Sensory Constraints on Volitional Modulation of Neural Activity in the Motor Cortex

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Biomedical Engineering

Carmen Fernández Fisac

B.S., Biomedical Engineering, Universidad Carlos III de Madrid
M.S., Neuroscience, University College London

Carnegie Mellon University
Pittsburgh, PA

August 2022

ACKNOWLEDGEMENTS

As I type the last words of this thesis, it is impossible not to look back at the six years I have spent at Carnegie Mellon University. I am overwhelmed by the phenomenal people that have accompanied me all those months, without whose support this thesis would never have come to be. Although these brief remarks cannot possibly do them justice, I will nonetheless stubbornly do my best to capture my gratitude.

I would like to begin by thanking my PhD advisor and chair of my thesis committee, Steve Chase. I am grateful to Steve for taking a chance on me in 2016 and for his continued guidance all these years, always offering up his time for extra meetings and approaching setbacks with a sense of humor. I especially enjoyed getting into cheerful—and often very amusing if I do say so myself—arguments with Steve, which also fueled my determination to see this project through. I would like to take this chance to acknowledge the funding sources that supported this research: the National Science Foundation (NSF CAREER Award IOS1553252, S.M.C.), the "la Caixa" Foundation (100010434 LCF/BQ/AN15/10380007, C.F.F) and the Carnegie Mellon Neuroscience Institute (CMU Presidential Fellowship, C.F.F.).

I am also grateful to Byron Yu, Sandra Kuhlman and Marc Schieber for being part of my thesis committee and sharing their scientific insights. More importantly, I am immensely thankful for the excitement they have shown for the findings I have produced during these years and their confidence—even when my own faltered—in my ability to take this research project to completion and continue to work in neuroscience beyond my PhD. I could not have wished for a better committee. Many thanks to other faculty and staff that have mentored and inspired me over the years, especially Gelsy Torres-Oviedo and the BME department, as well as the SCABBY gang for the invaluable training in critical thinking and scientific integrity.

My presence at Carnegie Mellon University would not have been possible without the backing provided by the "la Caixa" Fellowship, which supported my decision to conduct research in the intersection of neuroscience and engineering and constituted an endorsement during the PhD application process. Beyond the financial and logistical means it provided, I am incredibly grateful for the opportunity to join the exceptional community of fellow *caixeros* scattered all over the globe, of which I will always be a part.

I am thankful every day for the wonderful members of the Chase Lab, past and present: Kat, Adam, Xiao, Tze Hui, Darby, Brian, Jay, Lindsay, Abhinav and Sadhana. Their company made life better. Thanks for those chats about the brain—and about everything else—, for our recurring game nights and for the occasional rant about research. As we often remind each other: *"The good thing about grad school is that every day is a Saturday. The bad thing about grad school is that you work on Saturdays."* Among them, a special mention *has* to go to Kat, my roommate of three years and *labmate* of six, whose friendship and

expert comedic relief made the hard times easier to bear. I am beyond grateful to have ridden the PhD roller coaster alongside this tremendous human and it is only fitting that we should graduate within a month of one another.

Words cannot adequately describe my gratitude for the extraordinary Tartan Salsa community. It was early 2017 when we decided to found a new student organization in which people from different backgrounds could come together to learn salsa dancing. What started as a small group of students quickly turned into hundreds of members attending our weekly lessons. Through those dances, I have been lucky to connect with terrific individuals: Jen, Marilina, Alyson, Shena, Andresa, KK, Manuel, Moataz, Nora, Kedar, Silvino, Gus... A special thank you goes to Jen, who has been with me at the leadership of Tartan Salsa since the very beginning and whose caring and unconditional friendship has been an endless source of moral support in grad school. Witnessing our students' growth has been truly rewarding; we have helped so many newcomers as they worked hard on their skills and challenged their assumptions to become seasoned dancers eager to make salsa communities everywhere more welcoming and empowering. For every initiative that we have put together throughout the years, I continue to look around me in awe at my fellow Tartan Salsa teammates, these dedicated humans—all of them remarkable—that selflessly devote their very limited time and energy to keep our shared project alive. Working with them all to make Tartan Salsa possible has been one of the greatest joys in my life.

Also deserving a special mention are *los vinos* (Uxue, Raquel, Beñat, Angélica and Ginés), a group of (mostly) Spaniards brought together by Pittsburgh with whom the nights were filled with *croquetas* and old songs played on the guitar, and whose absence from Pittsburgh during the last couple of years has been sourly felt. I would also like to extend my gratitude to my friends and family back home, who—despite my long absences from Madrid and my complete inability to stay in touch—have a way of welcoming me back like I haven't been gone more than a day. To *las uc3mitas* and *los jalapeños*, two groups of close friends that have been there for me for over a decade. To my four loving grandparents and my extended family. To Angelita, my *mother from another mother*.

My final thank you is to those who have seen the very darkest of my days over these six years and who have helped brighten them. To Jared, whose warmth and partnership accompanied me for the majority of this adventure and who came to be—and will never be anything but—family to me. To our dog Freddie, whose contagious cheerfulness and timely nudges kept me alive and fed, especially during the final months of my PhD. And, above all, to my extraordinary family. My brother Jaime, the lighthouse that can guide me through any storm, who inspires me on a daily basis and on whom I trust and rely more than anyone in this world. My parents, who selflessly encouraged my every step, regardless of how far from home that would lead me, and whose voices I carry with me in everything I do. Simply saying that this dissertation would not exist without their unconditional support would be the understatement of a lifetime.

Thank you all.

ABSTRACT

A fundamental way in which we interact with the world around us is voluntary movement. The primary motor cortex (M1) is known for its central role in driving this intentional movement, and decades of research in the field of motor control has shown that individuals can learn to modulate even single neurons in M1 at will. Indeed, the fact that neural activity in M1 can be volitionally controlled makes it a powerful target for brain-computer interface (BCI) devices and their clinical applications.

Yet M1 also encodes non-volitional information. This brain area receives pronounced sensory inputs. It also contributes to sensory-evoked—as well as voluntary—motor responses. What does this duality in M1 imply for its volitional control? To what extent do non-volitional signals restrict our ability to voluntarily modulate M1?

This thesis provides insight into the constraints imposed by sensory context on M1 by explicitly probing voluntary control of individual neurons in this area under a variety of sensory conditions. To do this, three macaque monkeys were trained in a BCI paradigm that decoupled volitional modulation of M1 from specific aspects of sensory feedback. In these experiments, the firing rate of a single neuron—termed the *command neuron*—directly determined the position of a computer cursor along a one-dimensional axis. Altering the orientation and location of this movement axis and the posture in which the animal's arm rested created distinct sensory contexts for the subject without changing the neural requirements of the task. We leveraged this paradigm to assess monkeys' ability to modulate individual neurons in M1 under different sensory contexts.

We found that sensory context persistently affected volitional control of single neurons in M1. For all three subjects, the ability to perform the task significantly changed based on movement orientation: rotating the feedback axis could render the same neural task effortless or problematic. Axis location within the workspace also affected the ability to modulate

individual neurons, albeit to a lesser extent than orientation. Notably, the disparity in single-neuron control across sensory contexts was not resolved even after extended training in the task. We found that additional practice within a session or across multiple days was not sufficient to erase the interaction between movement orientation and the ability to voluntarily modulate individual neurons.

Overall, these findings suggest that sensory context *does* limit the degree to which M1 activity is under volitional control. This interplay between sensory and volitional signals, which manifested at the level of individual neurons within M1, may constitute an additional constraint on motor learning. The notion that sensory inputs may impose bounds on the voluntary modulation of M1 could also have direct implications for the development of clinical neuroprosthetic devices and other advances in neurotechnology.

RESUMEN

Una herramienta fundamental a la hora de interactuar con el mundo que nos rodea es el movimiento voluntario. La corteza motora primaria (M1) es conocida por su rol capital en la ejecución de movimiento intencionado y décadas de investigación en el campo del control motor han demostrado que es posible aprender a modular hasta neuronas individuales de esta área del cerebro a voluntad. Esa capacidad para controlar la actividad neuronal de la M1 de forma voluntaria la convierte en un blanco enormemente atractivo para el desarrollo de interfaces cerebro-máquina (BCI por sus siglas en inglés) y la aplicación clínica de estos dispositivos.

Sin embargo, la corteza M1 también codifica información no volitiva. Esta región recibe potentes señales sensoriales y contribuye también al desempeño de acciones motoras suscitadas no por una intención sino en respuesta a estímulos externos. ¿Qué implica esta dualidad de la M1 a la hora de controlar su actividad de forma intencional? ¿Hasta qué punto restringen esas señales no volitivas nuestra capacidad de modular voluntariamente la región M1?

Esta tesis doctoral constituye un estudio explícito del control voluntario de neuronas individuales ante distintas condiciones del entorno y contribuye a esclarecer los límites que el contexto sensorial impone sobre la corteza M1. Para ello, tres macacos Rhesus recibieron entrenamiento para controlar un modelo de interfaz cerebro-máquina cuyo diseño disocia la modulación volitiva de la M1 de la retroalimentación sensorial. En estas pruebas, la actividad de una sola neurona, a la que llamamos la *neurona comandante*, determina de forma directa la posición del cursor en pantalla a lo largo de un eje unidimensional. Este paradigma de BCI permite alterar el ángulo de ese eje de movimiento, su localización respecto al centro de la pantalla o la postura en la que descansa el brazo del animal, todo ello sin cambiar las propiedades neuronales de la tarea asignada al sujeto, lo que da lugar a una serie de con-

textos sensoriales distintivos que comparten requisitos neuronales idénticos. Esta flexibilidad de diseño posibilita la evaluación de la capacidad del sujeto para modular la actividad de neuronas individuales en la M1 dado un abanico de contextos sensoriales.

Los resultados descritos a lo largo de esta tesis indican que el contexto sensorial afecta de forma persistente el control volitivo de neuronas individuales en la corteza M1. Todos los animales manifestaron un cambio significativo en su habilidad para mover el cursor en función de la orientación del movimiento: solo rotar el eje sobre el que se mueve el cursor podía hacer la misma tarea neuronal resultar fácil o suponer un desafío. También la localización del eje de movimiento en la pantalla demostró influir sobre la capacidad para controlar la neurona comandante, si bien en menor medida que el ángulo del eje. Cabe señalar que un entrenamiento prolongado en el uso de la misma neurona comando no fue capaz de resolver esta discrepancia de control entre contextos sensoriales. Tanto en una sola sesión como a lo largo de múltiples días, la ampliación del periodo de práctica no resultó suficiente para eliminar esta interacción entre la orientación del movimiento y la capacidad para modular neuronas individuales de forma voluntaria.

En conjunto, estos resultados indican que el contexto sensorial sí limita el alcance de nuestro control voluntario sobre la actividad de la corteza M1. Esta interacción entre señales sensoriales y volitivas presente al nivel de neuronas individuales puede constituir una restricción adicional sobre el aprendizaje motor. La noción de que la información sensorial sea capaz de condicionar la modulación neuronal voluntaria puede asimismo tener implicaciones directas para el desarrollo de neuroprótesis clínicas y otros avances en neurotecnología.

To my parents, Concha and Curro, and to my brother Jaime.

¿Quiénes somos sin socios?

Contents

List of Figures

1

Background

1.1 Volitional Nature of the Primary Motor Cortex

The primary motor cortex (M1) is heavily involved in the generation of intentional movement (Porter and Lemon, 1993; Scott, 2003). Neural activity in this brain area has been shown to correlate with a range of action-related variables, such as muscle activation, movement kinematics and limb dynamics (Evarts, 1968; Georgopoulos et al., 1982; Todorov, 2000; Holdefer and Miller, 2002; Sergio and Kalaska, 2003; for reviews see Porter and Lemon, 1993; Scott, 2003). For instance, single cells in the motor cortex display robust directional tuning during movement execution (Georgopoulos et al., 1982; Georgopoulos et al., 1986; Moran and Schwartz, 1993). These ties to intention execution and voluntary movement make it possible to volitionally modulate M1 activity and even dissociate this modulation from movement, as detailed below.

Brain-computer interface (BCI) paradigms are frequently employed in the study of sensorimotor control (Golub et al., 2016). This is due to their built-in flexibility: they allow researchers to reduce the complexity of the studied system by involving only a subset of the neural population, which can be recorded from in real-time, and defining the exact mapping between that population's activity and the sensory feedback returned. For instance, when a

subject performs an arm reach it impossible to know with any certainty whether the neurons being recorded from are the only ones contributing to the generation of that movement or to quantify every aspect of the sensory feedback that the subject receives throughout the reach—a combination of visual, proprioceptive and haptic information. In a BCI reaching task, however, it is up to the experimenter to explicitly stipulate the relationship between neural activity and sensory feedback, from selecting which of the recorded neurons will directly affect the outcome of the task to specifying the precise sensory modalities involved in the behavior. By providing this kind of tractable environment to probe motor control, BCI setups constitute an instrumental tool to assess the extent to which neurons in M1 can be volitionally modulated.

Early operant conditioning studies showed that individuals can learn to reliably modulate even single neurons in M1 at will if provided direct feedback of firing rate (Fetz, 1969; Fetz and Baker, 1973). In these pivotal studies, Fetz and colleagues measured the activity of a single cell in M1 and trained monkeys to increase the firing rate of this neuron to receive a biscuit. This was made possible because animals were given *neurofeedback*, that is, sensory feedback that directly related to the firing of the neuron in question. For instance, the subject would be allowed to look at a meter to aid in the task: the needle would go to one side of the meter when the firing rate was low, and it would go to the other side when the subject increased the neuron's firing rate. What they found was that animals were able to find a strategy to modulate the selected neuron, whether neurofeedback was provided visually or when the needle was replaced by an auditory cursor. These experiments were the first to demonstrate animals were able to voluntarily modulate activity of single cells in M1 in the context of a goal-directed task, guided by otherwise arbitrary effectors such as the position of a needle on a meter or the pitch of an auditory cursor.

Further examination following those studies has shown that M1 activity can be flexibly dissociated from movement across a range of behaviors (Schieber, 2011). For instance, the strength of the correlation between M1 modulation and the bursts of muscle activity

that ordinarily accompany it can be flexibly adjusted by conditioning on electromyography feedback, making it possible to both magnify the neuron-to-muscle connection (Davidson et al., 2007) or decouple it entirely (Fetz and Finocchio, 1971, 1975).

Indeed, during BCI control subjects often stop moving the native limb altogether, even when the decoder is trained on overt movements (Taylor et al., 2002; Carmena et al., 2003). Moreover, when neural activity is used to drive electrical stimulation of paralyzed muscle tissue, subjects can learn to modulate individual neurons to control muscles regardless of their original correspondence with limb movement (Moritz et al., 2008). Similarly, subjects can volitionally control small ensembles of neurons in the motor cortex irrespective of prior directional tuning or co-modulation patterns (Moritz and Fetz, 2011; Law et al., 2014; Koralek et al., 2012; Clancy et al., 2014; Kennedy et al., 2000).

Together, the above body of literature suggests that there are few constraints—if any—on the volitional modulation of M1. *However...*

1.2 Sensory Signals in the Primary Motor Cortex

Non-volitional signals are also present in the motor cortex, and may be powerful enough to compete with voluntary modulation. Anatomical studies describe prominent somatosensory inputs arriving at M1 (Rosén and Asanuma, 1972; Sato and Svoboda, 2010; Hooks et al., 2013), which can trigger salient sensory-evoked responses in the motor cortex (Kleinfeld et al., 2002; Ferezou et al., 2007; Huber et al., 2012; Petrof et al., 2015; Stavisky et al., 2017b).

At the same time, there is evidence of cortical involvement in reflexive motor responses. The long-latency response (LLR) is the slowest component of reflexes, and precedes the earliest voluntary motor command in the reaction to feedback disturbances during movement. LLRs can be triggered by postural (Bedingham and Tatton, 1984; Pruszynski and Scott, 2012) or visual (Day and Lyon, 2000; Franklin and Wolpert, 2008) perturbations, and have been shown to engage M1 and other cortical circuits that are fundamental to voluntary

movement (Evarts and Tanji, 1976; Cheney and Fetz, 1984; Jacobs and Horak, 2007; Scott, 2004; Pruszynski, 2014; Scott et al., 2015). Furthermore, LLRs share some features that are characteristic of voluntary control: they are sensitive to the subject's intention and to task goals (Crago et al., 1976; Tanji and Evarts, 1976; Colebatch et al., 1979; Shemmell et al., 2009; Nashed et al., 2012), they adapt during learning (Franklin et al., 2014), and they can change depending on factors such as the physical properties of the limb and the environment (Kurtzer et al., 2014; Shemmell et al., 2009).

Sensory-evoked modulation is not limited to somatosensation: visual responses have also been shown to drive motor cortex. During BCI experiments, visual perturbations of the cursor can drive fluctuations in motor cortical activity unrelated to corrective cursor movements (Stavisky et al., 2017a), and visual feedback can influence BCI control depending on whether it is congruent or incongruent with proprioception (Suminski et al., 2010). In human BCI clinical trials, Collinger and colleagues have observed that the neural encoding of grasp changes depending on whether that grasp is performed in open air or around an object, to the point where it could compromise grasp control (Downey et al., 2017).

The extent to which these sensory-evoked responses might impact volitional control is therefore not known.

1.3 Reconciling the Duality

Together, these two streams of observations beg certain questions. To what degree is neural activity in M1 under volitional control? How is volitional control dependent on sensory context? Can the volitional and sensory signals that coexist in M1 be balanced without compromising behavioral performance?

To address these questions, we leveraged BCIs to design a paradigm that dissociates volitional and sensory drivers of neural activity. We trained three macaque monkeys to perform experiments in which the modulation of a single neuron's firing rate in M1 dictated the

position of a computer cursor along a one-dimensional axis. While the neural requirements were kept identical throughout a session, we isolated and altered aspects of the task's sensory feedback in order to assess their influence on the subject's ability to succeed at the task. We tested three manipulations of sensory context through this framework: axis orientation, movement location in the workspace and arm posture.

Our paradigm revealed that certain sensory features can impact the ability to volitionally modulate single-cell activity in M1. Specifically, we found that the orientation of the cursor axis substantially impacted the ability to voluntarily modulate a majority of the neurons tested across all animals: subjects would be able to easily and quickly modulate a given neuron to drive the cursor along the axis at one angle, yet immediately exhibit difficulty modulating it when the cursor axis was rotated, despite the neural task remaining the same across conditions. Changing the location within the screen where the cursor movement took place also had an impact on control, albeit to a lesser extent than orientation. Our preliminary examination of postural effects, while inconclusive, demonstrated the promise of this novel paradigm to study the influence of sensory modalities beyond vision, such as somatosensation. Crucially, providing the subjects with days of additional training did not rescue their difficulty in controlling the cursor under certain sensory contexts.

Combined, the findings described in this thesis suggest that the extent to which M1 activity is under volitional control is limited by sensory context.

2

Experimental Setup

All animal procedures were approved by the Institutional Animal Care and Use Committees of Carnegie Mellon University and the University of Pittsburgh.

2.1 Surgical Procedure

Three male rhesus macaques (*Macaca mulatta*) were each implanted with a recording microelectrode array (Blackrock Microsystems), which was inserted into the primary motor cortex (M1) using standard procedures (Suner et al., 2005; Velliste et al., 2008; Zhou et al., 2019). Each chronically-implantable array contained 96 electrodes in a grid configuration, with an electrode spacing of 400 µm and an electrode shaft length of 1.5 mm.

Arrays were visually placed on the caudal convexity of the precentral gyrus at the level of the spur of the arcuate sulcus of the hemisphere contralateral to the active arm, aiming for the proximal arm representation. Implant surgery occurred only after animals had been trained to perform reaching tasks in a virtual reality environment.

2.2 Neural Recordings

Microelectrode array voltage signals were amplified (1000x – 32000x), band-pass filtered (250 Hz – 8 kHz) and processed with a multichannel acquisition processor system (Plexon Inc., Monkeys N and A; Blackrock Microsystems, Monkey R). Waveform data was converted to spiking events by threshold crossing, and snippets that crossed the threshold were manually sorted online to isolate neurons for BCI decoding.

