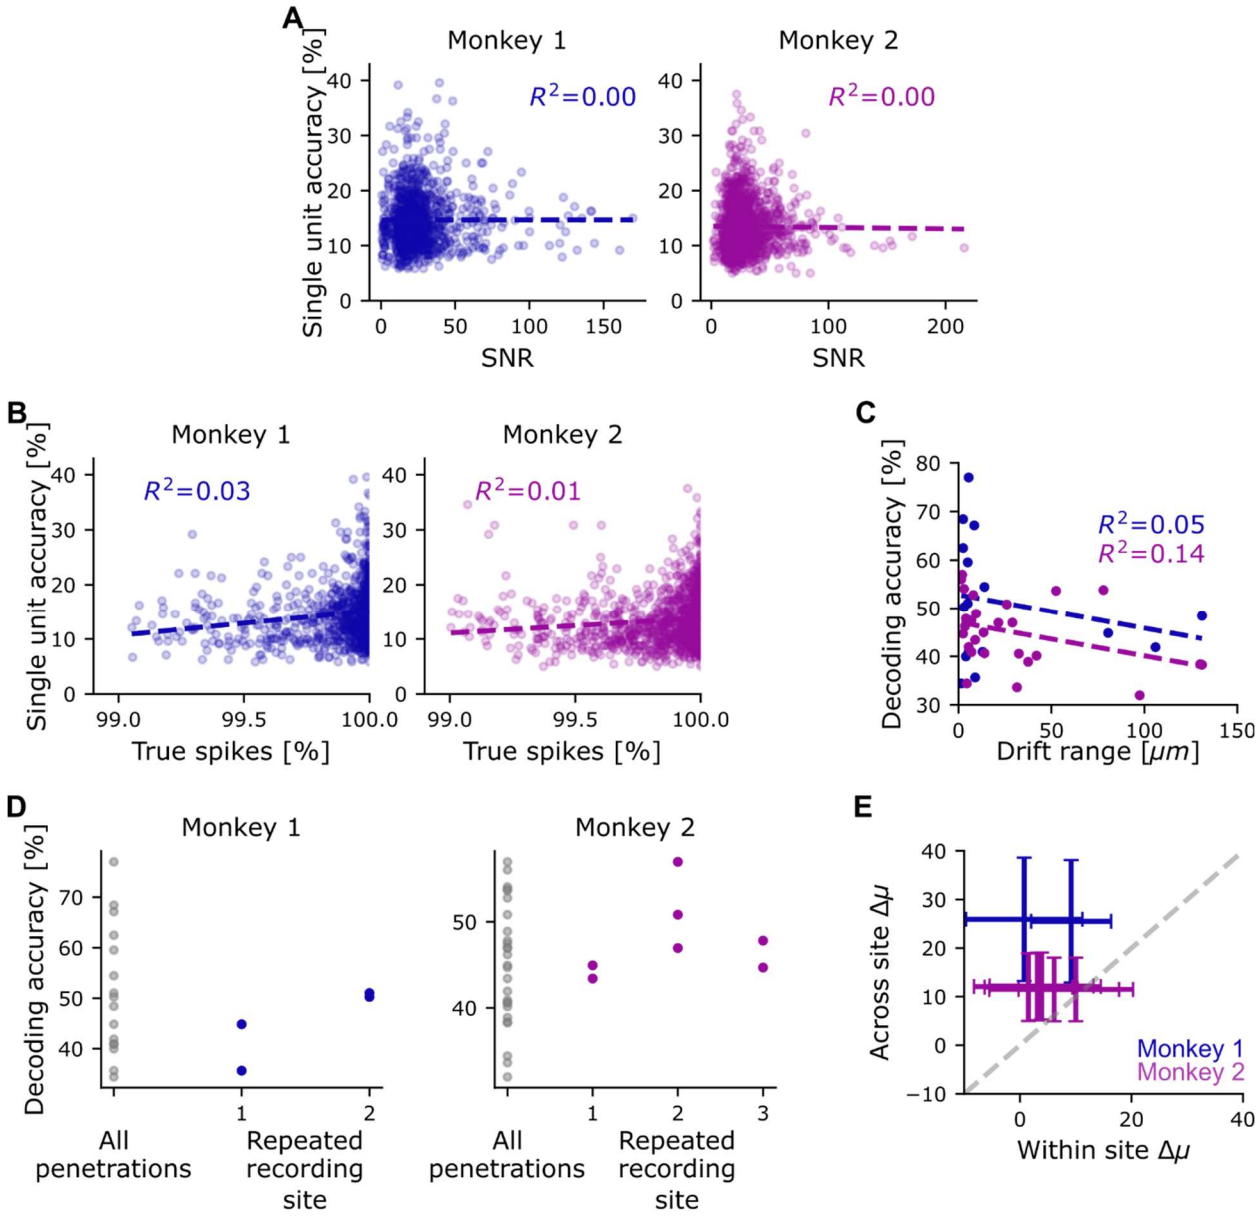

Figure S2.3. Decoding accuracy variability is not explained by recording quality. ***(A)*** *Relationship between single-unit decoding accuracy and the signal-to-noise ratio for each unit (See Unit quality metrics)* ***(B)*** *Relationship between the single-unit decoding accuracy and true spike percentage based on refractory period violations (See Unit quality metrics).* ***(C)*** *Relationship between the drift range computed by Kilosort 4 and decoding accuracy for each probe insertion.* ***(D)*** *Average decoding accuracy for each penetration at a recording site. Decoding accuracy from each repeated probe penetration (blue/purple) next to the decoding accuracy from all other probe penetrations for each monkey (gray).* ***(E)*** *The mean decoding difference for repeated penetrations compared to the mean decoding difference across sites. Error bars represent the standard deviation across unit group resampling (See Statistical analysis across multiple penetrations). The gray dashed line represents the identity line. Note: All colored dashed lines represent the linear regression fit.*

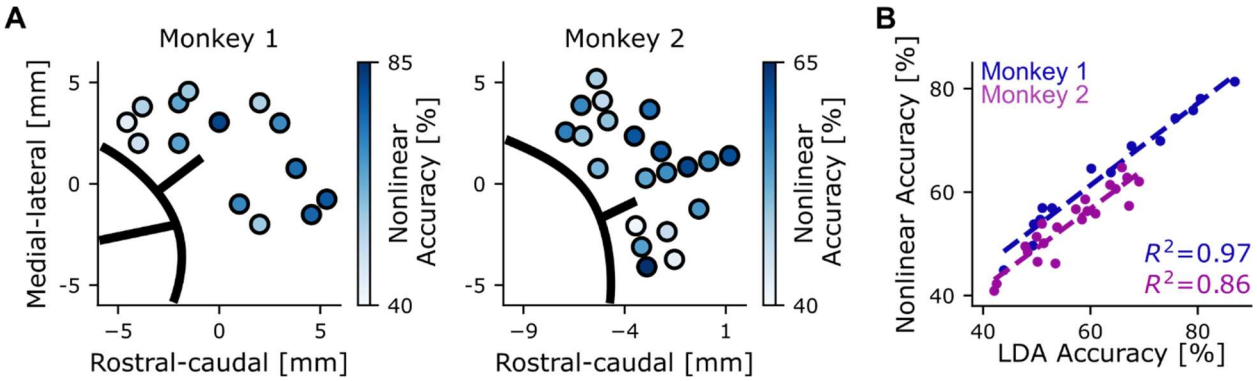

Figure S2.4. Decoding accuracy is correlated when computed with either linear or nonlinear methods. ***(A)*** *Site accuracy computed with a nonlinear decoding algorithm (See 2.2.5. Decoding analysis).* ***(B)*** *Relationship between nonlinear decoding accuracy and linear decoding accuracy (See 2.2.5. Decoding analysis). Dashed lines represent the linear regression fit.*

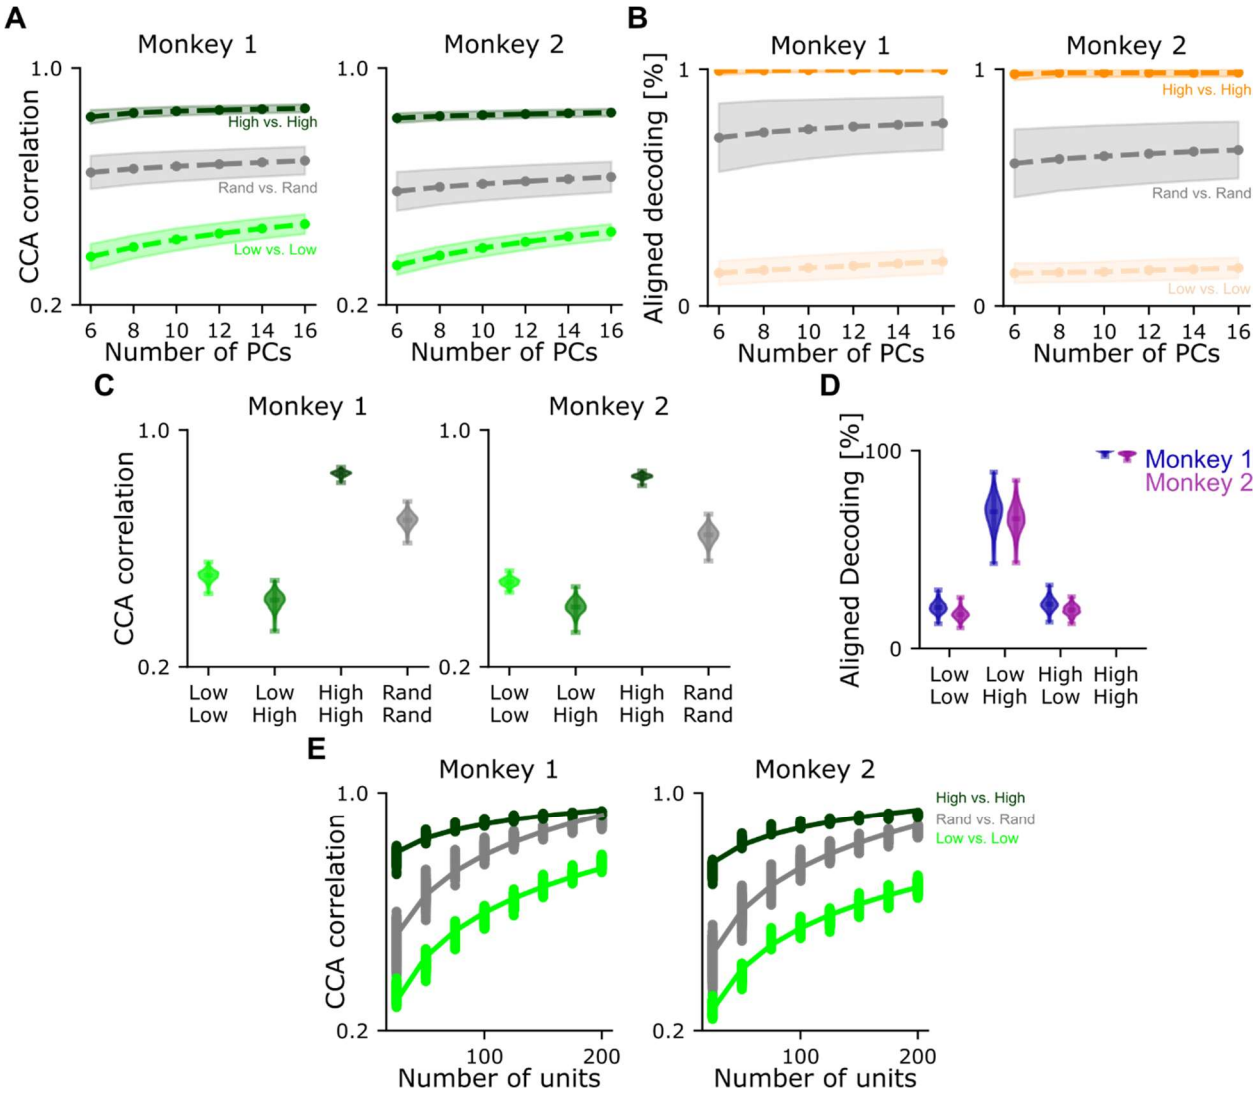

Figure S2.5. Dynamic correlations are specific to high task information populations across a range of principal components (PC). ***(A)*** *CCA correlation for task neural populations across a range of principal components (PC). Dashed lines represent the mean across random unit sampling and the shaded regions represent the 95% CI.* ***(B)*** *Same as (A) but for aligned decoding.* ***(C)*** *CCA correlation for task responsive neural populations while controlling for the explained variance. Each distribution represents the variability from all random unit resampling centered on the mean.* ***(D)*** *Formatted the same as (C), but for aligned decoding.* ***(E)*** *CCA correlation for task responsive neural populations across a range of neural population sizes. Solid lines indicate an exponential fit to the data. Each dot represents the CCA correlation from random unit resampling using the number of units indicated by the x-axis. For all panels, high and low decoding accuracy units were defined as the 100 units with the highest and lowest single-unit decoding accuracy, respectively (See 2.2.6. Neural population analysis and Fig. 2.4A).*

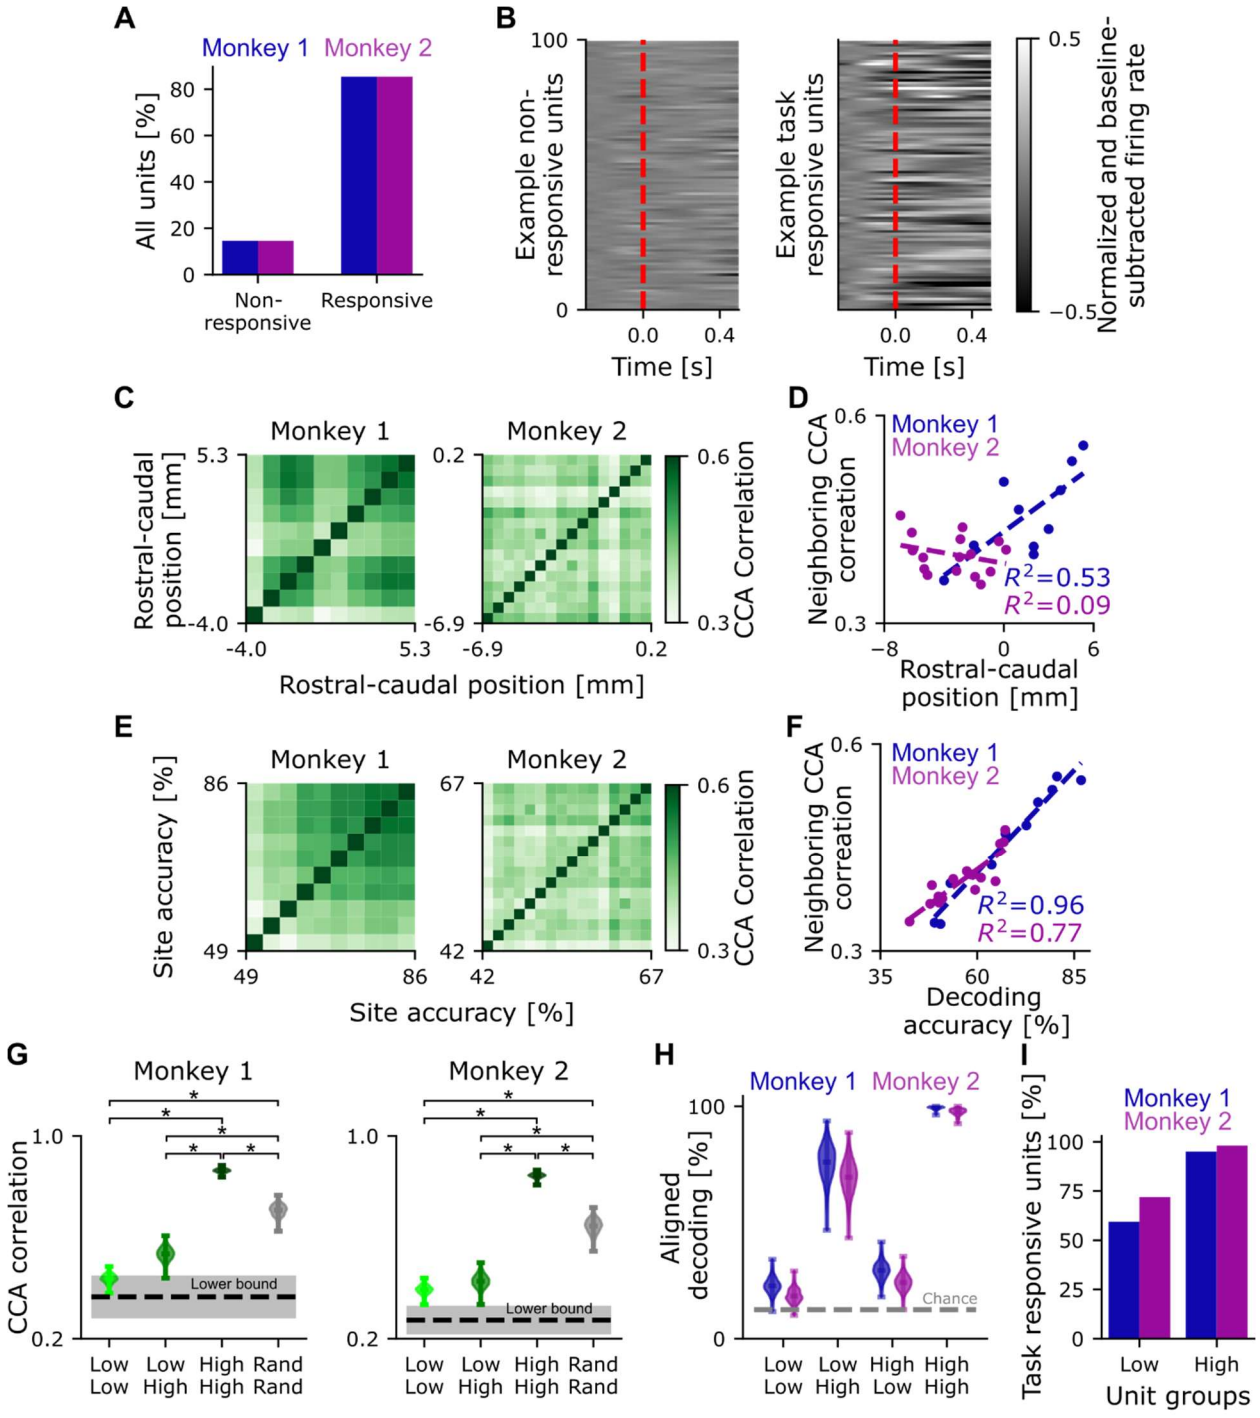

Figure S2.6. Dynamic correlations are specific to high task information populations when analysis is restricted to task-responsive units (A) *Percentage of all recorded units identified as not task responsive and task responsive for monkey 1 (blue) and monkey 2 (purple).* ***(B)*** *Trial averaged and baseline-subtracted firing rate as a function of time of 100 randomly selected example units from Monkey 1 that were identified as not task-responsive (left) or task-responsive (right; see 2.2.7. Statistical analysis). Firing rates were time-aligned to movement onset, indicated by the red dashed line.* ***(C)*** *Heatmap of CCA correlation between neural populations of only task-responsive units organized by rostral-caudal position for Monkey 1 (left) and Monkey 2 (right). Each pixel represents a pair of neural populations (N=45 for monkey 1; N=120 for monkey 2).* ***(D)*** *Neighboring CCA correlation based on rostral-caudal position for each distinct neural population (see 2.2.6. Neural population analysis) for monkey 1 (blue) and monkey 2 (pink). Each dot represents a recording site (N=10 for monkey 1; N=16 for monkey 2). The dashed line*

represents the linear regression. ***(E)*** *Heatmap of CCA correlation between distinct neural populations organized by the amount of task information for Monkey 1 (left) and Monkey 2 (right). Format as in (B).* ***(F)*** *Neighboring CCA correlation based on task information for each distinct neural population, formatted as in (C).* ***(G)*** *CCA correlation between distinct neural populations of task-responsive units defined by task information (Low and High) or randomly selected (Rand) for monkey 1 (left) and monkey 2 (right). The distribution for each comparison represents the variability from different random unit selections (N=1000). Each distribution depicts the CCA correlation results from all random task-responsive unit selections and is centered on the mean. The black dashed line represents the lower bound estimated from the inter-trial interval activity of task-responsive units (see 2.2.7. Statistical analysis) and the gray shading represents the 95% CI across random unit selections.* ***(H)*** *Aligned decoding for each neural population comparison. The distribution represents the variability of aligned decoding from different random task-responsive unit selection (N=1000). Each distribution depicts the aligned decoding results from all random task-responsive unit selections and is centered on the mean. The gray dashed line represents chance-level decoding.* ***(I)*** *Percentage of task responsive units in the high and low task information populations from Fig. 2.4.*

