THE UNIVERSITY OF CHICAGO

THE NATURAL VARIABILITY OF A DEXTEROUS MOTOR SKILL IS STABLY ENCODED IN THE CORTEX OF FREELY BEHAVING MICE

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES

AND THE PRITZKER SCHOOL OF MEDICINE

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL NEUROSCIENCE

BY

ELIZABETH ANN DE LAITTRE

CHICAGO, ILLINOIS

MARCH 2025

Abstract

The nervous system enables precise, skilled movements across diverse contexts, consistently performed day after day. However, the extent to which the brain's neural activity patterns are specific to each instance of a movement or behavior, and how similar these patterns are across days, remains unclear. To address this, we record calcium fluorescence activity in the motor cortex of freely moving mice performing a self-initiated, skilled reach-to-grasp task. The high trial counts and single-trial variability in this task allow for a rigorous statistical analysis of moment-to-moment movement encoding across matched behavioral sets over five days. We found that single neurons in motor cortex encode the single trial (reach-to-reach) details of paw, digit, and head movements, and this encoding is stable over days. This suggests that the stable contribution of individual cells in the motor circuit underlies the generation of skilled movements, even in complex sensory-driven tasks like reach-to-grasp.

Contents

List of Figures v

Acknowledgments vi

1. General Introduction 1

    1.1. The neural underpinnings of voluntary forelimb movements in mammals 1

    1.2. Neural encoding of naturalistic behavior 7

    1.3. Representational drift as a fundamental feature of neural coding 10

    1.4. Representational drift in the motor system 15

2. The natural variability of a dexterous motor skill is stably encoded in the cortex of freely behaving mice 18

    2.1. Introduction 18

    2.2. Results 20

    2.3. Discussion 39

    2.4. Methods 44

    2.5. Supplementary Figures 61

3. General Discussion 66

References 74

List of Figures

Figure 1. A dexterous reach-to-grasp task elicits variable yet spatiotemporally coordinated forelimb and digit movements 23

Figure 2. Each animal develops its own unique reaching pattern, yet consistently uses a similar range of movements across days. 25

Figure 3: Neurons in CFA M1 motor cortex modulate during reaching 27

Figure 4: Reach-locked average neuronal activity patterns (PETHs) are stable across days 29

Figure 5: Trial-by-trial variance in forelimb, digit, and head movements allows for accurate modeling of single-trial CFA M1 neuronal activity 35

Figure 6: Linear model encoding profiles of single cells are stable across days 38

Supplementary Figure 1: Mice exhibit stable task performance over the five days 61

Supplementary Figure 2: Overview of reach strategies for all mice 62

Supplementary Figure 3: Distributions of kinematic variables before and after covariate balancing procedure, for all 14 kinematic variables 63

Supplementary Figure 4: Predictor correlations, and count and proportion of significant linear model fits by dataset 64

Supplementary Figure 5: Linear model encoding profiles constructed from R-squared values are stable across days 65

Acknowledgments

The first and biggest acknowledgment goes to my advisor Jason MacLean, whose mentorship and encouragement have been vital throughout my PhD. Next, this project has been made better by countless insightful discussions with collaborators and colleagues, particularly Maayan Levy, Gabriella Wheeler Fox, Yuqing Zhu, Hal Rockwell, Chad Smith, Siwei Wang, Dalton Moore, Marina Sundiang, and Hank Grier; and with my thesis committee members Stephanie Palmer, Nicho Hatsopoulos, and Matt Kaufman. Special shoutouts go to JD Laurence-Chasen, Caleb Sponheim, Maayan Levy, Jen Ding, Will Liberti, and Matt Getz, all of whom have given me that invaluable kind of scientific feedback/advice that can only come from people who are both close colleagues and close friends.

The neuroscience community at UChicago is vibrant and engaged; I am grateful to have had the opportunity to do science here. My time at UChicago is peppered with memories of hallway conversations that started as a brief hello, then quickly turned into an extremely useful discussion of recording methodology or the most recent seminar. I would like to particularly thank Elena Rizzo, Stephanie Dubbeld, Lili Gonzalez Hernandez, and the rest of the neuroscience staff for extensive behind-the-scenes administrative work that is the foundation for this strong community.

My life in lab was bolstered by my life outside of lab, and the incredible people I got to spend time with in Chicago. Ellen Ketter, JD Laurence-Chasen, Maryn Carlson, Ben Ketter, and Caleb Sponheim have been there from the beginning, and Anna Gatdula, Madison Weigand, and Francis and Fred Ketter joined in soon after. You all are brilliant, thoughtful, caring, intentional

people. To each of you: I cannot fathom where I'd be without your wholehearted friendship. Thank you for everything. Let's cheers to how far we've come and to the adventures ahead.

Thank you to my sport communities for getting me out of lab and out of my head: the triathlon team (particularly Coach Linda, Anna Gatdula, and Matt Getz), my orca pod swim friends (particularly Hannah Martin, Nate Edwards, and Casey Kraus), my frisbee crew (particularly Jojo Amdur, Sabina Hartnett, Casey Kraus, and Sam Swank), and the 2020 DIY yoga crew.

Thank you to Maayan Levy for being a mentor, labmate, and friend, and for bringing me on so many adventures. Thank you to Hannah Martin for being the roommate/friend/yoga instructor/orca podmate/cat I needed in 2020 and beyond, and thank you to Sabrina Arif for rounding out our roommate trio and brightening our home with laughs and doodles. Thank you to Emily, Meike, Evan, Kourtney, Matt, Fernando, and Jimmy for welcoming me into the crew and for keeping me well fed and always laughing. Thank you to Wolfkat for being my constant partner-in-crime from afar.

Thank you Mom and Dad for fostering my curiosity, creativity, and general nerdy-ness from the beginning. Thank you to my immediate and extended family for encouraging me throughout this entire process, and also for keeping me grounded.

Lastly, thank you to my partner Steven for being an incredible support during these last few years. Thank you for feeding me, laughing with me, scheming with me, and most importantly for this thesis, for believing in me.

Chapter 1: General Introduction

Humans have an impressively flexible motor system. We can learn a wide range of movement sequences and patterns involving the coordination of many body parts and execute these movements with high precision. This impressive motor control is supported by complex neural circuitry consisting of many interconnected brain areas (Warriner et al., 2020; Asan et al., 2022; Nudo and Frost, 2007; Lemon, 2008). Understanding the neural underpinnings of voluntary movement will provide a vital foundation for understanding how it is impacted by disease, disorder, and damage, and for developing therapies for restoring function. In this work, I investigate the role of primary motor cortex (M1) in mice performing a complex, forelimb-intensive reach-to-grasp movement, and assess whether M1 encodes the low level execution details of this movement in a stable manner over days.

1.1 The neural underpinnings of voluntary forelimb movements in mammals

Voluntary forelimb movements and their neural correlates have mainly been studied in primates, although rodents have recently emerged as a model system for studying motor circuitry. A number of reasons have drawn motor neuroscientists to rodents, and to mice in particular. One, while manual dexterity in rodents is more limited than in primates, their overall motor repertoire is still extensive and flexible. Mice exhibit coordinated control of forelimb, paws, and digits during natural behaviors like manipulating food (Barrett et al., 2020; An et al., 2024; Whishaw et al., 2020). Additionally, they are capable of learning new complex and precise sequences of movements relatively quickly, making them excellent model systems for studying learned skills in particular. To study the control of the forelimbs, mice can be trained to do things

like press levers (Masamizu et al., 2014; Peters et al., 2014, 2017; Miri et al., 2017; Warriner et al., 2022), grasp and move joysticks (Mosberger et al., 2024; Wagner et al., 2019; Bollu et al., 2024), reach for water droplets (Galiñanes et al., 2018), and reach and grasp pellets or seeds (Guo et al., 2015; Wang et al., 2017; Becker et al., 2020). Two, rodents are relatively close evolutionary relatives of primates and as such their neural circuitry has many similarities, including a descending motor pathway from cortex to brainstem to spinal cord, interconnected with multiple feedback loops through basal ganglia and, separately, cerebellum (Nudo and Frost, 2007). Three, the genetic tools available for breaking down areal and cellular subunits of circuitry promise to elucidate mechanisms of neural computation within and across brain areas (for a review, see Warriner et al., 2020). Four, rodents are smaller and thus easier to house and handle than primates. As a model system, the rodent provides promising opportunities to extract principles of neural control of movement in a wide range of behaviors and contexts, at a level of biological detail currently unattainable in other species.

Here, I study a self-initiated, whole body reach-to-grasp movement in freely moving mice. This movement requires precision placement of the paw, which for freely moving animals is dependent on both whole body positioning and the activation of proximal muscles controlling the shoulder and the elbow. It also requires appropriate shaping and closing of the digits before and during contact with the object. Importantly, it requires the temporal coordination and coarticulation of these steps so that digit closure happens at the right paw placement. Finally, it requires activation of distal muscles controlling the digits to maintain grip on the seed while carrying it back to the mouse. All steps require online sensory feedback, both proprioceptive (where is my paw, what is the shape of my paw) and tactile (am I making contact with the object

and how). There are considerable similarities in the reach-to-grasp sequence between rodents and primates (Sacrey et al., 2009; Whishaw, Pellis, and Gorny, 1992; Guo et al., 2015; Becker et al., 2020; Whishaw et al., 2017). Although rodents are less dexterous than primates, they exhibit a number of coordinated movements during reaching and grasping behaviors that make this movement interesting to study from a systems neuroscience perspective. Rats adjust hand aperture when reaching for an object (a form of pre-shaping), and show individuation of the little finger when grasping (Whishaw, Dringenberg, and Pellis, 1992; Whishaw et al., 2010). Mice make coordinated use of their digits during reach-to-grasp when both freely moving and head-fixed (Whishaw et al., 2017). The fact that rodents demonstrate some but considerably less dexterity than primates is potentially an advantage rather than disadvantage, because the computations required to execute these movements are expected to be simpler and thus the neural circuitry required will potentially be easier to characterize. As I describe below, the neural basis of these movements is largely underexplored in rodents.

A significant foundation studying the neural correlates of reaching and grasping has been laid down in primates, focused on many interconnected motor and premotor cortical areas. Primary motor cortex (M1) is one area in this network. Because of its connectivity to brainstem and spinal cord through the pyramidal tract (Lemon, 2008), M1 is thought of as the final cortical stage of processing before output to downstream brain areas. Because of this downstream connectivity, M1 has very direct, short-latency influence on muscles via corticospinal neurons in mice (Miri et al., 2017; Ueno et al., 2018; Koh et al., 2024). Notably, in Miri et al., 2017, this influence of corticospinal projecting neurons was observed during a learned, forelimb-intensive behavior but not during locomotion, suggesting that this direct, short latency influence may be

specific to precision forelimb movements, while other circuits more directly control locomotion. In rodents, there are two motor cortical areas with extensive and distinct projections to cervical (forelimb) spinal cord: caudal forelimb area (CFA) and rostral forelimb area (RFA) (Carmona et al., 2024; Saiki-Ishikawa et al., 2024; Ueno et al., 2018; Hira et al., 2013). The shared and separate roles of CFA and RFA are still being worked out, though one inactivation study suggests that RFA may be more involved in grasping while CFA may be more involved in precision reaching (Wang et al., 2017). In both primates and rodents, primary motor areas are heavily interconnected with other cortical and subcortical areas, for example, with the cortico-basal-thalamic loop and intra-cortically (Anderson et al., 2010; Oswald et al., 2008; Hooks et al., 2013; Weiler et al., 2008; Carmona et al., 2024; Yamawaki et al., 2021), and through the brainstem to cerebellum (Wagner et al., 2019). This connectivity suggests involvement in sensory-driven tasks, and those where high-level task variables may need to be integrated with low level details of movement. The fact that M1 can have short latency influence on muscles during some behaviors and not others suggests that motor cortex can serve different roles in different contexts.

Lesion and inactivation studies have demonstrated that motor cortex is an important node in the multi-area circuit for the control of reach-to-grasp and other precision reaching behaviors in rodents (Guo et al., 2015; Wang et al., 2017; Miri et al., 2017; Galiñanes et al., 2018). It has also been demonstrated through lesion or inactivation that motor cortex can be required to learn a task, even when it is not required to perform that task later (Kawai et al., 2015; Hwang et al., 2019). Stimulation of CFA elicits movements of the forelimb, both proximal (shoulder and

elbow) and distal (wrist and digits), using either intracortical microstimulation or optogenetic activation (Tennant et al., 2011; Harrison et al., 2012; Omlor et al., 2018; Wang et al., 2017).

The study of neural coding in mouse motor cortex is built upon a foundation of studies in primates, which I very briefly review here. A large body of studies of M1 encoding for forelimb movements in primates have focused on the relationship of the activity of either single neurons or populations to various movement parameters during reaching, including kinematic features like direction (Georgopoulos et al., 1982; Schwartz et al., 1988) and unsigned speed (Moran & Schwartz, 1999), kinetic features like forces (Evarts, 1968) and directional changes in forces (Georgopoulos et al., 1992), and kinematics of the arm trajectory rather than the net change in position (Hocherman & Wise, 1991). Tuning for these parameters has been observed to change over the course of a movement (Churchland & Shenoy, 2007; Mason et al., 1998; Sergio et al., 2005; Hatsopoulos et al., 2007) and depending on the window considered, such as during preparatory versus movement periods (Churchland et al., 2010). Additionally, kinematics of the wrist and digits have been shown to explain variance in single neurons during grasping and reach-to-grasp (Vargas-Irwin et al., 2010; Saleh et al., 2010; Saleh et al., 2012). As an alternative to this representation-based approach, a growing body of studies have used the framework of dynamical systems to explain population activity in M1 (Churchland et al., 2012; Russo et al., 2018; see review in Shenoy et al., 2013). This line of work has highlighted differences between reaching movements and reach-to-grasp movements: for example, Suresh and colleagues observed that M1 population activity during grasping was much more tangled than during reaching, suggesting a reliance on inputs to generate the appropriate movement commands (Suresh et al., 2020).

In comparison, neural coding for reaching and grasping in mice is understudied. Thus far, most studies of neural coding in CFA and other cortical motor areas in mice have focused on the encoding of high level variables such as behavioral category (Dhawale et al., 2017a, Huber et al., 2012; Dombeck et al., 2009; Peters et al., 2017) and categorical reach direction (Galiñanes et al., 2018). More recently, motor cortical activity has been shown to encode muscle activations and kinematic details (position and velocity) of forelimb movements in mice trained to walk on a ladder or perform a climbing task (Omlor et al., 2018, Koh et al., 2024), suggesting that motor cortex may also be involved in the low level execution details of the constituent precision forelimb movements. These results suggest that the same would be true in a task requiring both precision reaching and dexterity, but this has yet to be described.

What role does motor cortex play in the sensory-driven reach-to-grasp behavior, which requires coordination of the whole body, from gross movements of all limbs and the trunk to fine control of the digits? Rats show individuated control of digits in a pellet grasp task, which never recovers after motor cortical stroke, implicating a role for motor cortex in dexterity in rodents (Alaverdashvili and Whishaw 2008). The neural correlates of digit control remain understudied in mice, however, despite careful behavioral descriptions demonstrating that mice make coordinated use of their digits in reach-to-grasp (Whishaw et al., 2017) and food handling (Barrett et al., 2020; An et al., 2024). Cortical encoding of whole body posture in rodents is similarly understudied. Notably, recording from the posterior parietal cortex and secondary motor cortex in rats freely behaving in an open arena, Mimica and colleagues found cells that were tuned to static posture of the head, neck, and back, measured in egocentric coordinates (Mimica et al., 2018). In rats standing on force sensors, Disse and colleagues observe responses

in hindlimb and trunk sensorimotor cortex during postural disturbances (Disse et al., 2023). It remains to be seen if either whole body posture or movements of the digits are encoded in forelimb primary motor cortex in the mouse in a task that combines across physical scales. Exploring what features of movement are encoded, including not only proximal forelimb but also whole body posture and details of digits, is the first step to breaking down the computations occurring there during complex movements like reach-to-grasp.

1.2 Neural encoding of naturalistic behavior

Building on foundational work with simpler tasks, the field has made a concerted effort to study naturalistic (or even natural behavior). By naturalistic behavior, I mean behavior that is less stereotyped, more flexible and adaptive to changing conditions, spatiotemporally complex (involving multiple parts of the body, potentially coordinated over a range of timescales), challenging, and potentially freely moving. Our understanding of this kind of behavior is limited. A number of tuning models described in the previous section have been applied to free behavior, with minimal success (Jackson et al., 2007; Aflalo & Graziano et al., 2006). There are many possible reasons why these models don't extend. For example, with a larger kinematic space, it is difficult to achieve a dense enough sampling of the free behavior to capture encoding reliably. Perhaps with additional sampling, we would find that tuning for allocentric kinematic and kinetic variables would indeed extend to explain sufficient variance in M1.

The purpose of constraining behavior in experiment is to remove variability of all other dimensions of movement as much as possible, to isolate the dimension in question. However, it is possible that when other parts of the body are allowed to move, or move in different ways, the

code changes. For example, M1 activity may multiplex encoding for forelimb with that for whole body movements, such that when a task requires animals to coordinate movements of multiple parts of their body, the relationship between neural activity and forelimb movements changes relative to conditions that only require forelimb movements. This could explain why findings in constrained behaviors do not extend to free behavior. Two findings from work in rodent motor cortex are consistent with this view. Omlor and colleagues observed that encoding of forelimb joint angles in sensorimotor cortex during a ladder walking task changes between ladders with regular and irregular rungs, suggesting that cortex might be more involved in the more challenging case of irregular rungs (Omlor et al., 2018). Warriner and colleagues observed that the population correlation structure of motor cortical activity differed across two lever tasks, pushing a precise amount, requiring activation of shoulder flexors only, versus holding a lever in place, requiring co-activation of proximal muscles (Warriner et al., 2022). Unfortunately, in both studies, the movements that animals make in the two different task paradigms are different and not controlled for in the encoding analyses. This means we cannot fully attribute the difference to changes in motor cortical encoding, though the magnitude of change across the two conditions in each study is notable. There is a general dearth of studies of cortical encoding of whole body posture in freely moving animals, and hence our understanding of how the encoding of forelimb movements may be modulated by movements of other parts of the body is also limited. Here we attempt to extend the field by studying a forelimb-intensive behavior in a freely moving animal, and consider how the movements of multiple body parts are encoded in motor cortex.