For all data sets included in this thesis, the subject's arms were lightly restrained during recordings. Subjects tended to exhibit minimal overt arm or hand movements while performing these BCI tasks; however this was not quantified during our experiments.

2.3 Virtual Reaching Environment

Before ever engaging in the single-neuron experiments detailed throughout this thesis (see *3 Dissociation of Sensory and Volitional Drivers of Neural Activity*), subjects had to learn to perform two core behavioral tasks involving a virtual-reality environment. Prior to being implanted with a recording array, each animal was first trained to execute hand reaches in this virtual environment. Once that stage had been mastered and the microelectrode array was in place, subjects were trained in a BCI task, in which they were required to modulate neural activity rather than perform an overt movement in order to interact with the virtual environment.

2.3.1 Hand-Control Center-Out Task

In our hand-control center-out paradigm, the subject sits in a non-human primate chair facing a mirror that reflects a 3D image from a stereoscopic computer monitor (Dimension Technologies). The display is mounted in a periscope-like design in front of the animal and incorporates a piece of black canvas immediately underneath. In this way, it is impossible

for subjects to see their limbs: the only visual feedback about arm movement is therefore through the computer monitor.

During the task, the movement of the arm ipsilateral to the recording site is restrained. Conversely, the contralateral arm is allowed to move freely, and through it the animal will be able to control cursor movement. To make this possible, the free arm is fitted with an infrared-emitting optical marker (Optotrak, Northern Digital) taped to the back of the hand. The 3D position of this infrared marker is tracked by 3 cameras placed in front of the subject, facing the active arm. Movement of the marker, and therefore the animal's hand, is captured by these cameras, fed to a computer, and returned in real time as the proportional movement of a spherical green cursor on the stereoscopic monitor.

On a typical hand-control trial, a blue spherical target is presented on the screen to indicate the possibility to start a new trial. For the trial to begin, the subject must move the tracked hand such that the green cursor is positioned over this central target, and keep the cursor within the target for a variable hold period (typically 500 ms). Once the hold period expires, the blue target disappears and reappears at one of the potential reach positions. The subject must then move the active arm to once again drive the green cursor to the blue target within a certain timeout period. Well-trained animals can perform this reaching movement in less than 600 ms, including reaction time. The subject is rewarded after holding the cursor within the target for a second hold period (also in the order of 500 ms).

In addition to the central target, there are a total of 26 possible peripheral targets, distributed in the 3D space within the virtual environment on the corners, faces, and edges of a cube and pushed outward to lie on a sphere, such that all the targets in the periphery are equidistant to the central target. This 3D reaching task typically includes two different types of reaches: center-out reaches require the subject to drive the cursor from the center of the workspace to one of the peripheral targets, whereas out-center reaches require the cursor to move from a target in the periphery to the central target.

For well-trained subjects, target and cursor radii are set to 6 mm and the workspace is scaled such that there are 85 mm from the center to the periphery, since this is as far as animals are typically able to reach. The workspace is carefully calibrated to ensure that the subject has free access to all the targets without the arm physically touching any obstacles in the vicinity.

2.3.2 Brain-Control Center-Out Task

In the brain-control—or just BCI, for short—center-out task, arm movement is unnecessary and the behavior of the cursor is instead governed by the real-time activity of a neural population. Specifically, subjects controlled the 2D velocity of a computer cursor by modulating the activity of the recorded neurons, and were tasked with hitting one of eight visual targets uniformly spaced around a circle of radius 85 mm.

Cursor movements were driven by a population vector algorithm (PVA) velocity decoder (Georgopoulos et al., 1986; Taylor et al., 2002; Chase et al., 2009). The implementation of this BCI decoder is described in detail in previous work from our group (Zhou et al., 2019). Briefly, spikes from each of the N recorded neurons were binned in 33-ms intervals, normalized by decoding parameters (baseline firing rate, tuning dynamic range), and smoothed online by averaging the last 5 time bins together. Each neuron's normalized firing rate values were used to scale a unit vector pointing in the neuron's preferred direction, and these scaled individual vectors were averaged to produce the cursor velocity.

Decoder parameters were obtained through a calibration period using a co-adaptive procedure (Chase et al., 2012) that cycled through the 8 possible targets and regressed the observed firing rates for each trial against target direction to recursively update estimates of the linear tuning while gradually decreasing the level of assistance provided to the subject's cursor control. After approximately 5 cycles (typically involving around 2 min of data collection), the subject was able to reliably control cursor movement, and decoding parameters were fixed for the remainder of the center-out session.

At the start of each trial, cursor position was reset to the center and a peripheral target appeared on the workspace. In a successful trial, the subject would drive the cursor to the target within 3 s and hold it in place for 100 ms to receive a reward.

A short session of the brain-control center-out task described here was run at the beginning of each single-neuron experiment in order to establish baseline firing rates across the neural population (see *3.1.1 Initialization of Target Firing Rates*).

3

Dissociation of Sensory and Volitional Drivers of Neural Activity

This chapter introduces the paradigm that enabled the findings reported in later chapters of this dissertation. Building on the foundational methodology discussed in chapter 2, here we provide detailed descriptions of the behavioral and neural aspects of this experimental setup.

We designed the **Feedback Alteration Single-cell Task** (FAST) to probe the influence that sensory context can exert on an individual's ability to voluntarily modulate M1 activity. This novel BCI paradigm separates the neural requirements of the task from certain features of sensory feedback, and allows us to evaluate a subject's ability to volitionally control the activity of a single neuron across a range of sensory conditions.

This is achieved by translating the firing activity of a single cell—termed the command neuron (CN)—into the real-time position of a computer cursor along a one-dimensional (1D) movement axis, while allowing for flexible manipulation of task feedback to test volitional modulation across sensory contexts (Figure 3.1).

To successfully complete a trial, subjects had to control the CN's firing rate to drive the

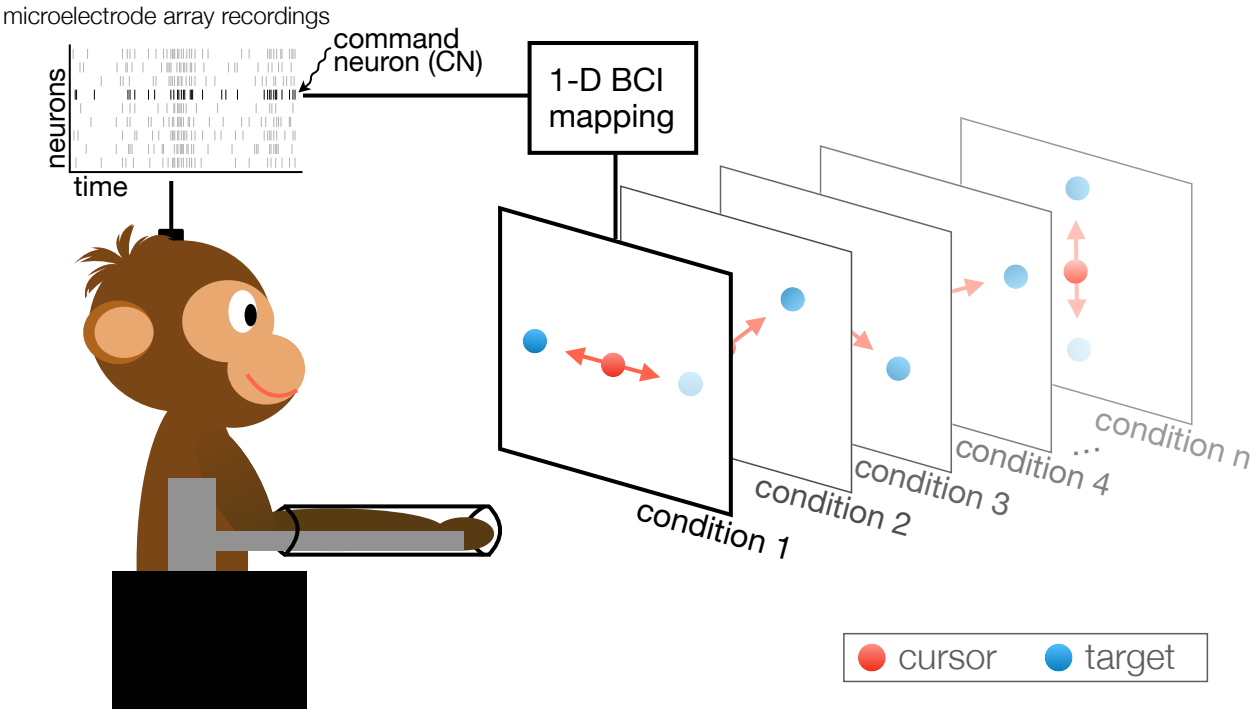

Figure 3.1 | The FAST paradigm dissociates the neural requirements of the task from its sensory consequences. A 1D BCI position decoder takes in the activity of a command neuron (CN) to drive a cursor along the movement axis on the display.

cursor along the axis on the display and sequentially hit a low firing rate target and a high firing rate target within 2 s. Trials were presented continuously, with no inter-trial interval.

Blocks of trials were characterized by the different sensory contexts provided; for instance, the orientation of the movement axis could vary across conditions (Figure 3.1). Specifically, a sensory condition in any FAST paradigm is fully defined by three parameters: orientation of the movement axis, coordinates of the axis origin and posture of the arm contralateral to the implanted hemisphere. On a typical block, subjects were given 4 min to continuously modulate the CN and complete as many trials as possible on the same feedback condition.

3.1 Task Flow

Cursor position along an invisible 1D axis was determined by the activity of a single command neuron (CN). Each trial required the subject to push the cursor to 2 targets appearing in succession on the movement axis: a *low* firing rate target, followed by a *high* firing rate target (Figure 3.2). Cursor and target radii were 4 mm each. No hold time was enforced at either target; subjects merely had to exhibit a firing rate that made the cursor pass through the target boundary for the hit to count.

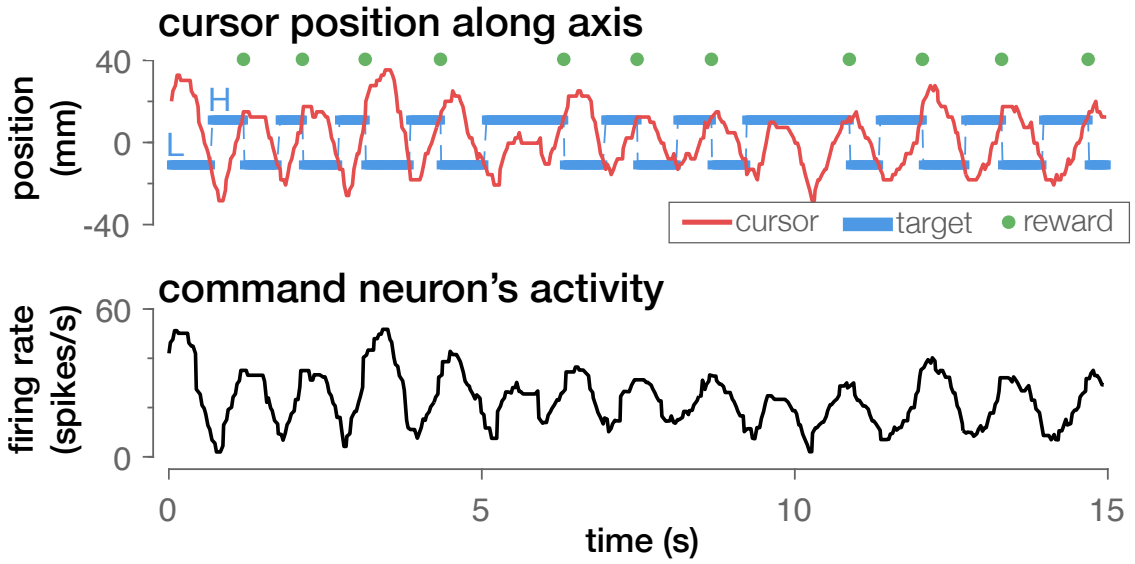

Figure 3.2 | The activity of the command neuron determines the position of the cursor along the axis. The CN's spike count is smoothed online and used as a control signal (black, bottom) to determine the position of the cursor (red, top) along the movement axis. In order to receive a reward (green dot), the subject must modulate the CN's firing rate to drive the cursor from the low (L, blue) to the high (H, blue) target in under 2 s. Data: session FOA06 (Monkey A).

Trials started with the low target, which remained on the screen until the end of the block unless it was hit by the cursor. Once the subject decreased the CN's firing rate enough to drive the cursor into the low target, the high target immediately appeared. The subject then had a 2-s window to drive the cursor into the high target before it disappeared. If the subject accomplished this, the high target disappeared and a liquid reward was made available, and the low target immediately reappeared to begin the next trial. If the subject

did not hit the high target within the 2-s window, the high target disappeared without a reward and the low target immediately reappeared, beginning the next trial.

During each block, subjects were free to complete as many trials as possible in the given sensory context. For Monkeys N and R, blocks of trials lasted 4 min with the exception of *focus* blocks (see *5.1 Orientation Dependence Remains Despite Within-Day Focused Training*). Monkey A's blocks ranged in duration (1.54 ± 0.57 min, mean ± 1 s. d.).

The CN's instantaneous firing rate was extracted from spike counts binned at 33-ms intervals and smoothed with a 20-bin boxcar filter to reduce cursor jitter. This filtered firing rate r_f was translated into a 1D cursor position s along the axis according to the linear map

$$s = r_f \, \frac{s_H - s_L}{r_H - r_L} \, ,$$

where r_H and r_L are the target rates that the CN's firing had to reach in order for the cursor to enter the high and low targets, respectively. The 1D positions of those targets along the movement axis, s_H and s_L, were always set to be 34 mm to either side of the axis origin.

The axis origin, axis orientation and arm posture fully defined the sensory context (see *3.2 Sensory Manipulations*), and were held fixed throughout a block of trials.

3.1.1 Initialization of Target Firing Rates

Before running any of the FAST paradigms each day, all subjects first performed a short, standard BCI center-out session to collect baseline firing statistics from the neural population, which were used to set the FAST parameters.

In this session, subjects used a PVA velocity decoder to control the 2D velocity of a computer cursor (see *2.3.2 Brain-Control Center-Out Task*). Note that, unlike the FAST paradigm, in this task the decoder includes the full neural population being recorded from in order to drive the cursor, rather than restricting it to a single neuron.

Following decoder calibration, subjects performed 162 ± 60 trials (i.e., approximately 20 trials per target) of brain-control center-out cursor control prior to the start of any given FAST experiment. At the end of the center-out session, we randomly selected one well-sorted neuron to be used as the command neuron in the FAST paradigm, with a bias towards selecting a new cell each day and choosing neurons that displayed a wide dynamic range. Finally, to compute the initial firing rate targets for the CN, we computed the 20th and 80th percentiles of the firing rates exhibited by the selected neuron across the entire 2D BCI center-out session.

3.1.2 Difficulty Progression

To more fully probe the limits on controllability, the paradigm run on Monkeys N and R included an automated schedule to gradually increase trial difficulty within each 4-min block as a function of recent task performance (Figure 3.3), similarly to the setup described by Law et al. (2014). The initial firing rate values for the target rates r_H and r_L were set, respectively, to the 80th and 20th percentile of the firing rates displayed by the CN during the BCI center-out task. If the animal was highly successful, these target rate values would be spread farther apart. The task enforced at all times a minimum difference between firing rate targets of 15 Hz and a minimum r_L value of 3 Hz.

The difficulty scaling was relative to the CN's firing statistics. Let D be the average firing rate between the 80th and 20th percentiles. Every 10 s, we computed the number of successes accrued over the past 20 s. If the number of successes was greater than 10 (i.e., greater than 30 successes/min), a full-step increase in difficulty was triggered, which increased the separation of the target rates by $\frac{1}{6}D$. That is, r_H was increased to $r_H + \frac{1}{12}D$ and r_L was decreased to $r_L - \frac{1}{12}D$. If the number of successes was not as high but was still greater than 7 (i.e., greater than 20 successes/min), a half-step increase in difficulty would instead be triggered, which increased the separation of the firing rate targets by $\frac{1}{12}D$. For the half-step increase, r_H increased to $r_H + \frac{1}{24}D$, and r_L decreased to $r_L - \frac{1}{24}D$. In this way,

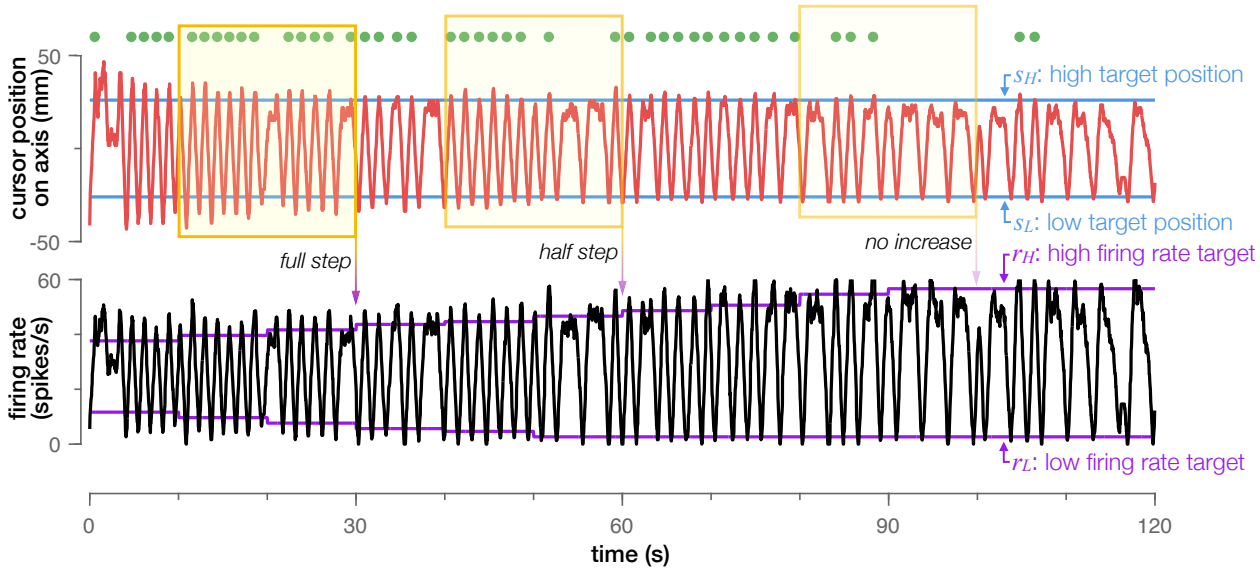

Figure 3.3 | Task difficulty gradually escalates as a function of behavioral performance. Cursor position (top) and firing rate (bottom) for one CN during the first 2 min of a block. Horizontal blue lines (top) show the positions, along the axis, of the high and low targets; the purple lines (bottom) show the firing rates required to hit those targets. Trial successes are indicated above cursor position (green). Task difficulty was given by the separation between the firing rate targets. Difficulty could increase by a half or a full step if the subject accrued more than 7 or 10 successes, respectively, within the most recent 20-s window. Each yellow box highlights an example of a 20-s window in which task performance triggered either a full-step increase, a half-step increase, or no increase in task difficulty. Data: session FON35 (Monkey N).

excelling at the task resulted in the target rates r_H and r_L being updated so as to increase the dynamic range required of the CN to complete a trial. Note that if the change would decrease r_L below 3 Hz, r_L was left at its last value and the entirety of the change was instead applied to r_H.

3.2 Sensory Manipulations

Sensory context in the FAST paradigm is fully defined by three parameters: the orientation of the movement axis, the location of the axis origin within the display and the passive posture of the arm contralateral to the implanted hemisphere (Figure 3.4). We define 8 possible orientations, 9 possible locations and 3 possible postures.