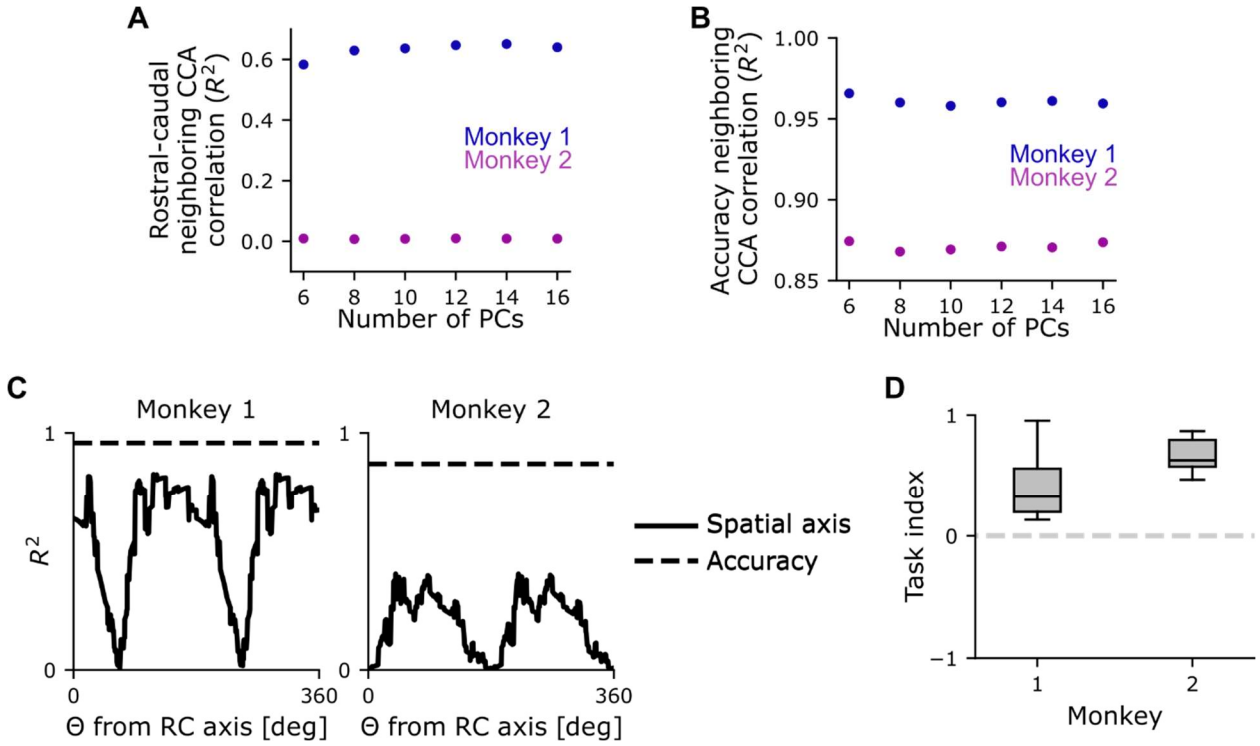

Figure S2.7. Spatial organization results are consistent across a reasonable range of principal components (PC) and across different spatial axes. ***(A)*** *R^2 of the relationship between neighboring CCA correlation and rostral-caudal position of recording sites across a range of principal components (PC).* ***(B)*** *Same as (A) except with recording sites organized by decoding accuracy.* ***(C)*** *Relationship of the R^2 from (A) except when computed with recording sites organized along a variety of different axis relative to the rostral-caudal (RC) axis (solid line) using 10 PCs. Dashed line represents the R^2 from (B) with 10 PCs.* ***(D)*** *Task index computed for each axis (C). Positive values indicate that decoding accuracy predicts similar dynamics better than spatial position (See Neural population dynamics between recording sites). Dashed line represents a task index of 0 where decoding accuracy and spatial axis predict correlated dynamics similarly well.*

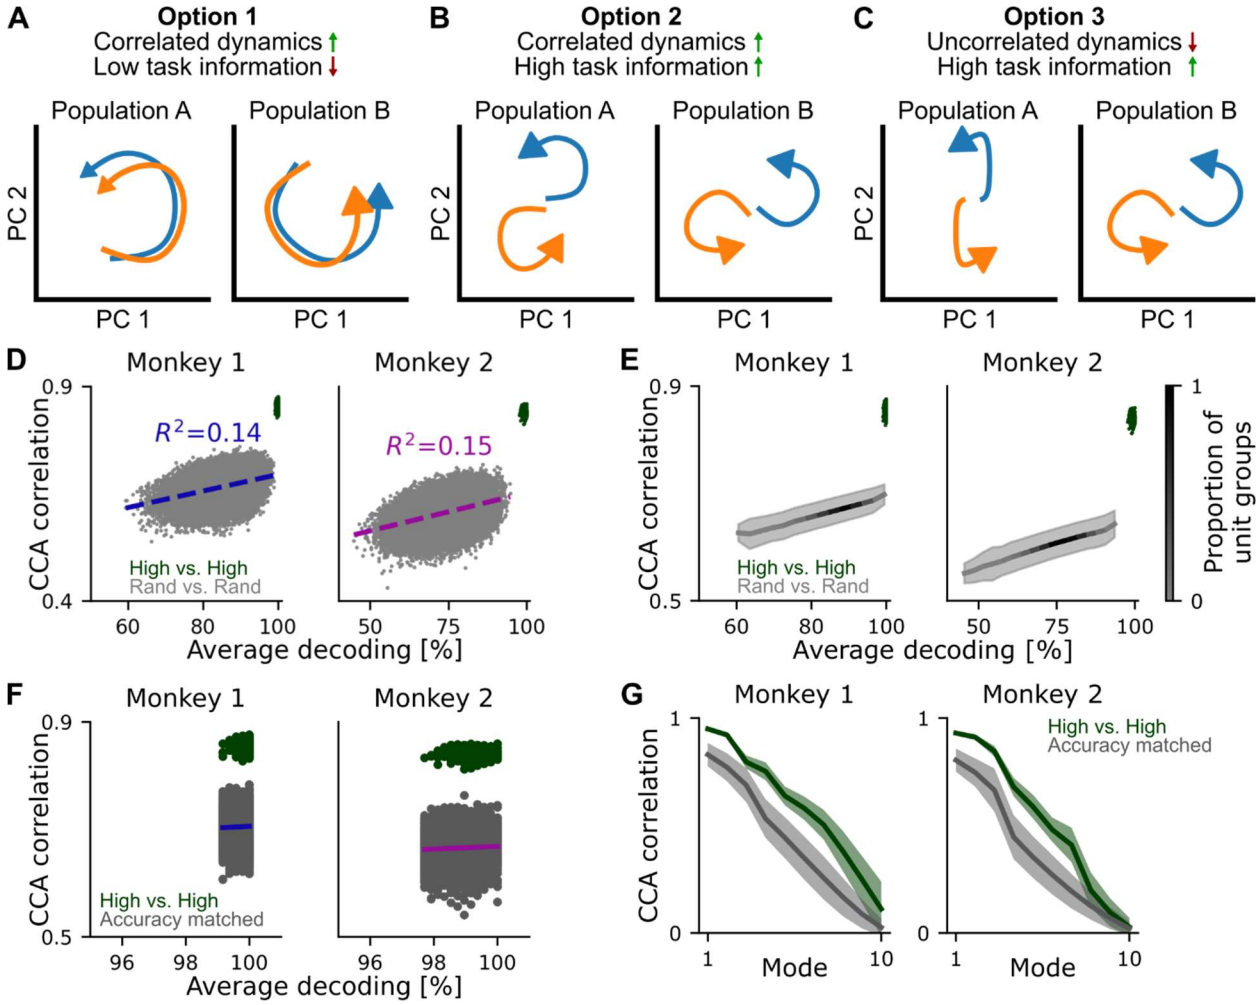

Figure S2.8. Correlated dynamics are not necessarily solely a consequence of high task information. (A) *Conceptual schematic diagram depicting correlated dynamics with low task information. Left: Population dynamics in PC space for neural population A. Each color represents the trajectory for a different task condition (i.e. target direction). Right: Same as Left, but for neural population B, composed of different units.* ***(B)*** *Same as (A) but for correlated dynamics with high task information* ***(C)*** *Same as (A) but for uncorrelated dynamics with high task information* ***(D)*** *CCA correlation between two distinct neural populations as a function of their average decoding accuracy. Gray dots represent a comparison between neural populations composed of randomly selected units (N=100,000). Green dots represent a comparison between neural populations composed of high task information units (N=1,000). Each dot is from different random unit sampling. The dashed line represents the linear regression.* ***(E)*** *Same as D except representing the distribution of results from unit groups. The solid line represents the mean, and the shaded region represents ±1σ. The gradient represents where units were distributed along the x-axis, normalized to the maximum number of groups.* ***(F)*** *Same as D but with random unit selection restricted to produce populations that match the range of decoding accuracies observed in the high-task information subpopulations (N=6536 for Monkey 1 and N=5562 for Monkey 2). The colored line represents the linear regression of the randomly selected neural populations.* ***(G)*** *CCA correlation as a function of CC mode. The shaded regions represent the 95% confidence interval for each mode across random unit selections.*

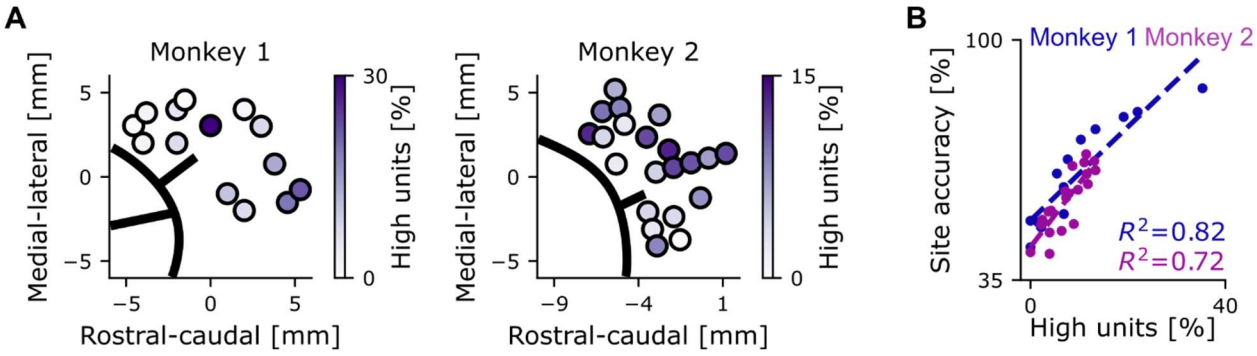

Figure S2.9. The spatial distribution of the high task information unit locations is correlated with the spatial distribution of population decoding accuracy. (A) *Spatial distribution of the recording location for high task information units for Monkey 1 (left) and Monkey 2 (right). Each circle represents a recording site and the x-y position indicates the coordinates within the chamber. The color at each recording site represents the percentage of units with high task information relative to the total number of units measured at each recording site.* ***(B)*** *Population decoding accuracy for a recording site (same analysis as Fig. 2.2G,H) as a function of the percentage of units with high task information recorded at that site for monkey 1 (blue) and monkey 2 (pink). Each data point represents a recording site (N=14 for monkey 1; N=21 for monkey 2); dashed lines represent a linear regression (R^2 = 0.82, $p < 0.0001$ for monkey 1; R^2 = 0.72, $p < 0.0001$ for monkey 1).*

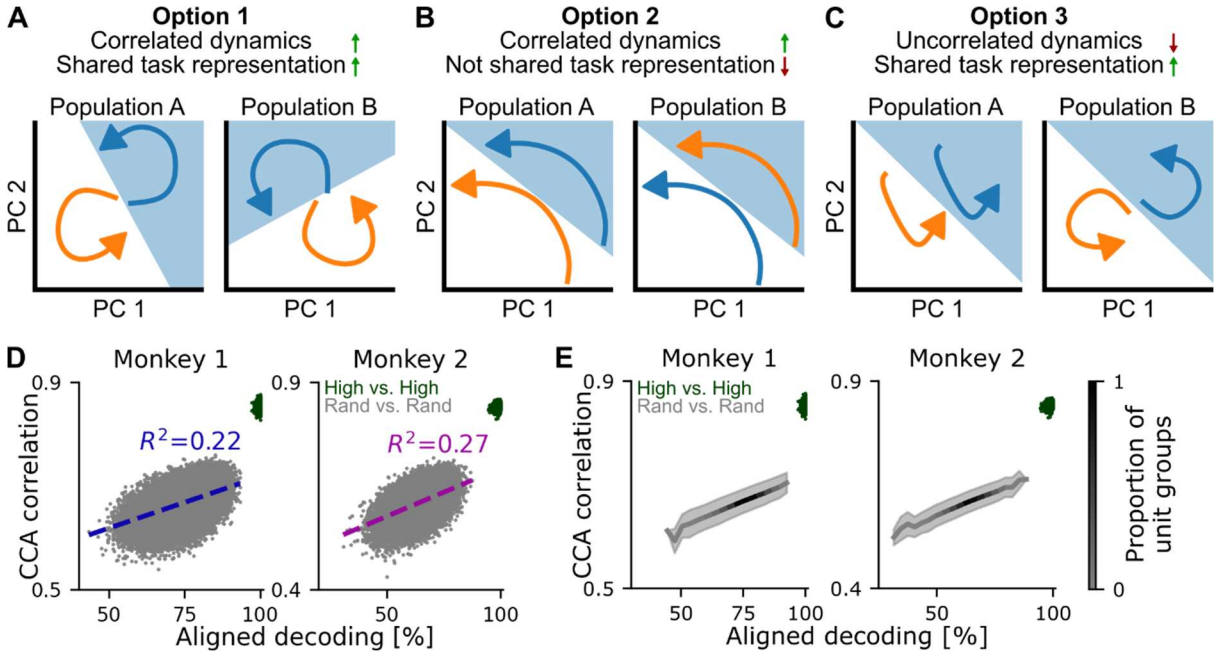

Figure S2.10. A shared task representation is not required by correlated dynamics. (A) *Diagram depicting correlated dynamics with a shared task representation. Left: Population dynamics in PC space for neural population A. Each color represents the trajectory for a different task condition (i.e. target direction). Right: Same as Left, but for neural population B, composed of different units. The blue shaded background represents a decision boundary between task conditions trained on Population A, then tested on Population B after alignment between Population A and B.* ***(B)*** *Same as A but for correlated dynamics without a shared task representation.* ***(C)*** *Same as A but for uncorrelated dynamics with a shared task representation.* ***(D)*** *CCA correlation between two neural populations as a function of their aligned decoding. Gray dots represent neural populations composed of randomly selected units (N=100,000). Green dots represent neural populations composed of high task information units (N=1,000). Each dot is from different random unit sampling. The dashed line represents the linear regression.* ***(E)*** *Same as D except representing the distribution of results from unit groups. The solid line represents the mean, and the shaded region represents ±1σ. The gradient represents where units were distributed along the x-axis, normalized to the maximum number of groups.*

3. The spatiotemporal structure of neural activity in motor cortex across motor computations

3.1. Introduction

Goal-directed movements largely consist of two computations: planning and execution. To pick up your coffee cup, you plan your movement, then execute the reach. Planning and execution are behaviorally linked as more time to prepare improves performance (Churchland et al., 2006). Yet planning and execution are also anatomically distinct. During the transition from planning to execution, information transitions from rostral to caudal frontal motor cortices (Cisek and Kalaska, 2005; Ouchi et al., 2025). Anatomically-agnostic population approaches suggest that planning and execution computations are coordinated. When neurons are pooled across frontal motor cortices, their activity evolves similarly in time, yet in orthogonal subspaces (Elsayed et al., 2016). It remains unclear how anatomically distinct neural populations in motor cortex transition between distinct, but linked, computations.

We hypothesized that the transition from planning to execution across spatially distributed neural populations could occur in one of three ways. First, distinct populations could be separately involved in planning and execution (Hypothesis 1; Figure 3.1A). Second, separate populations could be involved in planning and execution yet coordinate through an intermediate population that is involved in both computations (Hypothesis 2; Figure 3.1B). Third, the same neural population could be involved in both planning and execution with individual neurons reweighting their contribution when transitioning between computations (Hypothesis 3; Figure 3.1C).

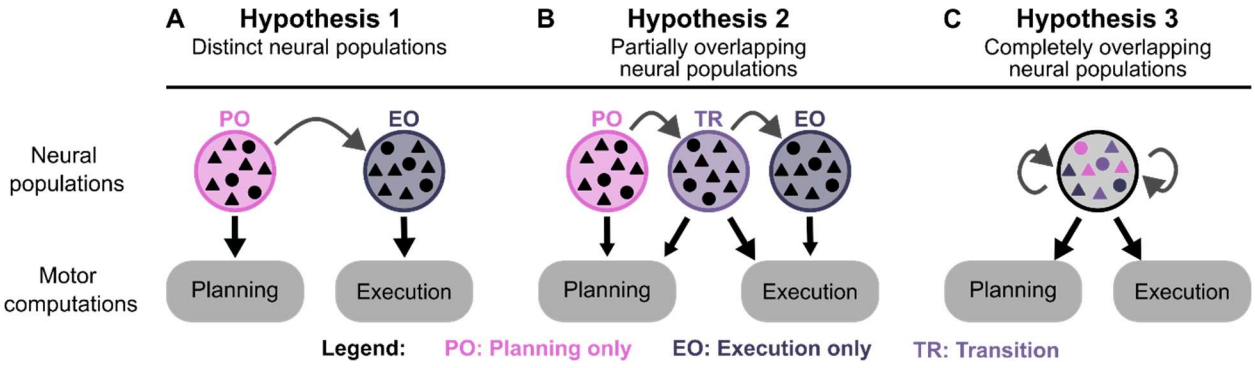

Figure 3.1. Proposed hypotheses for how distributed neural populations transition from planning to execution. ***(A)*** *The transition happens between distinct planning and execution neural populations.* ***(B)*** *The transition involves distinct planning and execution populations and a subset of units involved in both planning and execution.* ***(C)*** *All neurons are involved in both planning and execution.*

We addressed these hypotheses with spatially precise recordings of single-unit activity across frontal motor cortex using high density microelectrode arrays (Neuropixels). The recorded neural populations span from dorsal pre-motor cortex (PMd) to the arm region of primary motor cortex (M1). We recorded neural activity while monkeys performed a delayed center-out reaching task allowing us temporally separate neural activity associated with planning and execution. We quantified the contribution of individual neurons and neural populations to planning and execution with target-direction decoding and dynamics similarity analysis.

Single-unit decoding analysis revealed subpopulations of neurons involved only in planning or only in execution. As expected, planning-related neurons were primarily located in rostral motor cortex and execution-related neurons primarily located in caudal motor cortex (Cisek and Kalaska, 2005). We also observed a subset of units involved in both planning and execution, seen in previous work (Churchland and Shenoy, 2007), but found this population spatially distributed across frontal motor cortex. The results from our single-unit decoding analysis were consistent with either hypothesis 2 or 3, but not hypothesis 1, suggesting that spatially distributed neural populations do not coordinate across computations through distinct subpopulations.

To distinguish between hypothesis 2 and 3 we analyzed population dynamics because similar dynamics reflect linked computations. Hypothesis 2 predicts that a subset of units would have the most similar dynamics across computations because not all units are involved in both planning and execution. Hypothesis 3 predicts that all subsets of units would have similar dynamics across planning and execution because all neurons are involved in both computations. Consistent with hypothesis 2, we found that the subset of units most involved in both planning and execution also had the most similar dynamics between planning and execution even though the task encoding of individual units did not generalize.

Our results suggest that planning and execution are primarily linked by a subset of spatially distributed neurons. Furthermore, the transition between distinct but linked computations appear to be facilitated through this subset of units that has consistent dynamics during both computations.

3.2. Methods

3.2.1. Experimental procedures

The structure of neural activity across motor cortical computations was assessed using intracortical neural activity recorded from two monkeys (*Macaca mulatta*) while they performed a reaching task that allowed movements to be planned. Surgical, electrophysiological, and behavioral protocols are similar to those described in Chapter 2 and (Canfield et al., 2025). Briefly, monkeys performed a delayed center-out reaching task where hand movements were mapped to control the position of a cursor on a screen (Figure 3.2A). Trials began with the appearance of a center target. After a brief center hold, a peripheral target appeared 6.5 cm from the center target initiating the delay period. Subjects were required to keep the cursor in the center target during a delay period, the length of which was randomly drawn from a uniform distribution. All task parameters can be found in Table 2.1. All monkeys were experts at this task during all recording sessions and performed stereotyped movements across days (Figure 3.2B).