Naturalistic behaviors pose many interconnected challenges for experimental study. Throughout this manuscript we discuss these challenges and how we overcame them. I

summarize them here briefly. First, quantifying neural coding for behaviors that show more variance requires analytical approaches that don't assume stereotypy. As I discuss in Chapter 2 (the data chapter), trial averaging can still be useful for capturing the encoding of high-level behavioral or task variables, but these approaches ignore behavioral variance that could be relevant to the neural encoding. Second, statistical tools that do not assume stereotypy across trials, such as the linear regression over timepoints used in Chapter 2, require high trial counts and sufficient sampling of the behavioral space over which one would like to estimate coding. We describe in the discussion of Chapter 2 the steps we took to encourage variance in the animals' movements in our experiments. Finally, studying the neural encoding of this trial-to-trial variance requires a detailed description of behavior and task elements at a fine enough resolution to capture the variance of the behavior in the first place. And, interconnected with the previous reason, one needs to carefully quantify that variance for a large dataset, the challenge of quality at quantity. In our case, quantitatively describing the movements of freely moving animals requires pose estimation powered by computer vision, which is only recently feasible at scale. In particular, the inability to track individual digits for enough timepoints to generate large enough and statistically rich enough datasets has been a significant barrier to rigorous quantification of neural encoding for digit movements. Painstaking frame-by-frame manual tracking is not realistic for large datasets. Manipulanda such as levers and joysticks provide straightforward ways to record movement, but really only reflect the activation of proximal muscles of the arm (those that move the shoulder and elbow) and thus do not facilitate recording of wrist or digit movements. Quantification of whole body posture faces a similar challenge, and the study of both areas has been recently jump-started by the adoption of computer vision tools.

An important novel aspect of this study is the detail with which we described the movement of the head and the digits in freely moving animals, in addition to detailed description of the paw position and velocity.

1.3 Representational drift as a fundamental feature of neural coding

Improvements in recording technology and computational resources have allowed the recording of the same neural populations over days to weeks, allowing systems neuroscientists to assess whether the encoding properties of individual cells and populations remain fixed over time. Contrary to expectations, the neural encoding for various stimulus, task, and behavioral variables have been observed to change from recording session to recording session (some more recent examples are Marks and Goard, 2021; Geva et al., 2023; Wang et al., 2023; Deitch et al., 2021; Bauer et al., 2024; and the topic is reviewed in Micou and O'Leary, 2023; Driscoll et al., 2022; Rule et al., 2019; Clopath et al., 2017; Lütcke et al., 2013). This phenomenon, referred to as "representational drift," has been observed in many brain areas, in many species, using multiple recording methods, and has motivated significant experimental and theoretical work. In this section, I provide context and motivation for studying neural encodings over timescales longer than a single recording session. In the next section, I review in detail the study of representational drift in motor systems.

The study of representational drift originated in representationalist frameworks, which rely on quantifying a mapping from neural activity to some measured (usually external) variable. This mapping does not need to be limited to classic representationalist tools like tuning curves, PETHs, or even regression coefficients, but can include any measurable property of neural

activity associated with some external or internal variable, such as correlations between cells, principle components or latents of population activity, or a dynamical feature of population activity observed during a particular behavior/stimulus/etc. The motivating question at hand is whether whatever property of the nervous system that we are focusing on is stable or changing, and if changing, how? For simplicity, throughout this document, I will use the terms mapping or neural encoding to refer to some statistical regularity between neural activity and a stimulus, task, or behavior variable. Thus, in a study comparing neural encoding across time, the choice of mapping is just as important as in studies not looking across time. Any observation of drift or stability can only speak for that mapping. For example, in this work, I focus on two types of encoding in single cell activity in motor cortex: that for the average behavior (peri-event time histograms locked to the onset of reaching) and that for the timepoint-to-timepoint position and velocity of different body parts. These are two different ways to quantify the encoding of this behavior in motor cortex, which warrant separate attention to quantify whether each is drifting or not.

Why might we observe changes in neural coding from session to session? The first and most obvious reason is because of experimental confounds, which I discuss at the end of this section, along with considerations on how to avoid them. Changes in neural coding that are due to experimental confounds are not true representational drift, although if not accounted for, they can lead to the incorrect interpretation that an encoding is changing. One likely cause of drift is the inherent noise of the nervous system and the dynamic nature of biological tissue due to turnover of cellular components that influence neural excitability and synaptic transmission (reviewed in Faisal et al., 2008; Renart and Machens, 2014). For some of these processes, we

expect the effects on neural coding to be mean-reverting, meaning that comparison over a long enough timescale (days to weeks) would reveal short term fluctuations around a stable mean (as in Jensen et al., 2022). There are other potential causes, however, that we expect to be non-mean reverting (as in Rule et al., 2019), with statistical structure that could be informative about the underlying processes. These include continued consolidation of learned representations without an observable change in behavior, perhaps to increase efficiency, restructuring of representations to allow for incorporation of new information (continual learning, such as in Driscoll et al., 2017), and homeostatic plasticity to regulate properties of activity over time. Under the assumption that the space of possible neural activity patterns to represent a given stimulus or behavior is degenerate, many of these changes might occur without any apparent impact on cognition or behavior.

From a neural coding perspective, these sources of change are challenges to be overcome, and how downstream populations could reliably read out external variables from a drifting population is an active area of computational and theory work. For example, the encoding of natural movies in mouse primary visual cortex (V1) is unstable at the single neuron level over weeks to months (Marks and Goard, 2021), but a stable encoding can be found at the population level (Xia et al., 2021). Another example, stable readout at the level of the population despite drift in single neuron encoding has also been shown for the encoding of task and movement variables in mouse posterior parietal cortex during a virtual navigation task (Driscoll et al., 2017; Rule et al., 2020). Recent theory work has described how representational drift with different statistical properties could be compensated for by downstream populations without privileged knowledge of the external variable (Rule and O'Leary, 2022; Micou and O'Leary, 2024).

On the other hand, theory work has proposed potential benefits of drifting representations. Drift may make learned representations more robust by injecting stochasticity in activity or full dropout of individual neurons, which can prevent overfitting to particular experiences (Aitken et al., 2022). Mechanisms that allow neural circuitry to maintain dynamic representations might also allow the brain to more quickly and robustly incorporate new information or representations, potentially without losing older representations (Driscoll et al., 2017).

How can studying whether and how neural encodings change over time expand our understanding of the nervous system? It must first be noted that interpreting drift or stability inherits all the challenges of neural coding generally. For example, if we observe that a particular neural coding is stable, we cannot assume that information is being read out and used downstream, in the same way that if a particular kinematic variable explains significant neural variance that that activity is causally related to the behavior. Conversely, if we observe that a particular neural coding is unstable, we cannot assume that means the brain is not using that information. With all this in mind, characterizing the nature and dynamics of drift could be informative about biological mechanisms supporting neural encoding. For example, Driscoll and colleagues observe a particular distribution of drift rates across the population: small numbers of cells make large discrete jumps from session to session (Driscoll et al., 2017). These statistics narrow the possibility of underlying plasticity mechanisms that could be at work that either maintain or erode learned representations (Rule and O'Leary, 2022). Additionally, comparing the statistics of drift for different stimuli may reveal how those stimuli might be processed differently by the visual system. Recall that Marks and Goard observe that the encoding of

drifting gratings in mouse V1 is stable over weeks, while the encoding of natural scenes shows significant drift over the same time period (Marks and Goard, 2021). A stable signature of time in the natural movie could be found at the level of the population, however, suggesting that encoding of more complex stimuli could be more distributed across the V1 population than that for simpler stimuli like gratings (Xia et al., 2021).

There are a number of methodological considerations that must be taken into account when comparing neural encodings quantified in different sessions. When recording neurons via calcium imaging, as we do here, ensuring that you are recording from the same neurons in the same way requires robust protocols for returning to the same fields of view each day and strict criteria for matching cell bodies across sessions (Lutcke et al., 2013; Sheintuch et al., 2017). The same statistical considerations apply as when quantifying neural coding within a single session: you need to ensure the dataset contains sufficient sampling of the variable(s) you are quantifying coding for. To avoid confounds introduced by estimating coding from different ranges and sampling, steps must be taken to match distributions of the external variables over days, such as the covariate balancing procedure we do in this work. Day-to-day behavioral drift in movement has been shown to be correlated with slow neural drift (Jensen et al., 2022). If small changes of behavior are not accounted for from day to day, cells may appear more unstable than they actually are (Liberti, Schmid, et al., 2022). This is particularly important for studies involving any aspect of behavior or behavioral response, but can also impact studies of encoding of sensory variables alone. It has been shown that task-irrelevant movements (those not directly instructed by the task) can modulate neural activity across the brain (Musall, Kaufman et al., 2019; Stringer et al., 2019). Untracked changes in these behaviors can change the encoding of task-relevant

variables from session to session, not only for the encoding of movement but also stimulus variables (Sadeh and Clopath, 2022). A potential solution for this challenge is to track a variety of task-relevant and task-irrelevant variables and incorporate them into encoding models, as in Driscoll et al., 2017. Throughout this manuscript, we describe steps we took to avoid experimental confounds that could lead to the incorrect interpretation of drift, including performing high resolution tracking of movements of multiple body parts and covariate balancing before fitting encoding models.

1.4 Representational drift in the motor system

As described in the previous section, numerous studies of sensory and association areas have observed that neural encodings can change from session to session. In contrast, the encoding in motor areas thus far appears relatively stable, such as in rats performing a precision-timing lever press task (Jensen et al., 2022) or behaving freely in their homecages (Dhawale et al., 2017a), zebra finches performing directed or undirected song (Katlowitz et al., 2018, Liberti et al., 2016 (inhibitory neurons)), and macaques reaching to targets or controlling a cursor using a brain-computer interface (BCI) (Chestek et al., 2007; Ganguly and Carmena, 2009; Stevenson et al., 2011; Fraser and Schwartz, 2012; Flint et al., 2016; Gallego et al., 2020) (but also see Carmena et al., 2005; Rokni et al., 2007; and excitatory neurons in Liberti et al., 2016). Generally, studies of stability of encoding in the motor system have come out of two specific lines of work. In one, researchers observed a highly stereotyped behavior and seek to identify if it is supported by similarly stereotyped neural patterns of activity (as in Liberti et al., 2016; Katlowitz et al., 2018; and Jensen et al., 2022). In another, researchers seek to identify a stable

readout of movement intention for application in brain-computer interfaces (BCI) in humans (as in Chestek et al., 2007; Ganguly and Carmena, 2009; and Gallego et al., 2020). Naturally, these two differing subfields have different goals and thus different approaches for studying stability.

In the first line of work, the scientific question comes out of the rigid stereotypy of the behavior, the mating song of the male zebra finch or a heavily trained precision-timing lever press task in rats, a task designed to achieve the same stereotypy as birdsong. Thus, and understandably so, the neural encoding analyses employed in these studies have largely ignored behavioral variability within and across days, comparing instead the behavior-locked average responses of single neurons (peri-event time histograms, PETHs) (Liberti, et al., 2016; Katlowitz et al., 2018; Jensen et al., 2022) or time-averages of activity over all timepoints of a particular behavioral category (Dhawale et al., 2017a). Because of the specialized nature of the zebra finch song circuit, the stability of encoding for this behavior in particular is unlikely to speak to stability in motor systems broadly. It does however, provide evidence that the nervous system is capable of creating and maintaining stable circuits for stable behavioral output, despite the biological necessity of turnover of constituent cellular machinery.

In the second line of work, studying the encoding for single trial variance in the form of moment-to-moment kinematic details is more common. However, confidently tracking cells across sessions with electrophysiological recordings is difficult without continuous recording. Proposals have been made for matching units based on a collection of activity pattern metrics (Fraser and Schwartz, 2012), but using activity patterns to identify if cells are the same biases the sample towards stable cells, potentially obscuring an unstable portion of the population. To avoid this issue, researchers have described latent factors of population activity instead, which have

been observed to be stable over weeks to months (Gallego et al., 2020). However, the stability of encoding for the details of movements is largely under-explored at the single neuron level over periods longer than a single session. Identifying whether these stable latents correspond to stable or changing single neuron activity will narrow down possible network mechanisms responsible for maintaining stable motor encoding and output.

Finally, a concern raised about this entire collection of results is that the observed stable encodings only exist because the behaviors often studied in laboratory settings are highly stereotyped and simple. Perhaps these are special cases of "overtrained" behavior where the nervous system has allocated a protected, "crystallized" circuit or pattern of activity to generate these movements reliably. One way to address this concern is to perform longitudinal studies of a wider range of more naturalistic and flexible behaviors such as the reach-to-grasp. It remains unclear if this behavior, which requires coordinated, precision control of both gross and fine movements, would be stably encoded in mammalian motor cortex, and if so, whether all features of the behavior are stably encoded, including trial-by-trial variability in joint kinematics.

Chapter 2: The natural variability of a dexterous motor skill is stably encoded in the cortex of freely behaving mice

2.1 Introduction

Even well-practiced behaviors exhibit variability with each attempt. In basketball, for instance, no two shots are identical, as each is shaped by factors such as starting posture, fatigue, and interactions with defenders and teammates. Yet, players consistently adapt and successfully perform this precise skill across these dynamic contexts. While generating goal-directed movements under such conditions remains a significant challenge in robotics (Billard and Kragic, 2019), biological nervous systems excel at this. To understand how motor skills are consistently executed across varying contexts, it is crucial to investigate how movement variability is represented in neural circuits and whether the representation of learned, complex motor behaviors remains stable over time. This insight can shed light on the circuit-level precision required to adapt to perturbations while maintaining reliable motor performance over extended periods of time.

Prehension, the act of reaching and grasping objects, is a vital yet complex movement requiring precise coordination of the trunk, shoulder, arm, wrist, and fingers, guided by rapid sensory feedback processing. The need to interact flexibly with diverse objects further enhances its complexity. Recent work suggests that precision reaching in mice relies on online sensory feedback rather than pre-planned, ballistic movements (Becker et al., 2020). Motor cortex plays a central role in the brain-wide circuit controlling reach-to-grasp and other precise reaching behaviors in rodents, as shown by lesion and inactivation studies (Guo et al., 2015; Wang et al., 2017; Miri et al., 2017; Galiñanes et al., 2018). In mice, the caudal forelimb area (CFA) and

other cortical motor regions represent high-level categorical variables, such as behavioral category (Dhawale et al., 2017a; Peters et al., 2014, 2017; Huber et al., 2012; Dombeck et al., 2009) and reach direction (Galiñanes et al., 2018). Simultaneously, motor cortical activity captures low-level features of forelimb movements, including muscle activations and kinematics (e.g., position and velocity), during ladder walking (Omlor et al., 2018), climbing (Koh et al., 2024), and a precision reach-for-water task (Zhu et al., 2022). However, the extent to which CFA in mice encodes whole-body posture or digit movement details during sensory-driven reach-to-grasp behavior, which requires coordination of the whole body, from gross movements of all limbs and the trunk to fine control of the digits, is unclear. While there is evidence for coordinated digit use in mice (Whishaw et al., 2017; Barrett et al., 2020; An et al., 2024) and posture-related activity in rats' posterior parietal, secondary motor, and sensorimotor cortices (Mimica et al., 2018; Disse et al., 2023), it is unknown how these kinematic features are represented in CFA. Understanding how CFA encodes movement features—such as proximal forelimb actions, whole-body posture, and fine digit control—is essential for clarifying its role in complex behaviors like reach-to-grasp. If M1 contributes to the online control of these movements as part of a broader sensorimotor circuit, it is likely to encode their moment-to-moment details.

Neural encoding in sensory and association brain areas often shows session-to-session variability, a phenomenon termed representational drift (reviewed in Micou and O'Leary, 2023; Driscoll et al., 2022; Bauer and Rose, 2021; Chambers and Rumpel, 2017; Clopath et al., 2017; Lütcke et al., 2013). In contrast, motor areas exhibit more stable encoding over time (Jensen et al., 2022; Dhawale et al., 2017a; Katlowitz et al., 2018; Liberti et al., 2016; Chestek et al., 2007;

Fraser and Schwartz, 2012; Ganguly and Carmena, 2009; Gallego et al., 2020; Flint et al., 2016; but see Rokni et al., 2007 and Liberti et al., 2016). Most studies of the stability of motor encoding have relied on trial-averaged metrics, such as peri-event time histograms (PETHs), which can mask neural stability or instability by not accounting for behavioral variability within and across sessions. This raises the possibility that apparent neural instability could reflect changes in behavior rather than shifts in the mapping between neural activity and behavior (Liberti, Schmid, et al., 2022; Jensen et al., 2022).

The stability of moment-to-moment encoding for fine-grained movement features, such as joint position, velocity, or muscle activations, remains underexplored. This is partly due to the challenge of studying behaviors that are variable enough to reveal specific neural activity patterns for each instance of a movement, while maintaining sufficient control to ensure statistical rigor. To address this gap, we examine a complex whole-body, dexterous reach-to-grasp task that produces high trial counts with significant trial-to-trial variability within sessions. Despite this variability, individual animals consistently explore a similar kinematic movement space across sessions, enabling us to create matched datasets for directly comparing the encoding of movement variables—kinematics of the head, paw, and digits—across days.

2.2 Results

A dexterous reach-to-grasp task elicits variable yet spatiotemporally coordinated forelimb and digit movements.

To investigate the stability of the cortical encoding of moment to moment forelimb and head movements, we trained unrestrained mice on a challenging reach-to-grasp task. We trained

the mice to reach for a millet seed placed on a narrow pedestal positioned in one of three locations just outside the front of the arena (Fig 1A, B, C). Mice self-initiate the delivery of a seed from the automated system by breaking an infrared beam positioned at the back of the arena opposite the reach port (~15cm or 2 body lengths away). Due to the limited mobility of the mouse shoulder, a critical aspect of the task is positioning the body and head to ensure the pedestal is within a reachable angle through the reach port. After triggering the delivery of a new seed and returning to the reach port, mice performed a series of 1 to 10 individual reaches in quick succession until they either successfully retrieved the seed or knocked it off. Each reach is defined as a single out-and-back motion of the paw through the reach port, where the paw exits and re-enters the arena. A set of reaches performed in quick succession before moving to trigger the delivery of another seed is referred to as a bout. Throughout this paper, we treat each reach as a single trial. During a 40 minute session, a trained animal will perform anywhere from 789 to 2831 reaches and successfully retrieve and consume anywhere from 27 to 139 seeds (see Supp Fig 1).

The pedestal's diameter is similar to that of the seed, making it easy for the millet seed to be knocked off (Fig. 1B). Successfully retrieving the seed requires precise paw placement and highly coordinated, precisely timed digit control. The narrow pedestal prevents repositioning the grasp on the seed after initial contact, allowing the mouse only one attempt per contact. After grasping the seed, mice must maintain their grip while transporting it back into the arena, adding to the task's complexity. Performance on this task typically plateaus after about 10–14 days of training (Chen et al., 2014).