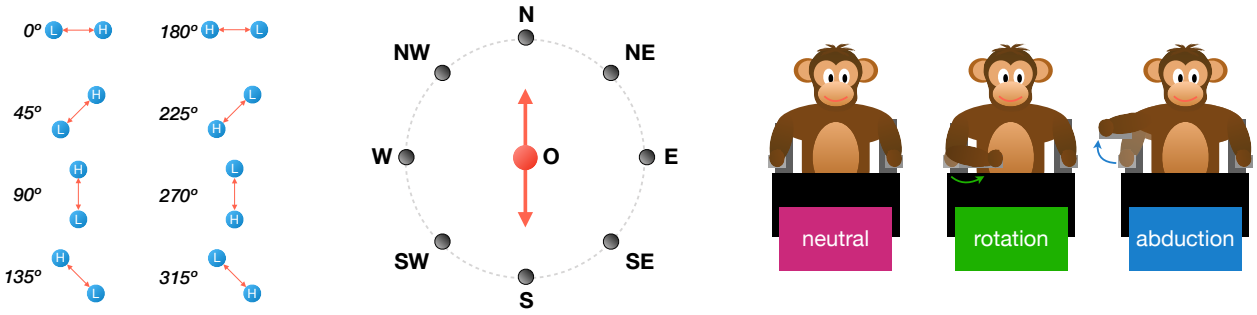

Figure 3.4 | Sensory context is defined by three features: orientation, location and posture. Orientation (left) represents the angle of the movement axis with respect to the horizontal (0^{o}), given the positions of the high (H) and low (L) targets. Location (middle) references the 2D coordinates of the axis origin (red) within the workspace. Posture (right) symbolizes the subject's arm positioning. Unless otherwise specified, conditions are assigned the central location (O) and neutral posture by default.

3.3 Quantification of Volitional Control

A critical step to study changes in a subject's ability to voluntarily modulate the command neuron is to quantify that volitional control. Our approach consisted in designing a metric that would integrate the key aspects of single-neuron modulation that *allow* the subject to succeed at the task. In other words, we sought to evaluate the elements necessary to *produce* a successful trial rather than measuring task performance itself. Naturally, both would be heavily correlated by definition. However, a closer examination of modulation across time would be able to paint a picture of the underlying mechanisms at play more nuanced than that given by a binary classification deeming a trial either a success or a failure.

This led to a crucial question: what makes volitional control *good*? Good volitional control of an individual neuron is characterized by an ability to hit a *wide* range of firing rate values, and an ability to *quickly* transition from one firing rate to another.

To push the subject's ability and comprehensively probe volitional control, in two subjects (Monkeys N and R) we gradually expanded the dynamic range required to hit the visual targets on the screen (Figure 3.3). While the high and low targets' visual positions along the movement axis remained constant throughout a block, the gain between firing rate and

cursor movement gradually shrank as a function of performance, increasing task difficulty. Thus, multiple consecutive rewards in close succession resulted in a separation of the targeted firing rates that the CN had to reach in order for the cursor to hit the targets on the screen. These target rates were reset to their original original values at the beginning of every new sensory context block.

3.3.1 Single-Neuron Controllability Computation

We designed *controllability* as a metric that captures the ease with which the subject can voluntarily modulate the firing rate of the CN across its dynamic range, given a sensory feedback condition.

Controllability was therefore computed as a function of two variables: the command neuron's normalized firing rate—to represent dynamic range—and the time derivative of its firing rate—to incorporate the ability to transition across firing rate values easily. Single-neuron firing rates z_t were obtained offline by smoothing spiking event timestamps with a Gaussian kernel ($\sigma_k = 150$ ms) and z-scoring these smoothed values over the entire FAST session for that day. The normalized time derivative of the firing rate $\dot{z}_t$ was computed by taking the finite difference of the smoothed firing rate over the sampling interval $h = 1$ ms, and scaling by the Gaussian kernel width σ_k:

$$\dot{z}_t = \frac{2\sigma_k}{h} \left(z_{t+h} - z_t\right) \quad .$$

In each block, we first identified the *peak minute*: the continuous minute of control that had the highest number of successful trials (Figure 3.5). We then subdivided this peak minute into 12 non-overlapping, 5-s bins. In each bin we computed a *control ellipse*: the 2-standard-deviation covariance ellipse between the CN's normalized firing rates z and their time derivative $\dot{z}$. The area of this ellipse captures both the dynamic range achieved by the CN during this 5-s bin and the speed with which the subject was able to modulate it. We thus quantified the controllability of the CN, given a sensory context, as the average area of these

control ellipses across the 12 bins of the peak minute. Blocks in which the subject failed to complete at least 5 successful trials within a continuous minute were discarded during

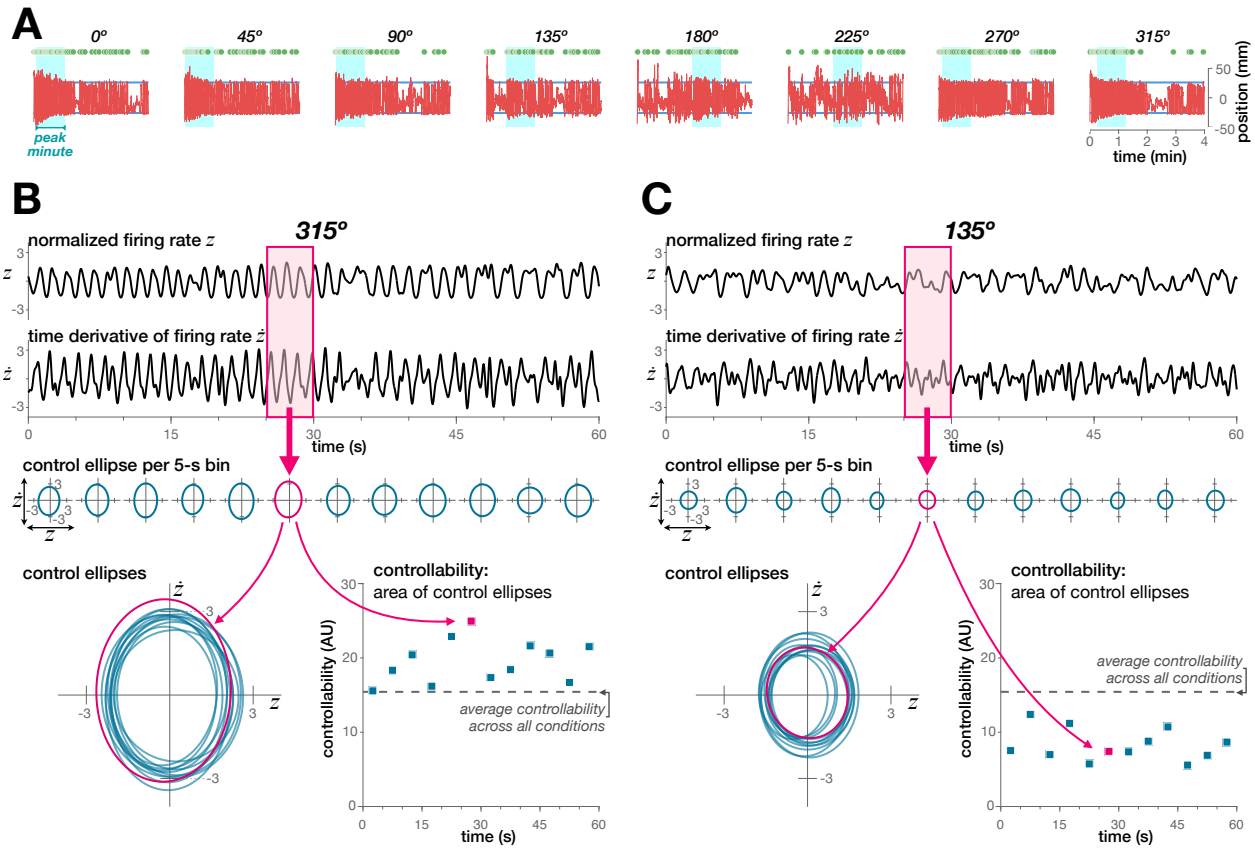

Figure 3.5 | Controllability quantifies the ease with which a neuron's firing rate can be modulated. (A) A block's peak minute (cyan shading) is defined as the 1-min, continuous interval of highest average success rate within the block. Cursor (red) and target (blue) positions produced by the same CN are shown for one 4-min block per condition. **(B)** Controllability is computed in 5-s non-overlapping bins as a function of the CN's firing rate and its derivative with respect to time. Normalized firing rate and its time derivative are plotted against time during the peak minute at the 315º condition (top). For each 5-s bin, we show the corresponding control ellipse: the 2-standard-deviation covariance ellipse between firing rate and its time derivative (middle, teal). Controllability is defined as the area of these control ellipses (bottom, left), plotted here for the peak minute of this 315º block against time (bottom, right). Dashed line indicates this CN's average controllability across all conditions. **(C)** Same as **(B)**, at the 135º condition. Firing rates close to baseline or resistant to change indicate difficulty controlling the CN's activity and produce low controllability values. At 315º, the subject was able to readily modulate this CN both widely and quickly, reflected respectively in the large fluctuations in firing rate and its derivative. At 135º, the animal was less adept at modulating the same CN. Data: session FON35 (Monkey N).

analysis. If a CN had more than one condition not represented because of this performance threshold, it was removed due to insufficient data. Monkey A had 16 blocks shorter than 1 min (46 ± 8 s); instead of selecting a peak minute, in these cases the full duration of the block was binned at 5-s intervals and used to compute controllability values.

We quantified the significance of the impact that sensory context had on a CN's controllability using one-way ANOVA to partition the total variation found in controllability into between-conditions variation and a within-conditions variation. The ratio between these components determines whether controllability means are significantly different across conditions. To assess the magnitude of this effect, we computed the fraction of variance explained by sensory context as:

$$V_{explained}(\%) = \left(1 - \frac{SSR/(N-K)}{SST/(N-1)}\right) \times 100 \quad ,$$

where SSR and SST denote respectively the residual and total sums of squares, K is the number of sensory contexts tested and N is the total number of 5-s bins across conditions in the experiment.

3.3.2 Alternative Metrics Explored

Finding a metric that captured the ease with which a subject is able to modulate a neuron was not trivial. We explored several other approaches to quantify volitional control, mainly focusing on performance proxies like target overshoot or reaction time. Ultimately, these were discarded in favor of our chosen metric, which we felt provided a more stable estimate of controllability.

Among them, the *difficulty-progression score* described below provided an intuitive framework to interpret the behavior observed during FAST experiments.

This strategy focused on task performance, commonly quantified through success rate. Can the subject readily drive the cursor to the required targets and be rewarded in the task? Average success rate can easily be computed as the number of trial successes accrued

within a certain window given its duration. The time resolution of such a metric is therefore determined by the bin size used to compute that average.

For two of our subjects, the automated difficulty scheduling provided an opportunity to further dissect volitional control (see *3.1.2 Difficulty Progression*). This opened the door to a finer assessment of performance based not only on success rate achieved at different difficulty levels but also the extent to which the animal was able to push the CN's dynamic range and continue to complete trials successfully as the task became more challenging.

Since multiple consecutive rewards in close succession resulted in an increase in task difficulty, we weighed the best performance a subject could achieve at a given level by the difficulty of the task at that point. Visually, this is equivalent to plotting maximum success rate achieved within each level against difficulty. We can quantify the proficiency with which a subject is able to volitionally modulate the CN in order to push through difficulty levels during a block by taking the area under that curve. We defined the difficulty-progression score as this area, easily computed by applying the trapezoidal rule to approximate the integral.

However, we found that an important limitation of this approach was the heavy weight that the metric assigned to level progression, shadowing other factors that should be equally decisive from a behavioral perspective—such as the achieved success rate across difficulties or the sharpness with which cursor movement was corrected. This issue was present regardless of whether difficulty progression was assessed in absolute firing rate terms or relative to the original dynamic range displayed by the CN.

Furthermore, this made the reliability of the metric vulnerable to fluctuations in the animal's motivation. Progression across difficulty levels is not independent from the length of the block: the longer the subject works at the task, the farther they could progress. This sensitivity to "work time" could therefore become problematic if the subject took a break in the midst of a block. One way that this challenge can be addressed is by modifying the

task flow to introduce a trial initiation step; this would mark periods of time in which the animal was not trying to perform the task, and could also be made to start and stop the block timer. Otherwise, controlling for this susceptibility would require closely monitoring motivation or attention changes, substantially reducing the practicality of this performance metric.

4

Impact of Sensory Context on Single-Neuron Controllability

We sought to assess the impact of sensory context on the ability to voluntarily modulate neurons within M1. To do this, we trained three rhesus macaques on an array of paradigms within the Feedback Alteration Single-cell Task (FAST) framework.

The set of FAST paradigms described in this thesis all hinge on one common element: a decoder that translates the firing activity of an individual command neuron into the position of a cursor along a 1D axis, regardless of the angle or the position in which that movement axis is displayed. By doing this, the FAST framework creates room for the flexible manipulation of task feedback in order to test a subject's ability to voluntarily modulate an individual neuron across a range of sensory conditions .

To quantify the volitional control of single cells in M1, we developed a *controllability* metric that captured both the magnitude and the speed with which the CN could be up- and down-modulated (see *3.3.1 Single-Neuron Controllability Computation*). Briefly, we first identify the 1-min, continuous interval of best performance within each block based on success rate, and divide this *peak minute* into 12 non-overlapping, 5-s bins. Good control is

characterized by large excursions in both firing rate and its derivative over this time period. In contrast, poor control is characterized by small deviations in these values (Figure 4.1). We captured these features with a *control ellipse*, formed by taking the area of a two-standard-deviation covariance ellipse between normalized firing rate and its time derivative. Larger areas subtended by the control ellipses, shown graphically as teal ovals in Figure 4.1, signal better volitional control. The overall controllability for a particular sensory context is defined

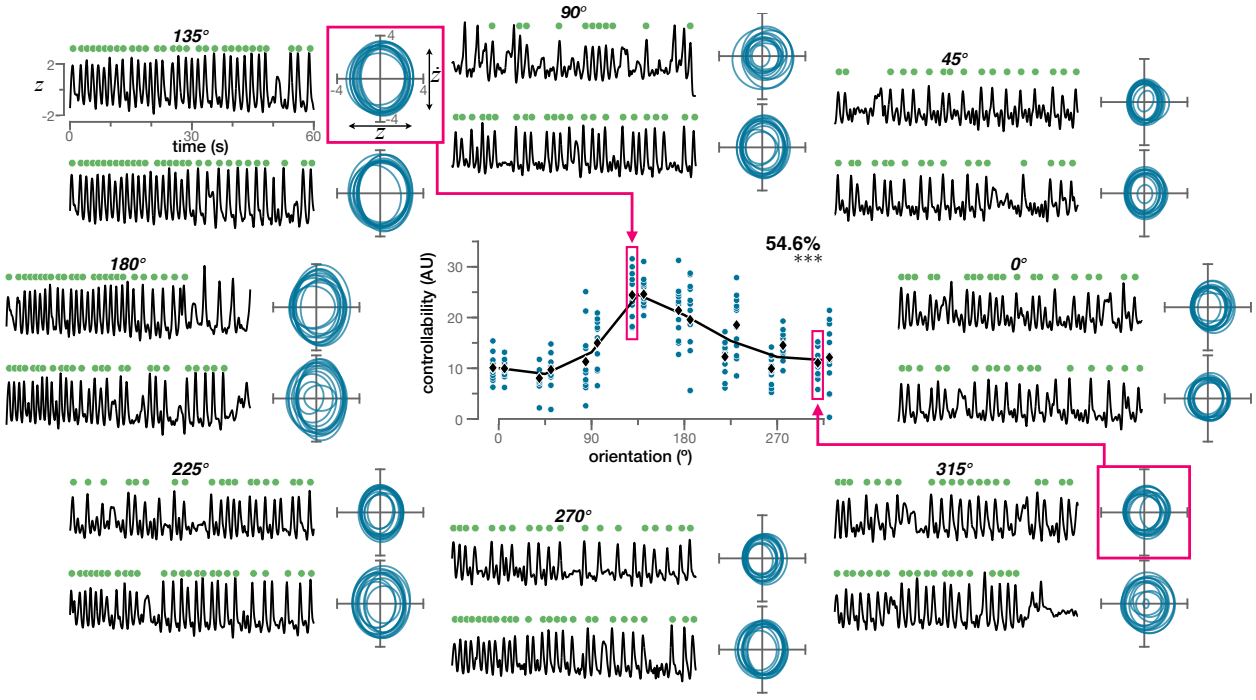

Figure 4.1 | Example showing modulation of controllability by orientation. Controllability throughout a session is determined by the modulation of the CN during the peak minute of each block. **Periphery:** Data per block are arranged in this figure based on the sensory condition. For each block, normalized firing rate (black) during the peak minute is plotted against time, with rewarded trials indicated by green dots. Beside it, the corresponding control ellipses (teal) for that peak minute are shown, each resulting from a 5-s bin in the peak minute. **Center:** Controllability, defined in each 5-s bin as the area of that bin's control ellipse, is plotted here against orientation. Controllability values from individual bins (teal) are grouped in the x-axis by condition, with the two block repeats at the same orientation slightly offset horizontally from one another for visibility. Average controllability of a block is indicated by a black diamond; a curve connects average controllability per condition. Of the variance present in this CN's controllability, the fraction that is explained by orientation is reported (%), along with significance (ANOVA, *** for $p < 0.001$). Data: session FON32 (Monkey N).

as the average of the control ellipse areas across the 12 bins within the peak minute, indicated as black diamonds in Figure 4.1 and throughout this thesis.

4.1 Movement Orientation Interacts with Controllability

We first tested whether the angle of the movement axis impacts the ability to volitionally control a given neuron. To probe this, we rotated the movement axis across conditions. If neurons are firmly tied to particular sensory contexts, then behavioral performance would change across directions. However, if voluntary modulation is truly dissociable from the effector, there should be little difference in controllability across conditions.

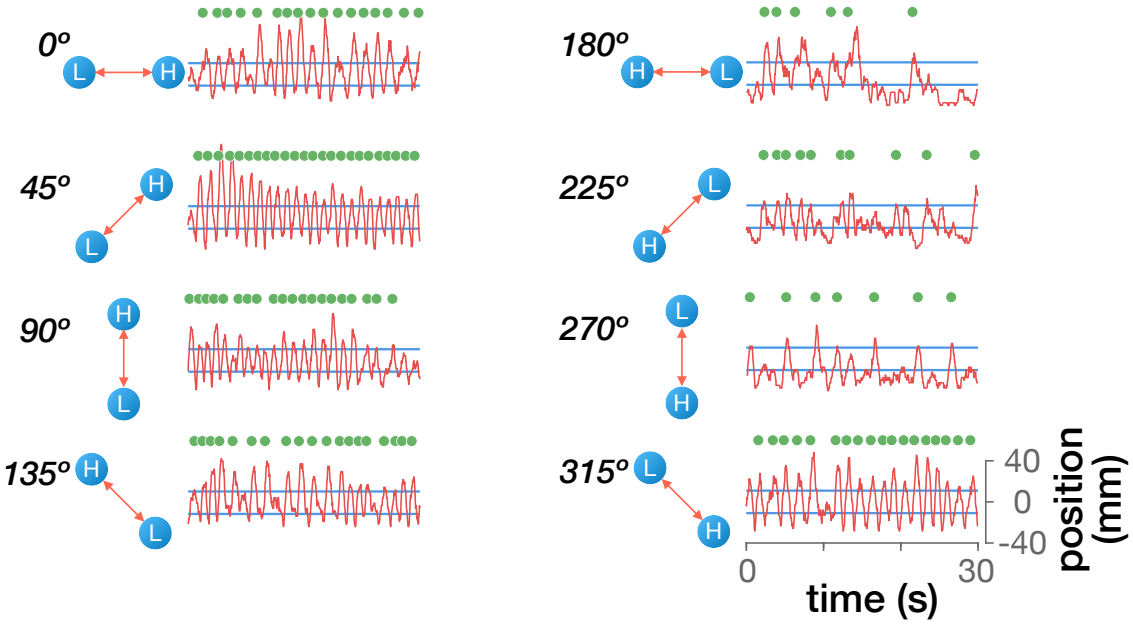

Figure 4.2 | Example of behavior in the FAST orientation paradigm. The CN's firing drives the cursor along axes with different angular orientations across blocks. Here we show the cursor position produced by a single CN during the first 30 s of each condition. Note the clear variability in performance across different axis orientations. Data: session FOA06 (Monkey A).