We quantified behavioral performance for each trial with reaction time. The reaction time was defined as the time between the go-cue and when the cursor crossed a speed threshold. Speeds greater than 4.75 standard deviations above the mean baseline cursor speed were considered significant. The cursor speed was estimated by the first derivative of the cursor positions, and digitally low pass filtered at 20 Hz (Ouchi et al., 2025). The threshold for significant cursor speed was defined for each trial based on the baseline cursor speed during that trial (from peripheral target onset to go-cue).

Neural activity was recorded from across frontal motor cortices with Neuropixels high-density laminar microelectrode arrays (IMEC, Leuven Belgium)(Jun et al., 2017). We recorded from 14 and 21 unique grid locations, or recording sites, on the cortical surface for Monkey 1 and 2 respectively (Figure 3.2D, E). We conducted these recordings on 16 and 30 non-consecutive days with multiple penetrations made at select grid locations. A majority of the data were collected using a single probe on each day except on four days with Monkey 2, where two probes were used to simultaneously record activity from different grid locations. Recording sessions were included in analyses using the following criteria: 1) they sampled neural activity from at least 2 mm along the Neuropixels probe, and 2) they contained at least 25 reaches to each target that that had a duration below the trial length threshold (see 2.2.3 Behavioral data analysis). All analysis was conducted using 25 successful reaches to each target.

The raw AP band signals were post-processed to sort action potentials (spike sorting) using Kilosort 4.0 with default parameters (Pachitariu et al., 2024). Each putative neuron (unit) was then automatically assessed using quality metrics derived from the International Brain Laboratory (Table 2.2) (International Brain Laboratory, 2024) and manual inspection.

3.2.2. Decoding analysis

We quantified task information during the reach planning and execution phases of the task with target-direction decoding analysis similar to that described in Chapter 2 and (Canfield et al., 2025). Briefly, firing rates were aligned to either peripheral target onset or reach onset and considered to be activity related to either planning or execution, respectively (Figure 3.2B). Reach onset was approximated by the time when the cursor left the center target.

We quantified if units were task-modulated during planning or execution by predicting target direction. Neural activity from 50ms before to 200ms after either target onset or reach onset were averaged across time for each unit and each trial. Linear decoding analysis was performed on the time-averaged and smoothed firing rates of either single-units or neural populations using linear discriminant analysis (LDA) implemented using the Python library Scikit-learn (Pedregosa et al., 2011) with the eigen solver and automatic shrinkage. Reported decoding accuracies were estimated by averaging the prediction accuracy from 4-fold cross validation. Each fold used 180 randomly selected trials to train the model and the remaining 60 trials for testing.

Population decoding analysis was applied to assess task information encoding during the planning or execution computations. To estimate task information during planning, we pooled units across all recordings, randomly selected 25 units, then performed decoding analysis with the activity aligned to target onset (i.e. the planning related activity). This procedure was repeated 1000 times with different randomly selected unit populations to estimate the 95% bootstrap confidence interval while controlling for population size. We applied the same process to execution related activity to estimate task information during execution with the same 1000 neural populations.

Population decoding analysis was also used to test the generalization of task information encoding across computations (i.e. from planning to execution). Decoders were trained on neural activity from either planning or execution, then tested on the neural activity from the other computation. Neural populations of 25 randomly selected units were used and the entire process was repeated with 1000 different neural populations to estimate the 95% confidence interval while controlling for population size.

The task-modulation of single-units was estimated with single-unit target-direction decoding analysis, then used to classify if each unit was modulated during planning or execution. Thresholds for modulation were defined as 0.7 and 1 standard deviation over the median single-unit decoding accuracy for Monkey 1 and 2, respectively. The specific threshold value was defined separately for planning and execution. Slightly different thresholds were used across monkeys to ensure a minimum number of units were classified in each group to enable the subsequent analysis. The results were not sensitive to the thresholds themselves. Units modulated during neither planning nor execution were defined as “not modulated” (NM) units. Units only modulated during planning were defined as “planning only” (PO) units and units only modulated during execution were defined as “execution only” (EO). Units modulated during both planning and execution were defined as “transition” (TR) units.

Population decoding of task modulated unit groups was quantified by applying the previously described population decoding analysis to the neural activity of randomly selecting 25 units within each of the NM, PO, EO, and TR population. Decoding analysis was repeated 1000 times with randomly selected neural populations from each unit group to estimate the 95% confidence interval.

3.2.3. Neural population analysis

We quantified dynamics similarity during and across the reach planning and execution with canonical correlation analysis (CCA) using a procedure similar to that described in Chapter 2 and (Canfield et al., 2025). Briefly, CCA was applied to Z-scored and smoothed firing rates following procedures based on previous work (Gallego et al., 2018). Principal Component Analysis (PCA) was applied to the neural activity around either planning or execution to project the neural population firing rates onto the top 10 principal components (PC) and estimate the latent dynamics (L) for each neural population. CCA was then applied to “align” the latent dynamics of each population by finding the new directions that maximize the pairwise correlation between latent trajectories, referred to as the “aligned” axes (Gallego et al., 2018). CCA produced a correlation value for each mode that represents the similarity of the latent trajectories between each neuronal population. We report the CCA correlation using the average computed across the top 4 modes to minimize the possibility of spurious correlations (Sussillo et al., 2015; Gallego et al., 2020).

CCA analysis was performed with planning and execution related activity, separately, to assess dynamics similarity of different unit groups within these computations. Within planning and execution, neural populations were formed by randomly selecting 25 units from each of the NM, PO, EO, and TR neural populations defined in *3.2.2 Decoding analysis.* CCA was performed between all combinations of these neural populations. When we compared the same group to itself as an internal control we applied CCA to two distinct populations of 25 units that group to control for population size. Unit group selection was repeated 1000 to estimate the 95% confidence interval.

CCA was also used to quantify the dynamics similarity of all combinations of unit groups across computations. Here, the planning-related neural activity of one group was compared to the execution-related activity of another group, for all unit group combinations. When we compared the planning-related activity to execution-related activity of the same unit group we used the same 25 units, quantifying how much the dynamics of those units change across populations. Unit group selection was repeated 1000 to estimate the 95% confidence interval.

Units with similar tuning during planning and execution were excluded from population analysis to ensure any observed dynamics similarity did not simply reflect consistent modulation across computations. Units with consistent tuning were defined as units whose across-event decoding accuracy (see *3.2.2 Decoding analysis*) was more than two standard deviations over the mean across event decoding accuracy across all units.

3.3. Results

We trained two monkeys to perform a delayed center-out reaching task, as described in Chapter 2, that allowed us to separate movement planning and execution-related activity (Figure 3.2A) even though planning and execution often occur simultaneously in practice (Lara et al., 2018b; Zimnik and Churchland, 2021). Monkeys were proficient in the task and produced highly stereotyped movements across days (Figure 3.2B). The monkey's performance, measured by reaction time, improved on trials when given more time to prepare, suggesting that planning occurred during the delay period (R^2: 0.27 for Monkey 1, and 0.36 for Monkey 2; Figure 3.2C). Therefore, we refer to time immediately after target onset as the planning period of the trial, and time immediately after reach onset as the execution period.

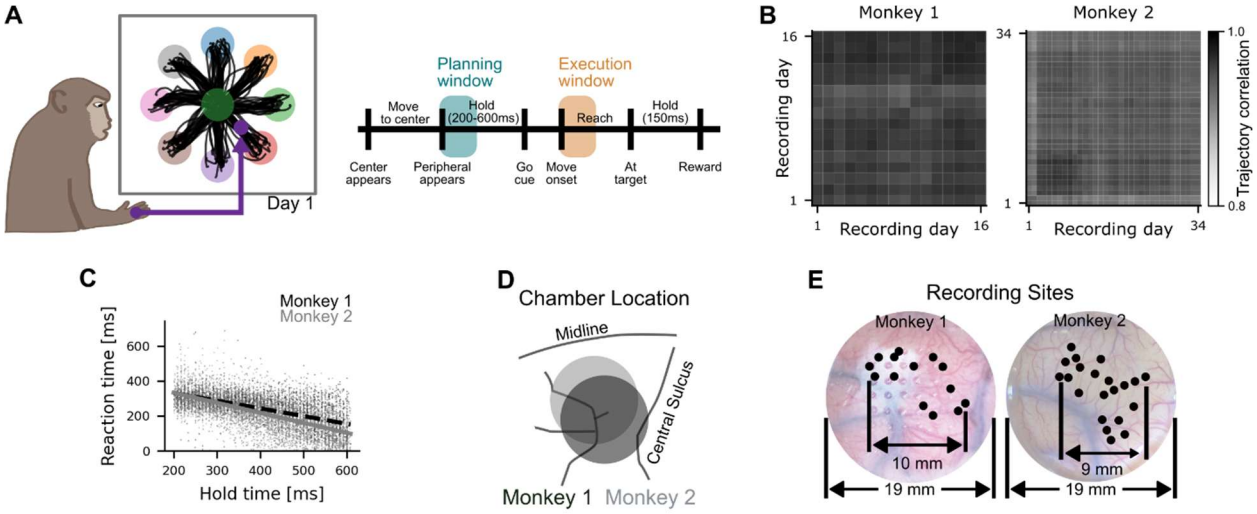

Figure 3.2. Large-scale recordings are made during highly stereotyped reaches that are planned. (A) *Schematic of the experimental setup. Hand position was tracked and mapped to the position of the cursor (Purple) on a visual display. Bottom: trial structure that for the delayed center out task that temporally separates planning (blue) and execution (orange) related activity* ***(B)*** *Average cursor trajectory correlation across days. Relationship of reaction time and cursor hold time. The dashed and solid lines represent the linear regression fit.* ***(D)*** *Chamber location.* ***(E)*** *Neuropixels probe insertion locations (black dots) on the cortical surface.*

We recorded single-unit extracellular activity from frontal motor cortices while monkeys performed planned reaches (Figure 3.2D). We used high-density laminar microelectrode arrays (Neuropixels) to make spatially precise recordings of neural populations from a large 3D cortical volume. Neural recordings covered a cortical surface area of 9.9 × 6.7 mm in Monkey 1 and 8.1 × 9.3 mm in Monkey 2 (rostral–caudal × medial–lateral; Figure 3.2D, E). Each probe penetration recorded activity from the cortical surface to a depth ranging from 2.0 - 3.8 mm. Collectively, we recorded 903 units from Monkey 1 and 1280 units from Monkey 2. The behavioral stereotypy, spatially precise recordings across a broad cortical region, and large quantity of units allowed us to compare how spatially distributed neural populations transitioned from movement planning to execution.

3.3.1. Units modulated during both planning and execution are spatially distributed across frontal motor cortices

We first tested hypothesis 1 by examining if any units were modulated during both planning and execution computations. Hypothesis 1 predicts that distinct neural populations are involved in planning and execution; therefore, units modulated during both planning and execution would refute this hypothesis.

We examined how units were modulated during planning and execution by estimating the task information during each computation. We quantified task information by computing how well individual units predicted target direction during each task phase, separately (see *3.2.2 Decoding analysis*). Task information was unevenly distributed across units during both planning and execution because this single-unit decoding distributions were non-normal for both planning (Kolmogorov-Smirnov test: $p < 0.0001$ for Monkey 1, $p < 0.0001$ for Monkey 2; Figure 3.3A) and execution (Kolmogorov-Smirnov test: $p < 0.0001$ for Monkey 1, $p < 0.0001$ for Monkey 2; Figure 3.3B). Also, the planning and execution distributions were each skewed right (Planning skewness ± S.E.: 0.83 ± 0.08 for Monkey 1, and 0.83 ± 0.07 for Monkey 2; Execution skewness ± S.E.: 1.11 ± 0.08 for Monkey 1, and 0.89 ± 0.07 for Monkey 2) more than their respective chance distribution that was computed by shuffling target labels because the standard error did not overlap (Planning skewness ± S.E.: 0.27 ± 0.08 for Monkey 1, and 0.34 ± 0.07 for Monkey 2; Execution skewness ± S.E.: 0.39 ± 0.08 for Monkey 1, and 0.27 ± 0.07 for Monkey 2). Units with a high single-unit decoding accuracy in planning or execution were defined as modulated during that computation (see *3.2.2 Decoding analysis*; Figure 3.3C).

We assessed if units were involved in both planning and execution by classifying them based on which motor computation they were modulated during (Figure 3.3D). Units modulated during neither planning nor execution were considered ‘not modulated’ (NM). Units modulated only during planning or execution were defined as ‘planning modulated only’ (PO) or ‘execution modulated only’ (EO), respectively. If units were modulated during both planning and execution they were defined as ‘transition’ (TR) units. If TR units do not exist, that would provide evidence supporting hypothesis 1. Alternatively, the existence of TR units would be consistent with hypothesis 2 and 3 (Figure 3.3E).

We found TR units, modulated during both planning and execution, and that they were spatially distributed from rostral to caudal frontal motor cortices. Their distribution spanned more than 5 mm in both monkeys and were found in both rostral and caudal parts of the chamber. The PO and EO units followed the expected spatial distribution where PO units were primarily found in rostral motor cortices and EO units were primarily found in caudal motor cortices (Cisek and Kalaska, 2005)(Figure 3.3F).

We quantified modulation for each unit group during both planning and execution, separately, with population decoding analysis (see *3.2.2 Decoding analysis*). Populations of PO (95% CI: [26.6, 50.3] for Monkey 1, and [30.3, 52.6] for Monkey 2) and TR (95% CI: [37.0, 60.5] for Monkey 1, and [43.1, 69.3] for Monkey 2) units were similarly modulated during planning, and each had significantly more modulation than NM (95% CI: [9.4, 22.1] for Monkey 1, and [9.6, 22.9] for Monkey 2) and EO units(95% CI: [10.8, 22.4] for Monkey 1, and [12.5, 26.3] for Monkey 2) because the confidence intervals (CI) did not overlap (Figure 3.3G). A similar pattern was apparent during execution, except here the EO (95% CI: [83.5, 96.4] for Monkey 1, and [76.5, 92.8] for Monkey 2) and TR (95% CI: [89.1, 96.9] for Monkey 1, and [82.8, 94.3] for

Monkey 2) populations had similar modulation that was significantly more than either NM (95% CI: [34.4, 59.2] for Monkey 1, and [26.4, 50.3] for Monkey 2) or PO (95% CI: [30.3, 52.6] for Monkey 1, and [32.6, 56.2] for Monkey 2) populations (Figure 3.3H). Our finding that TR units were similarly modulated during planning and execution as PO and EO units, respectively, suggests that their existence does not depend on a low threshold for unit modulation. For visualization purposes, the recorded accuracies in Figure 3.3G,H were normalized to the maximum of any randomly selected neural population to highlight the differences across groups. Reported statistics were computed using the unnormalized decoding accuracies.

The existence of TR units that had similar modulation to PO units during planning and EO units during execution refutes hypothesis 1 and is consistent with both hypothesis 2 and 3. Going forward, we narrowed our focus to distinguish between hypothesis 2 and 3.

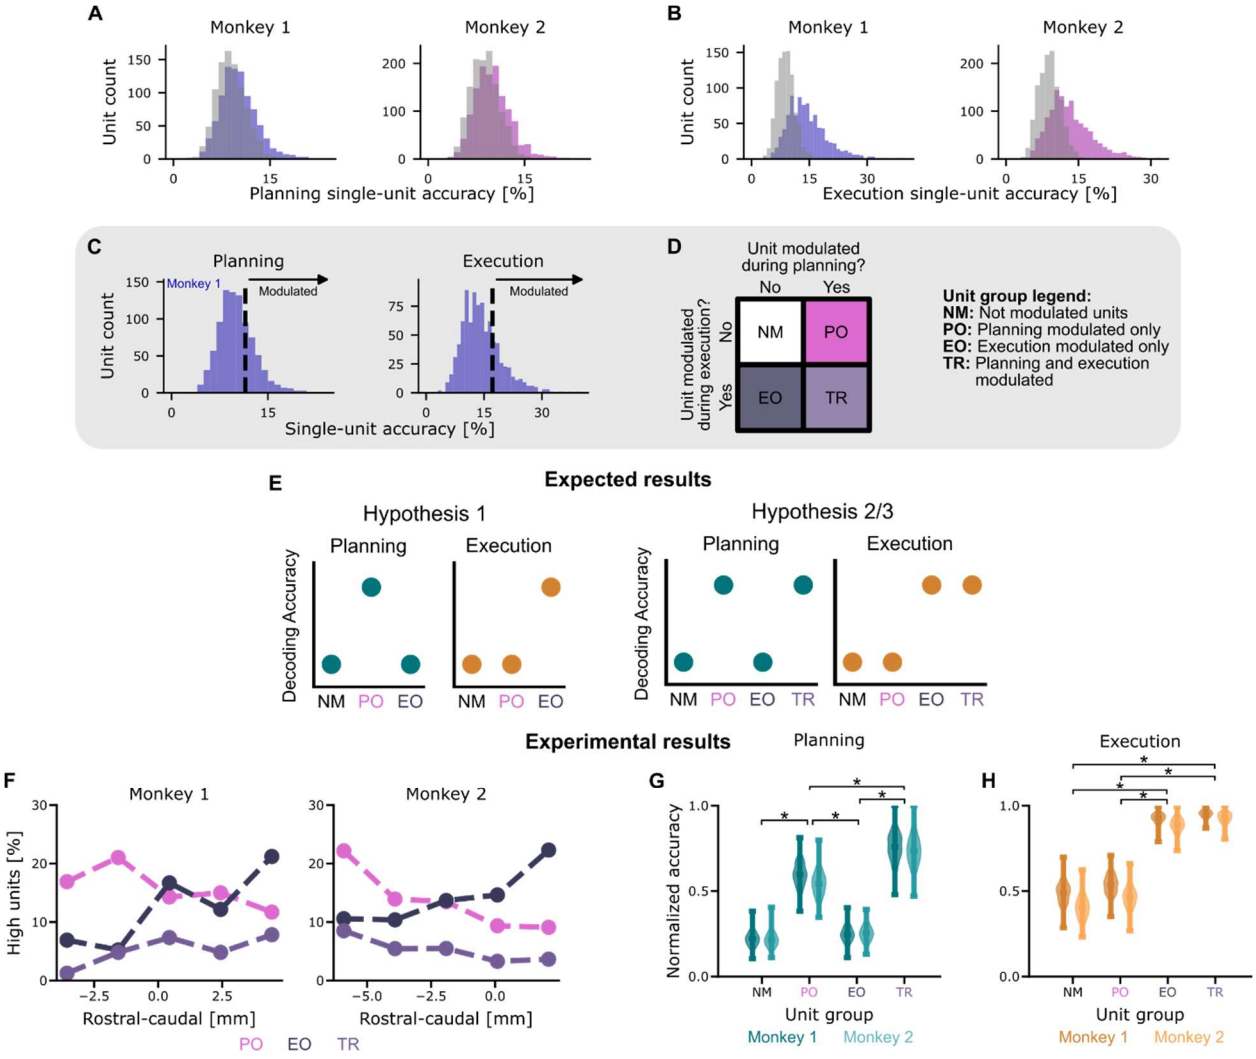

Figure 3.3. Units modulated during planning and execution are found across frontal motor cortex. (A) *Single-unit decoding accuracy distribution during planning overlaid with the chance distribution (Gray).* ***(B)*** *Same as A but during execution.* ***(C)*** *Diagram depicting how modulated units were defined based on single-unit decoding accuracy during planning (Left) and execution (Right). Units with a single-unit decoding accuracy greater than the threshold (dashed line) were considered modulated.* ***(D)*** *Diagram explaining how units were assigned to groups.* ***(E)*** *Expected decoding accuracy results for each proposed hypothesis (see Figure 3.1).* ***(F)*** *Spatial Distribution of unit groups along*

the rostral-caudal axis of frontal motor cortex where 0 approximately corresponds to the intersection of the arcuate sulcus and arcuate spur. ***(G)*** *Population decoding accuracy for each unit group normalized to the highest estimated decoding accuracy across all unit groups.*

3.3.2. Transition units have the most conserved dynamics across computations

The transition from planning to execution can be implemented in multiple ways. Previous work suggests that planning and execution related activity primarily evolve in orthogonal low-dimensional subspaces (Elsayed et al., 2016). Planning and execution are also linked through similar dynamics: the structure of dynamics during planning predicts the structure during execution (Elsayed et al., 2016). However, the link between planning and execution can either be driven by a subset of units (Hypothesis 2) or all units can contribute equally through a population mechanism such as reweighting (Hypothesis 3).