To quantitatively describe the animals' digits, paw, and head movements during reaching, we recorded high-speed video during all sessions and employed markerless keypoint tracking (DeepLabCut, Mathis et al., 2018; Nath et al., 2019) to reconstruct the kinematics of each reach (see Fig 1A, C, D). On the paw, we tracked the joints of the first three digits and the wrist; on the head, we tracked the nose tip, corner of the eye, and a prominent corner of the surgical implant (Fig 1A). From this keypoint tracking, we extracted the position and velocity of the paw and head centroids during reaching (Fig 1C); we also ran principle component analysis on the collection of distances of the paw and head keypoints from their respective centroids, to extract proxies for the shape and orientation of the paw and the orientation of the head (see Methods). The position and velocity of the paw centroid indicate activation of proximal forelimb muscles controlling the elbow and shoulder, while principal components of the conformation of the paw reflect the activation of distal muscles controlling the wrist and digits. Similarly, the position and orientation of the head reflect the whole-body posture of the animal and potentially the activation of shoulder and neck muscles.

Within a single 40 minute session, animals exhibit high reach-to-reach variability in the speed, shape, and duration of their digit, paw, and head movements during reaching (Figs 1E, F, G). This trial-to-trial variability persists even after the animals have been trained beyond the time needed to master the task, as evidenced by their stable success rates during the period of study. (Supplementary Fig 1). We considered this within-session variability advantageous, as it allows us to examine the moment-to-moment neural encoding of these movements (as in Omlor et al., 2018 and Becker et al., 2020).

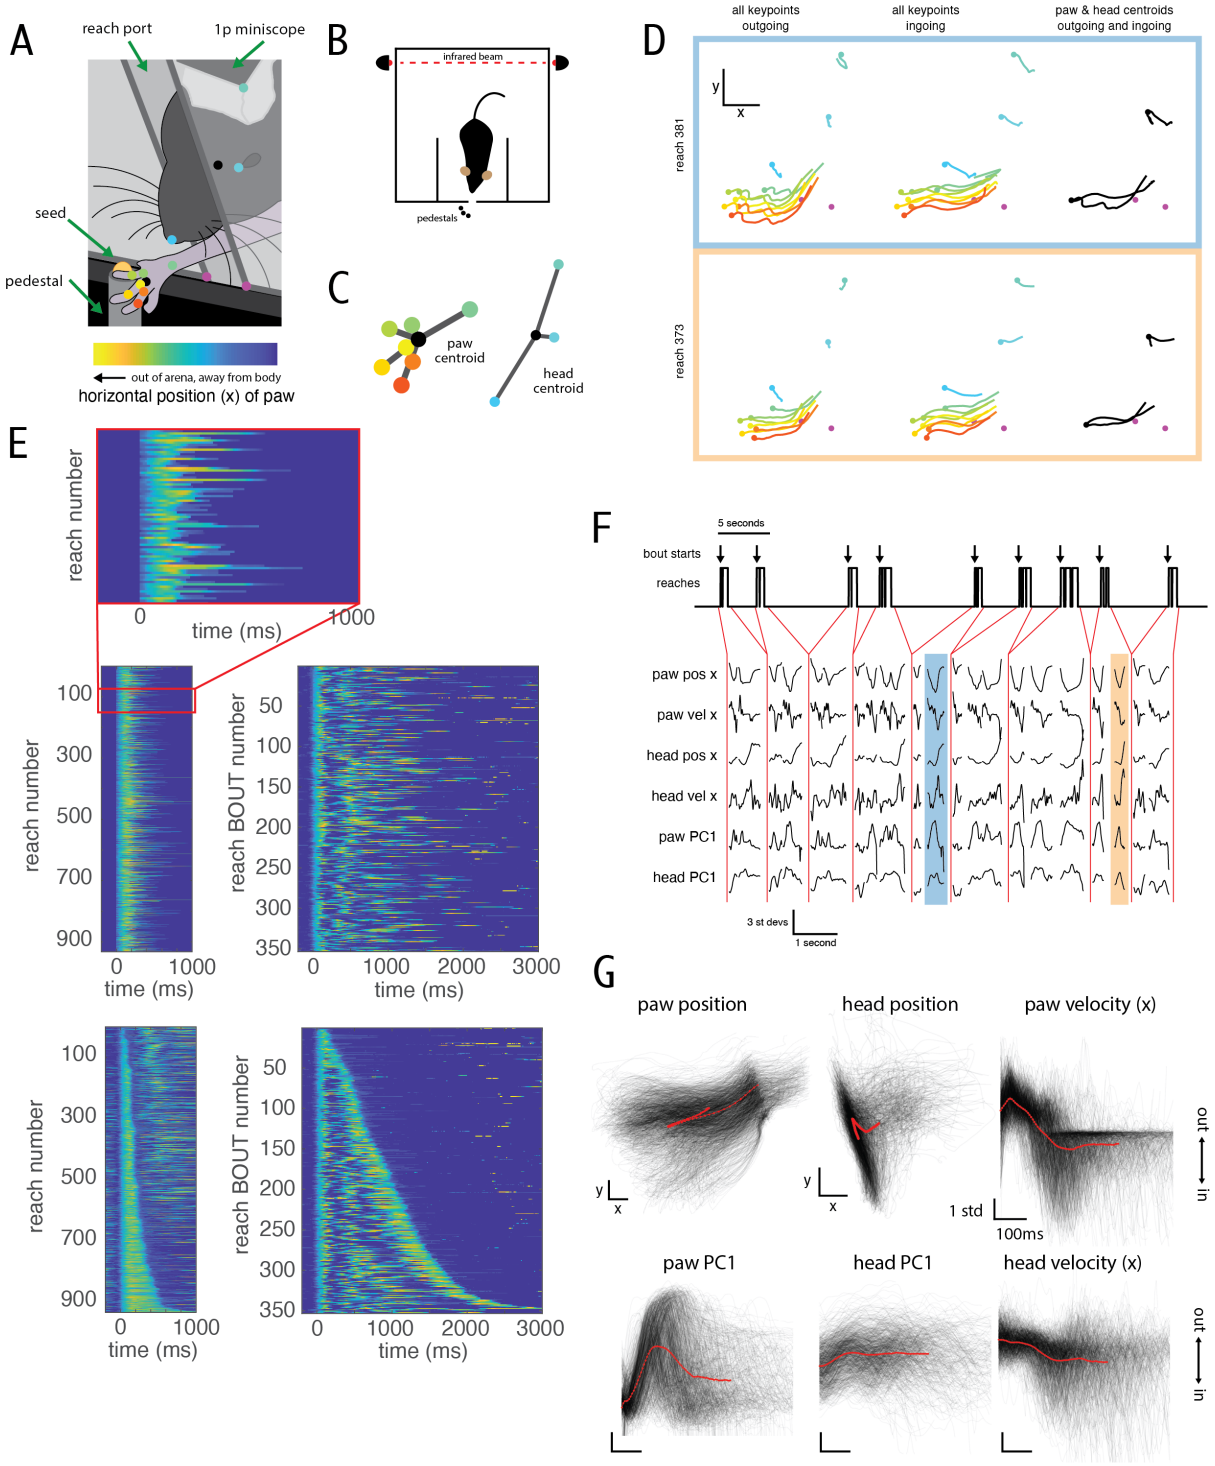

Figure 1: A dexterous reach-to-grasp task elicits variable yet spatiotemporally coordinated forelimb and digit movements.

A. Closeup view of mouse reaching, with 10 keypoints on the paw and head marked, from the viewing angle of the tracking camera. Colorbar at bottom indicates paw position in x relative to the view shown here; relevant for the heatmaps in E.

B. Schematic of mouse in arena, indicating location of pedestals, reachport, and infrared beam

Figure 1 (continued):

C. Closeup on keypoints on the paw and head, and their distances to the paw and head centroids (black dots). Colors match those in A.

D. Two example reconstructed reaches, where the colors match those in A and C. For all panels, dots indicate the endpoint (point of the reach that is furthest from the arena). Left column, all paw and head keypoints during the outgoing part of the two reaches. Middle column, same for the ingoing part of the reaches. Right column, the paw and head centroids only for both the outgoing and ingoing parts of the reaches. Magenta dots indicate the base corners of the reach port.

E. Paw position in the X dimension plotted for all reaches and bouts in one session for an example mouse (mouse 5). Colormap is under the schematic in A; yellow colors indicate further outside the arena, whereas blue colors indicate closer to or inside the arena. Top, reaches sorted in the order in which they occurred (left), and the same reaches resorted into bouts, again in the order in which they occurred (right). Bottom, the same reach and bout data sorted according to reach and bout duration.

F. Top, binary variable indicating when the animal is reaching (paw is outside the arena), with arrows at the start of bouts. Bottom, time courses of 6 of the 14 kinematic variables extracted via keypoint tracking, during the bouts plotted above.

G. All reaches in one session for mouse 5, for the same 6 variables as in G. Individual reaches are plotted in transparent gray, the average (time-locked to reach onset) is plotted on top in red. Scale bar for paw and head position correspond to 30 pixels in x and y; the view is the same as is schematized in A and for example reaches in D. Scale bar for paw velocity, head velocity, paw PC1, and head PC1 are 100ms and 1 standard deviation of each variable.

Each animal develops its own unique reaching pattern, yet consistently uses a similar range of movements across days.

To investigate whether the neural coding for this behavior changes or remains consistent over time, we analyzed data from sessions on five consecutive days late in training (starting on or after day 20), a period during which animals demonstrate stable success rates (Supplementary Fig. 1, success rates across 5 days). The within-session variability observed on a given day persists across the five days of study, with the variability centered around very similar means (two example animals in Fig 2A). The specific digit, paw, and head movements employed differs across animals, as each develops its own strategy for retrieving the seed (see Supplementary Fig. 2 for kinematic descriptions across all five mice). In some cases, the strategy involves rhythmic, consistent reaches, with individual reaches limited to a narrower range of magnitude and duration. In other cases, there is greater variability in reach magnitude, duration, and timing within bouts (see the temporal patterning of the two example mice in Fig 2B). For

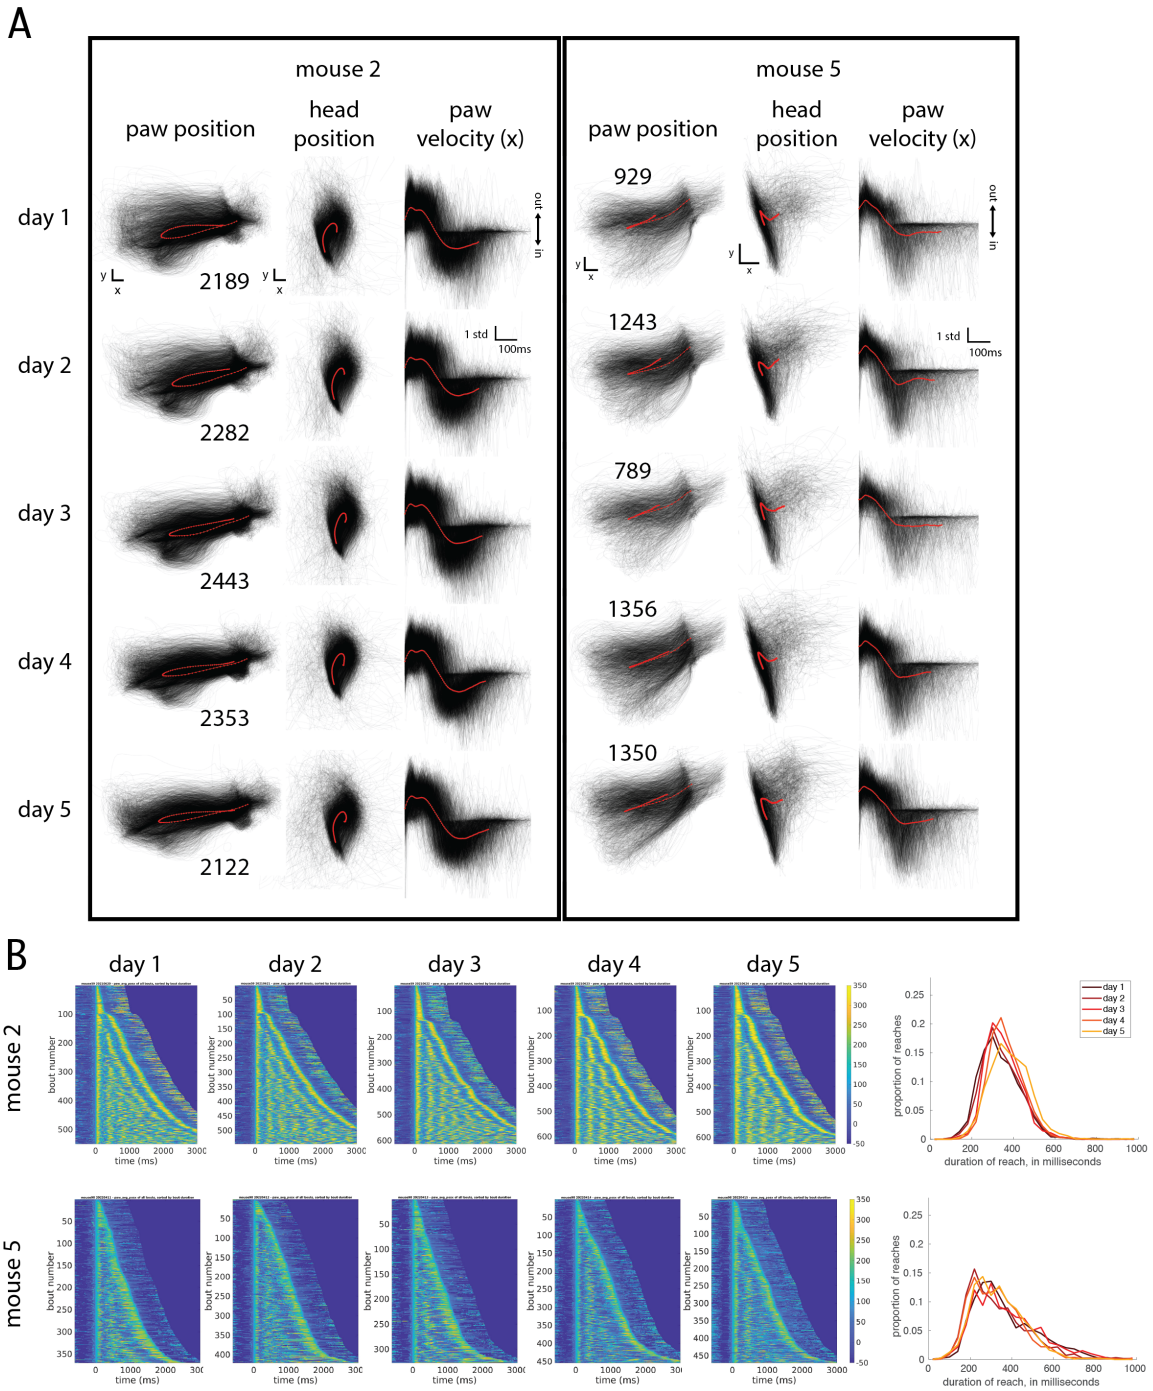

Figure 2: Each animal develops its own unique reaching pattern, yet consistently uses a similar range of movements across days.

A. All reaches in each of five sessions for two example mice, for 3 of the 14 kinematic variables extracted via keypoint tracking. As in Fig 1G, individual reaches are plotted in transparent gray, the average (time-locked to reach onset) is plotted on top in red. Numbers next to the paw position plots are the number of reaches in that session.

Figure 2 (continued):

B. For the same two example mice (mouse 2 and 5), for all five days: Left, paw position in the X dimension plotted the same way as in Fig. 1E. Right, histograms of duration of individual reaches.

each animal, the distributions of all measured kinematic properties and the corresponding success rates remain consistent across all five days. This stability is further validated by the covariate balancing procedure used for the linear models in later sections: after matching, the majority of reaching timepoints are retained across all days, indicating a similar joint distribution of kinematic variables over time (Supplementary Fig. 3; see Methods). This consistency in movement across sessions enables the use of this dataset to compare neural encoding of reach, grasp, and carry behaviors over multiple days.

Neurons in CFA M1 motor cortex modulate during reaching

We recorded the activity of populations of individual neurons in caudal forelimb area of mouse motor cortex (CFA, also referred to as forelimb M1) using an Inscopix miniscope to perform one photon calcium imaging in freely moving (non-headfixed) mice. We combined the miniscope with a prism to achieve an imaging plane perpendicular to the cortical surface (Andermann et al., 2013; Resendez et al., 2016). This setup enabled simultaneous imaging of putative cortical layers 2/3 and 5a, capturing hundreds of cells in a single field of view (244 +/- 50 neurons [138 347] (mean +/- st dev, range) per mouse per day). By consistently returning to the same field of view each day, we were able to identify and longitudinally track the activity of many of the same cells over multiple days.