In the FAST orientation experiments, each condition was therefore characterized by the angle of the movement axis, keeping the location of the axis origin fixed at the center of the

display and the subject's arm in the neutral posture throughout every session. We tested 8 orientations in 45º intervals (Figure 4.2). To control for motivation and other potential time-related confounds, in a typical FAST orientation session each orientation was presented in random order once, and then all 8 orientations were randomly presented a second time, for a total of 16 blocks of trials. Note that block shuffling was constrained to ensure that no condition was presented twice in a row and to prevent *opposite* orientations (i.e., angles separated by exactly 180º, which produce parallel movement axes with mirroring polarity) from appearing consecutively.

We found that the orientation of the movement axis had a substantial impact on controllability. An example of this effect can be seen in Figure 4.1, in which we show the peak minute for each block across the 8 axis orientations. When the movement axis was angled at 135º, control was smooth and fluid, as marked by the rapidly changing firing rate and large dynamic range, resulting in a high success rate. In contrast, control was more coarse when the axis was oriented at 315º: in multiple instances, the subject had to make several attempts to reach the target after failing to drive the cursor far enough to receive a reward on the first try. As a result, larger control ellipses stem from the 135º blocks than from the 315º blocks. The control ellipse areas for a given block correspond to the controllability values shown on the central chart of Figure 4.1. By examining the changes in controllability across movement angles, we can weigh the impact that sensory context has on the subject's ability to volitionally control the CN.

All three monkeys showed widespread changes in controllability as a function of movement orientation (Figure 4.3). A majority of the CNs showed significant differences in controllability across orientations (one-way ANOVA, $p < 0.05$): 96% for Monkey N (44 of 46 neurons), 100% for Monkey R (16 of 16 neurons), and 80% for Monkey A (8 of 10 neurons). We can quantify these changes in subjects ability to voluntarily modulate the CN's activity as a function of sensory context by computing the fraction of the total variance found in controllability that is explained by orientation. Disparities in controllability across blocks

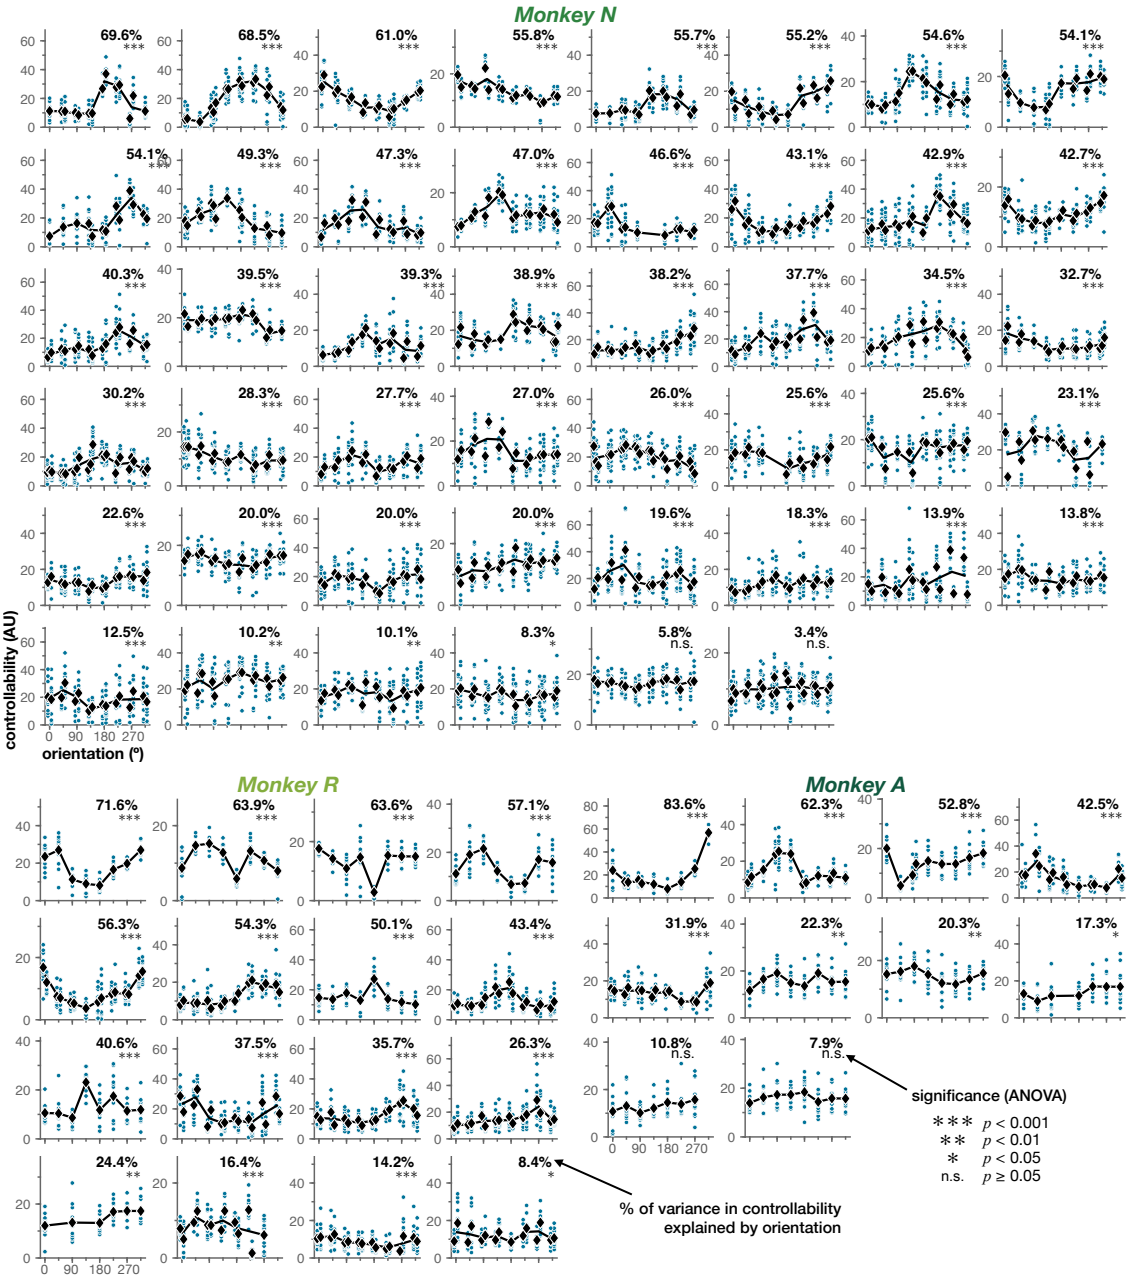

Figure 4.3 | Volitional control is modulated by the orientation in which the 1D visual feedback is provided. Single-neuron controllability changes with orientation (Monkey N: $n = 46$ neurons, Monkey R: $n = 16$ neurons, Monkey A: $n = 10$ neurons). Controllability values (teal) and block averages (black diamonds) are shown; curves connect average controllability per condition. Neurons are sorted by descending variance explained by orientation (%).

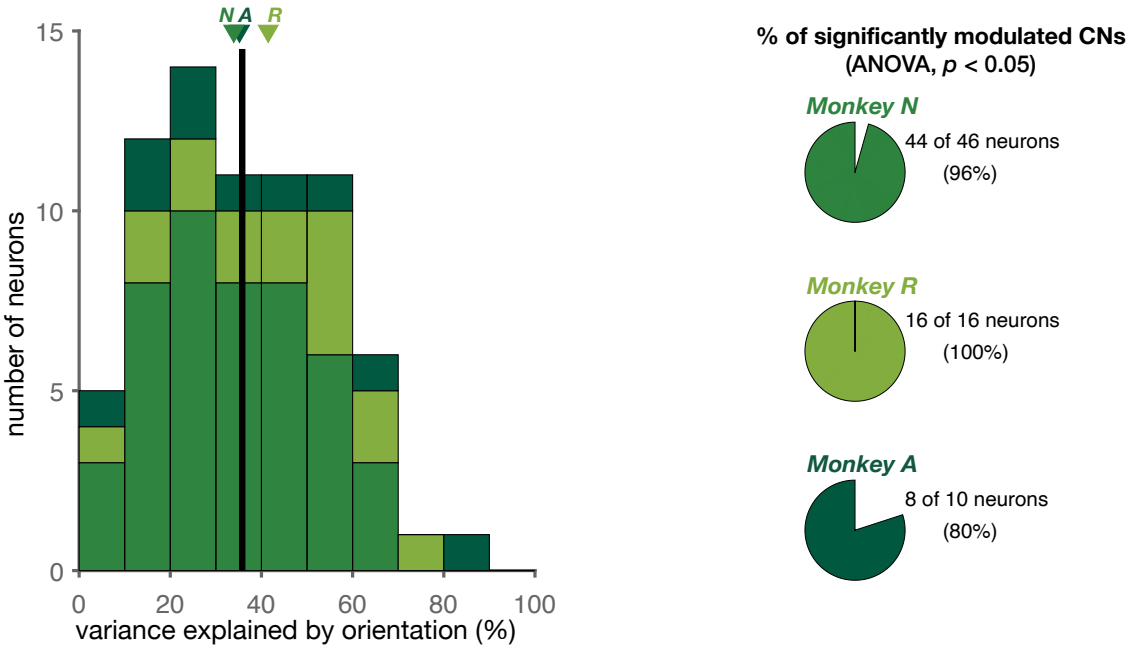

Figure 4.4 | Movement orientation accounts for a substantial portion of the variance found in controllability. Stacked histogram illustrates variance explained by orientation, indicating averages per subject (arrows) and overall (black line). Pie charts show fraction of CNs that displayed significant controllability differences across orientations (ANOVA, $p < 0.05$).

were largely accounted for by the changes in the angle of the movement axis (Figure 4.4), with an average across animals of 35.8% of variance in controllability being explained by orientation (Monkey N: 33.9%, Monkey R: 41.5%, Monkey A: 35.2%).

Controllability was modulated by movement angle across the population in a way that resembled the directional tuning found during brain-control center-out sessions (Figure 4.5), both in orientation and modulation depth. To examine each CN's tuning changes across paradigms, we linearly regressed controllability against axis orientation to define its FAST tuning direction and individually compared it to the directional tuning it displayed during the same day's BCI center-out session. The FAST and BCI tuning directions were strongly correlated for all animals ($\rho = 0.84$, Pearson's correlation). To further quantify tuning similarity, we considered the FAST-to-BCI angle: the difference in direction between the FAST and BCI tuning. Tuning directions did not seem to diverge markedly ($45.1 \pm 40.5^{\circ}$ on

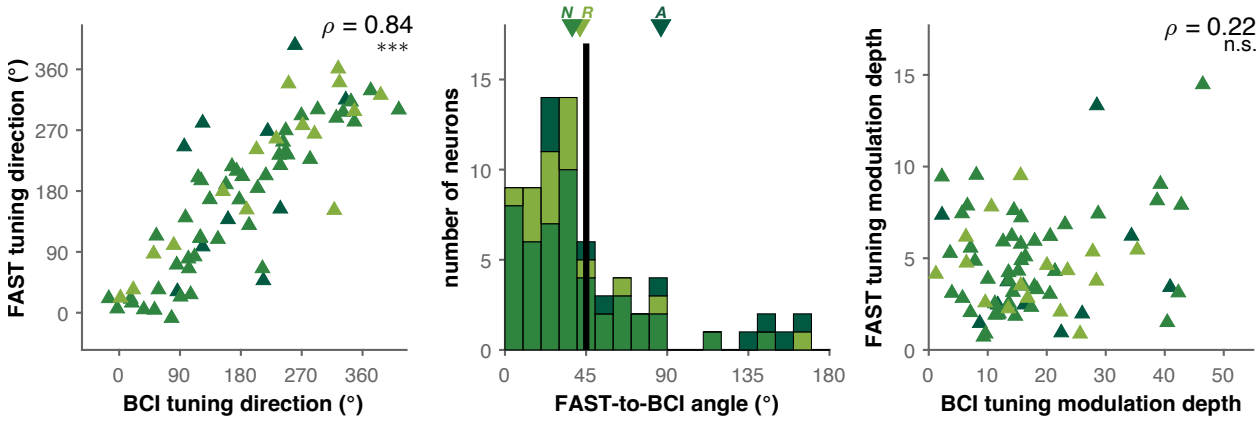

Figure 4.5 | Orientation interactions with controllability echoed BCI directional tuning. Left: Tuning directions corresponding to the FAST experiment and the brain-control center-out data collected on the same day, plotted against one another, reveal a strong correlation. **Center:** Histogram of the FAST-to-BCI angle, the separation in the tuning direction of each CN between paradigms. **Right:** Tuning modulation depth under FAST vs BCI, representing the strength of the relationship with direction. ρ indicates Pearson's correlation coefficient (*** if $p < 0.001$, *n.s.* if $p \geq 0.05$).

average, across subjects). Indeed, the FAST-to-BCI angle remained below 60° for a majority of CNs: 80% (37 of 46 neurons) for Monkey N, 81% (13 of 16 neurons) for Monkey R and 50% (5 of 10 neurons) for Monkey A. We compared the FAST and BCI tuning based on their modulation depth as well as their direction, but found no significant correlation in this aspect between the FAST and BCI settings (Pearson's correlation: $\rho = 0.22$, $p \geq 0.05$).

Since neural firing could be affected by fluctuations in motivation throughout the course of a session (Hennig et al., 2021), we added a control to test for temporal effects on controllability, caused by sensory conditions having been run block-wise. We controlled for time by fitting a linear regression between controllability and block number. We then used the residuals of this regression to re-compute the variance explained by sensory context. However, the dependence of controllability on orientation could not be explained by the passage of time across blocks (Figure 4.6).

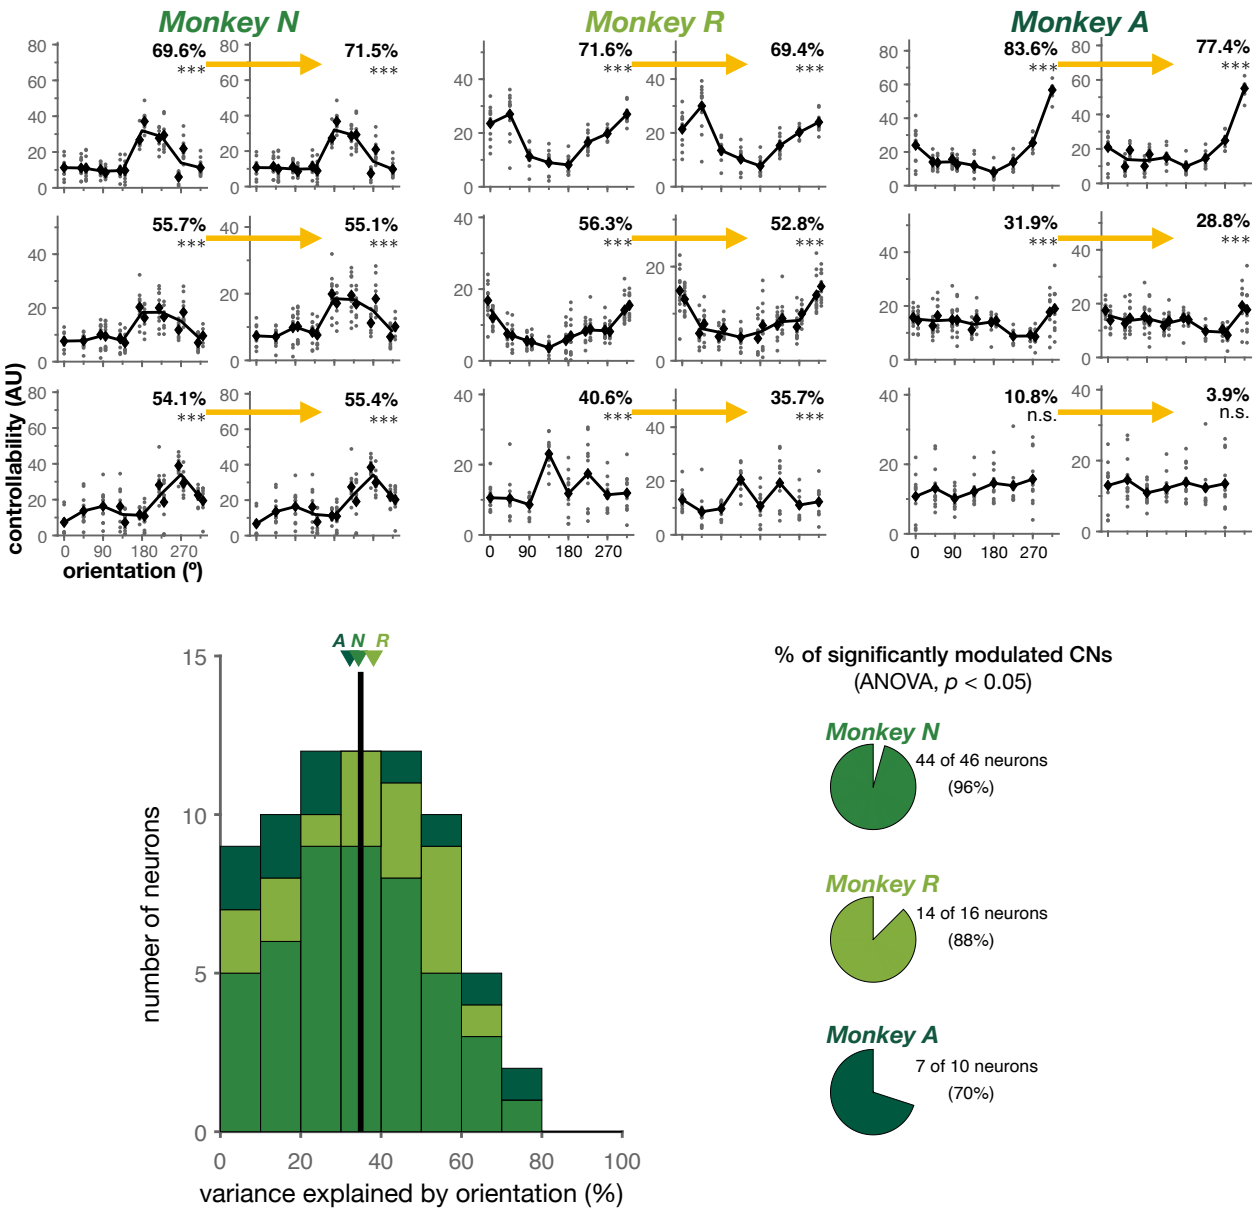

Figure 4.6 | The passage of time did not explain controllability modulation by orientation. Top: Controllability across orientations before (left of arrow) and after (right of arrow) factoring out block order through linear regression for 3 example CNs per subject. Controllability values (gray) and block averages (black diamonds) are shown; curves connect average controllability per condition. Variance explained by orientation is reported, along with significance (ANOVA): *** if $p < 0.001$, *n.s.* if $p \geq 0.05$. **Bottom:** Stacked histogram (left) of controllability variance explained by orientation after removing time. Pie charts (right) indicate fraction of CNs per subject that display significant controllability differences given orientation after factoring out block order (ANOVA, $p < 0.05$).

4.2 Controllability is to a Lesser Extent Influenced by Location

For our second manipulation, we sought to evaluate changes in controllability given by the overall location of the movement axis across the workspace. To do this, we trained 2 subjects (Monkey N, $n = 20$ neurons; Monkey R, $n = 3$ neurons) on the FAST location paradigm, in which we fixed axis orientation and instead varied the coordinates of the axis origin within the display (Figure 4.7).