We analyzed links between planning and execution by assessing the dynamics similarity of different units groups using canonical correlation analysis (CCA) (see *3.2.3 Neural population analysis*). We computed the CCA correlation across the previously defined unit groups (see *3.3.1*) during planning to compare their dynamics similarity. Hypothesis 2 predicts that both the PO and TR populations are involved in planning, but that the EO population is not. Therefore, the dynamics of PO and TR populations should be similar to each other and distinct from the EO population. Alternatively, hypothesis 3 predicts that all units are involved in the transition between computations predicting that the dynamics of all unit groups will be similar (Figure 3.4A). As an internal control we also split each unit group into distinct neural populations giving us the estimated the dynamics similarity within and across each group (see *3.2.3 Neural population analysis*; (Figure 3.4B). Consistent with hypothesis 2, we found that TR populations (95% CI: [0.45, 0.53] for Monkey 1, and [0.53, 0.6] for Monkey 2) had more similar dynamics than the EO populations (95% CI: [0.26, 0.39] for Monkey 1, and [0.28, 0.38] for Monkey 2) because the 95% Cis did not overlap (Figure 3.4C).

We repeated this analysis to estimate the dynamics similarity across and within unit groups during execution. Here, hypothesis 2 predicts that the EO and TR populations would have the most similar dynamics, while the PO populations would be distinct. Hypothesis 3 predicts that the dynamics of all unit groups would be similar (Figure 3.4D). We again found patterns consistent with hypothesis 2 in dynamics similarity between unit groups where the EO (95% CI: [0.59, 0.68] for Monkey 1, and [0.58, 0.66] for Monkey 2) and TR (95% CI: [0.68, 0.74] for Monkey 1, and [0.66, 0.72] for Monkey 2) populations had more similar dynamics than the PO population (95% CI: [0.34, 0.45] for Monkey 1, and [0.33, 0.46] for Monkey 2; Figure 3.4E,F).

Finally, we assessed dynamics similarity of each unit group across computations by comparing the planning and execution related activity of the same neural population. Hypothesis 2 predicts that only the TR population would have similar dynamics across computations while hypothesis 3 predicts that the dynamics of any neural population would be similar. We found that the TR (95% CI: [0.46, 0.56] for Monkey 1, and [0.51, 0.61] for Monkey 2) units had more similar dynamics across computations than the PO (95% CI: [0.31, 0.44] for Monkey 1, and [0.32, 0.45] for Monkey 2) or EO (95% CI: [0.29, 0.38] for Monkey 1, and [0.31, 0.42] for Monkey 2) populations, consistent with hypothesis 2 (Figure 3.4H,I). Units with consistent tuning across events were excluded from the dynamics similarity analysis suggesting that the link between planning and execution was not primarily driven by the grouping procedure (Fig. S3.1). Instead,

the similar dynamics between planning and execution possibly reflect coordination specific to transition units.

Collectively, these results are consistent with hypothesis 2 where distinct neural populations are involved in planning and execution, but the computations are linked by a subset of units. Therefore, the transition between these linked (but spatially distinct) computations may be facilitated through a subset of units with consistent dynamics during both computations.

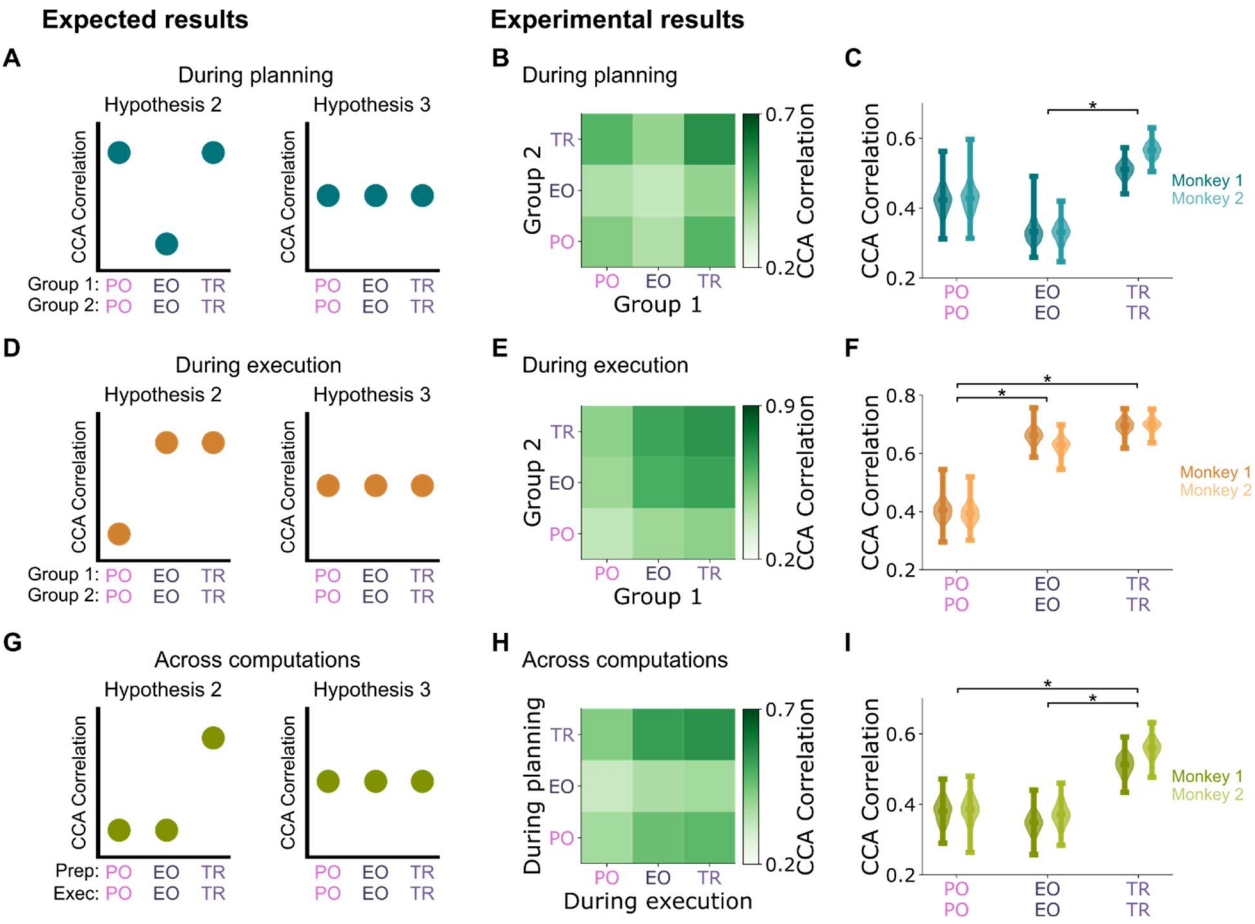

Figure 3.4. Transition units have the most conserved dynamics within and across computations. ***(A)*** *Expected CCA correlation for hypothesis 2 (left) and hypothesis 3 (right) within planning.* ***(B)*** *Average CCA between each modulated unit group.* ***(C)*** *CCA correlation distributions for select unit group comparisons, which are the x-axis tick labels.* ***(D, E, F)*** *Same as A, B, C, respectively, except during execution.* ***(G, H, I)*** *Same as A, B, C, respectively, except across planning and execution.*

3.4. Discussion

Limits in largescale recordings across frontal motor cortex have prevented investigations into how functionally and spatially distinct neural populations coordinate to transition between linked computations. We hypothesized that this transition could be implemented with distinct computation-specific populations, distinct populations with an overlapping subset, or with completely overlapping populations (Figure 3.1). We used high-density laminar microelectrode arrays to sample large neural populations across frontal motor cortex to assess how they

coordinate during and across movement planning and execution. Target decoding analysis revealed that units modulated only during planning or execution were primarily found in rostral and caudal motor cortex, respectively, as previously reported (Cisek and Kalaska, 2005). However, our large-scale sampling also revealed neurons modulated during both computations that were spatially distributed across frontal motor cortex (Figure 3.3F). These units were found to be similarly predictive of target direction during planning as planning modulated-only units and similarly predictive of target direction during execution as execution modulated-only units (Figure 3.4G). Combined, these results suggest that the transition from planning to execution is not implemented by distinct neural populations; at least some neurons are involved in both computations even though previous evidence suggests an anatomical distinction (Cisek and Kalaska, 2005; Ouchi et al., 2025). Dynamics similarity analysis further revealed links between planning and execution, as previously reported (Elsayed et al., 2016). However, not all neural populations had similar dynamics across computations. Neural populations composed of units modulated during both planning and execution had the most similar dynamics across computations. Our largescale sampling of neural populations provides new evidence that planning and execution are primarily linked by a subset of spatially distributed neurons. These results offer new perspectives on the structure underlying linked cortical computations.

Connection with other transition motifs

Our results suggest that distinct neural populations may coordinate differently during the transition between planning and execution computations. These results help reconcile previous findings that neurons modulated during planning and execution were found in spatially distinct areas (Cisek and Kalaska, 2005), yet produce coordinated activity patterns across this transition (Elsayed et al., 2016), by suggesting that a spatially distributed subset of neurons primarily links the coordinated activity patterns. Population-level perspectives also suggest that a condition-invariant signal (CIS) is involved in initiating the transition from planning to execution (Kaufman et al., 2016; Churchland and Shenoy, 2024). Our results expand these findings by suggesting that the CIS may trigger a transition that is implemented via computation-specific and computation-shared neural populations. Future investigations will be necessary to reveal how the CIS interacts with each population.

Anatomical structure

The potential relationship between population dynamics and connections between neurons (Sussillo et al., 2015) raises questions about the structural basis of the transition between motor computations. The transition from planning to execution is likely impacted by cortico-cortical and cortico-subcortical interactions because planning is facilitated by multiregional circuits (Inagaki et al., 2022). Frontal motor areas receive projections from upstream areas such as thalamus (Hooks et al., 2013) and supplemental motor areas (Ninomiya et al., 2019) providing the pathways for this facilitation. Information transfer from upstream areas and within cortico-cortical circuits is likely mediated by specific cell types because transcriptomically-defined and cortical projection cell types are known to carry distinct information (Inagaki et al., 2022). Additionally, transcriptomically-defined cell types undoubtedly play distinct roles during and across linked computations due to their unique electrophysiologic and connectivity properties (Tremblay et al., 2016; Lee et al., 2023). Future investigations with causal manipulations of specific cell types within and across areas will be critical to parse the anatomical pathways that facilitate planning and execution, and the cell types that mediate the computational transition.

Our results suggest that the information pathways in frontal motor cortex may have computational-shared and computation-specific motifs. If so, these motifs likely exist to confer a behavioral or computational advantage. One possibility is that the computation-specific populations allow parallelization while the computation-shared populations are critical to maintain the structure of information in the presence of changing computational demands. For example, the execution population may coordinate with the transition population during reaching, freeing the planning population to begin preparing for the next movement. This may be particularly important to perform sequential reaches or for naturalistic reaches with corrective sub-movements. However, it remains a question if the computation-shared population directly handles the information transfer between computations, is downstream of both computation-specific populations, or in parallel. Computational and causal manipulation experiments will be essential to parse the circuit architecture underlying the observed information transfer.

Implications for BCI

Brain-computer interfaces (BCI) are a promising motor rehabilitation therapy that map neural activity to movement, yet they almost completely disregard planning related activity. Planning related activity is important to consider because it can simultaneously occur with execution related activity when users are given little time to prepare (Lara et al., 2018b) or when multiple movements are performed in sequence (Zimnik and Churchland, 2021). BCI performance can likely be improved by accounting for planning related activity and computation-specific networks may help facilitate this. Furthermore, planning-related activity can likely be used to predict upcoming movements, thereby decreasing latency and improving performance for users. Understanding how to design BCIs that leverage our emerging understanding of cortical computations will be critical for their long-term improvement.

3.5. Appendix

3.5.1. Supplemental figures

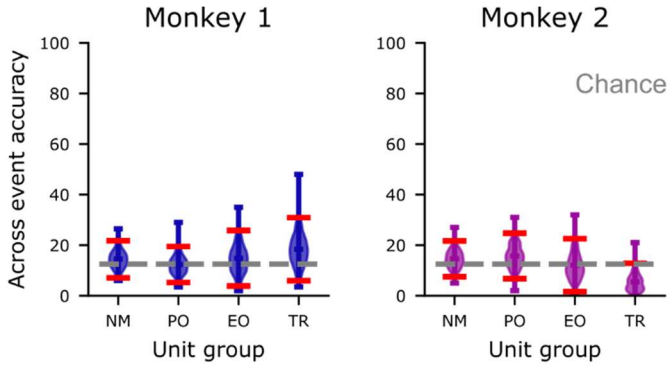

Figure S3.1 Tuning is not conserved from planning to execution for any unit group after units with conserved tuning had been excluded. *The gray dashed line represents chance decoding accuracy. Distributions are from repeating the decoding accuracy estimation with randomly selected neural populations to control for neural population size. Red lines represent the 95% confidence interval.*

4. The role of parvalbumin-expressing inhibitory neurons in visual sensitivity

4.1. Introduction

Cortical computations are implemented through the coordination of neurons. However, there are many different types of neurons. Neurons are often grouped based on whether they excite or inhibit downstream neurons, but excitatory and inhibitory neurons can be further split into specific cell-types based on their transcriptomic properties (Tremblay et al., 2016; Bakken et al., 2021; Lee et al., 2023). This level of classification is useful because these subtypes play unique roles in rodent sensorimotor computations (Chen et al., 2015; Mitani et al., 2018; Inagaki et al., 2022; Vendrell-Llopis et al., 2022) likely due to their distinct anatomical and electrophysiological properties (Tremblay et al., 2016; Lee et al., 2023). However, there is a need to study how specific cell-types facilitate cortical computation in behaving monkeys because they have diverse behavioral repertoires, their brains are more similar to humans (Van Essen et al., 2019; Bakken et al., 2021), and research pioneered in monkeys has directly translated to humans (Hwang et al., 2013; Orsborn et al., 2014; Sakellaridi et al., 2019; Silversmith et al., 2021).

The tools to record and manipulate specific cell-types in monkeys are in the early stages of development. Rodent cell-type specific experiments have been facilitated by combining optogenetics with transgenic animals, but transgenic monkeys are impractical. We achieved cell-type specific expression by injecting a replication deficient virus into the brain (El-Shamayleh and Horwitz, 2019). The virus contains an enhancer gene sequence that restricts opsin expression to our target cell population and opsins, such as channelrhodopsins, that allow the experimenter to control neural activity with light. Our experiments targeted parvalbumin-expressing (PV+) neurons because they produce widespread inhibition (Tremblay et al., 2016) and are a common cell-type found across cortex (Tremblay et al., 2016; Bakken et al., 2021).

We tested the efficacy of two virus candidates as tools to investigate the brain using histology, electrophysiology, and behavior. We focused on primary visual cortex (V1) because the spatial architecture and temporal dynamics are well understood. V1 follows a retinotopic structure where visual stimuli at specific locations in the visual field elicit localized responses in a predictable location in V1 (Talbot and Marshall, 1941; Gattass et al., 1981). Cortical responses to visual stimuli also occur at low-latency and persist while the visual stimulus is present (Talbot and Marshall, 1941). We used histology to test viral expression, electrophysiology to test if we can manipulate neural activity with light, and a contrast detection task to test if PV+ neuron activation can impact behavior by suppressing visual sensitivity.

We confirmed virus expression histologically and found the PV+ neuron specificity to be qualitatively consistent with previous reports (Vormstein-Schneider et al., 2020; Mich et al., 2021). Both viruses also allowed us to modulate neural activity with light in an awake, behaving monkey. We noticed behavioral deficits shortly after injecting one virus that prevented us from testing the role of PV+ neurons in visual sensitivity. However, we used the second virus to consistently and repeatably manipulate behavior through putative PV+ neuron activation. These results suggest that PV+ neurons play a direct role in visual sensitivity by modulating cortical information pathways. Our results highlight the need to target other cell types for cortical

computation experiments, and that a multi-step virus verification process will likely be necessary for success.

4.2. Methods

4.2.1. Histology to test virus expression

To test the expression of candidate viral vectors we made injections into rats and macaques. We produced viruses with the S5E2 (Vormstein-Schneider et al., 2020) and eHGT140h (Mich et al., 2021) enhancer sequences based on previous evidence suggesting that they drive selective expression in PV+ neurons. The enhancer sequences were combined with gene sequences to drive channelrhodopsin (ChR2) and fluorescent reporter (mCherry or eYFP) expression only in PV+ neurons. The two viral vectors we injected were:

1. AAV(php.eb) – S5E2 –ChR2 – mCherry
2. AAV(php.eb) – eHGT140h – ChR2 – eYFP

We injected virus 1 into the cortex of a macaque (*Macaca nemestrina*) and injected virus 2 into the cortex of a rat (*Rattus norvegicus*) due to limited monkey availability. Animals were perfused four weeks after the injection to allow time for protein expression. We sliced the brain in 50 µm sections and co-labelled cells for PV+ and the fluorescent reporter using immunohistochemistry (IHC).

4.2.2. Neural recordings

Neural activity was recorded from V1 in the left hemisphere of a female rhesus macaque (*Macaca mulatta*). We used extracellular tungsten microelectrodes (Alpha-Omega, Norcross, GA) to record electrical neural signals that were amplified, digitized, then recorded (Plexon, Dallas, TX). Electrodes were manipulated with a hydraulic microdrive (Narishige International USA, Inc., Amityville, NY). Offline spike sorting was performed manually with the Plexon offline sorter.

4.2.3. Injections into awake behaving macaque cortex

Macaque V1 is a structurally complex brain area. There is cortical gray matter (“Deep V1”) underneath superficial cortical areas (“Superficial V1”) and each responds to visual stimuli in different locations of the visual field (Gattass et al., 1981) (Figure 4.1A). We targeted our viral injections to locations in V1 where we could access both deep and superficial V1 with a single penetration track (Figure 4.1A). The goal of our injection strategy was to produce ChR2 expression in multiple parts of V1 that responded to visual stimuli in different visual field locations (Figure 4.1B).

We injected each virus after recording neural activity to identify the visual field locations (e.g. receptive fields) that elicited neural responses from the injection site. The monkey fixated at the center of the screen while an experimenter manually moved a bar, the visual stimulus, across different visual field locations (Figure 4.1B). There was an increase in recorded neural activity when the bar was in the receptive field of the injection site, providing us with an understanding of how the area of V1 that we were injecting into responded to visual stimulus. Multiple 1µL injections were made along each penetration track. Injection sites began 1mm below where neural activity was recorded (Below ‘Deep V1’), ended at the cortical surface, and were spaced every 0.5mm. The cannula was held at each injection site for 1 minute before proceeding. The

cannula was then held at the surface of the brain for 10 minutes before it was completely withdrawn.

We injected the two viruses in different regions of V1 in the same awake and behaving monkey. These regions were approximately 10 mm apart to minimize the possibility of overlap. Two penetration tracks, 2 mm apart, were made for each virus. The two viruses we used were the same used in histology experiments (see *4.2.1 Histology to test virus expression*). The injection of virus 1 was conducted with a Hamilton syringe and the injection of virus 2 was performed with an air compressor following established protocol (Kojima et al., 2021). We waited 4 weeks after each injection before conducting further experiments to allow time for expression.

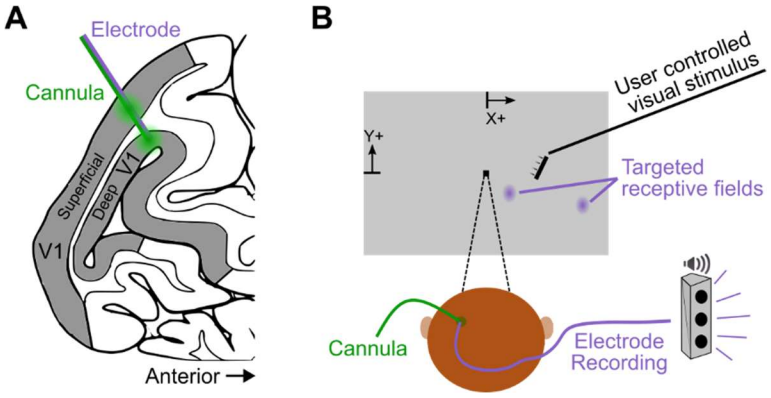

Figure 4.1. Virus injection procedure (A) *Representation of V1 targets for injection. Green circles indicate the targeted areas of cortex for injection.* ***(B)*** *Receptive field mapping during the injection. Purple dots represent the receptive field of corresponding to the injected cortex shown in A.*

4.2.4. Laser manipulation of neural activity

We returned to the injection site with an electrode and optical fiber (Thorlabs) to test for physiological levels of ChR2 expression. The electrode and optical fiber were independently manipulated by a dual-arm hydraulic micromanipulator (Narishige International USA, Inc., Amityville, NY) through a common guide tube. We first lowered the electrode alone to record neural activity and locate the receptive field corresponding to the part of V1 being recorded. Once we verified that the receptive field location of the recorded neurons matched the receptive field location of the injection, we slowly lowered the optical fiber to be approximately 250-500μm above the electrode tip. We manually pulsed the laser and listened for a strong auditory response in the recorded neural activity. Once found, we recorded the neural response to a sinusoidal light pattern ensuring that the recorded neural activity consistently responded to light.