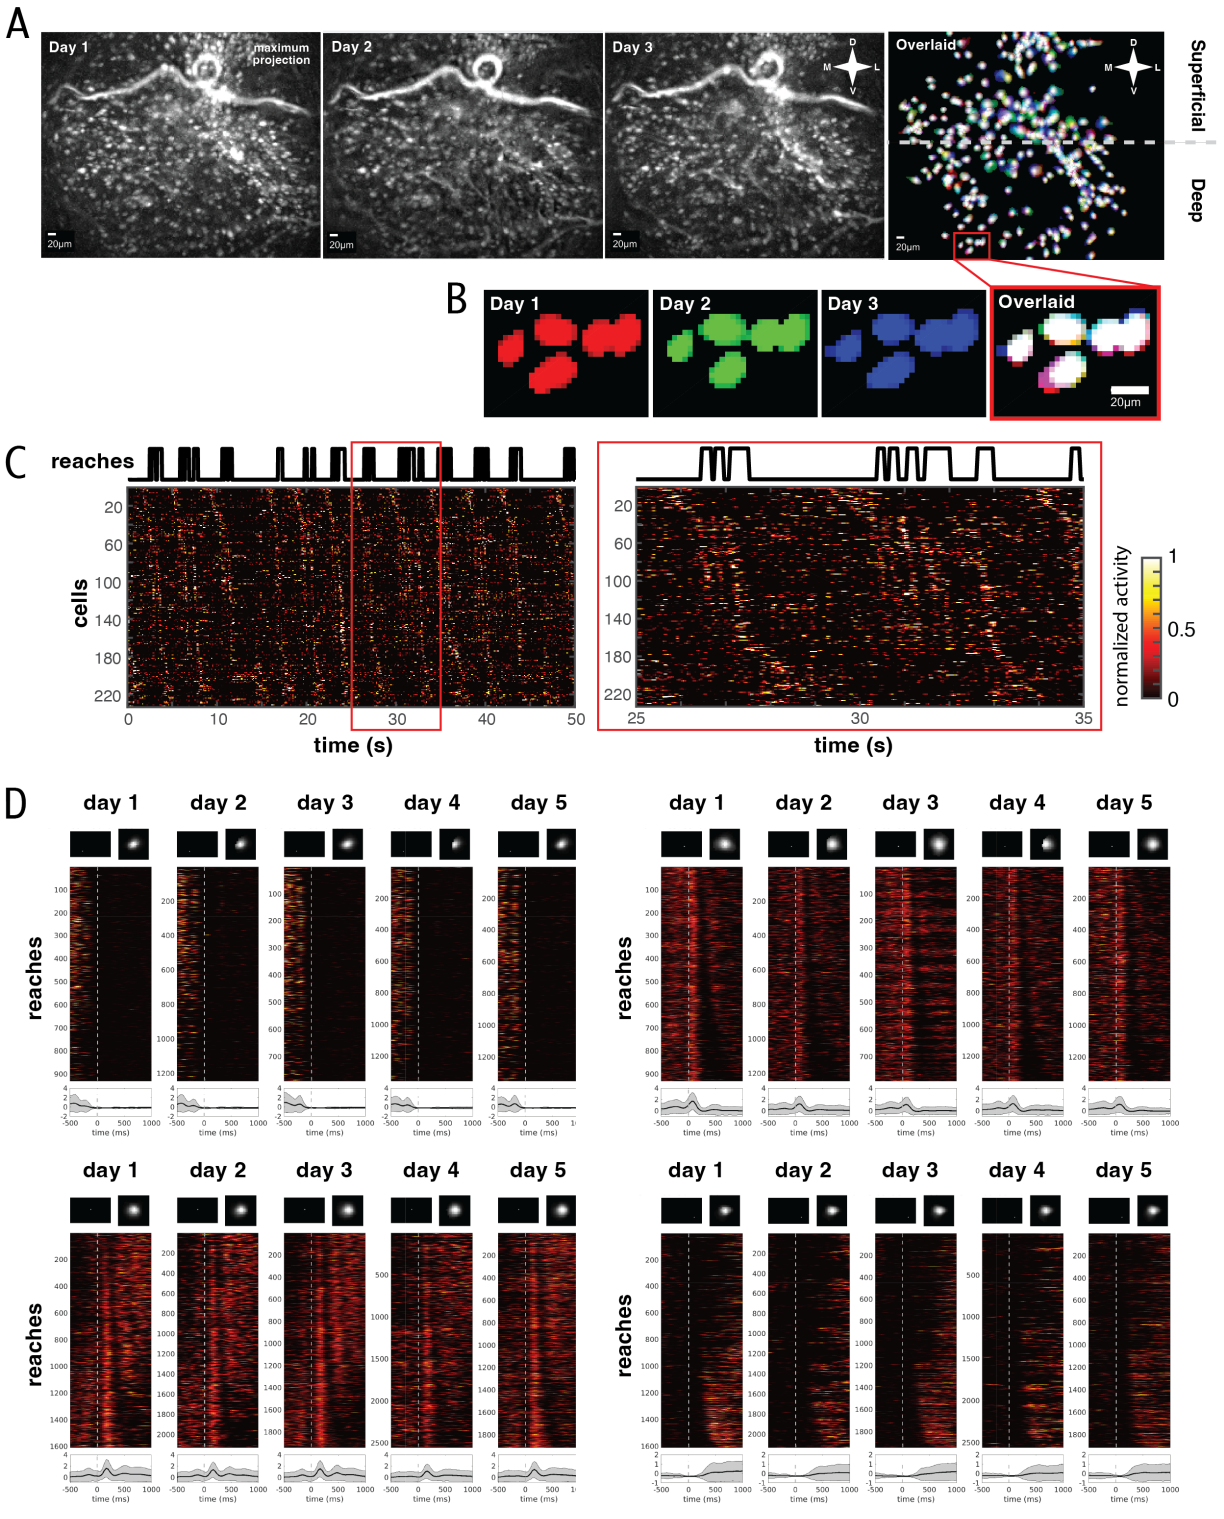

Figure 3: Neurons in CFA M1 motor cortex modulate during reaching

A. Far left, middle left, middle right: Maximum projections of the motion-corrected calcium imaging movie of the same field of view (mouse 4) on three subsequent days. Compass indicates dorsal-ventral and medial-lateral. Far right: overlay of footprints of individual cells identified on all three days, with day 1 in the red channel, day 2 green, and day 3 blue.

Figure 3 (continued):

B. Same as in A, but zoomed in on five cells found on all three days.

C. Top: Binary variable indicating when the mouse was reaching. Bottom: Heatmap of inferred spiking activity of all recorded cells from one field of view for 50 seconds of recording. Right inset: Temporal zoom to 10 seconds of recording. Each cell's activity has been normalized separately within this window: minimum value = 0, mean + 8*standard deviation = 1.

D. Single-trial inferred spiking activity of four example cells found on all five days. For each day, the top panel shows the spatial footprint of the cell (right) and its location within the field of view (left). The middle panel shows a heatmap of the single trial activity. Normalization is the same as in C. Trials are sorted according to the duration of reaches, as for the kinematics in Fig 1E: short reaches at the top, long reaches at the bottom. The bottom panel shows the average response across all trials on that day (PETH) as mean +/- st dev.

The majority of neurons exhibited significant modulation during reaching, as determined by an ANOVA comparing neural activity during reaching with that during other behaviors, including locomotion, grooming, food manipulation, and food consumption (5544 out of 6105 cell-days (90.8%), where each cell-day represents a single day's observation for a given neuron; $p = 0.05$ with Bonferroni correction for the number of cells in each field of view for that day; see Methods). Across the recorded population, neurons displayed diverse temporal response profiles, with activity observed before, during, and/or after reaching (Fig. 3C). Additionally, cells were active during the intervals between reach bouts, when the animal moved to the back of the arena to trigger the next seed delivery. In this analysis, each reach was treated as a trial. Neural activity showed notable variability across individual trials (Fig. 3C, D), reflecting the variability in the animals' movements. Trials are sorted on each day by reach duration, as in Fig 1E bottom left. Not all neurons were active during every reach, and even when a neuron was active, its activity magnitude varied significantly. Despite this variability, neurons demonstrated stable firing patterns across days (Fig. 3D), both in single-trial activity (middle) and in their trial-averaged response (bottom).

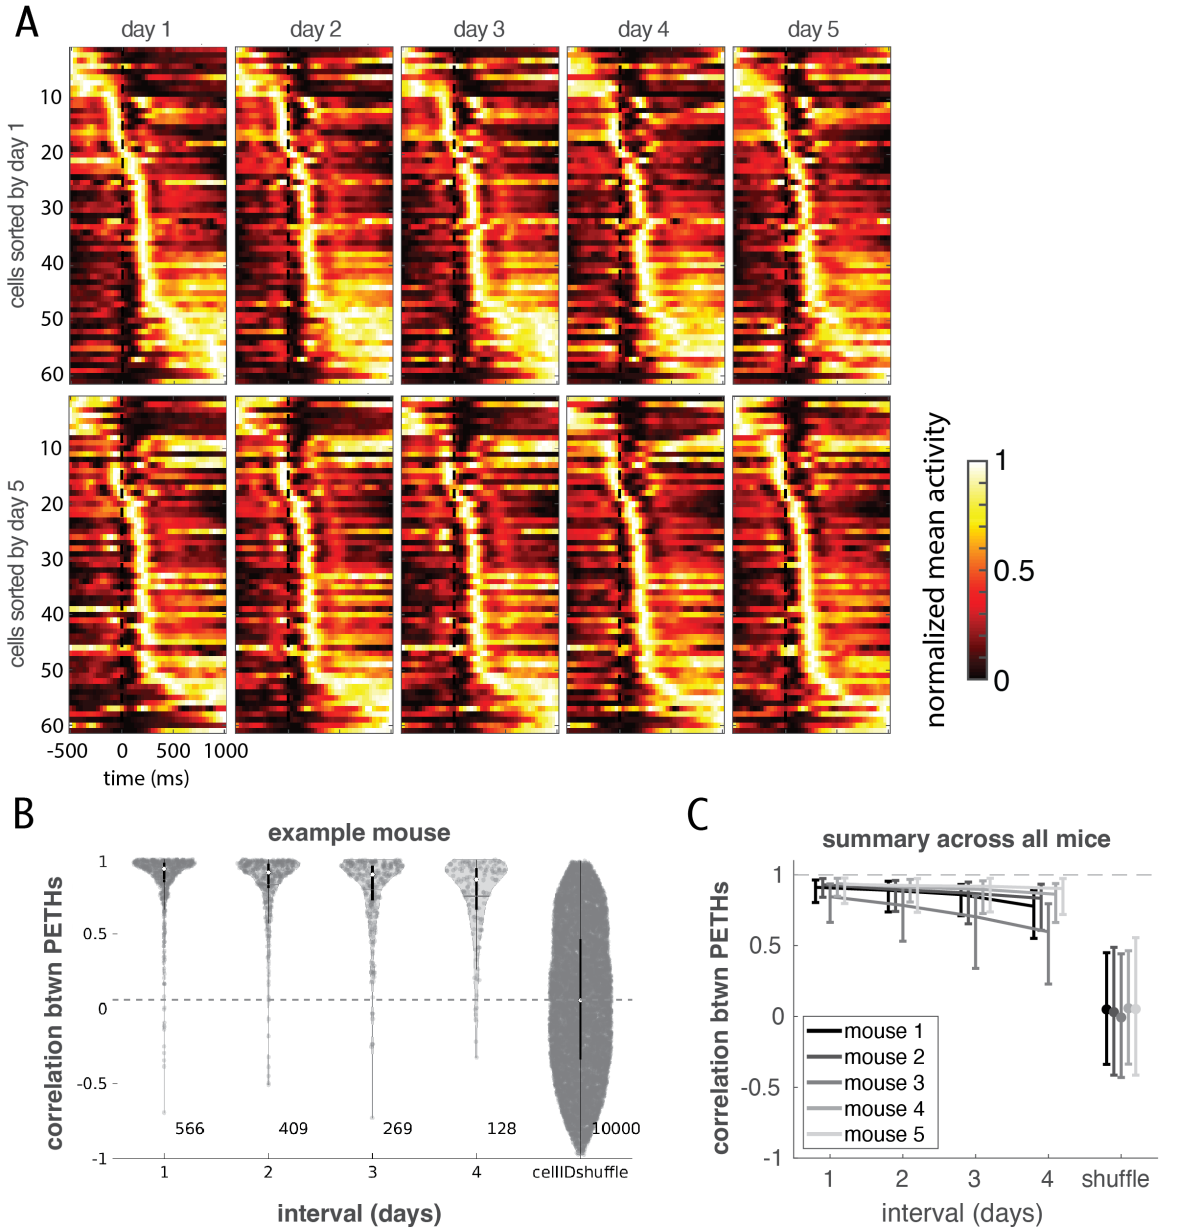

Figure 4: Reach-locked average neuronal activity patterns (PETHs) are stable across days

A. Normalized mean activity (PETHs) of neurons registered across all five days for an example mouse (mouse 4). Cells are sorted according to the time of their peak activity on either day 1 (top row) or day 5 (bottom row). Each cell's activity has been normalized separately within this window: minimum value = 0, maximum value = 1.

B. Correlation between PETHs for cells found on more than one day, as a function of the interval between the days, for an example mouse (mouse 4). Each datapoint is an observation of the same cell on two different days. Numbers to the bottom right of the violins indicate the number of observations included in that violin. For each violin, the white dot indicates the median of the distribution and the black box indicates interquartile range. Cell ID shuffle indicates the distribution of correlations between PETHs of two cells randomly chosen from any day (excluding any cell pairs that are registered as the same cell; see Methods). Gray dashed line indicates the median of this shuffle.

C. Correlation between PETHs for all mice. Datapoint is median, error bars are interquartile range across cells for that interval for that mouse.

Reach-locked average neuronal activity patterns (PETHs) are stable across days

We quantified the average response of each neuron using peri-event time histograms (PETHs) time-locked to the start of the reach. The variety of response timecourses observed in single trials (Fig. 3) was also evident in the PETHs (Fig. 4A). While many neurons showed peak activity during reaches, others were most active either before or after the reach. This suggests that these neurons may be modulated by other task components associated with the reach behavior, such as positioning the body at the reach port, locomotion to trigger the next seed delivery, or consuming seeds.

We found that, across the population, PETHs were highly similar across days. This is demonstrated visually by showing the PETHs for all cells tracked across all five days for an example animal (Fig 4A), sorted according to peak activity modulation on either day 1 or day 5. Regardless of sorting day, the rank order is maintained across all other days. To quantify similarity across days, we computed the correlation between each longitudinally tracked neuron's PETH on different days. For this analysis, we use the time window -100ms to +500ms around reach onset, which captures most reaching timepoints while excluding the majority of non-reaching timepoints. For each animal we examined the distribution of these correlation values as a function of the interval between days and compared it to a control distribution of correlations between PETHs for non-same cells (Fig 4B; see Methods for details of control distribution). We observed that the distribution of correlation values for each interval was clustered near the maximum correlation value of 1 and was far from the median of the control distribution, which was close to 0 (example animal in Fig 4B; all animals in Fig 4C). In other words, neurons are more similar to themselves over days than to other neurons. This indicates that the time course of

the average responses during and around reaching remained stable in this population across an interval of five days. These findings are consistent with observations for other behaviors (Katlowitz et al., 2018, Jensen et al., 2022), and suggests that encoding for dexterous movements in rodent motor cortex is also stable.

Trial-by-trial variance in forelimb, digit, and head movements enables accurate modeling of single-trial neuronal activity in CFA M1

Paw and head movements during reaching exhibit variability across multiple kinematic dimensions such as speed, duration, the timing of wrist pronation and of the digits opening and closing, as described in Fig 1 & 2. While PETHs quantify a neuron's response for the average movement, they do not reveal whether a neuron encodes the significant trial-by-trial variance in this behavior. To assess the neural encoding of the moment-to-moment details of reach-to-grasp kinematics, we constructed families of linear models, referred to as neural encoding models, to predict the activity of individual neurons based on the extensive set of kinematic variables we extracted using keypoint tracking (schematized in Fig 5A).

It is important to note that it is possible for the mean movement to be encoded, while the trial-by-trial variance around the mean is not. This would suggest that there is a general "reach" signal in motor cortex (indicating whether the animal is performing this behavior or not), without encoding the moment-to-moment details of the constituent movements during that behavior. While individual neurons exhibit significant trial-by-trial variance around their average responses (Fig 3D), it is possible that that variance does not co-vary with behavior. If movement

details are indeed encoded in motor cortex within a given session, it is also possible that the nature of that encoding is not the same across days, indicating instability (drift).

We first applied a covariate balancing subsampling procedure to the predictor sets for all five days, individually for each animal (see Methods and Supplementary Fig. 3). This step ensured that we compared the encoding of similar movements, submovements and their corresponding variability across days.

The majority of neurons on any given single day have a substantial amount of variance explained by the set of kinematic predictors compared to a timepoint-shuffled control. Specifically, 4883 out of 6105 cell-days (80.0%) showed significant variance explained, where each cell-day represents a single day's observation for a given neuron. This analysis included data from all days and all mice ($p = 0.05$, with Bonferroni correction for the number of cells in each dataset/field of view for that day; see Methods for details of the shuffle comparison).

This result alone should not be interpreted as definitive evidence of trial-by-trial variance encoding. All linear models in this study are timepoint-by-timepoint models, which treat each timepoint as an independent observation and regress across all observations in a session. This approach is standard for quantifying the neural encoding of movements, but it conflates encoding for the mean movement with encoding for the variability around the mean. We have already observed that individual cells exhibit time-varying average responses (PETHs in Fig 4A), which can lead to significant predictions in timepoint-by-timepoint models even in the absence of coding for variance around the mean. In other words, the time-varying average movement could explain neural activity above and beyond the timepoint shuffle control. To test for this possibility, we built another model for each cell, in which we predict the cell’s single trial

activity from trial-averaged movements (schematized in Fig 5B). In the predictor matrix of this trial-average model, we replace the timepoints of each reach with the reach-locked average, for each kinematic predictor separately. Approximately half of all recorded cells were significantly predicted by the trial average model (3097 out of 6105 cell-days, or 50.7%, where each day for a given cell is a separate observation (cell-day), using cells from all days for all mice; using the same shuffle comparison as for the single trial full model, see Methods for details), corresponding to 50.5% of each dataset (50.5% +/- 13.0% [17.4% 71.5%]; mean +/- std [range]). Most neurons with significant predictions in the trial-average model were also significantly predicted by the single-trial model (2972 out of 3097 cell-days, or 96.0%). Very few neurons were significantly predicted by only the trial-average model, indicating they encode only a general reach signal (125 out of 6105 cell-days, or 2.0%). In contrast, many neurons showed significant predictions in the single-trial model but not in the trial-average model, suggesting that they encode moment-to-moment details but not a general reach signal (1911 out of 6105 cell-days, or 31.3%).

Moreover, of the neurons significantly predicted by the single trial model, we found that the vast majority were better modeled by the single trial model than by the trial average model (Fig 5C-E; 4522 out of 4883 (92.6%) cell-days with significant single trial full model predictions, from all days for all mice). This finding suggests that these neurons encode single trial variance around the average reach. These results indicate that mouse motor cortex encodes not only the general “reaching” signal observed in the PETHs and trial average predictions but also the moment-to-moment details of movement. For the remainder of our analysis, we focus on the subset of neurons with significant single-trial model predictions. On average, this subset

comprises 79.4% of neurons in each dataset (79.4% +/- 16.3%, [46.8% 94.2%] (mean +/- sd, [range] across datasets, n=25 datasets), see Supp Fig. 4B,C for proportion and counts by mouse and dataset).