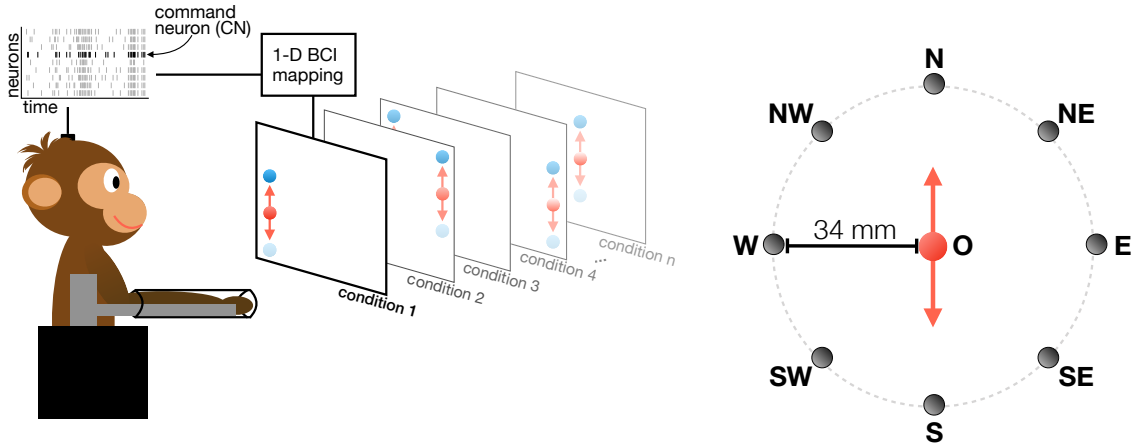

Figure 4.7 | The FAST location experiments assess the change in controllability when altering the 2D position of the axis origin. Each condition is defined by the location of the movement axis within the workspace and tested at a single axis orientation at a time (left). As an example, axis locations are shown for a 90° orientation (right).

Specifically, on each session we tested a series of locations at two distinct axis orientations. Conditions were thus defined as a location-orientation combination. Monkey N was tested across 10 location-orientation conditions, and was given two non-consecutive blocks at each. Conversely, Monkey R was tested across 18 location-orientation conditions but received only one block at each. All blocks in this paradigm had a duration of 4 min.

For Monkey N, we tested 5 axis locations: the center of the display (O) and four coordinates in the corners of the workspace, all of them at 34 mm from the center (Northwest, NW; Southwest, SW; Southeast, SE; Northeast, NE). At the start of each experiment, Monkey N

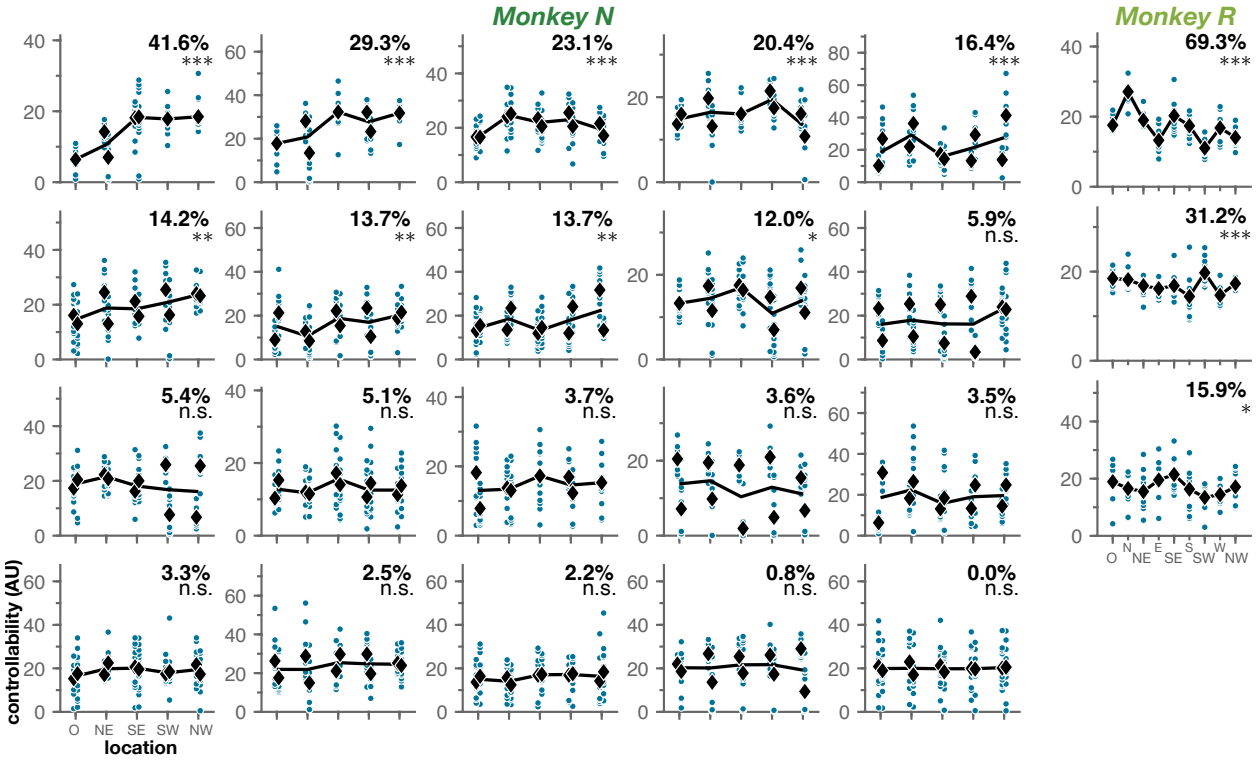

Figure 4.8 | Movement location can also interact with volitional control. Controllability of CNs across locations for Monkey N ($n = 20$ neurons) and Monkey R ($n = 3$ neurons). Controllability values (teal) and block averages (black diamonds) are shown; curves connect average controllability per condition. Neurons are sorted by descending variance explained by location, reported per CN along with significance (ANOVA): *** if $p < 0.001$, ** if $p < 0.01$, * if $p < 0.05$, *n.s.* if $p \geq 0.05$.

was given four arbitrary orientations at the central location (O). Of those, we selected two orientations (not opposites, separated by at least 90°) in which the subject performed well. Once selected, we first ran each orientation randomly intermixed across the 4 peripheral locations, and then ran the two orientations across all 5 locations in random order again. The presentation order was constrained so that both the orientation and the location changed across sequential blocks. For Monkey R, in addition to the locations described above, we tested 4 new locations following the cardinal directions (North, N; West, W; South, S; East, E), for a total of 9 locations—one in the center of the display and the others 34 mm away from it. In these experiments, Monkey R was first given one block at each of the 9 locations, randomly sorted, on one fixed axis orientation. After a full run through all locations

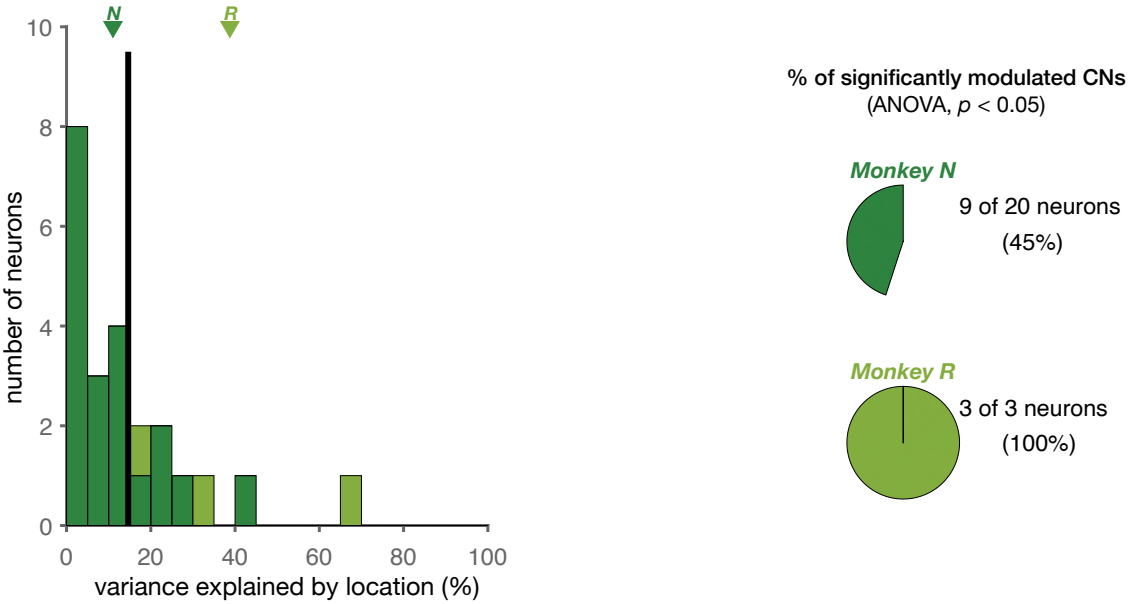

Figure 4.9 | Controllability was modulated by location more moderately than by orientation. Stacked histogram of the variance in controllability explained by location. Pie charts indicate fraction of CNs per subject that display significant controllability differences given location (ANOVA, $p < 0.05$).

given one angle, the axis was rotated to a second orientation, in which the subject was given another set of randomly sorted blocks across the 9 locations.

Both animals showed differences in controllability due to axis location (Figure 4.8). Yet, in contrast to orientation, overall a lower fraction of the CNs displayed a significant impact due to location (one-way ANOVA, $p < 0.05$). Regarding the magnitude of the effects, location had a modest effect on controllability compared to orientation (Figure 4.9), accounting for 14.6% of variance on average (Monkey N: 11.0%, Monkey R: 38.8%).

To control for any influence on the measured location dependence due to the rotation of the movement axis, we tested controllability across locations on a secondary orientation. We found that the impact of location on controllability was comparable across both orientations (Figure 4.10). Interestingly, within FAST location experiments location seemed to capture a smaller portion of the variance found in controllability as compared to the fraction of variability explained by orientation (Figure 4.11), although this difference was not significant

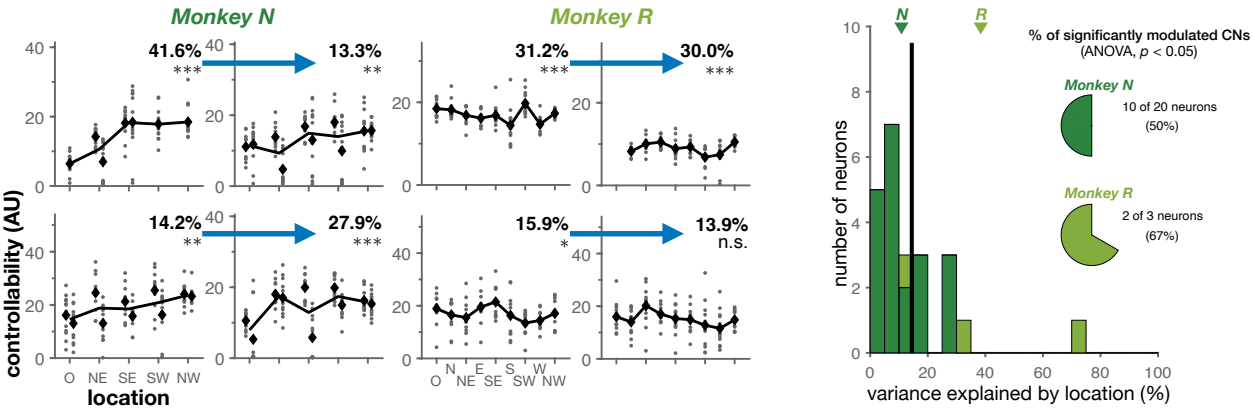

Figure 4.10 | Orientation changes do not erase the impact of location on controllability. Left: Controllability across locations given the primary (left of arrow) or secondary orientation (right of arrow) for 2 example CNs per subject. Controllability values (gray) and block averages (black diamonds) are shown; curves connect average controllability per condition. Variance explained by location is reported, along with significance (ANOVA): *** if $p < 0.001$, ** if $p < 0.01$, * if $p < 0.05$, *n.s.* if $p \geq 0.05$. **Right:** Stacked histogram of controllability variance explained by location at a secondary orientation. Pie charts indicate fraction of CNs per subject that display significant controllability differences given location (ANOVA, $p < 0.05$).

(paired-sample t-test, $p \geq 0.05$). Expanding this data set with additional experiments that test the controllability of a larger number of neurons could increase the statistical power of this comparison and shed light on the relationship between orientation and location.

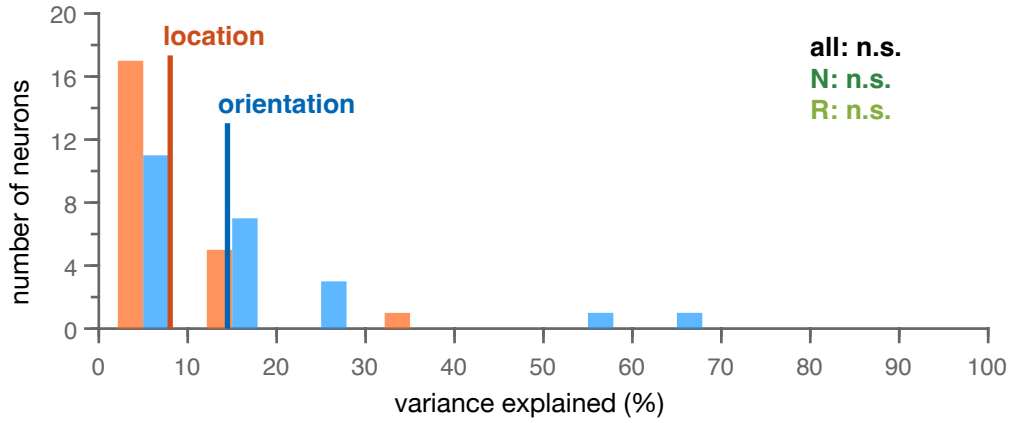

Figure 4.11 | Movement location may explain, on the same experiment, a lesser portion of controllability variance than orientation. Histogram of the variance in controllability present in FAST location experiments that is accounted for by either the location or the orientation label, irrespective of one another. Vertical lines indicate average variance explained by each label. Statistical testing for significant differences in variance explained by location versus orientation is reported (paired-sample t-test: *n.s.*, $p \geq 0.05$).

4.3 Preliminary Results on Posture-Orientation Interactions

With this last sensory manipulation, our goal was to examine whether changing the subject's resting posture could influence single-neuron controllability and its relationship with movement orientation. The FAST posture paradigm therefore constitutes an expansion of the FAST orientation experiments described above: here we incorporate postural manipulations to evaluate controllability across orientations given these changes in somatosensation.

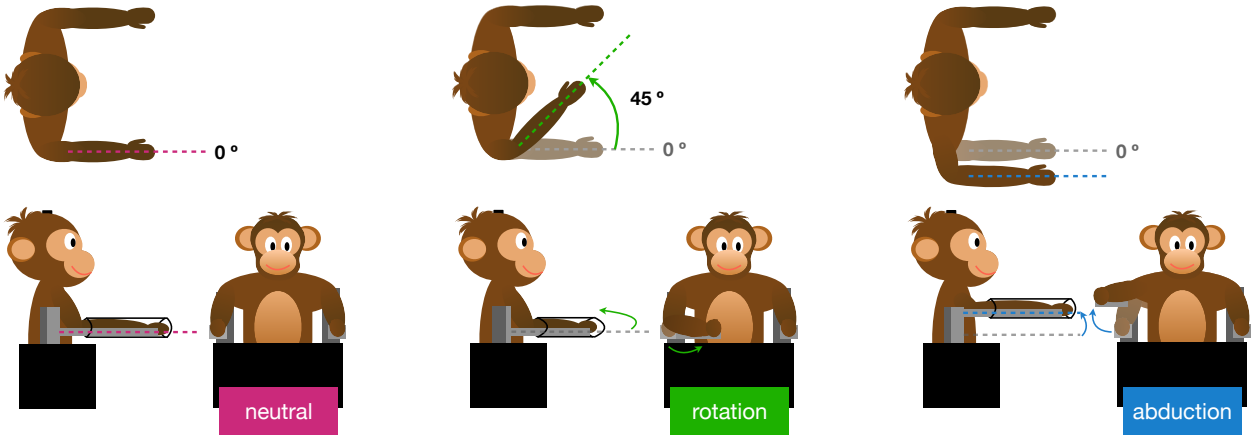

Figure 4.12 | Postural manipulations were designed to create maximal differences in shoulder posture. Three arm postures were evaluated: neutral (left), rotation through elbow flexion (center) and elevation through shoulder abduction (right). These postural manipulations are achieved through the configuration of an armrest and are hence passive: subjects do not have to actively engage their muscles to maintain each posture.

The three possible arm positions were designed to create a wide difference in shoulder posture while allowing the subject to be comfortable throughout the session (Figure 4.12). Compared with the neutral posture, the elbow extension posture resulted in an inward rotation of the forearm, while the shoulder abduction posture resulted in an elevation of the forearm, approaching shoulder height. Critically, these postural manipulations were passive; the animal's forearm was placed on an armrest at one of 3 configurations, where it was lightly restrained in position. Subjects were therefore able to relax their muscles, as opposed to being required to actively maintain a certain arm posture.

In the case of Monkey N ($n = 19$ neurons), every CN was tested across 3 different session, each characterized by a distinct posture. In other words, a FAST posture experiment required 3 days, since a single armrest configuration was used for the entirety of an individual session. Once the arm had been positioned in the desired posture for a given day, the paradigm's structure was comparable to a FAST orientation session: 8 orientation conditions were each presented twice in a non-consecutive order.

Conversely, Monkey R ($n = 12$ neurons) was exposed to all 3 postures on the same day; this, however, came at a cost. Since there is a hard limit on the length of time that a monkey can be expected to work on a behavioral task, each orientation was presented only once per arm posture, for a total of 24 orientation-posture conditions each given one 4-min block. While this mitigates the confounding effect of day-to-day variability inherent in Monkey N's version of the FAST posture paradigm, it also introduces a different confound. Since switching between arm configurations constitutes a small disruption of the task while the subject's arm is manually repositioned, it is not feasible for posture to vary from block to block like movement orientation can. This, combined with the inability to offer a second block at each orientation-posture condition, results in the posture label being inextricably linked to the passage of time within the session. To bypass this issue, Monkey R also received 4 sessions (marked by a pound sign on Figure 4.13 and Figure 4.14) in which block duration was reduced to 1 min. These shortened blocks made it possible to accommodate a second sweep across orientations at each posture, even after every combined condition had already been presented once.

The present FAST posture data set serves here as a preliminary exploration of the interaction between orientation and posture and their joint influence on single-neuron controllability. Despite day-to-day and within-session variability, controllability appeared to depend on orientation similarly at all three postures for a majority of neurons across both animals (Figure 4.13).