4.2.5. Behavioral Paradigms

During all behavioral paradigms the monkey was headfixed and placed 60 cm away from a screen. Trials were initiated when the monkey fixated within a 1° by 1.2° box around a dot at the center of the screen. Eye position was tracked with a scleral search coil.

Visually guided saccade task

We trained a monkey to perform a visually guided saccade task to estimate visual deficits (Figure 4.2). After fixation, a target appeared at varying locations in the monkey's peripheral vision. When the target appeared, the monkey was rewarded for making saccades to the target. We quantified the latency between target onset and saccade initiation and used it as a proxy for visual processing speed. We mapped the monkey's visual processing speed across the visual field by moving the target around the visual field.

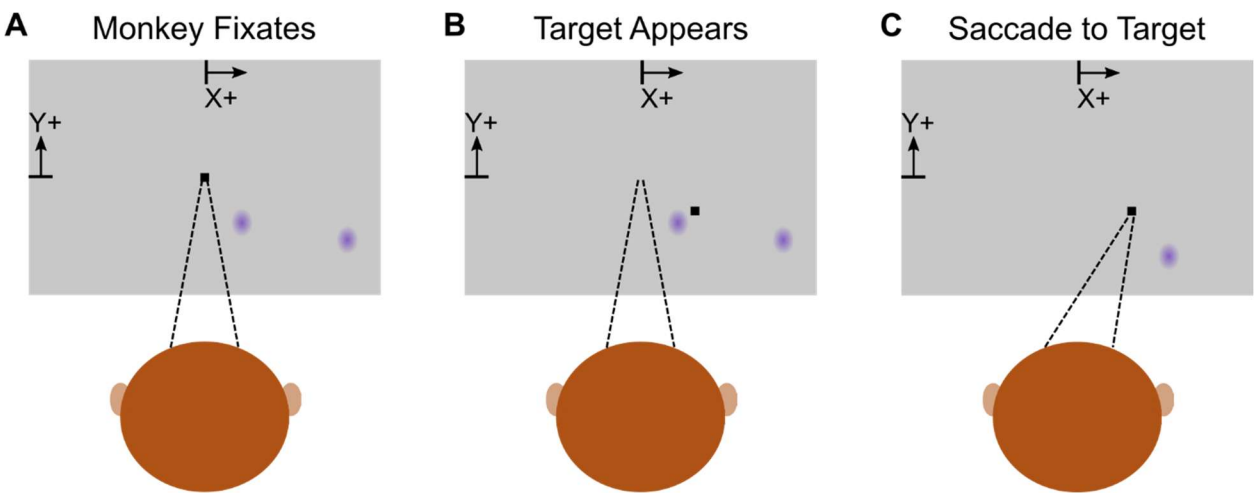

Figure 4.2. Visually guided saccade task diagram. (A) *The monkey initiates the trial by fixating.* ***(B)*** *During fixation a target appears elsewhere in the visual field that cues the monkey to saccade* ***(C)*** *The monkey is rewarded for making a saccade to the target (right). The purple dots in all panels represent the approximate receptive field locations found during virus injection (see Figure 4.1).*

Contrast detection task:

We trained a monkey to perform a contrast detection task to test how PV+ neuron activation impacted visual sensitivity (Figure 4.3). The monkey initiated the trial by fixating within an 1° by 1.2° box around a dot at the center of the screen. During fixation, a gabor stimulus appeared within the receptive field of the neurons expressing ChR2 on half the trials, then disappeared before fixation was broken. When the fixation point disappeared, two targets equally spaced appeared 2° away from the original fixation point along the horizontal axis. To indicate that the visual stimulus was present the monkey looked to the right, to indicate that the stimulus was absent the monkey looked to the left. Each trial resulted in one of four possibilities:

1. The visual stimulus was present and the monkey correctly identified it (“Hit”).
2. The visual stimulus was absent and the monkey correctly identified that (“Correct Reject”.
3. The monkey incorrectly indicated that the visual stimulus was present (“False alarm”).
4. The monkey incorrectly indicated that the visual stimulus was absent (“Miss”).

To ensure that the task was hard, but not impossible, the contrast of the visual stimuli was gradually stepped. If the monkey received rewards on three consecutive trials, the contrast of the visual stimulus was decreased, making the task harder. If the monkey was unrewarded on one trial, the contrast of the visual stimulus was increased making the task easier. After enough trials, the contrast stabilized at the monkey’s contrast threshold where the monkey was

rewarded on most of the trials and unrewarded on a few. When the monkey consistently achieved the same contrast threshold across multiple sessions and multiple days, we considered them well trained.

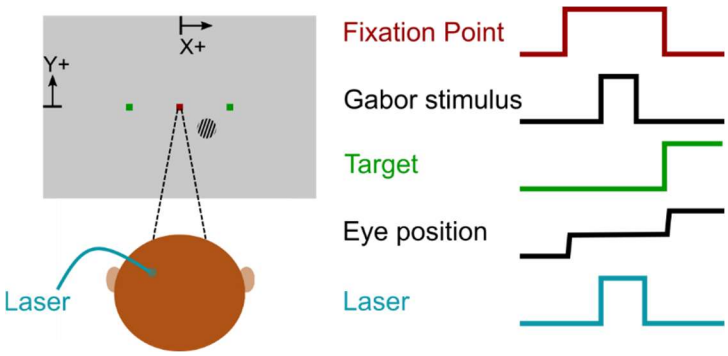

Figure 4.3. Contrast detection task summary. *Trials began when the monkey fixated on the center. During fixation the gabor stimulus could appear in the targeted receptive field. On some stimulus present trials the laser was also activated during stimulus presentation. After the stimulus disappeared, two targets appeared and the monkey indicated that the stimulus was present by looking right and that the stimulus was absent by looking left.*

When the monkey was well trained, we placed the visual stimulus in the receptive field of the ChR2-expressing neurons and provided optical stimulation simultaneously with visual stimulus. The laser was present on half of the trials that had the visual stimulus present and half of the trials that had the visual stimulus absent, but the contrast changed for laser and non-laser trials independently. Each trial within this task structure resulted in one of eight possibilities, Hits, Misses, CRs, and FAs for both laser and non-laser trials.

4.2.6. Optotagging

Optotagging is a valuable tool to identify neurons expressing ChR2, and thus a specific neuron subtype, without directly impacting behavior (Lima et al., 2009; Estebanez et al., 2017). When a neuron was recorded, we used intermittent square wave light pulses to stimulate the neuron. If the neuron responded to light onset with low-latency (~<5ms) we assumed that the neuron was expressing ChR2 and likely a PV+ neuron (Lima et al., 2009). Once a PV+ neuron was identified, we flashed an array of visual stimuli in its receptive field to identify the specific visual stimuli that drove its activity. We characterized any putative PV+ neurons with white noise color stimulus where random color stimuli patterns were presented in the RF of the recorded neuron. Then, we computed the spike-triggered average and spike-triggered variance for each unit using the stimuli patterns that immediately preceded each spike. Further details can be found in (De and Horwitz, 2021).

4.3. Results

We produced two candidate viruses to manipulate and record from PV+ neurons. These viruses used the S5E2 (Vormstein-Schneider et al., 2020) and eHGT140h (Mich et al., 2021) enhancer sequences to drive Channelrhodopsin-2 (ChR2) expression to modulate neural activity with light, and fluorescent reporter expression to histologically identify transduced neurons. The two viruses that we produced and tested were:

1. AAV(php.eb) – S5E2 –ChR2 – mCherry
2. AAV(php.eb) – eHGT140h – ChR2 – eYFP

4.3.1. Virus expression was verified histologically

We confirmed virus expression histologically with test injections in two anesthetized animals. Animals were perfused 4 weeks after injection, then we labelled neurons expressing parvalbumin and neurons expressing the fluorescent reporter. We found neurons expressing the fluorescent reporter for both viruses (Figure 4.4). For virus 1, mCherry-expressing neurons were labelled red and PV+ neurons were labelled green. For virus 2, eYFP-expressing neurons were labelled green and PV+ neurons were labelled red. For both viral vectors, the superposition of green and red labels (yellow and orange) indicate that viral expression was driven in a PV+ neuron. Any completely red neuron in Figure 4.4A or completely green neuron in Figure 4.4B suggests off-target expression. We visually confirmed that there are a minimal number of neurons with off-target expression. A full quantification of PV specificity was conducted in previous work (Vormstein-Schneider et al., 2020; Mich et al., 2021). Collectively, this provided confidence in the performance of our viral vector and we prepared to inject it into a behaving monkey.

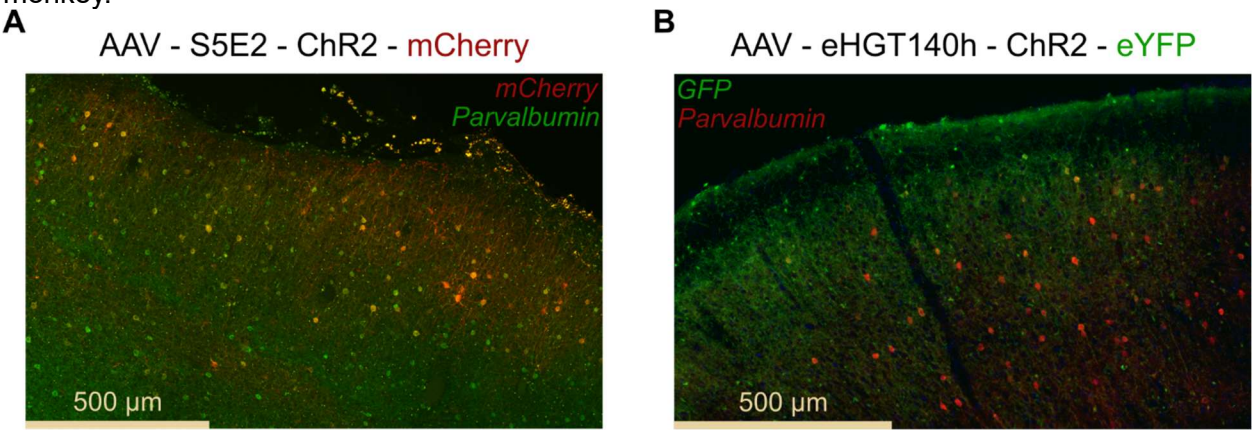

Figure 4.4. Histology from after virus injection confirms expression. (A) *Example tissue section after the injection of AAV-S5E2-ChR2-mCherry and counterstains for mCherry (red) and parvalbumin (green).* ***(B)*** *Example tissue section after the injection of AAV-eHGT140h-ChR2-eYFP and counterstain for green fluorescent protein (GFP; green) and parvalbumin (red).*

4.3.2. S5E2-ChR2 virus injection

Recording neural activity while optically stimulating the tissue at the injection site of Virus 1 (AAV – S5E2 – ChR2 – mCherry) revealed light-activated neurons. We pulsed the laser in a sinusoidal pattern and found spiking responses repeatably tied to power onset (Figure 4.5).

We noticed a lack of overall neural activity at the injection site suggesting tissue damage. With and without the laser, neurons were sparse. Lesions and tissue damage in V1 can lead to scotomas localized to the corresponding area in the visual field (Tong, 2003) allowing us to behaviorally test the retinotopic extent of the unhealthy tissue with a visually guided saccade task (see *4.2.5 Behavioral Paradigms*).

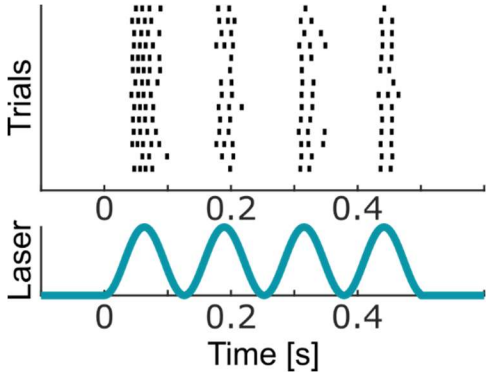

Figure 4.5. Example raster from a light responsive neuron transduced by the S5E2 virus (Virus 1). *Top: Each black dash represents an action potential, and each row is a trial. Bottom: Laser power modulation.*

The monkey performed a visually guided saccade task with targets that tiled the receptive fields of the injection sites and control sites on the other side of the screen. We expected that targets in the visual field areas corresponding to scotomas would be more difficult for the monkey to detect. We quantified performance by calculating the latency between target appearance and saccade initiation for each target (Figure 4.6), which gives us a proxy for visual processing speed. We found that saccade initiation took longer towards targets in the receptive field locations corresponding to the injection sites relative to control sites (Figure 4.6). The existing behavioral deficits and likely cortical damage at the injection sites made us unable to continue testing our ability to manipulate behavior by activating PV+ neurons with virus 1.

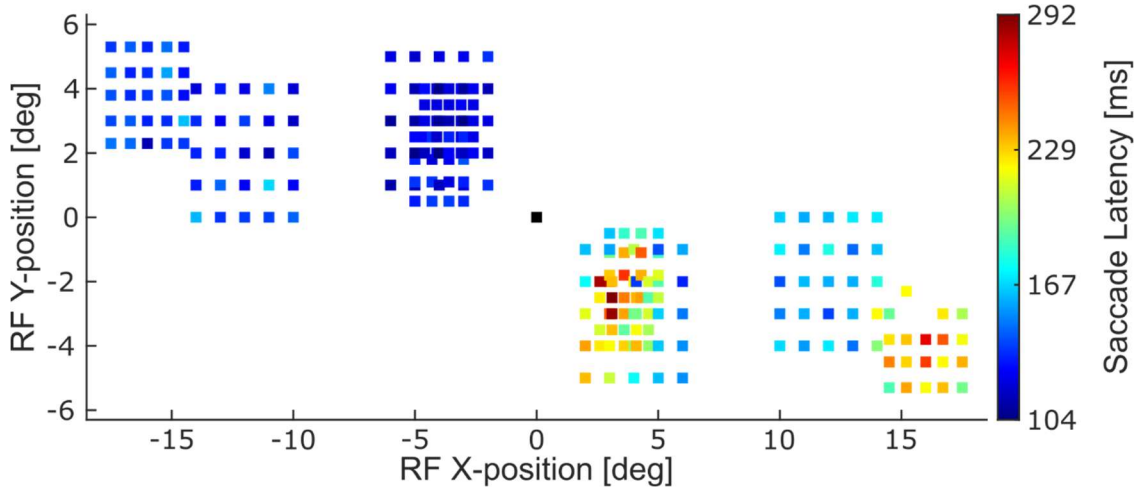

Figure 4.6. Behavioral deficits are localized to the targeted receptive fields. *Latency map between target onset and saccade initiation. Each dot represents a visual target and its location in the visual field.*

4.3.3. eHGT140h-ChR2 virus injection

Recording neural activity while optically stimulating the tissue at the injection site of Virus 2 (AAV – eHGT140h–ChR2 – mCherry) also revealed light-activated neurons (Figure 4.7). The light-activated neurons reliably responded to optical stimulation in a sinusoidal pattern consistent with physiological levels of ChR2 expression.

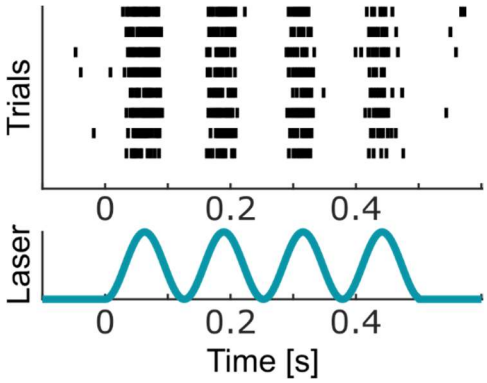

Figure 4.7. Example raster from a light responsive neuron transduced by the eHGT140h virus (Virus 2). *Top: Each black dash represents an action potential, and each row is a trial. Bottom: Laser power modulation.*

There was no evidence of neuronal damage at these injection sites so we proceeded to test if we could modulate behavior. We found a significant decrease in Hits and a corresponding increase in Misses when we placed the visual stimulus in the receptive fields of the ChR2-expressing neuron and introduced optical stimulation (Figure 4.8A). The contrast trajectory continued to increase on laser trials making the task easier, yet the monkey kept getting those trials wrong (Figure 4.8B). On laser trials the monkey made more mistakes with an easier task consistent with an inability to detect the visual stimulus. We immediately repeated the same experiment to verify that this result was not due to chance and found the same effect (Figure 4.8B). At control sites not expressing ChR2, we found an even number of Hits and Misses between laser and

non-laser trials (Figure 4.8C). The contrast trajectories for laser and non-laser trials were also similar. The monkey was correctly identifying the visual stimulus when it had the same contrast regardless of optical stimulation (Figure 4.8D). These results suggest that the behavioral effects of optical stimulation were specific to the sites expressing ChR2.

We repeated the experiment across days and penetrations to verify that the observed effect was repeatable. We compared the proportion of hits on laser trials to hits on non-laser trials for each session (Figure 4.8E). We found either the same or higher proportion of hits on non-laser trials compared to laser trials across all sessions suggesting that modulating PV+ neurons in V1 produces a behavioral effect in macaques by reducing visual sensitivity.

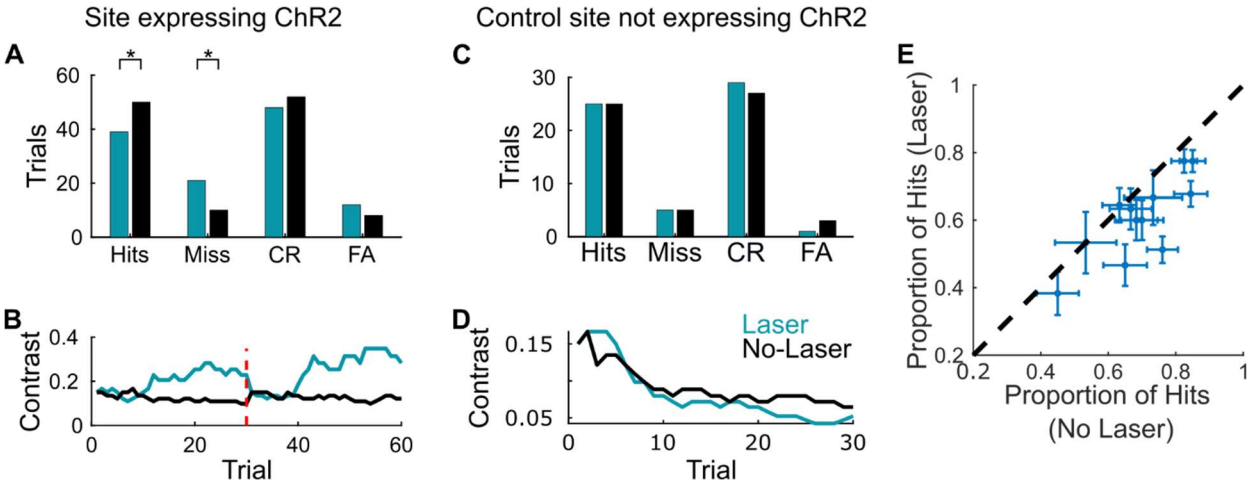

Figure 4.8. Optical stimulation of ChR2 expressing neurons reduces visual sensitivity. ***(A)*** *Results when optically stimulating a site expressing ChR2.* ***(B)*** *The contrast trajectory split by trials with and without light activation. The red line represents the point where the experiment was repeated and therefore the contrast was reset* ***(C)*** *Results when optically stimulating a site not expressing ChR2.* ***(D)*** *The contrast trajectory with and without the light activation.* ***(E)*** *Population results of optical stimulation. Each data point represents a different session. Error bars represent the 95% confidence interval. The black dashed line represents the identify line.*

4.3.4. Optotagging

Manipulating behavior by activating or silencing specific cell-types is a valuable tool to understand cortical computations. Another valuable application of optogenetics is optotagging, which can identify if the recorded neuron is the targeted cell-type by testing for ChR2 expression. Optotagging holds promise to reveal how different cell-types contribute to naturally occurring computations. We performed preliminary optotagging experiments in V1 as a proof-of-principle. After many days we found one putative PV+ neuron based on the low-latency (5ms) response to light stimulation (Figure 4.9). We characterized its response to visual stimuli through white noise color stimuli. We found no significant spike-triggered average (STA) result but found a significant result in the spike-triggered variance (STV) of the green channel, suggests that the PV+ neuron we were recording was a complex cell (Horwitz et al., 2007).