The R squared value from the single trial model represents the maximum amount of variance that the entire predictor set can explain for each individual neuron. Next we consider the contribution of individual predictors. Due to the nature of the task and the physics of the animals' bodies, many of our kinematic variables exhibit non-zero correlations (Supplementary Fig. 4A). Consequently, the coefficients in our single trial linear model are not unique and should not be interpreted as the maximum contribution of each individual predictor (Stevenson et al., 2018). A common approach to allow direct comparison of the coefficients is to orthogonalize the predictors to remove any correlations between them. However, this approach biases the analysis towards predictors that appear earlier in the design matrix, potentially underestimating the explanatory power of predictors that come later. To avoid this, we computed two metrics to quantify the contribution of individual kinematic predictors to explain a neuron's variance (schematized in Fig. 5A): the variance explained by each predictor separately (single variable R-squared values) and its unique contribution (delta-R-squared values). 1) Single Variable Model: For each predictor, we constructed a single-variable model using only the predictor in question. The R-squared value from the single-variable model reflects the maximum variance that that predictor can contribute to the full model. 2) Leave-One-Out Model: Starting with the full model, which predicts a neuron's activity using all kinematic predictors, we created a leave-one-out model by excluding one predictor at a time and refitting the model. The drop in R-squared when

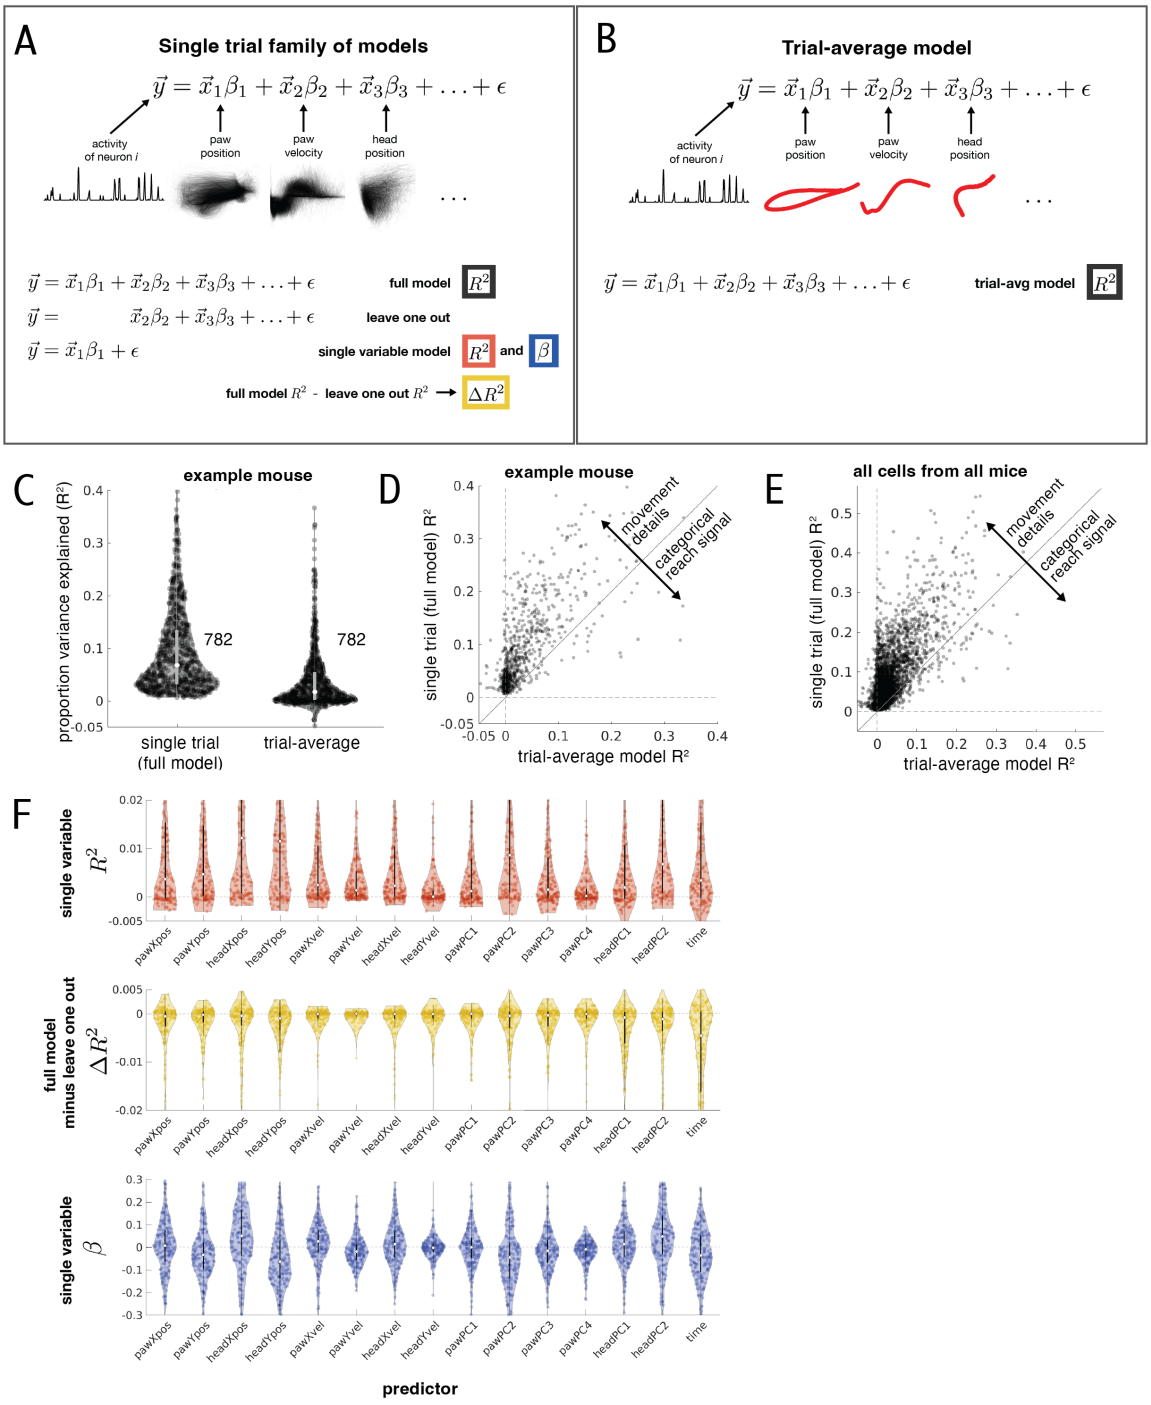

Figure 5: Trial-by-trial variance in forelimb, digit, and head movements allows for accurate modeling of single-trial CFA M1 neuronal activity

A. Schematic of trial-to-trial linear model families built to predict the activity of individual cells from kinematic features during reaching. The full trial-to-trial model contains all kinematic features and their observed values across trials. For each predictor, we build a leave-one-out model using all predictors except the one in question, and a single variable model, using only that predictor. The change in variance explained from the full model to the leave one out model indicates unique variance explained by that predictor.

Figure 5 (continued):

B. Schematic of the trial-average linear model, where each trial of the kinematics is replaced by the average timecourse for that kinematic predictor, averaged across reaches, locked to reach onset (see Methods).
C. Proportion of variance explained (R-squared) by the full single trial model or the trial average model for all cells with a significant full model prediction, for an example mouse (mouse 5). Numbers to the right of each violin indicate the number of observations included in the distribution. Each datapoint corresponds to a cell on one day. Cells registered across days have multiple dots if they have significant full trial-to-trial model predictions on multiple days. For each violin, the white dot indicates the median of the distribution and the black box indicates interquartile range.
D. Same data as in C, but with the full model and trial average model R squared values plotted against each other for each cell.
E. Same as D, but for all cells with significant full model predictions for all days for all mice.
F. Top: coefficients from single variable model fits for all cells in one example session. Each datapoint corresponds to a cell; each cell has one datapoint in each violin. Middle: same as top but for the variance explained (R-squared) for single variable model fits. Bottom: same as top but for the change in variance explained (delta R-squared) for full models to leave-one-out models for each predictor.

a predictor is removed, referred to as the delta-R-squared, indicates the unique variance that the predictor explains for that neuron. Together, these metrics provide a comprehensive view of both the unique and total explanatory power of individual predictors in modeling neuronal activity. Additionally, the coefficients from the single variable models can serve as tuning curves, revealing how the neuron's activity is modulated in relation to variations in the kinematic predictor.

Many of the predictors exhibit nonzero single variable R squared values and delta R squared values (Fig 5F top, middle), indicating that these movement variables are encoded by single cells in motor cortex. Additionally, distributions of single variable model coefficients (betas) are dispersed away from zero (Fig 5F bottom). Notably, these observations apply not only to paw-related predictors but also to head-related predictors. Head and paw variables are correlated (Supplementary Fig. 4A), which can account for some of the single variable R-squared values. However, several head variables have non-zero delta R-squared values for many neurons, indicating that the position and orientation of the head are encoded in forelimb motor cortical activity in addition to paw-related predictors. These findings are further evidence that

mouse motor cortex encodes both a general "reach" signal (observed in the PETHs and the trial average model predictions) and the moment-to-moment details of the reaching movements. These details include not only the position of the paw but its shape and orientation (posture of the digits) and the position and orientation of the head.

Linear model encoding profiles of single cells are stable across days

To assess whether the encoding of trial-by-trial variability is consistent across days, we compared components of the model fits across each day on which a neuron was recorded. A common approach to quantifying the stability of activity-behavior relationships is to assess how well a model trained on one day performs when applied to data from other days, often as a function of the number of days between the training and test sessions (e.g., Driscoll et al., 2017). When predictors are correlated, as in our study, there is an infinite set of possible coefficient combinations that can produce the same optimal fit (Stevenson et al., 2018). If the correlations between predictors, or between predictors and omitted (but relevant) variables, change even slightly from the training day to the test day, the coefficients derived from the training day may no longer accurately reflect the predictor relationships on the test day. This occurs even if the underlying relationships of interest—such as the link between individual predictors and the predicted variable (e.g., a neuron's fluorescence)—remain stable. This would result in a drop in performance of the model on the test day. If predictor correlations are not carefully accounted for, applying fitted models across days may mistakenly suggest that activity-behavior relationships for individual predictors are changing, when in fact, it is the relationships among the predictors that are shifting.

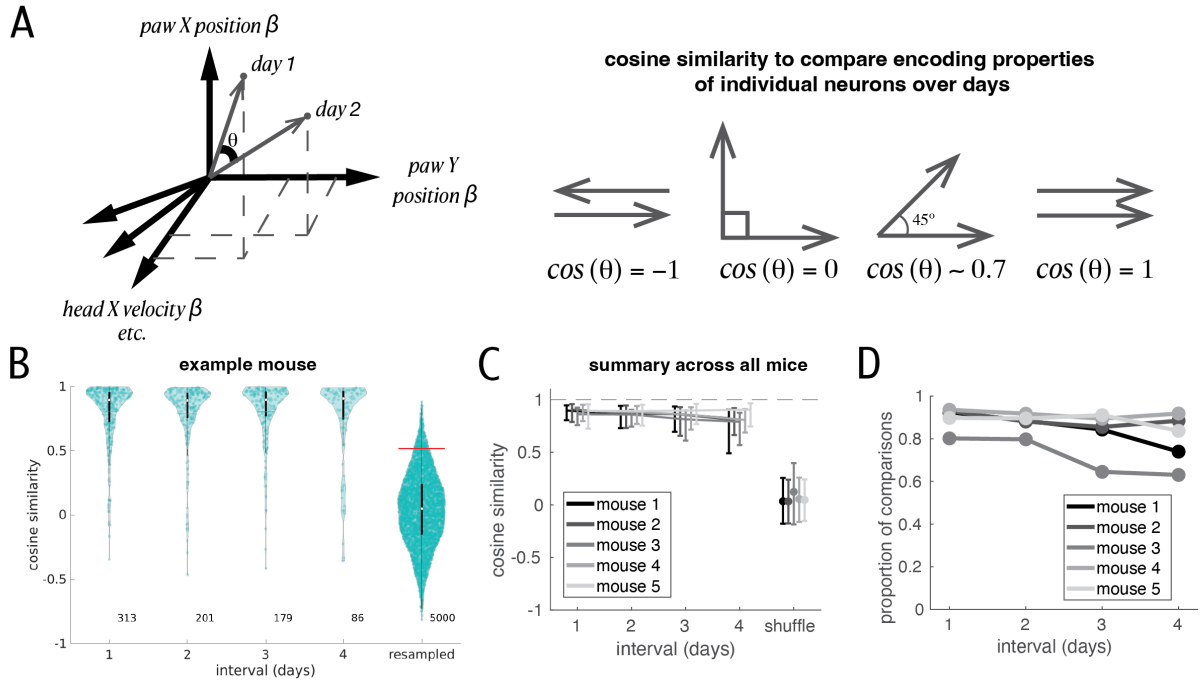

Figure 6: Linear model encoding profiles of single cells are stable across days

A. Schematic of method: using cosine similarity to compare vectors of encoding coefficients (betas) over days for individual cells

B. Cosine similarities of coefficient vectors as a function of the interval between observations for an example mouse (mouse 5). Each datapoint corresponds to a comparison of observations of a single cell on two different days. Numbers to the bottom right of the violins indicate the number of observations for that interval. For each violin, the white dot indicates the median of the distribution and the black box indicates interquartile range. Resampled shuffle control built by generating pseudo-vectors where each coefficient is drawn from the pool of coefficients observed across all neurons for that predictor within that day. Coefficients within the same pseudo-vector are drawn independently. The resampled shuffle distribution contains the same number of comparisons for each of the 10 possible pairs of the 5 days of study (5000 total, 500 per day pair). Red line indicates the 95th percentile cutoff of this distribution.

C. Same as B, but for all animals. Data is median, error bars are interquartile range across cells for that interval for that mouse.

D. Proportions of coefficient vector comparisons for a given interval that are statistically significant using the resampled shuffle distribution as a control.

To compare the contribution of the predictors across days, we instead compared the collection of single variable model coefficients (betas), referred to as a neuron's encoding profile. Specifically, for a given cell, we concatenated the coefficients for all predictors from one day into a vector and calculated the cosine similarity between a neuron's vectors from different days (Fig 6A, left). Cosine similarity is simply the cosine of the angle between the two vectors and

ranges from -1 to 1, indicating antiparallel and parallel vectors, respectively (Fig. 6A, right). This approach compares not the exact value of each coefficient over days, but instead the relationships between the coefficients for all predictors. Thus it provides a holistic comparison of a neuron's tuning on different days. Importantly, we are comparing the coefficients of the single variable models, rather than the coefficients of the full model, for all the reasons described in the previous paragraph. We only consider the stability of encoding of cells with significant full single trial model predictions; for cells with significant predictions on some days but not all days, we only compare encodings from days with significant predictions. We found that across intervals of up to four days apart, cosine similarity values are high for the vast majority of comparisons of same-cell encoding profiles (example mouse in Fig. 6B, all mice in Fig. 6C). Of the 4926 comparisons across all intervals for all mice, 4354 were significant (88.4%) relative to their respective resampled control distribution (Fig. 6D for breakdown by mouse and interval; see Methods for details of control distribution). We also performed the same comparison with encoding profiles made up of the collection of R-squared values (single-variable model R-squared values, full model R-squared, and delta R-squareds) and again observe significance for the majority of comparisons (3222 out of 4926 (65.4%); see Supplementary Fig. 5). These findings demonstrate that the encoding of moment-to-moment movement details in motor cortical activity remains largely stable at the single-neuron level across multiple days.

2.3 Discussion

Here, we investigated how the motor cortex encodes moment-to-moment movements of the paw, digits, and head during a self-initiated, skilled reach-to-grasp task, and whether this

encoding changes over days. Consistent with prior studies on the stability of motor systems (Jensen et al., 2022; Katlowitz et al., 2018; Liberti et al., 2016), our findings demonstrate that average responses (PETHs) in M1 remain highly stable over time. We extend previous work by recording motor cortex activity during the performance of a challenging motor skill characterized by many repetitions within a day and high reach-to-reach variability, features that make the skill more reflective of real-world motor behaviors. We find that neurons in the mouse motor cortex encode both general reaching signals and the detailed, moment-to-moment movements of the head, forelimb, and digits during reach-to-grasp behaviors. This encoding remains remarkably stable over days, providing a reliable neural foundation for producing skilled movements consistently over time.

Quantifying neural encoding during naturalistic behaviors with high trial-to-trial variability poses significant challenges which we overcame through a series of carefully designed approaches. First, we used linear encoding models that assume a consistent linear relationship between movement variables and neural activity within a day, rather than relying on assumptions of stereotyped movements across trials. This analytical flexibility allowed us to account for the inherent variability of naturalistic behaviors. Second, we designed the task to balance variability with sufficient trial counts. Key elements of the task, such as requiring animals to reposition themselves after each bout of reaches, randomizing seed positions across multiple locations, and omitting penalties for failed attempts, encouraged trial-to-trial variability while preventing the development of stereotyped automatized movement patterns. The task's inherent difficulty likely further contributed to the diversity in movements. Third, we captured these movements at a fine-grained resolution using high-speed, high-resolution cameras,

allowing us to precisely quantify trial-to-trial variability above the noise threshold. This detailed behavioral data was crucial for accurately identifying variations in movement. Finally, to ensure that observed changes in encoding were not confounded by behavioral shifts over days, we implemented a covariate balancing procedure before fitting linear models. This approach allowed us to quantify encoding within consistent regions of the high-dimensional kinematic parameter sets across days, effectively isolating neural changes from behavioral variability. Together, these strategies enabled us to rigorously estimate and compare neural encoding across days in a statistically robust manner. Our findings not only highlight the stability of motor cortex encoding during complex, variable behaviors but also underscore the value of incorporating detailed behavioral analyses into neural encoding studies. We hypothesize that similarly rigorous approaches to quantifying and incorporating behavior into encoding models, including movements not explicitly instructed by the task (Sadeh and Clopath, 2022; Musall, Kaufman et al., 2019), could reveal stable encoding throughout motor circuitry and potentially in sensory and association brain areas as well.

The peri-event time histogram (PETH) captures the time-varying average response of a neuron across repeated instances of a specific behavioral or task-related event. In this study, the PETH serves as a categorical descriptor of "reaching," since we aligned neural responses to the onset of each reach. Our analysis reveals a consistent spatiotemporal pattern of neural activation (PETH) during reaching that was conserved across days. For neurons reliably tracked over time, the vast majority exhibited stable average responses, consistent with findings in rats (Jensen et al., 2022) and zebra finches (Katlowitz et al., 2018; Liberti et al., 2016). The observed stability in the PETH for a task requiring dexterous digit control in mice under naturalistic conditions

reinforces the idea that the rodent motor cortex can sustain stable encoding of learned movements, irrespective of task complexity. Collectively these findings strongly suggest that stable behaviors arise from the consistent contributions of individual neurons within the motor circuit.