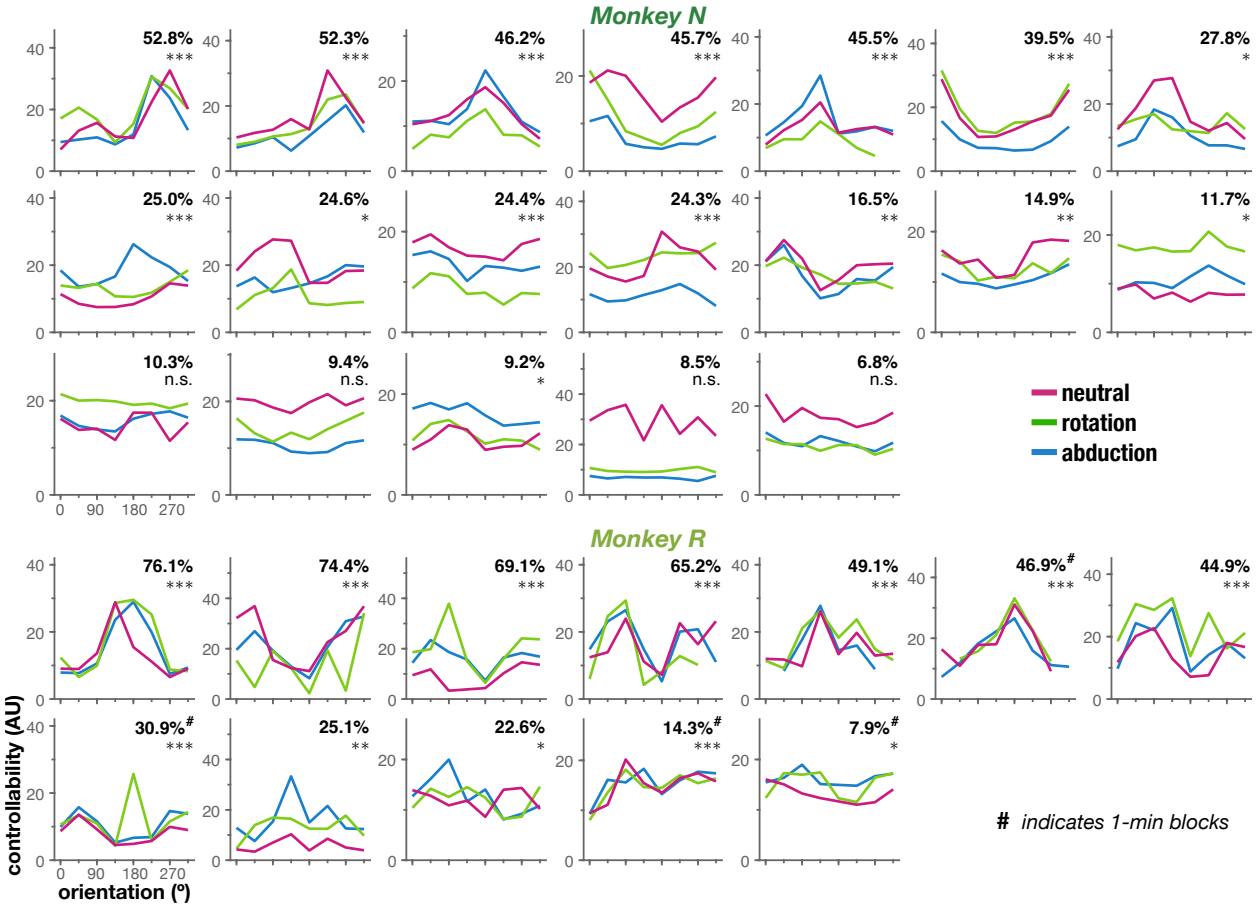

Figure 4.13 | Orientation effects on controllability seem largely conserved across postures. Controllability of CNs on each individual posture across orientations for Monkey N ($n = 19$ neurons) and Monkey R ($n = 12$ neurons). For each neuron, curves connect same-posture average controllability values across blocks. Neurons are sorted by descending average variance explained by orientation given posture, reported per CN along with the lowest significance among the three postures (ANOVA): *** if $p < 0.001$, ** if $p < 0.01$, * if $p < 0.05$, *n.s.* if $p \geq 0.05$. Monkey R's 1-min-block sessions are indicated by a pound sign (#).

The change in controllability as a function of posture is however difficult to interpret. In the case of Monkey R, who was exposed to all orientation-posture conditions within a single session, there seems to be a minor effect of posture overall, with significant effects appearing only for a small portion of neurons (Figure 4.14). For Monkey N, it is impossible to draw any further conclusions due to posture being inseparable from day-specific factors in these conditions.

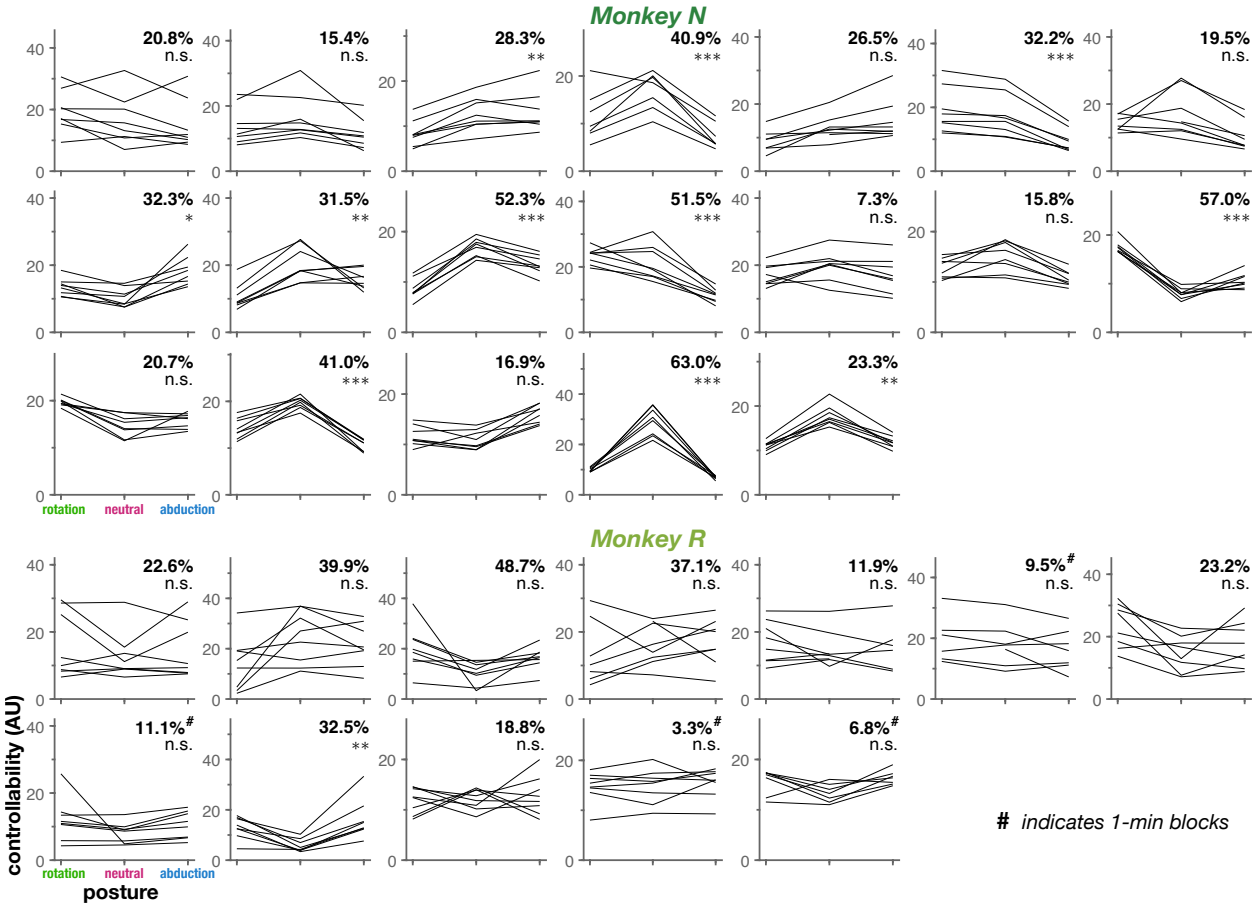

Figure 4.14 | Posture and day-to-day variability are inseparable in this implementation of the FAST paradigm. Controllability of CNs on each individual orientation across postures for Monkey N ($n = 19$ neurons) and Monkey R ($n = 12$ neurons). For each neuron, curves connect same-orientation average controllability values across blocks. Neurons are sorted by descending average variance explained by posture given orientation, reported for each CN. Stars indicate ANOVA significance level of the least significant posture effect among all orientations: *** if $p < 0.001$, ** if $p < 0.01$, * if $p < 0.05$, *n.s.* if $p \geq 0.05$. Monkey R's 1-min-block sessions are indicated by a pound sign (#).

Whether postural changes can modify the gain of these relationships is therefore a question that remains unanswered here, given the limitations of our initial experimental design. Nevertheless, the introduction of further controls in future implementations of the FAST posture paradigm may be able to shed light on the interplay between visual and somatosensory signals and their effect on voluntary modulation of M1.

5

Persistence of Sensory Effects on Controllability Through Time

The observations detailed until now do not rule out the possibility that, even though the neural requirements of the task were the same across sensory contexts, the subject may not have been aware of this. Perhaps the struggle to control the CN in certain conditions was in part due to the animal not realizing that the neural task remained unchanged.

We reasoned that, if this were the case, then providing subjects with additional training—whether focused at more troublesome contexts or generalized across all feedback conditions—would result in improved modulation control and a reduced interaction between sensory context and single-neuron controllability. In order to test this hypothesis, we designed two distinct experiments that leveraged the flexibility of the FAST framework to probe the efficacy of different training strategies.

5.1 Orientation Dependence Remains Despite Within-Day Focused Training

First, we introduced additional practice for a single condition in the middle of what would otherwise have been a standard FAST orientation experiment. This additional practice took place after every movement angle had been presented to the subject once, and consisted of two 15-min blocks in which the animal could gain more experience controlling the cursor under *one* specific condition.

By comparing controllability curves before and after those 30 min of training focused on a single orientation, we can evaluate whether extended practice in an individual sensory context results in better volitional control of the neuron and less interaction between volitional control and axis orientation.

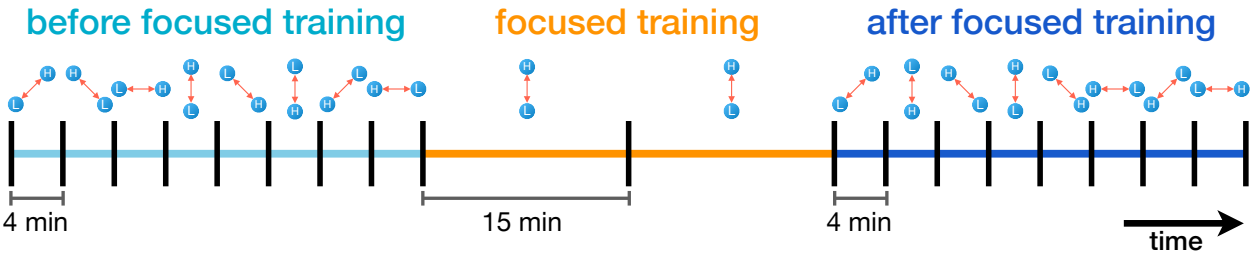

Figure 5.1 | The FAST focused training paradigm provides additional practice time in challenging conditions. An initial evaluation of controllability across all 8 conditions is followed by a extended period of training in a single orientation. Controllability across all sensory contexts is then assessed once again.

To do this, the focused training paradigm tested controllability across orientations twice: once before and once after the *focus* blocks (Figure 5.1). Similarly to the FAST orientation paradigm, each focused training experiment started with a single 4-min block at each of the 8 orientations. This served to establish a baseline level of controllability across all conditions. We then selected one of the orientations to be the focus of the additional training; typically, this was a condition in which the subject demonstrated an average performance in the first sweep across orientations. The focused training period started with one long block at

this orientation, lasting 15 min. As before, the difficulty ratcheting was active throughout this focus block, such that the firing rate targets would gradually separate if performance was strong. At the completion of this 15-min block, a new 15-min focus block at that same orientation started, which served to reset the difficulty level and make the task slightly easier again, renewing the monkey's motivation. Once this second 15-min block was completed, the

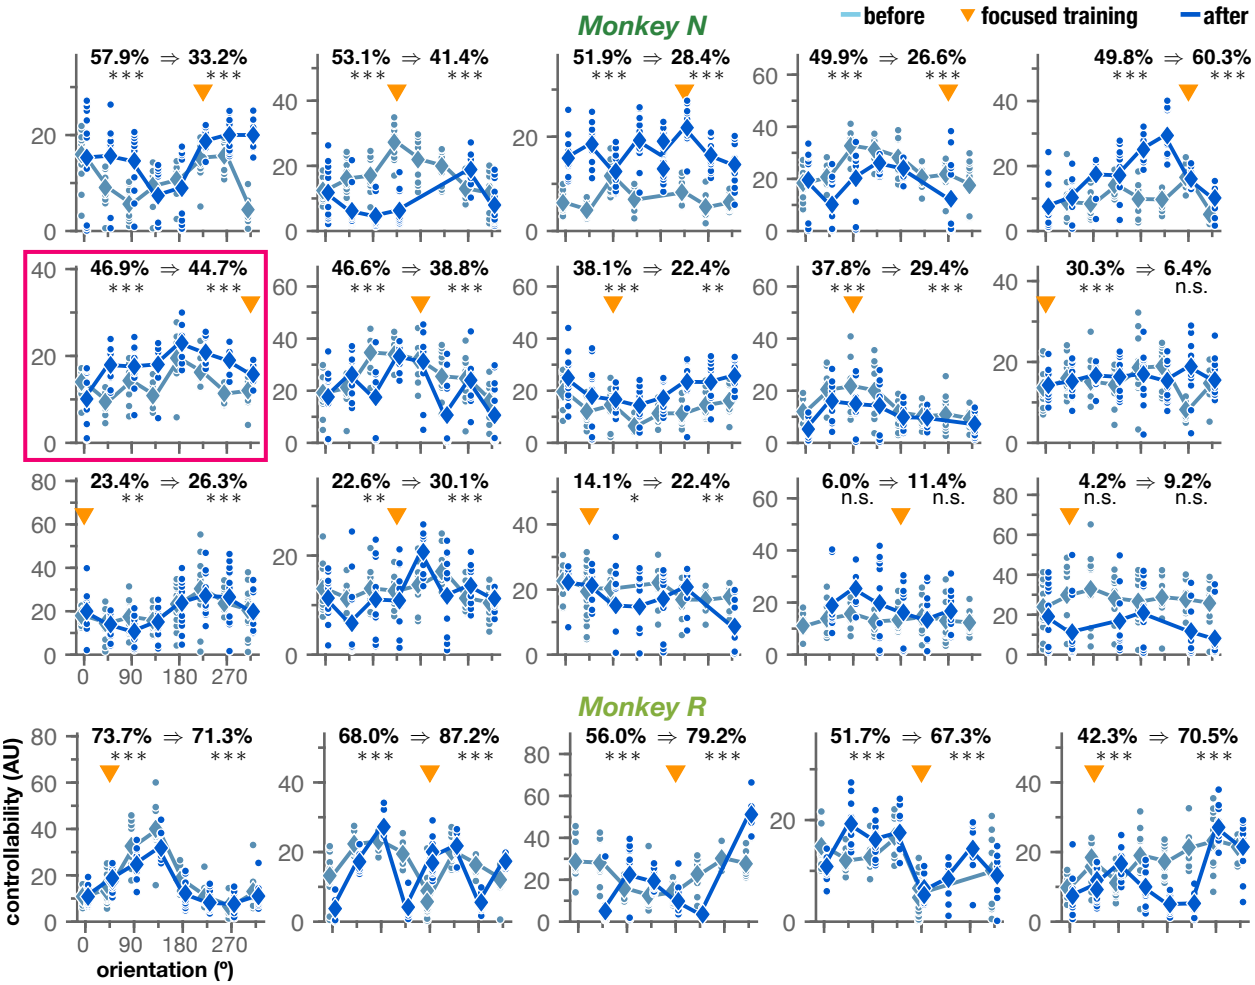

Figure 5.2 | Focused training on a "difficult" sensory context is not sufficient to extinguish orientation dependence. Controllability of CNs across orientations before and after focused training for Monkey N ($n = 15$ neurons), and Monkey R ($n = 5$ neurons). Controllability values (gray) and block averages (diamonds) are shown; curves connect average controllability per orientation before and after focus blocks. Variance explained by orientation is given as *before* ⇒ *after*, along with significance (ANOVA): *** if $p < 0.001$, ** if $p < 0.01$, * if $p < 0.05$, *n.s.* if $p \geq 0.05$. Neurons are sorted by descending variance explained before focused training.

subject was presented with a second set of 4-min blocks across all 8 orientation conditions.

Let us take a closer look at the example marked by a box in Figure 5.2: this CN displayed below-average controllability at the 315º condition during the baseline block. After the subject was given the two 15-min focus blocks at 315º to refine control in this sensory context, there was a visible improvement in controllability when the same orientation was presented again. However, this effect was not restricted to the trained condition exclusively: this prolonged training in the 315º orientation resulted in an overall increase in controllability across all sensory contexts.

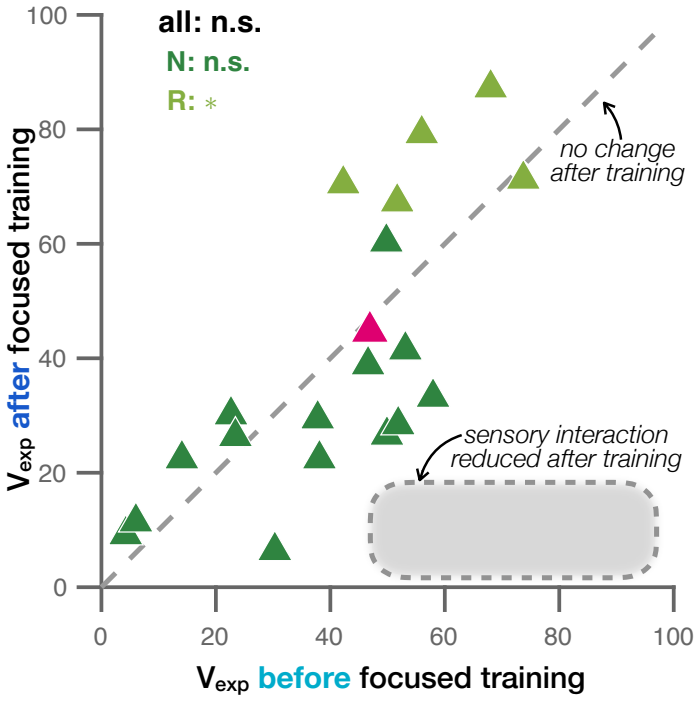

Figure 5.3 | Controllability showed no significant decrease in its sensitivity to orientation after focused training. Controllability variance explained by orientation after focused training is plotted against baseline before receiving the additional practice (dashed unity line shown for reference). Statistical significance is reported for comparisons within and across subjects (paired-sample t-test: * if $p < 0.05$, *n.s.* if $p \geq 0.05$). Overall, controllability did not change significantly after focused training ($p \geq 0.05$).

We found that prolonging practice in a condition that had proved challenging earlier in the session was not effective in neutralizing orientation dependence (Figure 5.2). Neither

subject showed large differences in controllability before versus after the focus blocks.

There was no overall significant change in variance explained by orientation before versus after the extended practice given during the focus blocks (Figure 5.3); in fact, in one monkey additional training actually *increased* the dependence of controllability on orientation.

5.2 Multiple-Day Training is Not Sufficient to Overcome Sensory Effects

To further evaluate the robustness of the orientation effect, we expanded our study of single-neuron controllability across multiple days. Could additional training spanning several sessions eliminate the interaction between sensory context and controllability?

The FAST multi-day training paradigm provided the subject with 5 days of practice at the orientation task to control a single CN (Figure 5.4). By extending training in all orientations over several days, subjects were allowed multiple sessions to control the same neuron and refine their ability to modulate its activity across conditions.

In a FAST multi-day training experiment, day 1 was identical to a typical FAST orientation session. We then held both the CN and the initial target rate values constant for days 2 through 5 and had the subject perform an equivalent FAST orientation session each day using those fixed parameters. In other words, aside from the order in which orientations were presented—which was randomized at the beginning of each day—the task was identical across all 5 days of the experiment. The CN was tracked across days by visual inspection, and confirmed offline through a neuron-tracking algorithm (Fraser and Schwartz, 2012).

We found that training for several days also did not eradicate the interaction of controllability with orientation (Figure 5.4). This can be appreciated in the boxed example: while average controllability increased across days, the interaction with orientation continued to be present throughout these sessions. The same example is highlighted in Figure 5.5.