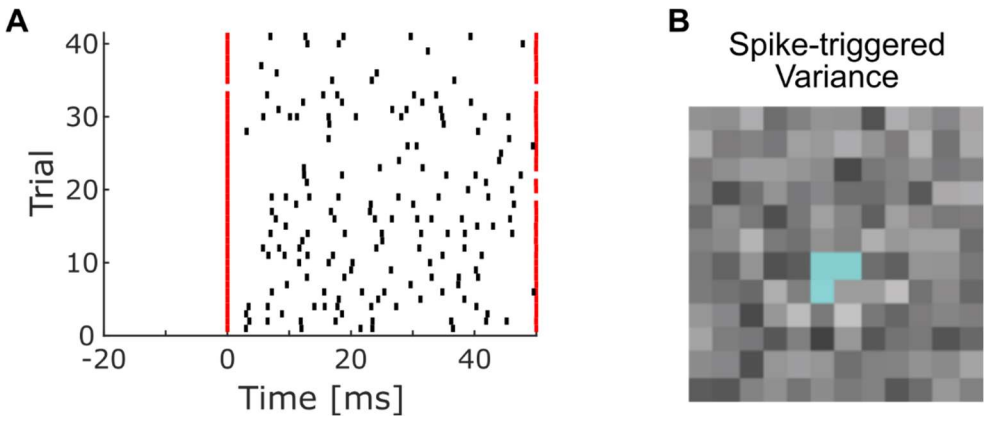

Figure 4.9. Potentially optotagged unit that responds to the laser with low latency. (A) *Raster of the unit's activity around laser onset. The red lines indicate laser on and laser off.* ***(B)*** *The significance map of the spike-triggered variance in the green color channel.*

4.4. Discussion

These experiments show that physiological levels of ChR2 expression can be achieved in a behaving monkey under the S5E2 and eHGT140h enhancers. These experiments also show that stimulating PV+ neurons in monkey V1 can induce behavioral changes, consistent with our hypothesis. It is tempting to compare the results from the two different enhancers but there were differences in the injection procedure making them not directly comparable.

It is notable that optical stimulation did not produce a significant behavioral result in every session that the monkey performed the contrast detection task. We expect that cortical magnification is one factor that plays a major role in this result. Cortical magnification is the phenomenon where locations near the fovea retinotopically map to larger areas of cortex compared to more eccentric visual field locations. The contrast detection experiments were conducted with RFs near the fovea, at approximately 2° of eccentricity. Therefore, it is possible that the optogenetic silencing impacted only a subset of the neurons whose receptive fields contained the visual stimuli allowing the monkey to effectively 'see' part of it.

These results represent an initial foray into leveraging optogenetic tools to understand the structure underlying cortical computations. The sparsity of putatively optotagged neurons highlights the need to combine cell-type specific optogenetic tools with high-density microelectrode arrays. As cell-type specific optogenetic tools continue to become more accessible for primate use, our understanding of structure across other cortical areas must improve to maximize their impact. We were able to test the behavioral effects of modulating PV+ neurons because the well understood spatiotemporal structure of V1 directed our hypothesis. These results highlight the value of understanding the spatiotemporal structure of cortex to form hypotheses about the role of specific cell types in other cortical computations.

5. Implications structure on BCI design

Motor brain-computer interfaces (BCI) provide an artificial means to control movements through newly created pathways. They record neural activity from portions of natural motor circuits, translate recorded neural activity into movement commands, and send those commands to a device. Sensory feedback of movement closes the control loop. In this section, we focus on intracortical motor BCIs in which sensory feedback is provided through native pathways, which represent the majority of existing research. An example is BCI control of a computer cursor using visual feedback. Historically, motor BCI design focused solely on maximizing the ability to predict a subject's intended movements. For example, algorithms to map neural activity into motor commands (the decoder) are commonly chosen on the basis of predictive accuracy in data sets where subjects performed or imagined moving as neural activity was recorded. This prediction is termed open-loop because the user does not have real-time feedback of decoded movements. Improvements in open-loop predictive performance do not necessarily translate into improved performance in closed-loop BCIs (Chase et al., 2009; Koyama et al., 2010). The difference between open-loop prediction accuracy and closed-loop BCI performance stems from the inherent integration of sensory and motor function. Creating artificial motor pathways also alters available sensory information, such as proprioception (the sense of the body's position and movement through space), and the link between sensory and motor information. That this does not catastrophically disrupt motor performance is likely due to mechanisms that continually update our sensorimotor systems to new environments and tasks. As a result, algorithms that optimally decode neural activity offline can have minimal benefits in closed-loop systems (Chase et al., 2009). The differences between open- and closed-loop systems change the nature of engineering problems for motor BCIs. BCI design cannot only consider prediction of motor intent—we must also consider how the full BCI sensorimotor system will be controlled and learned. Optimizing closed-loop motor BCIs will therefore require insight into the principles of sensorimotor control and learning in BCIs. In the following subsections, we summarize key observations about learning in motor BCIs and review how elements within a BCI may influence these processes.

5.1. Learning in Closed-Loop Motor Brain–Machine Interfaces

BCIs present the brain with two interrelated problems (Figure 5.1). The brain must first generate consistent neural activity in the features chosen for BCI control (readout) and then learn how a given pattern of neural activity relates to movement. To generate patterns of neural activity specific to the readout population, the brain may have to perform some form of credit assignment (Krakauer et al., 2019), where movement errors are attributed to particular neurons or connections. A growing body of research has explored learning in motor BCIs. These studies highlight a rich set of phenomena related to mapping and credit assignment problems with strong similarities to those of natural sensorimotor learning (Orsborn and Pesaran, 2017).

Perturbing BCI decoders has shed light on how the brain adapts to altered sensorimotor maps. Much like force-field and visuomotor perturbations (Bastian, 2008; Krakauer et al., 2019), these experiments change how the activity of readout neurons relates to movement variables. Perturbations that rotate neuron-behavior relationships produce deviations in movement trajectories (Jarosiewicz et al., 2008; Chase et al., 2012; Sadtler et al., 2014; Zhou et al., 2019) or goals (Hwang et al., 2013; Sakellaridi et al., 2019) that the brain can counteract in tens to hundreds of trials. Perturbations have been applied to both single neurons (changing the

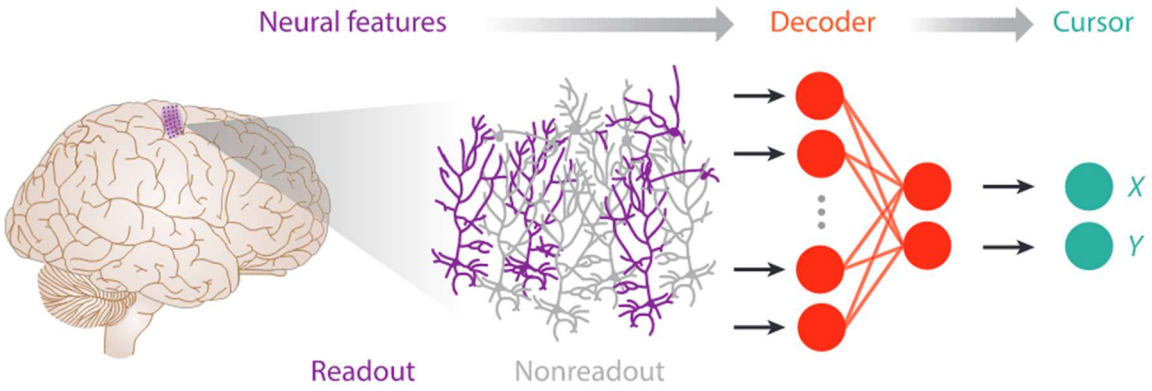

Figure 5.1. Motor brain–machine interfaces define multistep sensorimotor transformations. *This illustration highlights the transformation in an example interface where the firing of 100 neurons in the primary motor cortex (M1) controls the velocity of a two-dimensional cursor through a linear decoding algorithm. Of the many brain areas that contribute to movement, we chose to measure from one (M1). Our sensors measure only a small fraction of the thousands of neurons in M1. These 100 neurons now define the motor output (behavioral readout; purple), while all other neurons can only indirectly contribute to movement (nonreadout; gray). Finally, the decoding algorithm maps a pattern of neural activity in the 100 readout neurons into a two-dimensional velocity.*

direction in which a neuron moves the cursor) (Jarosiewicz et al., 2008; Chase et al., 2012; Hwang et al., 2013; Sakellaridi et al., 2019; Zhou et al., 2019) and neural populations (changing how neural firing patterns moves the cursor) (Sadtler et al., 2014). Behavioral adaptation to decoder perturbations resembles natural visuomotor adaptation, including the learning timescale and the presence of aftereffects following the removal of perturbations.

Neural plasticity during adaptation is distributed. Neural changes are consistent with subjects changing their aim (Jarosiewicz et al., 2008; Hwang et al., 2013; Golub et al., 2018; Sakellaridi et al., 2019), which involves learning a new association between the target location and the action chosen but does not require altering movements themselves. Reaiming produces changes in all neurons in a population (Hwang et al., 2013) but does not change the patterns of neural activity generated—only how they are associated with movement targets (Golub et al., 2018). While these global learning mechanisms dominate, if only a subset of neuron–movement relationships are perturbed, learning leads to a mixture of global and targeted changes (Jarosiewicz et al., 2008; Chase et al., 2012; Zhou et al., 2019). The presence of multiple learning mechanisms is consistent with the current understanding of computations driving visuomotor adaptation (Taylor et al., 2014).

The form of decoder perturbation also influences learning. A series of experiments used decoder perturbations structured so that the experimenters could control whether the perturbation required neural populations to change their correlation structure. Mappings where neuron correlation structures do not have to change are learned within a day ((Sadtler et al., 2014); discussed above). In contrast, perturbations where correlation structure must change require multiple days of training and lead to new neural activity patterns (Oby et al., 2019). Learning mechanisms in natural sensorimotor control are also influenced by the type of visuomotor perturbation. For instance, rotations in movement–vision relationships appear to lead to updates to an existing internal model of sensorimotor relationships, while mirror reversals lead to the creation of new models (Yang et al., 2021). Potential correspondence

between natural sensorimotor learning mechanisms and observations in BCI remain to be established, but existing research clearly demonstrates the presence of multiple learning mechanisms that are flexibly deployed (Orsborn and Pesaran, 2017).

Another line of research explores how the brain learns initial control of a BCI. These experiments define some fixed neuron–movement relationship in a task structure that resembles de novo skill learning in natural sensorimotor systems (Krakauer et al., 2019; Yang et al., 2021). These studies have found that the activities of single neurons (Fetz, 1969; Moritz and Fetz, 2011), neural populations (Ganguly and Carmena, 2009), and local field potentials (LFPs) (Wander et al., 2013) can all be operantly conditioned. When using neuron populations for control (most comparable to the decoder perturbation studies discussed above), learning to control a new decoder mapping takes multiple days (Ganguly and Carmena, 2009), and performance improves both with practice and after breaks in training (Gulati et al., 2014, 2017). Once learned, these mappings can be recalled and resist interference (Ganguly and Carmena, 2009), paralleling natural motor skills (Dayan and Cohen, 2011).

Distributed neural plasticity underlies the acquisition of control of a new decoder. Improvements in performance coincide with the formation of stable relationships between the activity of readout neurons and movement (Ganguly and Carmena, 2009). Neural patterns are variable initially, but the variance gradually decreases with learning (Athalye et al., 2017, 2018). Both patterns are similar to observations in natural skill acquisition (Peters et al., 2014; Dhawale et al., 2017). Experiments measuring activity across multiple brain areas but using only one for the BCI readout demonstrate that neural changes are distributed across cortical and subcortical regions (Koralek et al., 2012, 2013; Wander et al., 2013; Liu and Schieber, 2020). Blocking plasticity in the striatum can also extinguish the ability to learn a BCI decoder controlled by primary motor cortex (M1) activity (Koralek et al., 2012).

BCI learning may involve credit assignment computations on long (multiday) timescales. Plasticity when learning a novel decoder leads to differential activation of readout neurons in comparison to nonreadouts, even when they are immediately next to one another (Ganguly et al., 2011; Arduin et al., 2013; Clancy et al., 2014; Gulati et al., 2017). Similarly, while brain areas not used for readout, such as the striatum, are needed for learning, changes in striatal neurons are targeted to those that project to readout neurons (Koralek et al., 2013), consistent with plasticity targeted to a readout-specific neural circuit. Perturbing the decoder for only a subset of readout neurons produces changes specific to perturbed units that emerge over days (Zhou et al., 2019). These credit assignment–related computations may be sleep dependent (Gulati et al., 2017), consistent with their emergence over multiday training.

While many questions remain, these studies highlight that closed-loop BCIs engage a diverse range of learning mechanisms that parallel innate functions of natural sensorimotor systems. Learning is directly related to the sensorimotor transformations defined by the BCI system, which include all aspects of the artificial control loop from neural recording to sensory feedback. Optimizing BCI performance, then, will require considering how each element of a BCI system influences this transformation and learning computations.

5.2. Brain Areas

A key choice in a motor BCI is which part of the brain to record from. Differences in the computations performed by each area influence the ability to predict motor intention from neural

activity. Different brain areas also play different functional roles in sensorimotor control and learning computations (Cisek and Kalaska, 2005; Krakauer et al., 2019), which may influence closed-loop performance (Gallego et al., 2022).

5.2.1. Motor areas

Motor BCI research originated in M1 (Fetz, 1969). Anatomically, M1 is positioned to contribute to motor output: It is the cortical area with the largest proportion of cells that directly project to the spinal cord (Dum and Strick, 2004). Its diverse inputs also include peripheral sensory afferents, highlighting contributions to sensorimotor computations (Dum and Strick, 2004). How M1 represents movement variables is an area of active debate (Omrani et al., 2017). However, M1 activity accurately predicts a variety of kinematic (Carmena et al., 2003; Gallego et al., 2020) and other dynamic variables, including muscle activity (Cherian et al., 2011), in open-loop systems. In consequence, M1 is often used for continuous BCIs such as controlling moment-by-moment cursor velocity.

Premotor cortical areas have also been explored for BCI control. These areas are anatomically defined as frontal regions that send projections primarily to M1; many also send projections to the spinal cord (Dum and Strick, 2004). Therefore, they are positioned to support preparation and execution of movements. Premotor area computations are not fully understood, but they are distinct from those of M1. For instance, dorsal premotor cortex (PMd) activity precedes that of M1, suggesting contributions to planning (Cisek and Kalaska, 2005). This observation has motivated the use of premotor areas to control discrete BCI tasks, such as selecting sequential target pairs (Shanechi et al., 2012). However, PMd activity is also often combined with M1 for continuous control tasks (Gilja et al., 2012; Orsborn et al., 2014; Stavisky et al., 2015).

The parietal cortex, particularly the posterior parietal cortex (PPC), is another motor area used for BCI control (Mulliken et al., 2008). Anatomical connections position the PPC to transform sensory inputs to motor outputs: It predominantly receives input from sensory areas and sends dense projections to frontal motor areas (Dum and Strick, 2004; Whitlock, 2017). Motor-related activity in the PPC plans movement goals (Andersen and Buneo, 2002) in visual coordinates (Batista et al., 1999) rather than muscle activity; as a consequence, the PPC has been used primarily for discrete BCI tasks.

5.2.2. Nonmotor areas, noncortical areas, and multiple areas

Building upon research showing that neural activity in many cortical regions can be operantly conditioned (see *5.1 Learning in Closed-Loop Motor Brain–Machine Interfaces*), select studies have used brain areas outside the motor system for BCI readout. For example, rodents can control an auditory pitch cursor using activity in the primary visual cortex (V1) (Neely et al., 2018), learning with the same timeline and accuracy as using M1 activity (Koralek et al., 2013); this observation suggests common principles across brain areas. Interestingly, animals could learn to control a V1 BCI with the lights on or off, but could not generalize across conditions (Neely et al., 2018). Innate functions and connections between brain areas, then, likely influence learning processes and functionality. The distributed nature of motor-related signals in the brain (Musall et al., 2019), combined with the anatomical flexibility provided by learning, opens opportunities for clinical applications where people have damage to motor areas (e.g., cortical stroke) (Ganguly et al., 2009).

Subcortical areas have not been used for direct BCI device control. However, motivated by subcortical contributions to learning, researchers have used striatal signals to adapt cortical decoding algorithms (Mahmoudi and Sanchez, 2011). Such multiarea decoding strategies may be particularly valuable given that cortical–subcortical interactions are necessary for BCI learning (Koralek et al., 2012; Neely et al., 2018).

Leveraging multiarea activity may require additional insights into learning. Early BCI studies combined activity from many cortical areas (Carmena et al., 2003). Open-loop decoding suggested that M1 was most predictive of arm-movement kinematics, but after closed-loop training all areas contributed to BCI movement. Alternatively, extended training with electrocorticography (ECoG) BCIs across multiple areas resulted in the consolidation of control to a small region (Silversmith et al., 2021). The dynamic nature of closed-loop BCIs presents challenges in a priori selection of brain regions for motor output.

5.2.3. Learning differences across areas

Few studies have quantitatively compared closed-loop BCI performance between areas. Synthesis across the literature, however, highlights important qualitative differences when BCIs are controlled with different brain areas.

Learning dynamics differ in BCIs controlled with frontal versus parietal motor areas. Studies in M1 and PMd show that the brain can learn a wide range of mappings (Fetz, 1969; Ganguly and Carmena, 2009; Moritz and Fetz, 2011; Oby et al., 2019). This learning can occur within a session (Hwang et al., 2013; Sadtler et al., 2014) or across multiple days ((Ganguly and Carmena, 2009; Oby et al., 2019; Zhou et al., 2019); see *5.1 Learning in Closed-Loop Motor Brain–Machine Interfaces*). Rapid learning is dominated by global strategies like reaiming (Golub et al., 2018). Longer timescales involve changes targeted to neurons used for BCI control (readouts) (Ganguly et al., 2011; Koralek et al., 2012; Clancy et al., 2014; Zhou et al., 2019). In contrast, learning in BCIs using the PPC for control appears to be dominated by rapid learning mechanisms. Both humans and animals struggle to learn PPC BCI mappings that cannot be solved by reaiming strategies, and neural changes are global (Hwang et al., 2013; Sakellaridi et al., 2019). Differences in observed learning dynamics may stem from differences in computations performed by each area. The PPC is thought to represent primarily movement goals, which may constrain the types of learning computations it can perform to manipulations of goals. However, differences in BCI control that have been used in each area (discrete control in PPC; continuous in M1) also influence the timing and forms of feedback available, which in turn may also influence learning (see *5.4 Forms of Feedback*).

Brain areas make distinct contributions to natural sensorimotor learning and control. Differences in how learning occurs across the motor system may influence closed-loop BCIs.

5.3. Neural Features

From the brain region(s) chosen to record from, BCIs compute neural features used for device control. These features are influenced by the sensors used and the signal processing applied to measurements. Recent reviews summarized differences between types of neural measurements and implications for decoding (Lu et al., 2021). In the following subsections, we highlight commonly used neural features and discuss how they may influence closed-loop BCI control and learning.

5.3.1. Types of features and key differences.

Many BCIs use electrophysiologically measured action potentials. While early studies focused on well-isolated single neurons (Fetz, 1969; Ganguly and Carmena, 2009), signals that combine action potentials from different neurons are now common, such as unsorted spiking activity (multiunit) data, threshold crossings (e.g., (Pandarinath et al., 2017)), and spike-band power (Nason et al., 2020). These features blend action potentials at the sensor level, grouping neurons that are physically nearby. An alternative is neural features that group neurons based on structure within neural populations. For example, correlations among action potentials have been used to define “latent factors” (Sadtler et al., 2014). These features blend action potentials based on statistical relationships independent of physical proximity. Latent factor decoding may be done with multiunit data, leading to spatial and statistical blending (Trautmann et al., 2019). Action potentials measured with calcium imaging have also been used for BCIs (Clancy et al., 2014; Prsa et al., 2017; Mitani et al., 2018; Trautmann et al., 2021), which can provide cell type–specific (Mitani et al., 2018) and subcellular (Trautmann et al., 2021) neural features.

The other main category of features is electrophysiological field potentials, which capture neural activity across a larger spatial area and on slower timescales than action potentials (Buzsáki et al., 2012). The spatiotemporal properties of field potentials are influenced by sensor size, sensor placement (Lu et al., 2021), and signal frequency (low-frequency signals capture more global signals relative to higher frequencies). A wide variety of field potential measurements have been used in BCIs, including intracortical electrodes (e.g. to record LFPs) (Engelhard et al., 2013; Flint et al., 2013; So et al., 2014; Stavisky et al., 2015) and potentials electrodes on the cortical surface (e.g. ECoG) (Wander et al., 2013; Benabid et al., 2019; Silversmith et al., 2021). Field potential features are most commonly defined in the spectral domain (e.g., power in a range of frequencies). BCIs often use features across a wide spectral range (e.g., (So et al., 2014)), potentially blending neural activity across a variety of spatiotemporal scales.