While the PETH provides an estimate of a neuron's mean activity during reaching, it does not account for trial-to-trial variability in the behavior, which could offer additional insights into the dynamics of neural encoding. We used single-trial encoding models to investigate whether the significant trial-to-trial variability in movements is encoded in the motor cortex and whether this encoding remains stable over days. These models linked the activity of individual neurons to a carefully curated set of 14 kinematic variables, capturing the movements of specific body parts (e.g., position and velocity of the paw and head, as well as the conformation of the digits and the orientation of the paw and head). While more complex models could likely explain additional variance, they often come at the cost of interpretability. Despite using straightforward linear approaches, we found that moment-to-moment movement details during reach-to-grasp explained a statistically significant fraction of the variance in most recorded neurons. Furthermore, these relationships remained stable across multiple days, representing the longest duration over which such detailed single-cell encoding of moment-to-moment movement has been examined.

The finding that the motor cortex encodes single-trial variability in reach-to-grasp movements suggests a role in the online monitoring and control of this sensory-driven behavior, operating as part of broader sensorimotor loops involving other brain regions (Sauerbrei et al., 2020; Wagner et al., 2019). The linear models showed that the principal components of paw

markers, which reflect wrist and digit states, contributed explanatory power comparable to paw position, which reflects the state of proximal joints such as the shoulder and elbow. Although one causal manipulation study suggests that the rostral forelimb area (RFA) is more involved in grasping than the caudal forelimb area (CFA) (Wang et al., 2017), our finding of digit encoding supports CFA involvement. Further investigation is needed to clarify how digit-related signals in CFA and RFA might work together to coordinate grasping.

Interestingly, the position and orientation of the head also contributed explanatory power comparable to paw- and digit-related predictors in the encoding model. This could reflect two potential (and not mutually exclusive) explanations: first, head-related variables may be correlated with omitted kinematic variables that are more directly linked to neural activity (Stevenson et al., 2018). Alternatively, head position and orientation may themselves be directly relevant to computations in the motor cortex, providing crucial information about body positioning relative to the reach port and pedestal. This latter explanation implies a role for egocentric planning signals in guiding the paw to the pedestal, as observed in encoding of whole body posture in rats (Mimica et al., 2018). These findings highlight the importance of examining how multiple body parts are encoded across various brain areas, rather than limiting the focus to regions traditionally associated with specific movements. A comprehensive understanding of whole-body movement generation requires this broader perspective. Furthermore, to uncover how the brain produces the flexible, context-dependent, and coordinated movements essential for daily human activities, it is crucial to study comparable behaviors in animal models. The reach-to-grasp behavior analyzed here exemplifies such complexity, combining whole-body posture with precise forelimb and digit control. This task closely parallels primate reaching and grasping

behaviors (Sacrey et al., 2009; Whishaw et al., 1992; Guo et al., 2015; Becker et al., 2020; Whishaw et al., 2017) and serves as a powerful model for investigating the neural mechanisms underlying these intricate actions.

2.4 Methods

Mice

All experimental procedures were approved by the University of Chicago Institutional Animal Care and Use Committee and were performed in compliance with the Guide for the Care and Use of Laboratory Animals. Data were collected from Ai148D mice of either sex (n = 2 female, n = 3 male; Jax strain 030328). The line was originally obtained from Jackson Labs and bred in house. Mice were between 10 and 12 weeks old at the start of training.

Surgical procedures

Animals underwent three surgical procedures: AAV-Cre injections, prism implantation, and baseplate implantation. For all surgeries, anesthesia was induced with isoflurane (induction at 4%, maintenance at 1–1.5%). During AAV injection, two small burr-hole craniotomies were drilled to allow insertion of a micropipette and injection of virus. Injections were made at three depths (650, 400, and 200um below the surface) and in two locations targeting caudal forelimb area (0.25mm anterior and 1.25mm lateral to bregma, and again at 0.25mm anterior and 1.75mm lateral to bregma). After injections, the burr holes were sealed with dental cement and the skin sutured. Animals recovered for a minimum of seven days before prism implantation. During prism implantation, a 2mm diameter craniotomy was drilled, centered at 0.25mm anterior and 1.5mm lateral to bregma. A stereotaxically mounted scalpel was used to make a 1mm wide incision running along the medial lateral axis, centered at the same coordinates as the craniotomy; the incision ran to a depth of 800um. The prism was then inserted into the tissue along this incision to the same depth, with the angled part of the prism directed posteriorly,

meaning the resulting field of view is anterior to the site of insertion. The remaining surface of the tissue was covered with Kwil-sil and the prism was cemented in place. Animals recovered for a minimum of seven days before baseplate implantation. During baseplate implantation, a plastic baseplate was cemented above the prism and grin lens, placed to allow the best alignment of the microscope with the focal plane of the prism. This final procedure is quick and non-invasive. Animals recovered for a minimum of two days before the start of food restriction and habituation in the arena.

Behavioral training

The entire behavioral experience consisted of three parts: habituation, shaping, and training. During all phases, the mice were briefly (~1min) anesthetized with isoflurane so that the miniscope could be attached to their headplates. Mice were then placed in the arena and allowed to recover completely (defined by showing unencumbered, normal ambulatory movement), at least 5 minutes. The arena was 15cm x 15cm with a reach port (6mm wide opening that ran from the floor up ~8cm) located in the center of one of the walls. The wall with the reach port was slanted at a 60 degree angle relative to the floor to enable the mice to maneuver better with the miniscopes on their heads.

Food restriction, Habituation and Shaping

At least one full day after headplate placement, mice began food restriction, in which they received a limited amount of pellet food, titrated to maintain their weight between 85% and 90% of their baseline weight. Mice were food restricted for at least one week before habituation

began, sometimes longer in order to achieve a weight between 85% and 90% of baseline. Mice were single housed throughout food restriction, habituation, and training. Mice were habituated to the training arena and the miniscope for at least seven days before shaping began. During habituation, mice were placed in the behavioral arena for twenty to thirty minutes, once per day. At the start of habituation, twenty millet seeds were placed on the ground inside the arena right next to the reach port, to encourage an association between the reach port and food.

During shaping, mice were again placed in the behavioral arena for twenty to thirty minutes, once per day. Instead of being presented inside the arena, seeds were presented in two places outside the arena: approximately ten seeds were presented on the ground outside the arena immediately in front of the reach port. These were reachable by the animal's tongue. Seeds were also presented in a low square trough outside the arena near the reach port. The trough was positioned a few millimeters laterally to the center of the reach port, such that the animal would only be able to successfully reach the trough with its contralateral forelimb. All mice were shaped and trained to use their left paws. The goal of shaping was to train the animal to associate reaching its forelimb outside the arena with getting food. Functionally, during shaping, the mice would scoop or knock seeds out of the trough and onto the ground, then lick them up with their tongues. Any large magnitude forelimb movement through the reach port was considered a reach. In order to move from shaping to training, mice needed to make at least 50 reaches with their left forelimb in a window of twenty minutes, and at least 80% of their reaches in that window needed to be with their left forelimb. This shaping procedure is based on that described in Chen et al., 2014. Shaping was performed daily until mice achieved these criteria. A given shaping session

was not stopped when animals reached these requirements, and often on the last shaping session before training, mice reached 75-100 times with their left forelimbs. If a mouse achieved this criteria in a shaping session, training began the next day. If, on the first training day, the animal did not show interest in reaching for the seeds on the pedestal, or reached very few times, the training session was converted to another shaping session, to reinforce the association between using their forelimbs and getting seeds. The animal was held to the same requirements as before for moving on to training.

We found that on the first day of shaping there was a critical moment requiring manual involvement. Sometimes, the mice initially reach out of the port with the correct limb but do not make contact with the seed trough or do not make strong enough contact to knock any seeds out. In the shaping phase the mice are quick to give up reaching entirely if their initial reaches are unsuccessful. We started manually knocking seeds out of the trough the first one or two times that the animal reached outside the arena with their left paw. This was generally enough.

Training

During training, seeds were only presented to the mouse on the seed delivery pedestals, which could only hold one seed each. There were two different kinds of seed delivery methods used for the mice included in this study. Mice only ever experienced one seed delivery method throughout their training.

For two of the mice, there were three permanent seed delivery pedestals positioned in an arc outside the reach port. Their positioning was offset to the mouse's right such that the pedestals were not reachable with the right forelimb. Seeds were delivered manually by the experimenter. Mice were required to turn away from the seed port and walk towards the back of the arena in order for a new seed to be delivered, thus initiating a new "trial". Only one seed was placed per trial. The order of seed placement was random but counterbalanced such that there was approximately the same number of deliveries per pedestal in a given session (matlab code).

For one of the mice, the same three permanent seed delivery pedestals were used, but seeds were only placed on one of the pedestals (the center one, equivalent to the lateral pedestal in Fig 1C).

For the final two mice, there was a single mobile seed delivery pedestal which was moved to two different positions throughout the session by a linear actuator. Their positioning was offset to the mouse's right such that the pedestals were not reachable with the right forelimb. For these mice, seeds were delivered automatically: the pedestal was lowered into a large trough of seeds and then lifted again, having caught a seed (see website for details on the automated seed delivery machinery). The order of the two delivery positions was random in a given session (python arduino code), a Bernoulli random variable. To trigger the delivery of another seed, mice had to break an infrared beam at the back of the arena.

For the manual seed delivery, if mice turned away to trigger another seed delivery, but the seed was present, the seed was left in position three times before that seed was removed and a new

seed was placed. For the automated seed delivery, there was no timeout for unsuccessful reaches and no requirement that the seed be removed from the pedestal before triggering another seed delivery.

On a given training day, the mice were in the arena for between 75 and 90 minutes. They performed the task in chunks, where each chunk consisted of 3-5 minutes of seed presentations followed by 2-4 minutes of break (no seeds delivered regardless of behavior). There were 8-12 chunks per session. In total, each mouse reached (performed the behavior, was allowed to have seeds delivered) for up to 40 minutes per day. Training was performed at approximately the same time each day for a given mouse. Mice were trained for a minimum of 24 days. Imaging took place every training day. For each mouse, five sessions from late in training were used in this study, generally starting between day 20 and 22.

Behavioral video recording and kinematic reconstruction

Videography and keypoint extraction

We took high speed, color video of the mouse during reaching, at a frame rate of 200Hz, using a Flir 16S2C Blackfly camera. All kinematic data used in this paper is captured in 2D, with a camera that is at an approximately 30 degree angle from head on, offset to the mouse's left to best view the left forelimb and paw. Additionally, the camera view is angled approximately 20 degrees from above. We used DeepLabCut to extract a variety of kinematic features from this video data (Mathis et al., 2018; Nath et al., 2019). After extracting keypoints using DeepLabCut, kinematic data was further processed with custom MATLAB software. Points that were

identified as poorly labeled (because of large distances traveled from the previous timepoint, or because markers that are anatomically close together were too far apart) were linearly interpolated over. Data was then smoothed: missing frames were linearly interpolated over, we smoothed using a butterworth filter (4th order, low-pass filter with 50Hz cutoff), then removed the interpolated (originally missing) frames.

Alignment across days

The position of multiple non-moving features of the arena (referred to as static markers) were extracted from the videos and used to align kinematic data across days. We adjusted for translation only: the x and y position of a static marker was subtracted from all the kinematic variables. Generally this translation was minimal because the cameras were fixed in place relative to the arena. We also extracted the position of the pedestal and the presence of a seed for each reach automatically from the video data.

Calculating kinematic variables of interest

We calculated the centroid of the paw as the average position of seven markers (wrist, digits 1-3 bases and middles). The centroid of the head was calculated as the average of all three markers on the head (nose, front corner of the eye, and corner of the microscope). See Fig 1A and C for visualization of these markers. To calculate a low dimensional estimate of paw shape and orientation, we subtracted the paw centroid position from the seven markers on the digits and wrist, in x and y separately, then performed pca on these 14 dimensions (seven markers, in x and y). We projected data onto the first 4 dimensions, then included it in the linear model. We

performed the same procedure for the three markers on the head to capture an estimate of head orientation, keeping 2 principal components of the 6 total to include in the linear model. We pooled aligned data across all five days before performing PCA, such that the PCA values were meaningful relative to each other across days. We calculate velocity from the position data using the following approach: we first calculate the difference between position in consecutive timepoints and divide that by the amount of time between timepoints (generally 5ms, but not always). We then treat that as the velocity at the halfway point between the observed timepoints and use linear interpolation to infer the velocity at the observed timepoints.

Interpolation

For the linear encoding model analysis, kinematic data (200Hz) was interpolated to the timepoints of the neural data (30Hz) using the MATLAB function interp1.

Segmenting reaches

Reaches are defined by paw position relative to the reach port: the paw was considered outside the arena when it was to the left of the corner of the right wall of the reach port (right magenta dot in Fig 1A), and inside the arena when it was to the right. The onset of a reach was identified as the moment the paw left the arena and similarly the end of the reach was identified as the timepoint the paw re-entered the arena. By this definition, reaches often consist of one out and back motion, but could include multiple out and back motions if the paw did not fully re-enter the arena in between. Reaches were clustered into multiple reaches in quick succession, followed by a longer break during which the animal left the reach port to trigger the delivery of another

seed (see Fig 1F top for illustration). Bout starts were identified using the inter-reach interval: an inter-reach interval of greater than 1 second separated two bouts. For some mice in some sessions the threshold between bouts was shorter (down to 700ms), depending on the distribution of inter-reach intervals.

The three possible pedestal positions are shown in Fig 1B: medial, lateral, and far lateral. For one of the mice (mouse 3), seeds were presented in only one of the pedestal positions (lateral). For two of the mice, seeds were presented in any of the three positions, in random, counterbalanced order. For the far lateral position the animals showed no success and it was deemed too far lateral for them to reach from the port. The remaining two animals were only trained on the medial and lateral positions. For all animals, only reaches when the seed was on the medial and lateral pedestals were included to calculate the success rates. All reaches to all pedestals were included in the neural coding analyses, as they all can be used to study the motor cortical encoding of different kinematic features.

Performance on the task

A successful reach is one where the animal grasped the seed and brought it into the arena. Success rate is defined as the proportion of reaches that are successful out of all reaches made with a seed present on the pedestal in the medial or lateral position (see previous section). Reaches where a seed is not present at the beginning of the reach are not included in the total. Success was scored manually by watching the video recording of each reach (yeah).

Inclusion of reaching data in analysis

Success versus fail: Due to the nature of the task, successful reaches naturally exhibit lower variance across multiple kinematic features compared to failed reaches. For example, the paw must pass within a certain distance of the pedestal during the reach for the animal to grasp the seed. In all neural coding analysis, we include both successful and failed reaches, as both sets of movements can be used to study the motor cortical encoding of different kinematic features.

Timepoints during reaching: We had the least occlusion and the most reliable tracking of the paw and digit markers when the paw was outside of the arena so we focus our analysis of neural coding on the animal's movements there. Hence, only timepoints during reaching were included in the linear model.

One-photon calcium imaging

Image acquisition

One-photon calcium imaging was performed using an Inscopix nVista 3.0 miniscope. During recording sessions, the lightweight (2g) miniscope was attached to a baseplate, which itself was chronically attached to the animal's skull and centered over the chronically implanted grin lens and prism. Data was collected in a single imaging plane at 30Hz. Field of view corresponds to approximately 900 x 650 micrometers, captured in 1236 x 1040 pixels. Each recording session lasted approximately 40 minutes, broken up into 3 to 5 minute recording chunks, separated by breaks of 2 to 4 minutes. Imaging was continuous during a recording chunk, but paused during

breaks. When analyzing imaging data from a given session, all data from that session was concatenated and analyzed together.

Locating the same field of view over days

The same field of view was imaged across all days of training. Because the prism is chronically implanted into the neural tissue, the only degree of freedom in identifying the field of view is depth (because of the 90 degree rotation from the prism and its orientation, more "shallow" corresponds to more posterior, whereas deeper into the tissue corresponds to more anterior). On the first day of training, the depth with the most active cells was selected while the animal was moving around the arena but not performing the task. Then, some small blood vessels were identified and used as fiduciary markers to find the same plane on subsequent days. Imaged FOVs are estimated between 50 and 100 micrometers deep in the tissue from the surface of the prism.

Pre-processing of calcium imaging data

Data was pre-processed using Inscopix's proprietary software (Inscopix Data Processing Software, IDPS), including downsampling by averaging (4x downsampling means a 4x4 square of pixels is averaged to one pixel), spatial filtering, and motion correction. We exported the resulting stacks of images for cell identification using Caiman CNMFE (Zhou et al., 2018; Giovannucci et al., 2019) to extract spatial footprints and single-cell fluorescence time series. We took the inferred spike data (S output from Caiman), thresholded at 1 standard deviation above

the mean for each neuron on each day separately, then convolved this signal with a gaussian of 30ms standard deviation. This is the neural activity signal used for all analysis.

Registering cells across days

We registered cells across days using CellReg (Sheintuch et al., 2017), with a strict, fixed requirement that across two consecutive days, weighted spatial footprints had to have a spatial correlation of 0.9 to be considered the same cell.

Quantification and statistical analysis

Unless otherwise stated, numbers refer to the following. For description and analysis of behavior, n refers to the number of mice studied, generally n = 5 mice. For description and analysis of neural activity, n refers to the number of neurons recorded from a field of view on one day in one mouse. For comparison of neural encoding properties across days, n refers to the number of pairwise across-day comparisons made, totaled across all cells found on two days that interval of days apart; note that most cells are found on more than two days and thus have more than one comparison made for them.

Analyses were performed on all reaching data, regardless of the success of the reach. See section on kinematic reconstruction for definition of reaching.

Peri-event time histograms (PETHs)

We calculated PETHs from snippets of neural activity where we've locked t = 0 to the start of the reach or the start of the first reach in a bout. We focused on the time window from 100ms before reach onset to 500ms after reach onset. For visualization and for comparison across days, PETHs were normalized to the lowest value = 0, highest value = 1 within this time window.

Linear encoding models of single neurons

On each day separately, we fit a linear model where we predicted the activity of each cell based on measured behavioral and task variables. We used an instantaneous encoding model; neural activity at time t was predicted from behavior and task variables at time t. No time-shifted versions of any predictors were included. Behavioral variables were all interpolated to 30Hz (see kinematic processing section).

Preparing predictor sets/design matrices:

Treatment of predictors: The predictors included are the position and velocity in x and y of the paw and the head (8 variables), the first four PCs of the paw markers, and the first two PCs of the head markers, and time within the session (not including breaks between recording chunks; time variable is equivalent to the index of the calcium imaging data). Only timepoints during reaching (when the paw was outside the arena) were used in the model. Because we want to compare the encoding for the same absolute values of the predictors, we Z-score on all days together rather than on each day separately. For each mouse separately, we treated each predictor as follows: We concatenated all reach timepoints for that predictor across all five days, then z-scored. We then

separated the data back into days and performed the covariate balancing procedure described below.