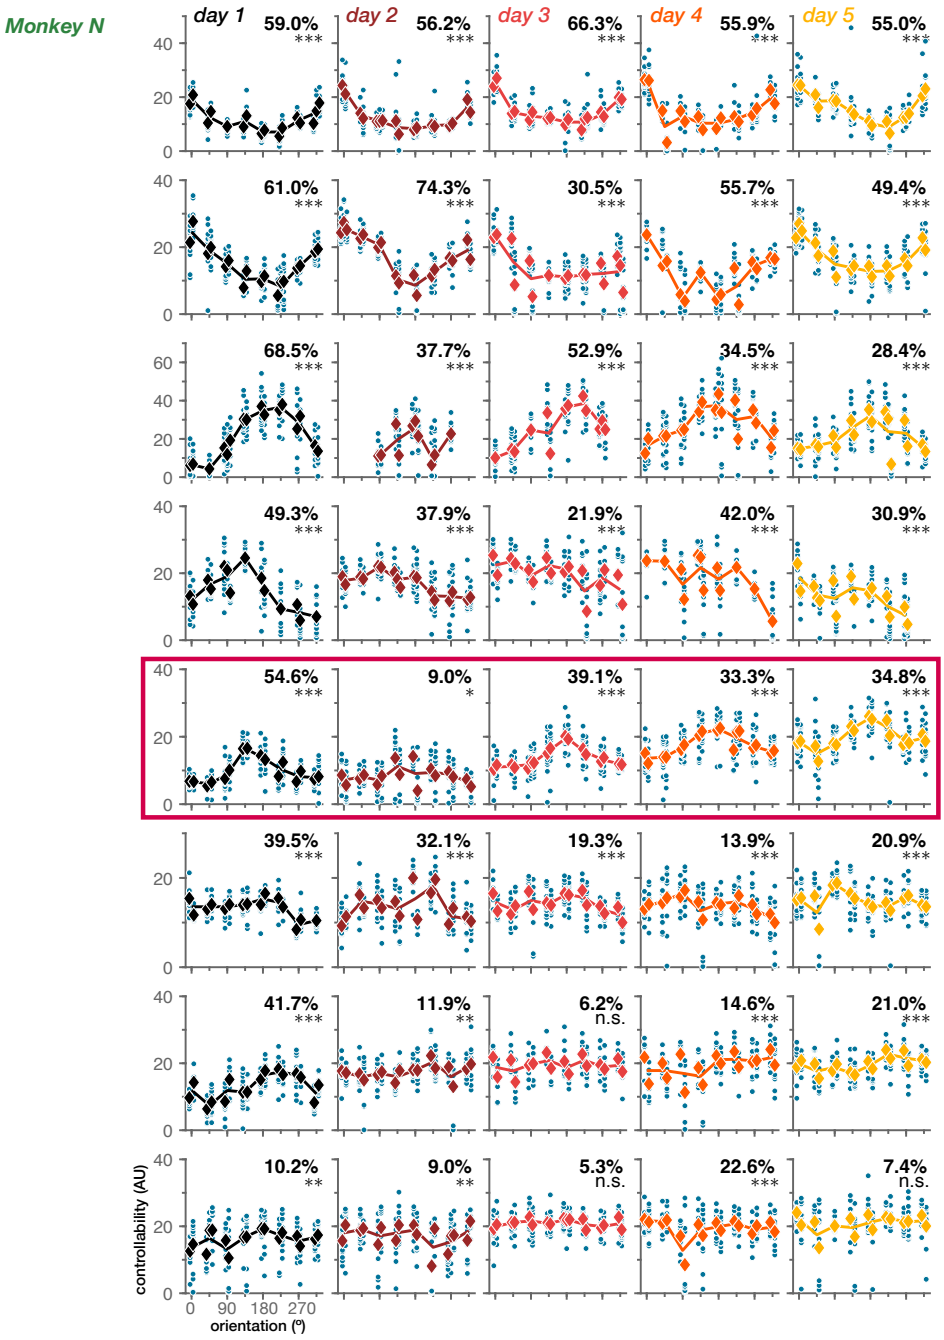

Figure 5.4 | Training over several days was not sufficient to extinguish orientation dependence. Controllability of CNs across days for Monkey N ($n = 8$ neurons). Each row represents a single CN, used to control the cursor on 5 sequential days of training. Average controllability per block and condition-average curves are color-coded based on training day. Variance explained by orientation is reported per day, along with significance (ANOVA): *** if $p < 0.001$, ** if $p < 0.01$, *n.s.* if $p \geq 0.05$. Neurons are sorted by descending variance explained, averaged across days.

Across neurons, multi-day training did not erase sensory interactions. While we observed a slight weakening of the orientation effect during the first 2 days of repeated training with a given CN, additional training in the following days did not further attenuate the impact of orientation, and the interaction between orientation and controllability never disappeared (Figure 5.5).

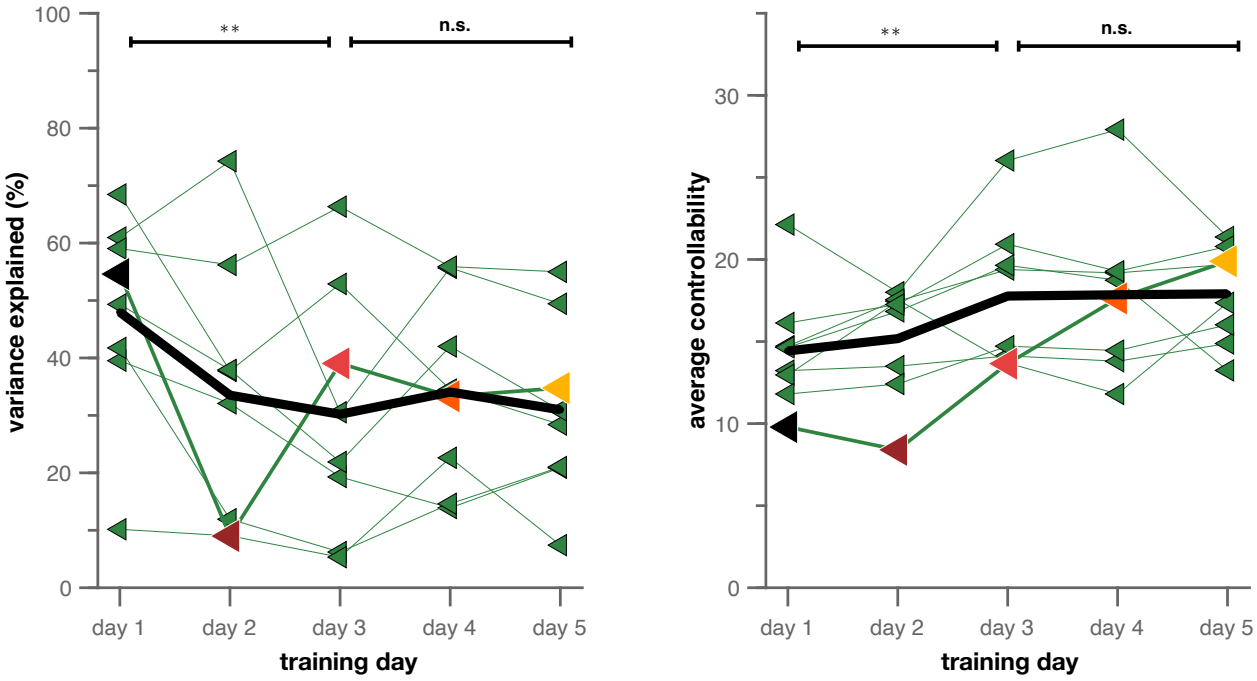

Figure 5.5 | The reduction in orientation dependence slowed down throughout multi-day training. Changes in variance explained (left) and average controllability (right) across days. Black curve represents average trend. Significance is reported for day 1 to day 3 comparisons and day 3 to day 5 comparisons (paired-sample t-test: ** if $p < 0.01$, *n.s.* if $p \geq 0.05$).

6

Repercussion of Sensory Constraints on Voluntary Neural Activity

This thesis investigates whether the ability to volitionally modulate the activity of single neurons in the primary motor cortex can be affected by changes in sensory context, that is, by the manner in which the feedback of the cell's firing rate is provided to the subject. We found that sensory context can indeed impose constraints on the degree to which individual neurons can be volitionally controlled.

Across animals, controllability was heavily impacted by the orientation of cursor movement: subjects showed consistent differences in success when attempting to modulate individual neurons at various movement angles. Movement location within the workspace had an influence on single-neuron controllability as well; location effects were, however, less pronounced than those of orientation. Finally, we found that the interaction between sensory context and volitional controllability was robust to additional training. Providing animals with extended practice within the same session at an orientation that they had found problematic earlier was not sufficient to overcome these differences in controllability, and neither was granting a subject multiple days to improve control of a single neuron.

6.1 Implications for Natural and Neuroprosthetic Motor Learning

How might sensory constraints fit in with learning? We observed that even several days of additional training did not lead to a sizable decrease in sensory interaction. Of course, it could be the case that the animals *could* have learned the task but never did.

Previous research has identified key factors that govern whether a BCI mapping should be learned quickly or slowly. Typically, a BCI mapping that only requires existing patterns of neural activity to be mapped to new movements are learned fast, showing substantial improvements on a single day of training (Sadtler et al., 2014; Golub et al., 2018; Hennig et al., 2018). In contrast, BCI mappings that require the formation of new neural activity patterns can be harder to learn and may necessitate incremental training across several days in order to facilitate the expansion of the neural repertoire (Oby et al., 2019). It is also possible to learn a perturbed BCI mapping that disrupts only a subset of the neurons involved in cursor control, which has been associated with a fast-and-global neural activity change in the population, combined with a gradual-and-local adjustment to the relative contributions of individual neurons predominantly in the perturbed subset (Jarosiewicz et al., 2008; Zhou et al., 2019).

In the FAST paradigm, trial success is determined exclusively by the firing rate of a single neuron and imposes no demands on the rest of the population: the subject is free to take advantage of any neural activity pattern at any time to carry out the task. In principle, changing sensory context should therefore be quite readily learned, since it is by definition well within the existing neural repertoire (Sadtler et al., 2014; Golub et al., 2018). Instead, we find no improvement within one session, and hardly any after multiple days. This interplay between sensory and volitional signals seems to be a novel, added constraint that might limit learning.

The notion that sensory context can impose bounds on our ability to voluntarily modulate M1 can have direct implications for the development of clinical neuroprosthetic devices. If the activity of individual neurons can be tied to particular sensory outcomes, these relationships could have an effect on how readily a BCI mapping can be learned, analogous to the involvement of existing versus novel neural activity patterns we discussed above. In other words, building decoders that are consistent with the sensory interactions in the population could constitute an additional factor in dictating how learnable a BCI mapping can be. Further study of this sensorimotor crossroads in M1 could enable decoders to take advantage of native neural pathways and dodge configurations that prove detrimental for BCI control.

Overall, our results suggest that volitional modulation of M1 can be conditional to certain constraints exacted by sensory integration in the motor cortex. This would be consistent with a model of cortex that contemplates sensory and volitional signals impinging on the same neurons, therefore resulting in a mixed response at the level of individual cells—within M1 and beyond. Such a framework would be consistent with the observation of volitional modulation across diverse areas of the brain, including primary sensory areas (Cerf et al., 2010; Koralek et al., 2012; Clancy et al., 2014; Neely et al., 2018; Chauvière and Singer, 2019; Andersen et al., 2019; Patel et al., 2021; Gallego et al., 2022; Fukuma et al., 2022; Jeon et al., 2022), as well as explaining sensory-driven activity in motor areas.

If both volitional and sensory signals coexist across the entire cortex, perhaps the key difference that distinguishes sensory and motor areas is precisely the degree to which they are under volitional control versus sensory-obligate modulation.

This possibility merits further exploration.

6.2 Follow-Up Studies

6.2.1 Population-Level Contributions to Single-Neuron Controllability

The results described in this thesis concern the command neuron and the behavioral performance in the task. However, the neural data sets collected during the experiments detailed throughout incorporate high-dimensional neural recordings. In other words, the use of microelectrode arrays allows us to simultaneously record the activity of a neural population of around 100 cells in addition to the command neuron. While beyond the scope of this thesis, it would be of interest to conduct population-level analyses that explore the role of those neurons not directly involved in the behavioral task.

Specifically, studying the dimensionality of the neural subspace occupied when subjects perform this task as compared to hand- and brain-controlled center-out paradigms may provide insight into the neural strategies employed to gain control of cursor movement and what their weaknesses may be. Indeed, identifying their limitations may explain the interaction between controllability and sensory context. It should also be interesting to dissect how neural states may vary as a function of other parameters, such as the escalation of task difficulty, the presence of compounded manipulations—those combining orientation and location changes or orientation and posture changes—and the continuation of training within and across days.

Interestingly, some studies have found that population dynamics in the neural space (for a review on latent dynamics, see Vyas et al., 2020) can change across contexts (Dubreuil et al., 2022). It is possible that the altered interaction between sensory feedback and volitional modulation might explain the existence of different flow fields across contexts.

6.2.2 Deeper Assessment of Location Effects on Controllability

Movement location seemed to weigh less heavily on controllability. This could be partly attributable to the magnitude of the displacement across location conditions, which was limited in our paradigm to accommodate all axis orientations and still guarantee that cursor movement remained within the confines of the display. The lesser impact of movement location we observed is, however, consistent with previous reports that cursor position in a BCI setting had only minor correlates in motor cortical activity and did not interfere with performance during neuroprosthetic control (Stavisky, Rezaii, et al., 2018).

Nevertheless, it would be of interest to design an alternative version of the FAST location paradigm in which the spatial separation of the different conditions was expanded. This could be achieved by narrowing the focus of the experiment to a single movement orientation: reducing the scope of orientations evaluated would create room to push the peripheral locations outwards within the workspace without compromising the amplitude of cursor movement available.

Moreover, this single-orientation design could incorporate the strengths of both versions of the FAST location experiments described in this work. One version of the task, given to Monkey N, provided the subject with repeated presentations of each individual condition, which facilitated the elimination of critical confounds like time and motivation. The other version of the task, given to Monkey R, forfeited the repetition of each condition in order to accommodate a finer mesh of locations, almost doubling the number of conditions in which controllability could be tested. By keeping axis angle constant across all conditions, this location-only FAST paradigm would therefore achieve a more thorough examination of movement location effects on controllability by combining larger spatial manipulations, finer lattice assessment and multiple instances of each condition.

6.2.3 Other Factors Potentially Imposing Additional Constraints

In the work detailed throughout this thesis, we manipulated feedback primarily within the bounds of a single sensory modality: vision. Our results here therefore invite one to consider what other forms of sensory information might influence volitional modulation in the motor cortex.

Postural changes, for instance, have been shown to sway directional tuning of individual neurons in M1 (Scott and Kalaska, 1995, 1997; Sergio and Kalaska, 2003; Velliste et al., 2014); it is likely that the impact of posture could extend to controllability as well. Our preliminary experiments combining postural and visual manipulations proved to be too noisy to confirm or discard this possibility. While it was impossible to ascertain whether the gain differences observed were due in part to posture changes, it would be interesting to more closely examine the influence of somatosensory feedback on single-neuron controllability. Further iterations of the FAST posture paradigm that incorporate controls for day-to-day neural variability could provide valuable insight in this respect.

Not far away from the vestibular system, auditory signals have also been shown to evoke activity in the motor cortex. More specifically, human participant studies (Wilson et al., 2004; Cheung et al., 2016) have found that speech sounds can elicit activity in motor areas involved in speech production—much like mirror neurons respond to observed, as well as executed, movement (Tkach et al., 2007). Strikingly, a group led by Shenoy and Henderson reported that speech production could drive neural activity in the arm and hand area of the dorsal motor cortex enough that the spoken words could be accurately decoded (Stavisky, Kao, Nuyujukian, et al., 2018; Stavisky et al., 2019). Although their subsequent study suggested that speech production did not interfere with BCI cursor control in the case of a population decoder (Stavisky et al., 2020), it is unclear whether this might hold true at the level of individual neurons.

Individual neurons in M1 have also been observed to change their firing as a function of

object-related information during reach-to-grasp naturalistic movement (Rouse and Schieber, 2016). Similarly, whether a reach involved an object or not led to differences in M1 activity that disrupted BCI control of a robotic arm if not explicitly accounted for in decoder design (Downey et al., 2017).

Alongside these, other variables not directly linked to sensory stimuli could potentially influence single-neuron controllability as well. Fluctuations in internal states—from attention to emotional condition—have been observed to modulate ongoing cognitive processes like decision-making, motor learning or sensory perception (Accolla and Carleton, 2008; Kennedy, 2011; Ruff and Cohen, 2019; Steinmetz et al., 2019; Cowley et al., 2020; Hennig et al., 2021; Smoulder et al., 2021), and can be detrimental for BCI decoding performance (Perge et al., 2013; Dunlap et al., 2020). Further examination would be required to determine whether these factors could also be limiting our ability to volitionally modulate individual neurons in M1.

6.2.4 Sensory-Driven Modulation During Movement Observation

Passively observing familiar movements has been shown to be able to trigger activity in motor cortical areas as well as through the execution of the action (Gallese et al., 1996; Hari et al., 1998; Rizzolatti and Craighero, 2004; Cisek and Kalaska, 2004). This includes the primary motor cortex, where modulation and directional tuning at the level of individual neurons does not change between observation and execution (Tkach et al., 2007).

Could we predict a neuron's controllability at different sensory conditions based on its observation-evoked modulation? To answer this, one could present subjects with a passive observation task that resembles the FAST paradigm and record from the neural population. Comparing a neuron's response in this task to the same neuron's controllability across sensory contexts could shed light on the coexistence of sensory and volitional signals within the primary motor cortex.

6.2.5 Response Dissociation in a Cursor Jump Perturbation

A cursor jump perturbation task could be a powerful tool to study whether neurons show a distinctive volitional or reflexive preference. This kind of paradigm can dissociate volitional, slow responses from reflexive, faster responses within M1 (Cross et al., 2019). In combination with FAST experiments, a cursor jump perturbation task could allow us to assess whether that potential preference between volitional and reflexive responses may correlate with single-neuron controllability across the population, as well as with the degree of adaptation in preferred direction in the context of a visuomotor rotation.