Most closed-loop BCIs use a single feature type, but combinations of features have also been explored. A common combination is action potentials and LFPs, which can be recorded simultaneously from intracortical electrodes (Hwang and Andersen, 2009; Stavisky et al., 2015).

5.3.2. Engineering considerations

Neural features differ in their temporal stability, affecting BCI usability and learning. Single-neuron signals are prone to drift with current electrophysiology technologies. Population-level features (Degenhart et al., 2020; Gallego et al., 2020) and field potentials (Flint et al., 2013; Silversmith et al., 2021) are more stable, potentially through different mechanisms. Population-level features appear to capture computations performed by the population that remain stable even if the contributing neurons change (Gallego et al., 2020). Field potentials are thought to capture similar neural activity over time as a result of physical properties of the measurements. Nonstationary neural features reduce decoding accuracy and require methods to adapt decoders (e.g., (Degenhart et al., 2020)); however, frequent changes in neural features and decoders can reduce or prevent learning (Ganguly and Carmena, 2009; Orsborn et al., 2014). Long-term ECoG BCI use leads to high performance and user learning (Benabid et al., 2019; Silversmith et al., 2021); population feature BCI studies have reported primarily maintenance of performance (Degenhart et al., 2020). Instability in features, or in the measured neural activity contributing to them, could affect learning processes such as credit assignment, reducing long-

term learning. More research is needed to understand what forms of feature variability influence learning.

Neural features are measured with different temporal resolutions, which affect closed-loop control and learning. The temporal resolution of a feature is tied to estimation methods. For instance, while action potentials are fast (millisecond resolution), most BCIs use the action potential rate, which is estimated by taking averages over time bins, constraining temporal resolution (Shanechi et al., 2017). This limitation can be overcome by using methods that directly estimate the underlying rate at millisecond resolution (Shanechi et al., 2017). The temporal resolution of features influences the rate of movement control and delay between neural activity and movement. In closed-loop systems, maximizing control rates (Shanechi et al., 2017) and minimizing delays (Willett et al., 2013) improve performance. Given the importance of sensorimotor timing for learning-related computations (Wolpert and Flanagan, 2001), timing differences in neural features may also affect learning, though this hypothesis remains to be explored.

The number of neural features used will influence both decoding performance and learning in BCIs. Open-loop decoding always benefits from adding information, but the relationship between the number of features and closed-loop performance is less well characterized and will depend on the features used. For instance, features that incorporate population dynamics can retain performance after losing many neurons due to redundancy (Kao et al., 2017). Learning is also influenced by the number of neurons used, though this effect has been explored only with small numbers (tens) of neurons (Clancy et al., 2014; Law et al., 2014). How these findings extend to features often used in clinical BCIs is unclear. Using a large number of neural features may make it easier for the brain to find effective BCI solutions by increasing control redundancy (Law et al., 2014), but how the number of features might influence learning processes like credit assignment is unknown.

5.3.3. Biological considerations

How neural features relate to the underlying functional neural circuitry must be considered. These considerations include the physical locations of neurons, cell types, and circuit connectivity—properties that are not completely independent. For instance, connectivity patterns differ across classes of inhibitory neurons (Tremblay et al., 2016).

Neural features differ in how they blend activity across neurons. Studies in which subjects control arbitrary BCI decoders suggest that increasing the physical distance between readout neurons negatively influences learning (Law et al., 2014). How features that combine activity across large distances, such as population-based features, affect processes like credit assignment learning is unknown. A potentially related question concerns the impact of feature spatial resolution. BCI performance typically improves when features are spatially localized. For instance, action potentials provided better closed-loop performance than did local motor potentials from LFPs (i.e. low-frequency timeseries), despite comparable offline predictive power (Stavisky et al., 2015). Whether these differences stem from the relative temporal or spatial resolution of these features has not been determined. Insights into the relevant scales of plasticity during motor learning will likely inform BCI feature design.

The relationship between features and anatomy may also influence learning. For instance, while cortical layers play distinct roles in natural motor learning (Makino et al., 2016), neural features

used in BCIs often ignore layer boundaries, and layer-specific differences have been explored only in open-loop decoding (Markowitz et al., 2011). Interestingly, several studies have demonstrated distinct roles for cell types in BCIs. For example, bursting neurons adapt differently from nonbursting neurons during BCI learning (Garcia-Garcia et al., 2020). The cell types used for online BCI control also influence neural activity changes (Mitani et al., 2018) and learning rates (Vendrell-Llopis et al., 2022). Differences in both the anatomical and firing-rate properties between cell types contribute to these effects (Vendrell-Llopis et al., 2022). How anatomical information can be leveraged to improve closed-loop BCIs remains an open question.

Neurons ultimately control movements through the activity of neuronal populations coordinated via anatomical and functional connections. Many neural features used in BCIs combine neural activity independent of anatomical connectivity. Features based on correlations in population activity, such as latent population dynamics, may provide one way to create features that reflect functional relationships. Decoder perturbations that do not require changes in population correlation structures are more rapidly learned than those that do not (Sadtler et al., 2014; Oby et al., 2019). This finding highlights the potential importance of designing BCI features that capture functionally relevant circuit properties, which will require additional insight into how neural populations perform motor control and learning computations.

5.4. Forms of Feedback

Feedback closes the sensorimotor loop. This feedback can take many forms, including visual, tactile, auditory, and artificial. In BCIs, some form of feedback is necessary for goal-directed movement, real-time feedback control, and learning (Fetz, 1969; Koralek et al., 2013; Shanechi et al., 2017). The importance of feedback motivates research to provide artificial sensation. Here, we briefly discuss how feedback influences motor control and learning.

The timing of feedback relative to actions is critical for sensorimotor control and learning. In the natural motor system, learned internal models are thought to allow the brain to anticipate the sensory consequences of motor commands and attenuate them to highlight expected errors. Disrupting the timing between movement and sensory consequences reduces this attenuation (Blakemore et al., 1999), thus influencing learning computations. While temporal delays occur in all BCIs, discrete BCIs typically introduce significant delays between neural activity and motor outcomes, which may contribute to differences in forms of learning observed in discrete versus continuous BCIs (e.g., compare (Hwang et al., 2013) versus (Ganguly et al., 2011)). The relative timing between movements and sensory feedback must be considered when optimizing BCIs.

The form of feedback must also be considered. Natural motor learning is driven by multiple types of error feedback, including sensory prediction errors (expected versus actual sensory consequences), reward prediction errors (expected versus actual outcomes), and task error (actual movement versus the goal). Each error signal contributes to different aspects of learning (Krakauer et al., 2019). Discrete feedback of movement outcomes, the only feedback available in discrete BCIs, largely eliminates sensory prediction error signals, while continuous BCIs provide all forms of feedback. The richer feedback signals available in continuous BCIs may contribute to differences in learning between the two. Indeed, BCI tasks that can be learned with continuous feedback cannot be learned with discrete task error feedback alone (success or failure) (Koralek et al., 2012). Artificial stimulation of brain areas associated with reward can

facilitate learning (Eaton et al., 2017; Athalye et al., 2018), highlighting the contributions of reward prediction errors. Deeper insights into how sensory feedback signals contribute to control and learning in BCIs will place important constraints on the design of both motor and sensory devices.

6. Conclusions and open questions

6.1. Conclusions and contributions

This thesis investigates the structure of cortical computations from spatial, temporal, and transcriptomic perspectives. These perspectives are not independent; they each play an interconnected role in the mechanistic implementation of computations. However, it is difficult to study all perspectives simultaneously. First, we must separately understand the spatiotemporal structure underlying a given computation and the unique properties of specific cell types (e.g. anatomical, electrophysiological). Then, we can hypothesize how the observed spatiotemporal patterns are mechanistically implemented by specific cell-types (i.e. transcriptomic structure). Here I advance our understanding of the spatiotemporal structure of computations in motor cortical areas, where it is less well established. Then I shift towards investigating the transcriptomic structure in visual cortex where the spatiotemporal structure is well established.

To investigate the spatiotemporal structure of motor computations I leveraged next-generation approaches to record large neural populations that were both spatially distributed and sampled with high spatial resolution. Decoding analyses revealed heterogeneity across both the cortical surface and depths, with significantly more task information at some locations than others. Population dynamics analysis also revealed heterogeneity in the temporal structure of task-related population activity across the cortical surface. Neurons that were most informative about the task were heterogeneously distributed in space yet also had activity patterns that evolved coherently in time. This work highlights the complex spatiotemporal structure of frontal motor cortices and that populations of task-modulated units display the most coherent activity patterns. These results suggest that a subset of neurons that are spatially distributed across frontal motor cortex may facilitate individual motor cortical computations.

Next, I examined how the spatiotemporal structure of motor cortical computations transitions between distinct, but linked, computations: planning and execution. I expanded on the results from Chapter 2 by finding distinct computation-specific neural populations and a spatially distributed computation-shared neural population within which neural activity patterns were conserved across computations. This subpopulation coexisted with the computation-specific subpopulations linking the two distinct computations.

I turned to visual cortex to begin investigating the transcriptomic structure of cortical computations. The spatiotemporal structure of primary visual cortex is well understood, which allowed a precise and experimentally testable hypothesis to be formed about the role of specific cell-types. I first used histology and electrophysiology to test and validate two viruses that targeted parvalbumin-expressing (PV+). However, the injection of one virus produced behavioral deficits requiring the focus to shift to the other virus for behavioral experiments. Selectively activating PV+ neurons impaired visual stimulus detection suggesting that they play an important role in modulating visual sensitivity. Additional work will be needed to use these optogenetic tools to study the transcriptomic structure of cortical computation in other brain areas.

6.2. Open questions

The experiments and investigations described in this thesis advance our understanding of how structure relates to cortical computations. Further investigations into other cortical computations

and how other cell types contribute will be critical to develop a comprehensive picture of cortical structure. To maximize the translational benefit of this research it will also be necessary to investigate how our understanding of cortical structure can be used to improve rehabilitation therapies, such as BCIs.

The studies in this thesis focus on target direction, one form of task information. It remains unclear how our findings generalize to other forms of task information because the relationship between motor cortex and movements is not fully understood. Neurons in motor cortex also encode other forms of task information (e.g. force, kinematics) with no clear spatiotemporal relationship (Fetz, 1994; Kakei et al., 1999). Directly testing how the spatiotemporal structure defined by target direction relates to other forms of task information will be important to help create a unified picture of how motor cortex relates to movement. Even without a unified understanding of motor computations, understanding the spatiotemporal structure of individual forms of task information will still be valuable for BCIs to target implants and design decoders.

The cortical computations investigated in this thesis underlie select behaviors that were performed in controlled environments; how the results generalize to other computations remains an open question. Future work will be necessary to reveal if a similar cortical structure supports other movements and the sensory input necessary to produce all goal-directed movements. Furthermore, it will be necessary to study how different cortical computations coordinate during goal-directed movement production. This thesis focused on the transition from movement planning and execution, but motor computations are also likely influenced by visual, tactile, and proprioceptive sensory inputs. Recent work studying primates in ethologically relevant environments (Moore et al., 2024; Ma et al., 2025) shows promise because of the rich and naturalistic behavior. However, addressing these questions will likely require parallel advances in wireless recording technology and the ability to monitor behavior.

It will also be necessary to investigate the patterns of cortical coordination within and across the computations that support learning. Our results are consistent with previous work suggesting that specific spatiotemporal activity patterns are formed during well-learned movements (Peters et al., 2014; Moore et al., 2024). It remains unclear how the brain converges on the specific activity pattern associated with each movement, but the computations likely involve subcortical areas (Ghanayim et al., 2025; Ramot et al., 2025) and specific cell types within cortex (Donato et al., 2013; Chen et al., 2015; Vendrell-Llopis et al., 2022; Blumenstock et al., 2025).

While a cell-type agnostic perspective of cortical computations is valuable, a mechanistic understanding requires determining how different cell types contribute to the spatiotemporal structure. This thesis focuses on PV+ neurons, a specific inhibitory neuron subtype, in visual cortex but specific cell-types likely play different roles in cortical computations due to their distinct electrophysiological and connectivity properties (Tremblay et al., 2016; Lee et al., 2023). Further work will be necessary to investigate how different cell types contribute to cortical computations. Most cell types also exist across cortex (Tremblay et al., 2016; Bakken et al., 2021) with inhibitory neurons being particularly similar from a transcriptomic perspective (Tasic et al., 2018), providing a pathway for their computational principles to be conserved across areas. Examining the computational roles of specific cell types may reveal unifying principles even though the spatiotemporal structure of neural activity varies between cortical areas.

As our understanding of structure in cortical computations develops, it is important to use our findings to improve translational therapies in parallel. The results in this thesis motivate new lines of investigation for future improvements. Device performance may improve with advanced mapping techniques that allow implants to target specific neurons and neural populations. These results also highlight the need to account for the range of ways that neural population activity evolves in time because the temporal structure can directly impact device performance. Collectively, these results provide promising new avenues of research and evidence for new considerations when designing BCIs to maximize their performance and translatability.

7. References

Amirikian B, Georgopoulos AP (2003) Modular organization of directionally tuned cells in the motor cortex: Is there a short-range order? Proc Natl Acad Sci USA 100:12474–12479.

Andersen RA, Buneo CA (2002) Intentional maps in posterior parietal cortex. Annu Rev Neurosci 25:189–220.

Arduin P-J, Fregnac Y, Shulz DE, Ego-Stengel V (2013) “Master” Neurons Induced by Operant Conditioning in Rat Motor Cortex during a Brain-Machine Interface Task. Journal of Neuroscience 33:8308–8320.

Asanuma H, Rosén I (1972) Topographical organization of cortical efferent zones projecting to distal forelimb muscles in the monkey. Exp Brain Res 14:243–256.

Athalye VR, Ganguly K, Costa RM, Carmena JM (2017) Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control. Neuron 93:955-970.e5.

Athalye VR, Santos FJ, Carmena JM, Costa RM (2018) Evidence for a neural law of effect. Science 359:1024–1029.

Bakken TE et al. (2021) Comparative cellular analysis of motor cortex in human, marmoset and mouse. Nature 598:111–119.

Banerjee A, Dean HL, Pesaran B (2010) A Likelihood Method for Computing Selection Times in Spiking and Local Field Potential Activity. Journal of Neurophysiology 104:3705–3720.

Bastian AJ (2008) Understanding sensorimotor adaptation and learning for rehabilitation. Current Opinion in Neurology 21:628–633.

Batista AP, Buneo CA, Snyder LH, Andersen RA (1999) Reach Plans in Eye-Centered Coordinates. Science 285:257–260.

Benabid AL et al. (2019) An exoskeleton controlled by an epidural wireless brain–machine interface in a tetraplegic patient: a proof-of-concept demonstration. The Lancet Neurology 18:1112–1122.

Best MD, Suminski AJ, Takahashi K, Brown KA, Hatsopoulos NG (2016) Spatio-Temporal Patterning in Primary Motor Cortex at Movement Onset. Cerebral Cortex 27 Available at: https://academic.oup.com/cercor/article/doi/10.1093/cercor/bhv327/3056241.

Blakemore S-J, Frith CD, Wolpert DM (1999) Spatio-Temporal Prediction Modulates the Perception of Self-Produced Stimuli. Journal of Cognitive Neuroscience 11:551–559.

Blumenstock S, Arakelyan D, Del Grosso N, Schneider S, Shao Y, Gjoni E, Klein R, Dudanova I, Komiyama T (2025) Optogenetic restoration of neuron subtype-specific cortical activity ameliorates motor deficits in Huntington’s Disease mice. Available at: http://biorxiv.org/lookup/doi/10.1101/2025.02.07.637155.

Buzsáki G, Anastassiou CA, Koch C (2012) The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. Nat Rev Neurosci 13:407–420.

Canfield RA, Ouchi T, Fang H, Macagno B, Smith LI, Scholl LR, Orsborn AL (2025) The spatiotemporal structure of neural activity in motor cortex during reaching. :2025.10.17.683171 Available at: https://www.biorxiv.org/content/10.1101/2025.10.17.683171v1.

Carmena JM, Lebedev MA, Crist RE, O'Doherty JE, Santucci DM, Dimitrov DF, Patil PG, Henriquez CS, Nicolelis MAL (2003) Learning to Control a Brain–Machine Interface for Reaching and Grasping by Primates Idan Segev, ed. PLoS Biol 1:e42.

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

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

Chehade NG, Gharbawie OA (2023) Motor actions are spatially organized in motor and dorsal premotor cortex. eLife 12:e83196.

Chen SX, Kim AN, Peters AJ, Komiyama T (2015) Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. Nat Neurosci 18:1109–1115.

Cherian A, Krucoff MO, Miller LE (2011) Motor cortical prediction of EMG: evidence that a kinetic brain-machine interface may be robust across altered movement dynamics. J Neurosci 106:564–575.

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. Nature 487:51–56.

Churchland MM, Shenoy KV (2007) Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. Journal of Neurophysiology 97:4235–4257.

Churchland MM, Shenoy KV (2024) Preparatory activity and the expansive null-space. Nat Rev Neurosci 25:213–236.

Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV (2006) Neural Variability in Premotor Cortex Provides a Signature of Motor Preparation. J Neurosci 26:3697–3712.

Cisek P, Kalaska JF (2002) Modest Gaze-Related Discharge Modulation in Monkey Dorsal Premotor Cortex During a Reaching Task Performed With Free Fixation. Journal of Neurophysiology 88:1064–1072.

Cisek P, Kalaska JF (2005) Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. Neuron 45:801–814.

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

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

Cunningham JP, Yu BM (2014) Dimensionality reduction for large-scale neural recordings. Nat Neurosci 17:1500–1509.

Dadarlat MC, Canfield RA, Orsborn AL (2023) Neural Plasticity in Sensorimotor Brain–Machine Interfaces. Annu Rev Biomed Eng 25:51–76.

Dayan E, Cohen LG (2011) Neuroplasticity subserving motor skill learning. Neuron 72:443–454.

De A, Horwitz GD (2021) Spatial receptive field structure of double-opponent cells in macaque V1. Journal of Neurophysiology 125:843–857.

Degenhart AD, Bishop WE, Oby ER, Tyler-Kabara EC, Chase SM, Batista AP, Yu BM (2020) Stabilization of a brain-computer interface via the alignment of low-dimensional spaces of neural activity. Nat Biomed Eng 4:672–685.

Dhawale AK, Smith MA, Ölveczky BP (2017) The Role of Variability in Motor Learning. Annu Rev Neurosci 40:479–498.

Donato F, Rompani SB, Caroni P (2013) Parvalbumin-expressing basket-cell network plasticity induced by experience regulates adult learning. Nature 504:272–276.

Dum R, Strick P (2002) Motor areas in the frontal lobe of the primate. Physiology & Behavior 77:677–682.

Dum R, Strick P (2004) Motor Areas in the Frontal Lobe: The Anatomical Substrate for the Central Control of Movement. In: Motor Cortex in Voluntary Movements: A Distributed System for Distributed Functions, 1st ed., pp 44. Boca Raton: CRC Press.

Eaton RW, Libey T, Fetz EE (2017) Operant conditioning of neural activity in freely behaving monkeys with intracranial reinforcement. Journal of Neurophysiology 117:1112–1125.

Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP (2016) Reorganization between preparatory and movement population responses in motor cortex. Nat Commun 7:13239.

El-Shamayleh Y, Horwitz GD (2019) Primate optogenetics: Progress and prognosis. Proc Natl Acad Sci USA 116:26195–26203.

El-Shamayleh Y, Kojima Y, Soetedjo R, Horwitz GD (2017) Selective Optogenetic Control of Purkinje Cells in Monkey Cerebellum. Neuron 95:51-62.e4.

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

Estebanez L, Hoffmann D, Voigt BC, Poulet JFA (2017) Parvalbumin-Expressing GABAergic Neurons in Primary Motor Cortex Signal Reaching. Cell Reports 20:308–318.

Evarts EV (1968) Relation of pyramidal tract activity to force exerted during voluntary movement. Journal of Neurophysiology 31:14–27.

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

Fetz EE (1994) Are movement parameters recognizably coded in the activity of single neurons? In: Movement Control, 1st ed. (Cordo P, Harnad S, eds), pp 77–88. Cambridge University Press. Available at: https://www.cambridge.org/core/product/identifier/CBO9780511529788A013/type/book_part.

Finkelstein A, Daie K, Rózsa M, Darshan R, Svoboda K (2026) Connectivity underlying motor cortex activity during goal-directed behaviour. Nature 649:416–422.

Flint RD, Wright ZA, Scheid MR, Slutzky MW (2013) Long term, stable brain machine interface performance using local field potentials and multiunit spikes. J Neural Eng 10:056005.