Covariate balancing across days to allow for direct comparison of linear fits: In order to directly compare properties of our linear model fits across days, we must ensure that each predictor has a similar distribution to itself over days, and also that the joint distribution of predictors is similar. Without this step, we cannot guarantee that observed instability is not a result of the animals exploring a slightly different part of kinematic space each day. We use a greedy matching algorithm to balance the joint distribution of all behavioral predictors (the time predictor is not included). All single day results presented in this paper are covariate balanced across all five days for that animal. After covariate balancing, each predictor is again z-scored within day before fitting the models.

Train-test split

We used a train-test split of 70% train, 30% test. We did the following procedure to ensure that each train-test split contained a balanced sampling of timepoints from across the session. For all the timepoints included in a day for a given mouse, we split the data into 100 chunks. We then grouped every 10th chunk together, resulting in 10 sets. For example, chunks 1, 11, 21, etc. went into set 1, while chunks 2, 12, 22, etc went into set 2. We then chose 7 random sets to use as the train set, and held out 3 chunks as the test set. We repeated this procedure for 10 cross-validations for all model fits. All linear fit values are reported as averages across the 10 cross-

validations. Note that due to the train-test split, r-squared values can sometimes be positive. Within a day, the train-test splits for each cross-validation were the same for all cells.

Fitting single trial models

Single trial models refers to the standard instantaneous linear model fit where we predict neural activity at time t from behavioral and task variables at time t. We fit a family of models for each cell. Full models were fit using all predictors. We also fit a single variable model for each predictor that included only that predictor. We also fit a leave one out model for each predictor that included all predictors except for the one in question. Delta-r-squared values are computed as the drop in r-squared from the full model to the leave one out model for the predictor in question, as schematized in Fig 5a.

All reported r-squared and delta-r-squared values are averaged across 10 cross-validations of the train-test split. Delta-r-squareds are calculated within each cross-validation, then averaged across cross-validations. Due to the train-test split, the Delta R-squared (leave one out models) can sometimes be positive, indicating that removing the predictor in question improved the test R-squared values. This typically occurs for small R-squared values or poor fits, where the model is fitting to noise in the train set and thus cannot generalize to the test set.

Significance testing of full model r-squared values

To test whether a full model fit was better than expected by chance, we did a significance test for each neuron on each day. We shuffled the timepoints of the neural activity vector (predicted

variable) and refit the full model. We repeated this procedure 1000 times. If the observed r-squared value was greater than 95% of time-shuffled r-squared values, we deemed the fit for that neuron on that day to be significant.

Fitting trial average models

For each predictor separately, we calculated a time-varying average movement during reaching, locked to the onset of reach. This is very similar to calculating a PETH, but where we average across snippets of a kinematic variable, instead of snippets of neural activity. We built a new design matrix where for each predictor we replaced each reach with a copy of that predictor's average. We only replaced reach timepoints: for a reach that is 10 timepoints long, we replaced those 10 timepoints with the first 10 timepoints of the reach-locked average, whereas for a reach that is 20 timepoints long, we replaced those 20 timepoints with the first 20 timepoints of the reach-locked average. Only reach timepoints were used for the models, as for the single trial models. All trial average models were fit using the same train-test procedure described for single trial models above.

Cosine similarity of r-squared values and coefficients

To quantify the similarity of the tuning profile of a given cell across two days, we compared two sets of linear model fit components: r-squared values and coefficients. For each set separately, we did the following comparison. For each cell, we make a vector out of all of the r-squared values for all models: the cross-validated r-squared and delta-r-squared for all predictor variables and the r-squared for the full model (13 predictors*2 + 1 r-squared for full model = 27 dimensional

vector). We then calculate the cosine of the angle between this vector for the cell on day 1 and this vector for the cell on day 2. Because of our test-train split, any r-squared value can be positive or negative, so these vectors are not guaranteed to occupy the same hyper-quadrant. Thus, the cosine of the angle varies between -1 (anti-parallel) and 1 (parallel). We also perform the same procedure for the coefficients from all the single variable predictors for a given cell. The coefficients are unbounded, so the cosine of the angle between possible coefficient vectors also varies between -1 and 1. This approach to quantifying similarity of tuning properties ignores the absolute magnitude of tuning for each predictor and instead compares the magnitude of the r-squared values for the predictors relative to each other.

Control shuffle: For each specific interval between days, we generated a resampled control distribution. If the observed cosine similarity exceeded 95% of this distribution, the individual comparison was considered significant. It's important to note that using one neuron in the population as a control for another is not an appropriate test because most cells are correlated with each other during motor behavior. Therefore, this analysis focused on the number of cells that showed significant similarities in their encoding profiles across different intervals (see the last panel in Figure 6). A large proportion of comparisons across a given day interval are significant - line plot. Note this is comparisons, not cells - an individual cell may show up in multiple comparisons for the same interval, e.g. 1-day interval for a cell that shows up on day 1, 2, and 3.

2.5 Supplementary Figures

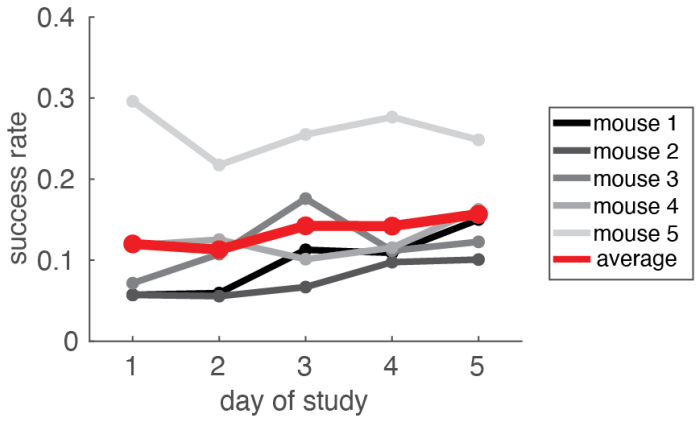

Supplementary Figure 1: Mice exhibit stable task performance over the five days. Success rate on each day for each animal, calculated as the proportion of reaches where the mouse successfully retrieved the seed, out of all reaches with seeds present. Average is calculated across mice.

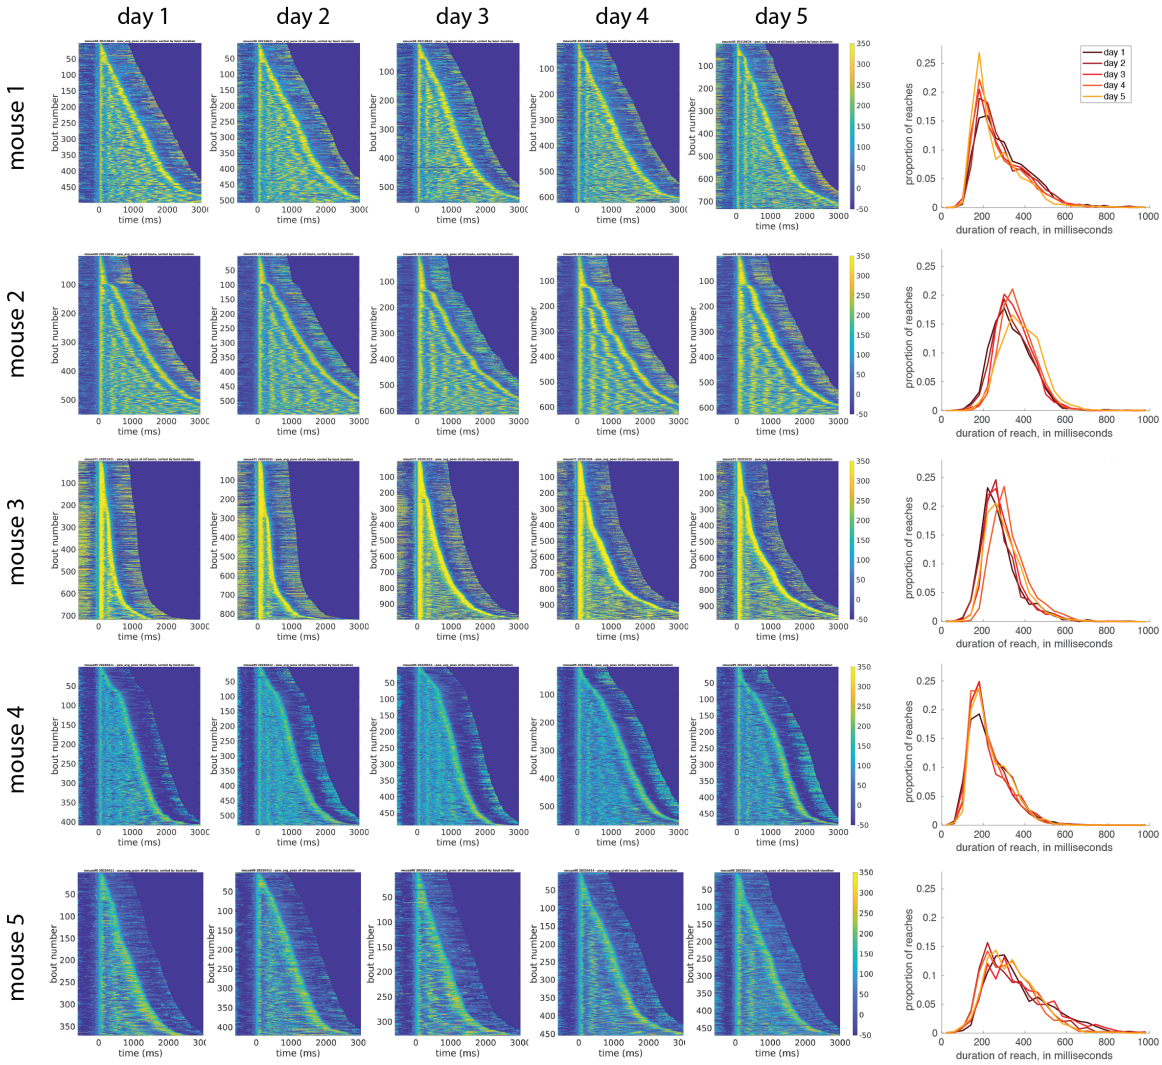

Supplementary Figure 2: Overview of reach strategies for all mice
Same as in Fig 2B, but for all five mice. Left, heatmaps of x position of the paw for all bouts, sorted by bout duration, using the same color scheme as Fig 2B. Right, probability density distributions of time durations of reaches for each day.

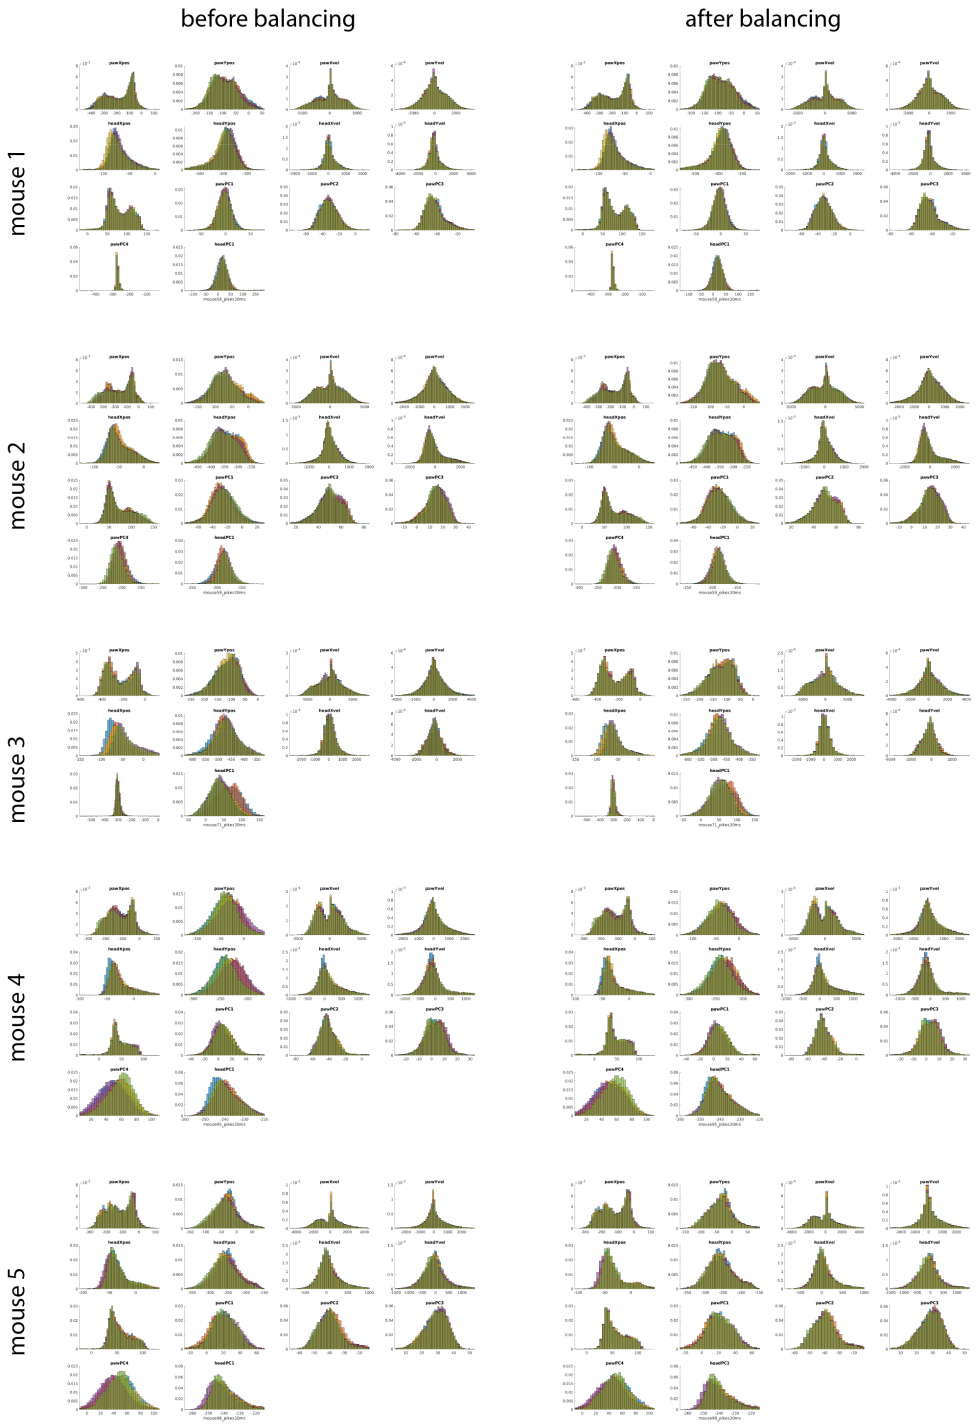

Supplementary Figure 3: Distributions of kinematic variables before and after covariate balancing procedure, for all 14 kinematic variables.
See methods for details of covariate balancing procedure. Principal components of the paw were not included for mouse 3 in any analysis because of poor tracking of the digits (see methods).

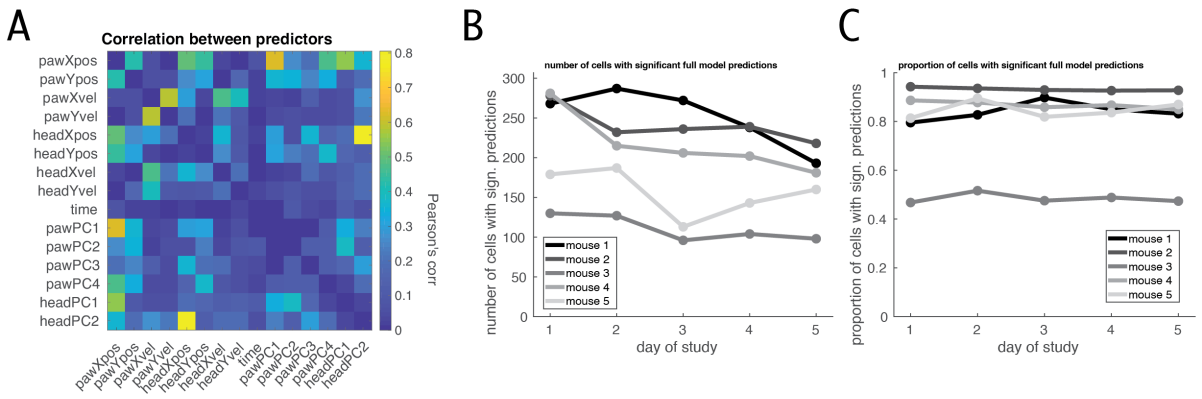

Supplementary Figure 4: Predictor correlations, and count and proportion of significant linear model fits by dataset.

A. Pearson's correlation between predictors for one example session (mouse 5, day 2).

B. Number of cells with significant linear model fits for the single trial full model, on each day for each mouse.

C. Same as in B but as the proportion of cells in the session.

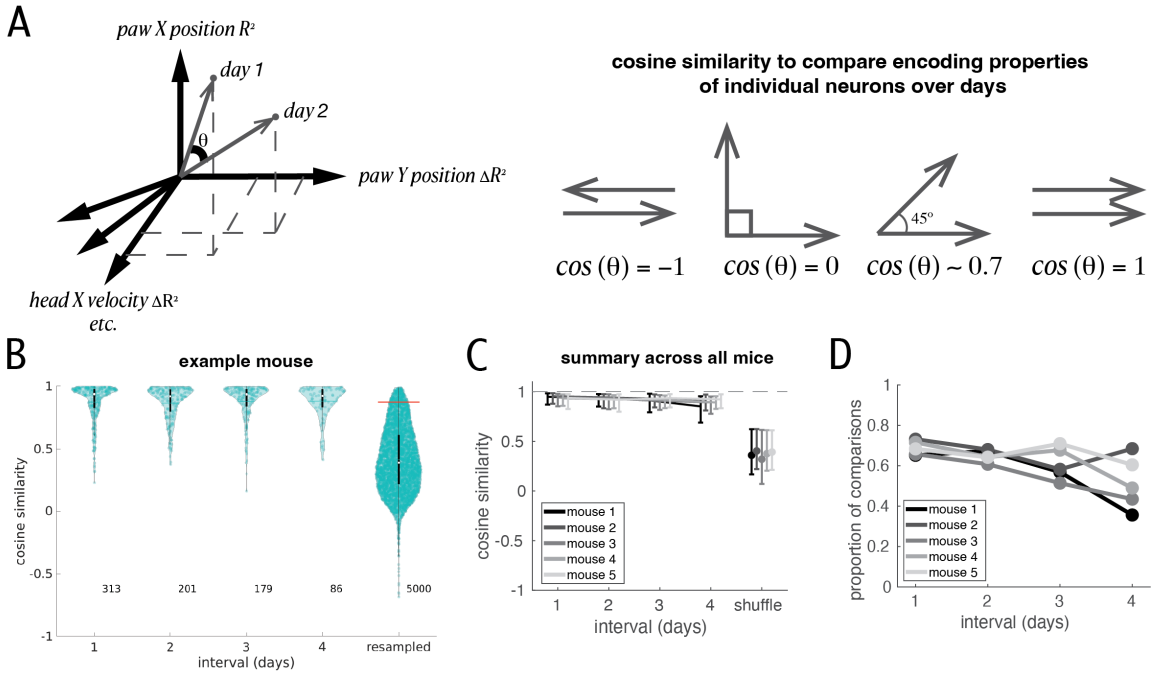

Supplementary Figure 5: Linear model encoding profiles constructed from R-squared values are stable across days.