Bibliography

Accolla, R., & Carleton, A. (2008). Internal body state influences topographical plasticity of sensory representations in the rat gustatory cortex. *Proceedings of the National Academy of Sciences*, *105*(10), 4010–4015. https://doi.org/10.1073/pnas.0708927105

Andersen, R. A., Aflalo, T., & Kellis, S. (2019). From thought to action: The brain–machine interface in posterior parietal cortex. *Proceedings of the National Academy of Sciences*, *116*(52), 26274–26279. https://doi.org/10.1073/pnas.1902276116

Bedingham, W., & Tatton, W. (1984). Dependence of EMG Responses Evoked by Imposed Wrist Displacements on Pre-existing Activity in the Stretched Muscles. *Canadian Journal of Neurological Sciences / Journal Canadien des Sciences Neurologiques*, *11*(2), 272–280. https://doi.org/10.1017/S0317167100045534

Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., Patil, P. G., Henriquez, C. S., & Nicolelis, M. A. L. (2003). Learning to Control a Brain–Machine Interface for Reaching and Grasping by Primates. *PLoS Biology*, *1*(2), e42. https://doi.org/10.1371/journal.pbio.0000042

Cerf, M., Thiruvengadam, N., Mormann, F., Kraskov, A., Quiroga, R. Q., Koch, C., & Fried, I. (2010). On-line, voluntary control of human temporal lobe neurons. *Nature*, *467*(7319), 1104–1108. https://doi.org/10.1038/nature09510

Chase, S. M., Kass, R. E., & Schwartz, A. B. (2012). Behavioral and neural correlates of visuomotor adaptation observed through a brain-computer interface in primary

motor cortex. *Journal of Neurophysiology*, *108*(2), 624–644. http://jn.physiology.org/content/108/2/624.abstract

Chase, S. M., Schwartz, A. B., & Kass, R. E. (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*(9), 1203–1213. https://doi.org/10.1016/j.neunet.2009.05.005

Chauvière, L., & Singer, W. (2019). Neurofeedback Training of Gamma Oscillations in Monkey Primary Visual Cortex. *Cerebral Cortex*, *29*(11), 4785–4802. https://doi.org/10.1093/cercor/bhz013

Cheney, P. D., & Fetz, E. E. (1984). Corticomotoneuronal cells contribute to long-latency stretch reflexes in the rhesus monkey. *The Journal of Physiology*, *349*(1), 249–272. https://doi.org/10.1113/jphysiol.1984.sp015155

Cheung, C., Hamilton, L. S., Johnson, K., & Chang, E. F. (2016). The auditory representation of speech sounds in human motor cortex. *eLife*, *5*. https://doi.org/10.7554/eLife.12577

Cisek, P., & Kalaska, J. F. (2004). Neural correlates of mental rehearsal in dorsal premotor cortex. *Nature*, *431*(7011), 993–996. https://doi.org/10.1038/nature03005

Clancy, K. B., Koralek, A. C., Costa, R. M., Feldman, D. E., & Carmena, J. M. (2014). Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat Neurosci*, *17*(6), 807–809. http://dx.doi.org/10.1038/nn.3712

Colebatch, J. G., Gandevia, S. C., McCloskey, D. I., & Potter, E. K. (1979). Subject instruction and long latency reflex responses to muscle stretch. *The Journal of Physiology*, *292*(1), 527–534. https://doi.org/10.1113/jphysiol.1979.sp012869

Cowley, B. R., Snyder, A. C., Acar, K., Williamson, R. C., Yu, B. M., & Smith, M. A. (2020). Slow Drift of Neural Activity as a Signature of Impulsivity in Macaque Visual and Prefrontal Cortex. *Neuron*, *108*(3), 551–567. https://doi.org/10.1016/j.neuron.2020.07.021

Crago, P. E., Houk, J. C., & Hasan, Z. (1976). Regulatory actions of human stretch reflex. *Journal of Neurophysiology*, *39*(5), 925–935. https://doi.org/10.1152/jn.1976.39.5.925

Cross, K. P., Cluff, T., Takei, T., & Scott, S. H. (2019). Visual Feedback Processing of the Limb Involves Two Distinct Phases. *The Journal of Neuroscience*, *39*(34), 6751–6765. https://doi.org/10.1523/JNEUROSCI.3112-18.2019

Davidson, A. G., Chan, V., O'Dell, R., & Schieber, M. H. (2007). Rapid changes in throughput from single motor cortex neurons to muscle activity. *Science*. https://doi.org/10.1126/science.1149774

Day, B. L., & Lyon, I. N. (2000). Voluntary modification of automatic arm movements evoked by motion of a visual target. *Experimental Brain Research*, *130*(2), 159–168. https://doi.org/10.1007/s002219900218

Downey, J. E., Brane, L., Gaunt, R. A., Tyler-Kabara, E. C., Boninger, M. L., & Collinger, J. L. (2017). Motor cortical activity changes during neuroprosthetic-controlled object interaction. *Scientific Reports*, *7*(1), 16947. https://doi.org/10.1038/s41598-017-17222-3

Dubreuil, A., Valente, A., Beiran, M., Mastrogiuseppe, F., & Ostojic, S. (2022). The role of population structure in computations through neural dynamics. *Nature Neuroscience*, *25*(6), 783–794. https://doi.org/10.1038/s41593-022-01088-4

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

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

Evarts, E. V., & Tanji, J. (1976). Reflex and intended responses in motor cortex pyramidal tract neurons of monkey. *Journal of Neurophysiology*, *39*(5), 1069–1080. https://doi.org/10.1152/jn.1976.39.5.1069

Ferezou, I., Haiss, F., Gentet, L. J., Aronoff, R., Weber, B., & Petersen, C. C. (2007). Spatiotemporal Dynamics of Cortical Sensorimotor Integration in Behaving Mice. *Neuron*, *56*(5), 907–923. https://doi.org/10.1016/j.neuron.2007.10.007

Fetz, E. E., & Baker, M. A. (1973). Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. *Journal of Neurophysiology*, *36*(2), 179–204. https://doi.org/10.1152/jn.1973.36.2.179

Fetz, E. E. (1969). Operant Conditioning of Cortical Unit Activity. *Science*, *163*(3870), 955–958. https://doi.org/10.1126/science.163.3870.955

Fetz, E. E., & Finocchio, D. V. (1971). Operant Conditioning of Specific Patterns of Neural and Muscular Activity. *Science*, *174*(4007), 431–435. https://doi.org/10.1126/science.174.4007.431

Fetz, E. E., & Finocchio, D. V. (1975). Correlations between activity of motor cortex cells and arm muscles during operantly conditioned response patterns. *Experimental Brain Research*, *23*(3). https://doi.org/10.1007/BF00239736

Franklin, D. W., & Wolpert, D. M. (2008). Specificity of Reflex Adaptation for Task-Relevant Variability. *Journal of Neuroscience*, *28*(52), 14165–14175. https://doi.org/10.1523/JNEUROSCI.4406-08.2008

Franklin, D. W., Franklin, S., & Wolpert, D. M. (2014). Fractionation of the visuomotor feedback response to directions of movement and perturbation. *Journal of Neurophysiology*, *112*(9), 2218–2233. https://doi.org/10.1152/jn.00377.2013

Fraser, G. W., & Schwartz, A. B. (2012). Recording from the same neurons chronically in motor cortex. *Journal of Neurophysiology*, *107*(7), 1970–1978. https://doi.org/10.1152/jn.01012.2010

Fukuma, R., Yanagisawa, T., Nishimoto, S., Sugano, H., Tamura, K., Yamamoto, S., Iimura, Y., Fujita, Y., Oshino, S., Tani, N., Koide–Majima, N., Kamitani, Y., & Kishima, H. (2022). Voluntary control of semantic neural representations by imagery with conflicting visual stimulation. *Communications Biology*, *5*(1), 214. https://doi.org/10.1038/s42003-022-03137-x

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

Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, *119*(2), 593–609. https://doi.org/10.1093/brain/119.2.593

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J.Neurosci.*, *2(11)*(11), 1527–1537. https://doi.org/citeulike-article-id:444841

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal Population Coding of Movement Direction. *Science*, *233*(4771), 1416–1419. https://doi.org/10.1126/science.3749885

Golub, M. D., Chase, S. M., Batista, A. P., & Yu, B. M. (2016). Brain–computer interfaces for dissecting cognitive processes underlying sensorimotor control. *Current Opinion in Neurobiology*, *37*, 53–58. https://doi.org/http://dx.doi.org/10.1016/j.conb.2015.12.005

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

Hari, R., Forss, N., Avikainen, S., Kirveskari, E., Salenius, S., & Rizzolatti, G. (1998). Activation of human primary motor cortex during action observation: A neuromagnetic study. *Proceedings of the National Academy of Sciences*, *95*(25), 15061–15065. https://doi.org/10.1073/pnas.95.25.15061

Hennig, J. A., Golub, M. D., Lund, P. J., Sadtler, P. T., Oby, E. R., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Yu, B. M., & Chase, S. M. (2018). Constraints on neural redundancy. *eLife*, *7*. https://doi.org/10.7554/eLife.36774

Hennig, J. A., Oby, E. R., Golub, M. D., Bahureksa, L. A., Sadtler, P. T., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Chase, S. M., & Yu, B. M. (2021).

Learning is shaped by abrupt changes in neural engagement. *Nature Neuroscience*, *24*(5), 727–736. https://doi.org/10.1038/s41593-021-00822-8

Holdefer, R. N., & Miller, L. E. (2002). Primary motor cortical neurons encode functional muscle synergies. *Experimental Brain Research*, *146*(2), 233–243. https://doi.org/10.1007/s00221-002-1166-x

Hooks, B. M., Mao, T., Gutnisky, D. A., Yamawaki, N., Svoboda, K., & Shepherd, G. M. G. (2013). Organization of Cortical and Thalamic Input to Pyramidal Neurons in Mouse Motor Cortex. *Journal of Neuroscience*, *33*(2), 748–760. https://doi.org/10.1523/JNEUROSCI.4338-12.2013

Huber, D., Gutnisky, D. A., Peron, S., O'Connor, D. H., Wiegert, J. S., Tian, L., Oertner, T. G., Looger, L. L., & Svoboda, K. (2012). Multiple dynamic representations in the motor cortex during sensorimotor learning. *Nature*, *484*(7395), 473–478. https://doi.org/10.1038/nature11039

Jacobs, J. V., & Horak, F. B. (2007). Cortical control of postural responses. *Journal of Neural Transmission*, *114*(10), 1339–1348. https://doi.org/10.1007/s00702-007-0657-0

Jarosiewicz, B., Chase, S. M., Fraser, G. W., Velliste, M., Kass, R. E., & Schwartz, A. B. (2008). Functional network reorganization during learning in a brain-computer interface paradigm. *Proceedings of the National Academy of Sciences*, *105*(49), 19486–19491. https://doi.org/10.1073/pnas.0808113105

Jeon, B. B., Fuchs, T., Chase, S. M., & Kuhlman, S. J. (2022). Existing function in primary visual cortex is not perturbed by new skill acquisition of a non-matched sensory task. *Nature Communications*, *13*(1), 3638. https://doi.org/10.1038/s41467-022-31440-y

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

Kennedy, P., Bakay, R., Moore, M., Adams, K., & Goldwaithe, J. (2000). Direct control of a computer from the human central nervous system. *IEEE Transactions on Rehabilitation Engineering*, *8*(2), 198–202. https://doi.org/10.1109/86.847815

Kleinfeld, D., Sachdev, R. N., Merchant, L. M., Jarvis, M. R., & Ebner, F. F. (2002). Adaptive Filtering of Vibrissa Input in Motor Cortex of Rat. *Neuron*, *34*(6), 1021–1034. https://doi.org/10.1016/S0896-6273(02)00732-8

Koralek, A. C., Jin, X., Long II, J. D., Costa, R. M., & Carmena, J. M. (2012). Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature*, *483*(7389), 331–335. https://doi.org/10.1038/nature10845

Kurtzer, I., Crevecoeur, F., & Scott, S. H. (2014). Fast feedback control involves two independent processes utilizing knowledge of limb dynamics. *Journal of Neurophysiology*, *111*(8), 1631–1645. https://doi.org/10.1152/jn.00514.2013

Law, A. J., Rivlis, G., & Schieber, M. H. (2014). Rapid acquisition of novel interface control by small ensembles of arbitrarily selected primary motor cortex neurons. *Journal of Neurophysiology*, (585), 1528–1548. https://doi.org/10.1152/jn.00373.2013

Moran, D. W., & Schwartz, A. B. (1993). Motor cortical representation of speed and direction during reaching. *J.Neurophysiol.*, *70*(1), 2676–2692. http://www.ncbi.nlm.nih.gov/pubmed/10561437

Moritz, C. T., & Fetz, E. E. (2011). Volitional control of single cortical neurons in a brain-machine interface. *Journal of neural engineering*, *8*(2), 025017. https://doi.org/10.1088/1741-2560/8/2/025017

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

Nashed, J. Y., Crevecoeur, F., & Scott, S. H. (2012). Influence of the behavioral goal and environmental obstacles on rapid feedback responses. *Journal of Neurophysiology*, *108*(4), 999–1009. https://doi.org/10.1152/jn.01089.2011

Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M., & Carmena, J. M. (2018). Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. *Neuron*, *97*(6), 1356–1368. https://doi.org/10.1016/j.neuron.2018.01.051

Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., Chase, S. M., & Batista, A. P. (2019). New neural activity patterns emerge with

long-term learning. *Proceedings of the National Academy of Sciences*, *116*(30), 15210–15215. https://doi.org/10.1073/pnas.1820296116

Patel, K., Katz, C. N., Kalia, S. K., Popovic, M. R., & Valiante, T. A. (2021). Volitional control of individual neurons in the human brain. *Brain*, *144*(12), 3651–3663. https://doi.org/10.1093/brain/awab370

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

Petrof, I., Viaene, A. N., & Sherman, S. M. (2015). Properties of the primary somatosensory cortex projection to the primary motor cortex in the mouse. *Journal of neurophysiology*, *113*(7), 2400–7. https://doi.org/10.1152/jn.00949.2014

Porter, R., & Lemon, R. (1993). *Corticospinal Function and Voluntary Movement*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198523758.001.0001

Pruszynski, J. A. (2014). Primary motor cortex and fast feedback responses to mechanical perturbations: a primer on what we know now and some suggestions on what we should find out next. *Frontiers in Integrative Neuroscience*, *8*. https://doi.org/10.3389/fnint.2014.00072

Pruszynski, J. A., & Scott, S. H. (2012). Optimal feedback control and the long-latency stretch response. *Experimental Brain Research*, *218*(3), 341–359. https://doi.org/10.1007/s00221-012-3041-8

Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*, *27*(1), 169–192. https://doi.org/10.1146/annurev.neuro.27.070203.144230

Rosén, I., & Asanuma, H. (1972). Peripheral afferent inputs to the forelimb area of the monkey motor cortex: Input-output relations. *Experimental Brain Research*, *14*(3), 257–273. https://doi.org/10.1007/BF00816162

Rouse, A. G., & Schieber, M. H. (2016). Spatiotemporal Distribution of Location and Object Effects in Primary Motor Cortex Neurons during Reach-to-Grasp. *Journal of Neuroscience*, *36*(41), 10640–10653. https://doi.org/10.1523/JNEUROSCI.1716-16.2016

Ruff, D. A., & Cohen, M. R. (2019). Simultaneous multi-area recordings suggest that attention improves performance by reshaping stimulus representations. *Nature Neuroscience*, *22*(10), 1669–1676. https://doi.org/10.1038/s41593-019-0477-1

Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., & Batista, A. P. (2014). Neural constraints on learning. *Nature*, *512*(7515), 423–426. https://doi.org/10.1038/nature13665

Sato, T. R., & Svoboda, K. (2010). The Functional Properties of Barrel Cortex Neurons Projecting to the Primary Motor Cortex. *Journal of Neuroscience*, *30*(12), 4256–4260. https://doi.org/10.1523/JNEUROSCI.3774-09.2010

Schieber, M. H. (2011). Dissociating motor cortex from the motor. *The Journal of physiology*, *589*(Pt 23), 5613–24. https://doi.org/10.1113/jphysiol.2011.215814

Scott, S. H. (2003). The role of primary motor cortex in goal-directed movements: Insights from neurophysiological studies on non-human primates. *Current Opinion in Neurobiology*, *13*(6), 671–677. https://doi.org/10.1016/j.conb.2003.10.012

Scott, S. H. (2004). Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, *5*(7), 532–545. https://doi.org/10.1038/nrn1427

Scott, S. H., Cluff, T., Lowrey, C. R., & Takei, T. (2015). Feedback control during voluntary motor actions. *Current Opinion in Neurobiology*, *33*, 85–94. https://doi.org/10.1016/j.conb.2015.03.006

Scott, S. H., & Kalaska, J. F. (1995). Changes in Motor Cortex Activity During Reaching Movements With Similar Hand Paths but Different Arm Postures. *Journal of Neurophysiology*, *73*(6), 2563–2567. https://doi.org/10.1152/jn.1995.73.6.2563

Scott, S. H., & Kalaska, J. F. (1997). Reaching Movements With Similar Hand Paths But Different Arm Orientations. I. Activity of Individual Cells in Motor Cortex. *Journal of Neurophysiology*, *77*(2), 826–852. https://doi.org/citeulike-article-id:2242397

Sergio, L. E., & Kalaska, J. F. (2003). Systematic Changes in Motor Cortex Cell Activity With Arm Posture During Directional Isometric Force Generation. *Journal of Neurophysiology*, *89*(1), 212–228. https://doi.org/10.1152/jn.00016.2002

Shemmell, J., An, J. H., & Perreault, E. J. (2009). The Differential Role of Motor Cortex in Stretch Reflex Modulation Induced by Changes in Environmental Mechanics and Verbal Instruction. *Journal of Neuroscience*, *29*(42), 13255–13263. https://doi.org/10.1523/JNEUROSCI.0892-09.2009

Smoulder, A. L., Pavlovsky, N. P., Marino, P. J., Degenhart, A. D., McClain, N. T., Batista, A. P., & Chase, S. M. (2021). Monkeys exhibit a paradoxical decrease in performance in high-stakes scenarios. *Proceedings of the National Academy of Sciences*, *118*(35). https://doi.org/10.1073/pnas.2109643118

Stavisky, S. D., Kao, J. C., Nuyujukian, P., Pandarinath, C., Blabe, C., Ryu, S. I., Hochberg, L. R., Henderson, J. M., & Shenoy, K. V. (2018). Brain-machine interface cursor position only weakly affects monkey and human motor cortical activity in the absence of arm movements. *Scientific Reports*, *8*(1), 16357. https://doi.org/10.1038/s41598-018-34711-1

Stavisky, S. D., Kao, J. C., Ryu, S. I., & Shenoy, K. V. (2017a). Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions. *Neuron*, *95*(1), 195–208. https://doi.org/10.1016/j.neuron.2017.05.023

Stavisky, S. D., Kao, J. C., Ryu, S. I., & Shenoy, K. V. (2017b). Trial-by-Trial Motor Cortical Correlates of a Rapidly Adapting Visuomotor Internal Model. *The Journal of Neuroscience*, *37*(7), 1721–1732. https://doi.org/10.1523/JNEUROSCI.1091-16.2016

Stavisky, S. D., Rezaii, P., Willett, F. R., Hochberg, L. R., Shenoy, K. V., & Henderson, J. M. (2018). Decoding Speech from Intracortical Multielectrode Arrays in Dorsal "Arm/Hand Areas" of Human Motor Cortex. *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 93–97. https://doi.org/10.1109/EMBC.2018.8512199

Stavisky, S. D., Willett, F. R., Avansino, D. T., Hochberg, L. R., Shenoy, K. V., & Henderson, J. M. (2020). Speech-related dorsal motor cortex activity does not interfere with iBCI cursor control. *Journal of Neural Engineering*, *17*(1), 016049. https://doi.org/10.1088/1741-2552/ab5b72

Stavisky, S. D., Willett, F. R., Wilson, G. H., Murphy, B. A., Rezaii, P., Avansino, D. T., Memberg, W. D., Miller, J. P., Kirsch, R. F., Hochberg, L. R., Ajiboye, A. B., Druckmann, S., Shenoy, K. V., & Henderson, J. M. (2019). Neural ensemble dynamics in dorsal motor cortex during speech in people with paralysis. *eLife*, *8*. https://doi.org/10.7554/eLife.46015

Steinmetz, N. A., Zatka-Haas, P., Carandini, M., & Harris, K. D. (2019). Distributed coding of choice, action and engagement across the mouse brain. *Nature*, *576*(7786), 266–273. https://doi.org/10.1038/s41586-019-1787-x

Suminski, A. J., Tkach, D. C., Fagg, A. H., & Hatsopoulos, N. G. (2010). Incorporating Feedback from Multiple Sensory Modalities Enhances Brain-Machine Interface Control. *Journal of Neuroscience*, *30*(50), 16777–16787. https://doi.org/10.1523/JNEUROSCI.3967-10.2010

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

Tanji, J., & Evarts, E. V. (1976). Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *Journal of Neurophysiology*, *39*(5), 1062–1068. https://doi.org/10.1152/jn.1976.39.5.1062

Taylor, D. M., Tillery, S. I. H., & Schwartz, A. B. (2002). Direct Cortical Control of 3D Neuroprosthetic Devices. *Science*, *296*(5574), 1829–1832. https://doi.org/10.1126/science.1070291

Tkach, D., Reimer, J., & Hatsopoulos, N. G. (2007). Congruent Activity during Action and Action Observation in Motor Cortex. *Journal of Neuroscience*, *27*(48), 13241–13250. https://doi.org/10.1523/JNEUROSCI.2895-07.2007

Todorov, E. (2000). Direct cortical control of muscle activation in voluntary arm movements: a model. *Nature Neuroscience*, *3*(4), 391–398. https://doi.org/10.1038/73964

Velliste, M., Kennedy, S. D., Schwartz, A. B., Whitford, A. S., Sohn, J.-W., & McMorland, A. J. C. (2014). Motor Cortical Correlates of Arm Resting in the Context of a Reaching Task and Implications for Prosthetic Control. *Journal of Neuroscience*, *34*(17), 6011–6022. https://doi.org/10.1523/JNEUROSCI.3520-13.2014

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

Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation Through Neural Population Dynamics. *Annual Review of Neuroscience*, *43*(1), 249–275. https://doi.org/10.1146/annurev-neuro-092619-094115

Wilson, S. M., Saygin, A. P., Sereno, M. I., & Iacoboni, M. (2004). Listening to speech activates motor areas involved in speech production. *Nature Neuroscience*, *7*(7), 701–702. https://doi.org/10.1038/nn1263

Zhou, X., Tien, R. N., Ravikumar, S., & Chase, S. M. (2019). Distinct types of neural reorganization during long-term learning. *Journal of Neurophysiology*, *121*(4), 1329–1341. https://doi.org/10.1152/jn.00466.2018