Gallego JA, Makin TR, McDougle SD (2022) Going beyond primary motor cortex to improve brain-computer interfaces. Trends Neurosci 45:176–183.

Gallego JA, Perich MG, Chowdhury RH, Solla SA, Miller LE (2020) Long-term stability of cortical population dynamics underlying consistent behavior. Nat Neurosci 23:260–270.

Gallego JA, Perich MG, Naufel SN, Ethier C, Solla SA, Miller LE (2018) Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. Nat Commun 9:4233.

Ganguly K, Carmena JM (2009) Emergence of a Stable Cortical Map for Neuroprosthetic Control Ashe J, ed. PLoS Biol 7:e1000153.

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

Ganguly K, Secundo L, Ranade G, Orsborn A, Chang EF, Dimitrov DF, Wallis JD, Barbaro NM, Knight RT, Carmena JM (2009) Cortical representation of ipsilateral arm movements in monkey and man. J Neurosci 29:12948–12956.

Garcia-Garcia MG, Marquez-Chin C, Popovic MR (2020) Operant conditioning of motor cortex neurons reveals neuron-subtype-specific responses in a brain-machine interface task. Sci Rep 10:19992.

Gattass R, Gross CG, Sandell JH (1981) Visual topography of V2 in the macaque. J Comp Neurol 201:519–539.

Gatter KC, Powell TPS (1978) The intrinsic connections of the cortex of area 4 of the monkey. Brain 101:513–541.

Georgopoulos A, Kalaska J, Caminiti R, Massey J (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci 2:1527–1537.

Georgopoulos AP, Merchant H, Naselaris T, Amirikian B (2007) Mapping of the preferred direction in the motor cortex. Proc Natl Acad Sci USA 104:11068–11072.

Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal Population Coding of Movement Direction. Science 233:1416–1419.

Ghanayim A, Benisty H, Cohen Rimon A, Schwartz S, Dabdoob S, Lifshitz S, Talmon R, Schiller J (2025) VTA projections to M1 are essential for reorganization of layer 2-3 network dynamics underlying motor learning. Nat Commun 16 Available at: https://www.nature.com/articles/s41467-024-55317-4.

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

Golub MD, Sadtler PT, Oby ER, Quick KM, Ryu SI, Tyler-Kabara EC, Batista AP, Chase SM, Yu BM (2018) Learning by neural reassociation. Nat Neurosci 21:607–616.

Gulati T, Guo L, Ramanathan DS, Bodepudi A, Ganguly K (2017) Neural reactivations during sleep determine network credit assignment. Nat Neurosci 20:1277–1284.

Gulati T, Ramanathan DS, Wong CC, Ganguly K (2014) Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. Nat Neurosci 17:1107–1113.

Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature 442:164–171.

Hooks BM, Mao T, Gutnisky DA, Yamawaki N, Svoboda K, Shepherd GMG (2013) Organization of Cortical and Thalamic Input to Pyramidal Neurons in Mouse Motor Cortex. J Neurosci 33:748–760.

Horwitz G, Chichilnisky EJ, Albright T (2007) Analysis of non-linear cone signal combination in V1. Journal of Vision 7:3–3.

Hubel DH, Wiesel TN (1959) Receptive fields of single neurones in the cat's striate cortex. The Journal of Physiology 148:574–591.

Huntley GW, Jones EG (1991) Relationship of intrinsic connections to forelimb movement representations in monkey motor cortex: a correlative anatomic and physiological study. Journal of Neurophysiology 66:390–413.

Hwang EJ, Andersen RA (2009) Brain Control of Movement Execution Onset Using Local Field Potentials in Posterior Parietal Cortex. Journal of Neuroscience 29:14363–14370.

Hwang EJ, Bailey PM, Andersen RA (2013) Volitional Control of Neural Activity Relies on the Natural Motor Repertoire. Current Biology 23:353–361.

Inagaki HK, Chen S, Daie K, Finkelstein A, Fontolan L, Romani S, Svoboda K (2022) Neural Algorithms and Circuits for Motor Planning. Annu Rev Neurosci 45:249–271.

International Brain Laboratory (2024) Spike sorting pipeline for the International Brain Laboratory. :18689452 Bytes Available at: https://figshare.com/articles/online_resource/Spike_sorting_pipeline_for_the_International_Brain_Laboratory/19705522/4.

Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB (2008) Functional network reorganization during learning in a brain-computer interface paradigm. Proc Natl Acad Sci USA 105:19486–19491.

Jun JJ et al. (2017) Fully integrated silicon probes for high-density recording of neural activity. Nature 551:232–236.

Kakei S, Hoffman DS, Strick PL (1999) Muscle and Movement Representations in the Primary Motor Cortex. Science 285:2136–2139.

Kao JC, Ryu SI, Shenoy KV (2017) Leveraging neural dynamics to extend functional lifetime of brain-machine interfaces. Sci Rep 7:7395.

Karpowicz BM, Ali YH, Wimalasena LN, Sedler AR, Keshtkaran MR, Bodkin K, Ma X, Rubin DB, Williams ZM, Cash SS, Hochberg LR, Miller LE, Pandarinath C (2025) Stabilizing brain-computer interfaces through alignment of latent dynamics. Nat Commun 16:4662.

Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. Nat Neurosci 17:440–448.

Kaufman MT, Seely JS, Sussillo D, Ryu SI, Shenoy KV, Churchland MM (2016) The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. eneuro 3:ENEURO.0085-16.2016.

Kojima Y, Ting JT, Soetedjo R, Gibson SD, Horwitz GD (2021) Injections of AAV Vectors for Optogenetics in Anesthetized and Awake Behaving Non-Human Primate Brain. JoVE:62546.

Koralek AC, Costa RM, Carmena JM (2013) Temporally Precise Cell-Specific Coherence Develops in Corticostriatal Networks during Learning. Neuron 79:865–872.

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

Koyama S, Chase SM, Whitford AS, Velliste M, Schwartz AB, Kass RE (2010) Comparison of brain–computer interface decoding algorithms in open-loop and closed-loop control. J Comput Neurosci 29:73–87.

Krakauer JW, Hadjiosif AM, Xu J, Wong AL, Haith AM (2019) Motor Learning. In: Comprehensive Physiology, 1st ed. (Terjung R, ed), pp 613–663. Wiley. Available at: https://onlinelibrary.wiley.com/doi/10.1002/cphy.c170043.

Kunigk NG, Schone HR, Gontier C, Hockeimer W, Tortolani AF, Hatsopoulos NG, Downey JE, Chase SM, Boninger ML, Dekleva BD, Collinger JL (2024) Motor somatotopy impacts imagery strategy success in human intracortical brain-computer interfaces. Available at: http://medrxiv.org/lookup/doi/10.1101/2024.08.01.24311180.

Lara AH, Cunningham JP, Churchland MM (2018a) Different population dynamics in the supplementary motor area and motor cortex during reaching. Nat Commun 9:2754.

Lara AH, Elsayed GF, Zimnik AJ, Cunningham JP, Churchland MM (2018b) Conservation of preparatory neural events in monkey motor cortex regardless of how movement is initiated. eLife 7:e31826.

Law AJ, Rivlis G, Schieber MH (2014) Rapid acquisition of novel interface control by small ensembles of arbitrarily selected primary motor cortex neurons. Journal of Neurophysiology 112:1528–1548.

Lee BR et al. (2023) Signature morphoelectric properties of diverse GABAergic interneurons in the human neocortex. Science 382:eadf6484.

Lima SQ, Hromádka T, Znamenskiy P, Zador AM (2009) PINP: A New Method of Tagging Neuronal Populations for Identification during In Vivo Electrophysiological Recording Nitabach MN, ed. PLoS ONE 4:e6099.

Liu Z, Schieber MH (2020) Neuronal Activity Distributed in Multiple Cortical Areas during Voluntary Control of the Native Arm or a Brain-Computer Interface. eNeuro 7:ENEURO.0376-20.2020.

Lu H-Y, Lorenc ES, Zhu H, Kilmarx J, Sulzer J, Xie C, Tobler PN, Watrous AJ, Orsborn AL, Lewis-Peacock J, Santacruz SR (2021) Multi-scale neural decoding and analysis. J Neural Eng 18:045013.

Ma X, Rizzoglio F, Bodkin KL, Miller LE (2025) Unsupervised, piecewise linear decoding enables an accurate prediction of muscle activity in a multi-task brain computer interface. J Neural Eng 22:016019.

Mahmoudi B, Sanchez JC (2011) A Symbiotic Brain-Machine Interface through Value-Based Decision Making Bongard J, ed. PLoS ONE 6:e14760.

Makino H, Hwang EJ, Hedrick NG, Komiyama T (2016) Circuit Mechanisms of Sensorimotor Learning. Neuron 92:705–721.

Markowitz DA, Wong YT, Gray CM, Pesaran B (2011) Optimizing the decoding of movement goals from local field potentials in macaque cortex. J Neurosci 31:18412–18422.

Matsumura M, Sawaguchi T, Oishi T, Ueki K, Kubota K (1991) Behavioral deficits induced by local injection of bicuculline and muscimol into the primate motor and premotor cortex. Journal of Neurophysiology 65:1542–1553.

Mich JK et al. (2021) Functional enhancer elements drive subclass-selective expression from mouse to primate neocortex. Cell Reports 34:108754.

Mitani A, Dong M, Komiyama T (2018) Brain-Computer Interface with Inhibitory Neurons Reveals Subtype-Specific Strategies. Current Biology 28:77-83.e4.

Mollazadeh M, Aggarwal V, Davidson AG, Law AJ, Thakor NV, Schieber MH (2011) Spatiotemporal Variation of Multiple Neurophysiological Signals in the Primary Motor Cortex during Dexterous Reach-to-Grasp Movements. J Neurosci 31:15531–15543.

Moore DD, MacLean JN, Walker JD, Hatsopoulos NG (2024) A dynamic subset of network interactions underlies tuning to natural movements in marmoset sensorimotor cortex. Nat Commun 15:10517.

Moritz CT, Fetz EE (2011) Volitional control of single cortical neurons in a brain–machine interface. J Neural Eng 8:025017.

Mulliken GH, Musallam S, Andersen RA (2008) Decoding Trajectories from Posterior Parietal Cortex Ensembles. Journal of Neuroscience 28:12913–12926.

Musall S, Kaufman MT, Juavinett AL, Gluf S, Churchland AK (2019) Single-trial neural dynamics are dominated by richly varied movements. Nat Neurosci 22:1677–1686.

Musk E, Neuralink (2019) An Integrated Brain-Machine Interface Platform With Thousands of Channels. J Med Internet Res 21:e16194.

Nason SR, Vaskov AK, Willsey MS, Welle EJ, An H, Vu PP, Bullard AJ, Nu CS, Kao JC, Shenoy KV, Jang T, Kim H-S, Blaauw D, Patil PG, Chestek CA (2020) A low-power band of neuronal spiking activity dominated by local single units improves the performance of brain–machine interfaces. Nat Biomed Eng 4:973–983.

Neely RM, Koralek AC, Athalye VR, Costa RM, Carmena JM (2018) Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. Neuron 97:1356-1368.e4.

Ninomiya T, Inoue K, Hoshi E, Takada M (2019) Layer specificity of inputs from supplementary motor area and dorsal premotor cortex to primary motor cortex in macaque monkeys. Sci Rep 9:18230.

Oby ER, Degenhart AD, Grigsby EM, Motiwala A, McClain NT, Marino PJ, Yu BM, Batista AP (2025) Dynamical constraints on neural population activity. Nat Neurosci.

Oby ER, Golub MD, Hennig JA, Degenhart AD, Tyler-Kabara EC, Yu BM, Chase SM, Batista AP (2019) New neural activity patterns emerge with long-term learning. Proc Natl Acad Sci USA 116:15210–15215.

Omrani M, Kaufman MT, Hatsopoulos NG, Cheney PD (2017) Perspectives on classical controversies about the motor cortex. Journal of Neurophysiology 118:1828–1848.

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

Orsborn AL, Pesaran B (2017) Parsing learning in networks using brain–machine interfaces. Current Opinion in Neurobiology 46:76–83.

Ouchi T, Scholl LR, Rajeswaran P, Canfield RA, Smith LI, Orsborn AL (2025) Mapping eye, arm, and reward information in frontal motor cortices using electrocorticography in non-human primates. J Neurosci:e1536242025.

Pachitariu M, Sridhar S, Pennington J, Stringer C (2024) Spike sorting with Kilosort4. Nat Methods 21:914–921.

Pandarinath C, Nuyujukian P, Blabe CH, Sorice BL, Saab J, Willett FR, Hochberg LR, Shenoy KV, Henderson JM (2017) High performance communication by people with paralysis using an intracortical brain-computer interface. Elife 6:e18554.

Pandarinath C, O'Shea DJ, Collins J, Jozefowicz R, Stavisky SD, Kao JC, Trautmann EM, Kaufman MT, Ryu SI, Hochberg LR, Henderson JM, Shenoy KV, Abbott LF, Sussillo D (2018) Inferring single-trial neural population dynamics using sequential auto-encoders. Nat Methods 15:805–815.

Park MC, Belhaj-Saïf A, Gordon M, Cheney PD (2001) Consistent Features in the Forelimb Representation of Primary Motor Cortex in Rhesus Macaques. J Neurosci 21:2784–2792.

Passingham RE, Perry VH, Wilkinson F (1983) The long-term effects of removal of sensorimotor cortex in infant and adult rhesus monkeys. Brain 106:675–705.

Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay É (2011) Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research 12:2825–2830.

Penfield W, Boldrey E (1937) Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. Brain 60:389–443.

Perich MG, Gallego JA, Miller LE (2018) A Neural Population Mechanism for Rapid Learning. Neuron 100:964-976.e7.

Pesaran B, Nelson MJ, Andersen RA (2010) A Relative Position Code for Saccades in Dorsal Premotor Cortex. J Neurosci 30:6527–6537.

Peters AJ, Chen SX, Komiyama T (2014) Emergence of reproducible spatiotemporal activity during motor learning. Nature 510:263–267.

Prsa M, Galiñanes GL, Huber D (2017) Rapid Integration of Artificial Sensory Feedback during Operant Conditioning of Motor Cortex Neurons. Neuron 93:929-939.e6.

Ramot A, Taschbach FH, Yang YC, Hu Y, Chen Q, Morales BC, Wang XC, Wu A, Tye KM, Benna MK, Komiyama T (2025) Motor learning refines thalamic influence on motor cortex. Nature 643:725–734.

Riehle A, Requin J (1989) Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. Journal of Neurophysiology 61:534–549.

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

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

Safaie M, Chang JC, Park J, Miller LE, Dudman JT, Perich MG, Gallego JA (2023) Preserved neural dynamics across animals performing similar behaviour. Nature 623:765–771.

Sakellaridi S, Christopoulos VN, Aflalo T, Pejsa KW, Rosario ER, Ouellette D, Pouratian N, Andersen RA (2019) Intrinsic Variable Learning for Brain-Machine Interface Control by Human Anterior Intraparietal Cortex. Neuron 102:694-705.e3.

Schieber MH, Hibbard LS (1993) How Somatotopic Is the Motor Cortex Hand Area? Science 261:489–492.

Seely JS, Kaufman MT, Ryu SI, Shenoy KV, Cunningham JP, Churchland MM (2016) Tensor Analysis Reveals Distinct Population Structure that Parallels the Different Computational Roles of Areas M1 and V1 Latham PE, ed. PLoS Comput Biol 12:e1005164.

Semedo JD, Zandvakili A, Machens CK, Yu BM, Kohn A (2019) Cortical Areas Interact through a Communication Subspace. Neuron 102:249-259.e4.

Sergio LE, Hamel-Pâquet C, Kalaska JF (2005) Motor Cortex Neural Correlates of Output Kinematics and Kinetics During Isometric-Force and Arm-Reaching Tasks. Journal of Neurophysiology 94:2353–2378.

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

Shanechi MM, Orsborn AL, Moorman HG, Gowda S, Dangi S, Carmena JM (2017) Rapid control and feedback rates enhance neuroprosthetic control. Nat Commun 8:13825.

Silversmith DB, Abiri R, Hardy NF, Natraj N, Tu-Chan A, Chang EF, Ganguly K (2021) Plug-and-play control of a brain–computer interface through neural map stabilization. Nat Biotechnol 39:326–335.

So K, Dangi S, Orsborn AL, Gastpar MC, Carmena JM (2014) Subject-specific modulation of local field potential spectral power during brain–machine interface control in primates. J Neural Eng 11:026002.

Sohal VS, Zhang F, Yizhar O, Deisseroth K (2009) Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. Nature 459:698–702.

Stavisky SD, Kao JC, Nuyujukian P, Ryu SI, Shenoy KV (2015) A high performing brain–machine interface driven by low-frequency local field potentials alone and together with spikes. J Neural Eng 12:036009.

Sussillo D, Churchland MM, Kaufman MT, Shenoy KV (2015) A neural network that finds a naturalistic solution for the production of muscle activity. Nat Neurosci 18:1025–1033.

Talbot SA, Marshall WH (1941) Physiological Studies on Neural Mechanisms of Visual Localization and Discrimination*. American Journal of Ophthalmology 24:1255–1264.

Tasic B et al. (2018) Shared and distinct transcriptomic cell types across neocortical areas. Nature 563:72–78.

Taylor JA, Krakauer JW, Ivry RB (2014) Explicit and implicit contributions to learning in a sensorimotor adaptation task. J Neurosci 34:3023–3032.

Tong F (2003) Primary visual cortex and visual awareness. Nat Rev Neurosci 4:219–229.

Trautmann EM et al. (2021) Dendritic calcium signals in rhesus macaque motor cortex drive an optical brain-computer interface. Nat Commun 12:3689.

Trautmann EM, Stavisky SD, Lahiri S, Ames KC, Kaufman MT, O'Shea DJ, Vyas S, Sun X, Ryu SI, Ganguli S, Shenoy KV (2019) Accurate Estimation of Neural Population Dynamics without Spike Sorting. Neuron 103:292-308.e4.

Tremblay R, Lee S, Rudy B (2016) GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits. Neuron 91:260–292.

Van Essen DC, Donahue CJ, Coalson TS, Kennedy H, Hayashi T, Glasser MF (2019) Cerebral cortical folding, parcellation, and connectivity in humans, nonhuman primates, and mice. Proc Natl Acad Sci USA 116:26173–26180.

Vendrell-Llopis N, Fang C, Qü AJ, Costa RM, Carmena JM (2022) Diverse operant control of different motor cortex populations during learning. Current Biology 32:1616-1622.e5.

Veuthey TL, Derosier K, Kondapavulur S, Ganguly K (2020) Single-trial cross-area neural population dynamics during long-term skill learning. Nat Commun 11:4057.

Vormstein-Schneider D et al. (2020) Viral manipulation of functionally distinct interneurons in mice, non-human primates and humans. Nat Neurosci Available at: http://www.nature.com/articles/s41593-020-0692-9.

Wander JD, Blakely T, Miller KJ, Weaver KE, Johnson LA, Olson JD, Fetz EE, Rao RPN, Ojemann JG (2013) Distributed cortical adaptation during learning of a brain–computer interface task. Proc Natl Acad Sci USA 110:10818–10823.

Whitlock JR (2017) Posterior parietal cortex. Current Biology 27:R691–R695.

Willett FR, Avansino DT, Hochberg LR, Henderson JM, Shenoy KV (2021) High-performance brain-to-text communication via handwriting. Nature 593:249–254.

Willett FR, Suminski AJ, Fagg AH, Hatsopoulos NG (2013) Improving brain–machine interface performance by decoding intended future movements. J Neural Eng 10:026011.

Williams JJ, Rouse AG, Thongpang S, Williams JC, Moran DW (2013) Differentiating closed-loop cortical intention from rest: building an asynchronous electrocorticographic BCI. J Neural Eng 10:046001.

Wolpert DM, Flanagan JR (2001) Motor prediction. Curr Biol 11:R729-732.

Yang CS, Cowan NJ, Haith AM (2021) De novo learning versus adaptation of continuous control in a manual tracking task. Elife 10:e62578.

Yang X, Zhou T, Zwang TJ, Hong G, Zhao Y, Viveros RD, Fu T-M, Gao T, Lieber CM (2019) Bioinspired neuron-like electronics. Nat Mater 18:510–517.

Yazdan-Shahmorad P, Gibson S, Lee JC, Horwitz GD (2024) Preferential transduction of parvalbumin-expressing cortical neurons by AAV-mDLX5/6 vectors. Front Neurosci 17 Available at: https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2023.1269025/full.

Zhou X, Tien RN, Ravikumar S, Chase SM (2019) Distinct types of neural reorganization during long-term learning. Journal of Neurophysiology 121:1329–1341.

Zimnik AJ, Churchland MM (2021) Independent generation of sequence elements by motor cortex. Nat Neurosci 24:412–424.