A. Schematic of method: using cosine similarity to compare vectors of encoding coefficients (betas) over days for individual cells. This is the same analysis as described in Fig 5, but comparing the linear encoding profiles constructed from the collection of R-squared values for each cell (see Methods).

B. Cosine similarities of coefficient vectors as a function of the interval between observations for an example mouse (mouse 5). Each datapoint corresponds to a comparison of observations of a single cell on two different days. Numbers to the bottom right of the violins indicate the number of observations for that interval. For each violin, the white dot indicates the median of the distribution and the black box indicates interquartile range. Resampled shuffle control built by generating pseudo-vectors where each coefficient is drawn from the pool of coefficients observed across all neurons for that predictor within that day. Coefficients within the same pseudo-vector are drawn independently. The resampled shuffle distribution contains the same number of comparisons for each of the 10 possible pairs of the 5 days of study (5000 total, 500 per day pair). Red line indicates the 95th percentile cutoff of this distribution.

C. Same as B, but for all animals. Data is median, error bars are interquartile range across cells for that interval for that mouse.

D. Proportions of coefficient vector comparisons for a given interval that are statistically significant using the resampled shuffle distribution as a control.

Chapter 3: General Discussion

Ultimately, the goal in motor neuroscience is to understand how the nervous systems of freely moving and naturally behaving animals (including humans) generate movements across many timescales. These kinds of movements are inherently variable and spatiotemporally complex, often involving the coordination of many body parts. In this project, I characterized the motor cortical encoding of such a task, a whole-body reach-to-grasp, at the single cell level. I then quantified the stability of this encoding across a span of five consecutive days. All together, these results support the view of mouse motor cortex as a brain area that is indeed involved in the fine-grained, low-level execution details of head, forelimb, and digit movements in reach-to-grasp behaviors.

Our results corroborate previous findings of behavioral category encoding in rodent motor cortex (Dhawale et al., 2017a; Peters et al., 2014, 2017; Huber et al., 2012; Dombeck et al., 2009). The vast majority of neurons modulate during reaching relative to non-reaching timepoints (encompassing behaviors such as locomotion, grooming, and consuming seeds), and approximately half of the neurons carry a general reaching signal in the form of a significant prediction by the trial average model. In addition, many neurons encode the moment-to-moment details of paw, digit, and head kinematics. Mouse motor cortex has already been shown to encode forelimb muscle activation during head-fixed climbing (Koh et al., 2024) and forelimb kinematics during head-fixed ladder walking (Omlor et al., 2018), two tasks that require precision paw placement and a grasping component. Our observation that motor cortex also encodes forelimb kinematics during a freely moving reach-to-grasp is further evidence that mouse motor cortex may play a role in the execution of skilled forelimb movements. Further

work is needed to explore how both the high level behavioral category signals and low level execution detail signals are multiplexed within single cell and population activity.

We observe that single cells in CFA encode movements of the paw, head, and digits. What role might these signals play in M1 during reach-to-grasp? Our linear model analysis maps neural activity to kinematics happening at the same timepoint; it is an instantaneous encoding model. We do not differentiate if neural activity most strongly encodes movements that happen before (e.g. sensory signals) or movements that happen after (e.g. motor signals). The fact that M1 can have short latency effects during a lever press task (Miri et al., 2017) and climbing (Koh et al., 2024), suggests that encoding of the reach-to-grasp that we see could at least partially be direct movement commands. Based on depth, our recordings span putative L2/3 and L5a. Both could influence behavior directly via their strong projection to L5b (Weiler et al., 2008), the layer that sends the most projections to subcortical regions, including notably to interneurons in the cervical spinal cord (Carmona et al., 2024; Anderson et al., 2010; Economo et al., 2018). L5a also contains corticostriatal-projecting neurons (Anderson et al., 2010); the activity of those cells could influence behavior indirectly via cortico-striato-thalamic loops. Both layers also receive tactile somatosensory input from the lemnisco-cortical pathway via projections from hand primary somatosensory cortex (Yamawaki et al., 2021) as well as input from RFA and other frontal motor areas (Ueta et al., 2014), positioning these neural populations well to integrate sensory feedback with ongoing movement plans, for example for error correction. Future work with this dataset could interrogate the temporal relationship between neural activity and movement by identifying the time lags at which behavior best explains neural activity, which could provide clues about the causal directionality.

In our linear encoding models, we relate the activity of individual neurons to a carefully chosen, carefully measured set of 14 kinematic variables (position and velocity in x and y and the shape and orientation of the paw and orientation of the head). Adding many more predictors, such as extracted motion energy or compressions (SVD) of our behavioral videos (as in Musall, Kaufman, et al., 2019 and Stringer et al., 2019) could increase the predictive power of our models. We chose to stick to our set of measured kinematic variables to maintain interpretability and tie our results to the movements of specific body parts.

The amount of variance of single neurons that can be predicted from this measured set is, for 80.0% of neurons (cell-days), a statistically significant amount of variance. For the vast majority of those neurons, this amount, although statistically significant, is a fraction of the overall single-trial variance of the cell. What is the source of all that other variance? Some portion of it is undoubtedly due to noisy biological processes at the molecular, cellular, and synaptic level (Faisal et al., 2008; Renart and Machens, 2014). Given the strong connectivity within and across layers in M1, some portion of variance could reflect the internal dynamics of the local cortical network (MacLean and Yuste, 2005). Such activity may be necessary for generating motor cortical output but not necessarily be linearly correlated with kinematics (Shenoy et al., 2013). Some portion of variance might be explained by movements, sensory input, or cognitive variables that we did not measure, either related or unrelated to the task (Musall, Kaufman, et al., 2019). Finally, forelimb, digit, and head kinematics may be encoded in motor cortex in a non-linear way or not at the timescale we considered. Our approach to quantifying neural encoding is the simplest possible allowing for moment-to-moment encoding (linear models with individual timepoints as observations). It is possible that if we used a more

complex encoding model that we could explain more variance. Ultimately, we observe that the majority of cells in motor cortex have statistically significant amounts of explained variance with our single trial models, and that the majority of these cells maintain the nature of their mapping to these specific features over days. We interpret these findings to indicate that, during reach-to-grasp, motor cortex encodes the movements of the paw, digits, and head but likely encodes many other things as well.

In the task we describe here, animals show significant trial-to-trial variability within a session even across a time period where their success rate is stable. Classically, studies of motor skill learning have focused on a reduction in variability of goal-directed movements as a metric to define if a skill has been learned (Shmuelof et al., 2012; Sternad, 2018; Dhawale et al., 2017b; Dayan and Cohen, 2011). The free throw shot in basketball is an excellent example of a motor skill where this metric works well: the external environment is fixed and players are allowed to set up the same way each time, and an increase in success is coupled with an increase in stereotypy. However, a different kind of learned motor skill is one that remains flexible even after a plateau in performance. This is exemplified by a player's ability to make baskets from many places on the court under many different conditions. These shots are built from the same core movement elements as free throw shots, but require slight adjustments and different neural control signals in order to be successful, such as when the player needs to step around a defender or when their muscles are fatigued at the end of a game. If the movements a player makes while shooting are not all exactly the same across a range of shots, we do not consider them to be still learning or poor at the task. In fact, the opposite is true; the ability to make baskets in many ways is highly desirable in players. Ultimately, we gauge a player's level of mastery on their overall

success rate, as we have done here in our study of reach-to-grasp in the mouse. Notably, across days, we do not observe an apparent continued reduction in variability or increase in speed that might suggest the animals are still learning or optimizing their movements. Given that our goal was to quantify stability of encoding, the most important behavioral metrics were success rate and this similarity in the distributions of movement across days. We note that our covariate balancing procedure would not have been effective if the within-subject distributions of kinematic variables were not already largely the same across days, i.e. we would have ended up with very few matched samples to fit our models with.

Taking a broader perspective, characterizing variability in skilled movements can be informative about both neural control processes and behavioral strategies. Undoubtedly, some movement variability is due to noise in the nervous system and musculature. The basketball analogy above highlights one non-noise-based source: learned behaviors can show variability because the task demands variability. In such a task, what makes one successful is the ability to adapt a core set of movements to the specific demands of each instance. Additionally, learned behaviors can show variability because of degeneracy in the space of successful movements, either at the level of gross movement strategy (e.g. classic overhead free throw shot versus the maligned but reliable underarm "granny" shot) or at the level of fine-grained precision. Nervous systems have limited resources and thus may not prioritize unnecessary stereotypy in goal-directed movements if it is not required by the task. It has been proposed that precision of movement might be optimized only over a limited set of dimensions that are highly task relevant, those that are under the most pressure in a reinforcement learning context (Sternad et al., 2018), leaving significant variability in many other dimensions. Finally, learned behaviors may show

variability because it conveys some benefit, potentially making the behaviors more readily adaptable (Dhawale et al., 2017b). Indeed, Dhawale and colleagues demonstrated that rats are capable of regulating their movement variability in response to changing task demands (Dhawale et al., 2019). Identifying which dimensions of movement variability are under active regulatory control and whether that variability is encoded by the nervous system could reveal principles of movement control more broadly.

To our knowledge, this is the longest duration over which the single cell encoding of this level of movement detail has been considered. Our approach of combining calcium imaging in mice with careful quantification of the fine details of behavior allowed us to describe the stability with which individual cells encode these low-level movement features. Because of this goal of characterizing stability at the single cell level, we used an encoding approach, predicting the activity of individual neurons, rather than a decoding approach where the norm is to predict an external variable from population activity. Although it is possible to isolate the contribution of individual neurons to population decoders, our encoding approach allowed us to focus on movement encoding of individual neurons in a more straightforward way. Our findings of stable single neuron activity suggest that the stable latents observed in macaque motor cortex could result from stable single neuron activity as well (Gallego et al., 2020).

Why is the encoding of movement apparently so stable relative to that for many kinds of sensory stimuli? Some of this dramatic difference in stability across brain regions is likely due to systematic differences in the way we study neural encoding of movements versus sensory stimuli. Separately, some of it may be due to inherent differences in neural circuitry and hence in processing. Many motor areas of the brain receive a constant, structured training signal about the

efficacy of movement commands in the form of somatosensory feedback. This tight sensorimotor loop may be responsible for maintaining the reliable relationship between neural activity and movement that the field has thus far observed in motor areas. This explanation, however, could also be used to argue for the converse: having a tight feedback loop could potentially grant some tolerance for drifting representations, particularly for movements or tasks that are repeated often and do not require high precision on each trial. In tasks like this, there would frequently be updated "training data" with which to adjust the internal mapping from neural activity to movement, such that mappings may drift over a long timescale like days to weeks without a decrement in performance. Further study of the role of sensory feedback in maintaining skilled movements and the stability of its encoding in motor areas could shed light on this issue.

Another potential explanation is based on naturalism and task complexity. It has been hypothesized that task complexity would be inversely correlated with drift rates, because greater demand on neural computational resources would limit the number of unused degrees of freedom in neural activity (Micou and O'Leary, 2023). Hence, we might expect a more stable encoding of a task that involves coordination of multiple body parts over multiple timescales than for a task that isolates movements to one body part. And the same for tasks that require flexible versus stereotyped movements. Similarly, in tasks that require animals to learn an arbitrary mapping from sensory stimuli to movement response, we might expect that tasks with simpler mappings, in the form of fewer or less complex stimuli or response options, would require fewer degrees of freedom in neural activity and thus have more dimensions in neural activity space in which to drift. And finally, for encodings of sensory stimuli that are passively viewed, we might not expect there to be any stable coding given that the stimulus has no trained behavioral relevance

and frequently little ethological relevance. As discussed in the introduction, the instability of sensory representations may in fact provide benefits. Our nervous systems are capable of mapping features of experiences (including arbitrary sensory stimuli) to specific movements and ultimately to behavioral outcomes. Sensory representations that are continually updating to incorporate new information may underlie this ability (Driscoll et al., 2017, 2022).

References

Aflalo, T. N., & Graziano, M. S. A. (2006). Partial tuning of motor cortex neurons to final posture in a free-moving paradigm. *Proceedings of the National Academy of Sciences*, *103*(8), 2909–2914. https://doi.org/10.1073/pnas.0511139103

Aitken, K., Garrett, M., Olsen, S., & Mihalas, S. (2022). The geometry of representational drift in natural and artificial neural networks. *PLOS Computational Biology*, *18*(11), e1010716. https://doi.org/10.1371/journal.pcbi.1010716

Alaverdashvili, M., & Whishaw, I. Q. (2008). Motor cortex stroke impairs individual digit movement in skilled reaching by the rat. *European Journal of Neuroscience*, *28*(2), 311–322. https://doi.org/10.1111/j.1460-9568.2008.06315.x

An, X., Matho, K., Li, Y., Mohan, H., Xu, X. H., Whishaw, I. Q., Kepecs, A., & Huang, Z. J. (2024). A cortical circuit for orchestrating oromanual food manipulation. bioRxiv. https://doi.org/10.1101/2022.12.03.518964

Andermann, M. L., Gilfoy, N. B., Goldey, G. J., Sachdev, R. N. S., Wölfel, M., McCormick, D. A., Reid, R. C., & Levene, M. J. (2013). Chronic Cellular Imaging of Entire Cortical Columns in Awake Mice Using Microprisms. *Neuron*, *80*(4), 900–913. https://doi.org/10.1016/j.neuron.2013.07.052

Anderson, C. T., Sheets, P. L., Kiritani, T., & Shepherd, G. M. G. (2010). Sublayer-specific microcircuits of corticospinal and corticostriatal neurons in motor cortex. *Nature Neuroscience*, *13*(6), 739–744. https://doi.org/10.1038/nn.2538

Asan, A. S., McIntosh, J. R., & Carmel, J. B. (2022). Targeting Sensory and Motor Integration for Recovery of Movement After CNS Injury. *Frontiers in Neuroscience*, *15*, 791824. https://doi.org/10.3389/fnins.2021.791824

Barrett, J. M., Raineri Tapies, M. G., & Shepherd, G. M. G. (2020). Manual dexterity of mice during food-handling involves the thumb and a set of fast basic movements. *PLOS ONE*, *15*(1), e0226774. https://doi.org/10.1371/journal.pone.0226774

Bauer, J., Lewin, U., Herbert, E., Gjorgjieva, J., Schoonover, C. E., Fink, A. J. P., Rose, T., Bonhoeffer, T., & Hübener, M. (2024). Sensory experience steers representational drift in mouse visual cortex. *Nature Communications*, *15*(1), 9153. https://doi.org/10.1038/s41467-024-53326-x

Bauer, J., & Rose, T. (2021). Mouse vision: Variability and stability across the visual processing hierarchy. *Current Biology*, *31*(19), R1129–R1132. https://doi.org/10.1016/j.cub.2021.08.071

Becker, M. I., Calame, D. J., Wrobel, J., & Person, A. L. (2020). Online control of reach accuracy in mice. *Journal of Neurophysiology*, *124*(6), 1637–1655. https://doi.org/10.1152/jn.00324.2020

Billard, A., & Kragic, D. (2019). Trends and challenges in robot manipulation. *Science*, *364*(6446), eaat8414. https://doi.org/10.1126/science.aat8414

Bollu, T., Whitehead, S. C., Prasad, N., Walker, J., Shyamkumar, N., Subramaniam, R., Kardon, B., Cohen, I., & Goldberg, J. H. (2024). Motor cortical inactivation impairs corrective submovements in mice performing a hold-still center-out reach task. *Journal of Neurophysiology*, *132*(3), 829–848. https://doi.org/10.1152/jn.00241.2023

Carmena, J. M., Lebedev, M. A., Henriquez, C. S., & Nicolelis, M. A. L. (2005). Stable Ensemble Performance with Single-Neuron Variability during Reaching Movements in Primates. *Journal of Neuroscience*, *25*(46), 10712–10716. https://doi.org/10.1523/JNEUROSCI.2772-05.2005

Carmona, L. M., Thomas, E. D., Smith, K., Tasic, B., Costa, R. M., & Nelson, A. (2024). Topographical and cell type-specific connectivity of rostral and caudal forelimb corticospinal neuron populations. *Cell Reports*, *43*(4), 113993. https://doi.org/10.1016/j.celrep.2024.113993

Chambers, A. R., & Rumpel, S. (2017). A stable brain from unstable components: Emerging concepts and implications for neural computation. *Neuroscience*, *357*, 172–184. https://doi.org/10.1016/j.neuroscience.2017.06.005

Chen, C.-C., Gilmore, A., & Zuo, Y. (2014). Study Motor Skill Learning by Single-pellet Reaching Tasks in Mice. *Journal of Visualized Experiments : JoVE*, *85*. https://doi.org/10.3791/51238

Chestek, C. A., Batista, A. P., Santhanam, G., Yu, B. M., Afshar, A., Cunningham, J. P., Gilja, V., Ryu, S. I., Churchland, M. M., & Shenoy, K. V. (2007). Single-Neuron Stability during Repeated Reaching in Macaque Premotor Cortex. *Journal of Neuroscience*, *27*(40), 10742–10750. https://doi.org/10.1523/JNEUROSCI.0959-07.2007

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

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., & Shenoy, K. V. (2010). Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron*, *68*(3), 387–400. https://doi.org/10.1016/j.neuron.2010.09.015

Churchland, M. M., & Shenoy, K. V. (2007). Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. *Journal of Neurophysiology*, *97*(6), 4235–4257. https://doi.org/10.1152/jn.00095.2007

Clopath, C., Bonhoeffer, T., Hübener, M., & Rose, T. (2017). Variance and invariance of neuronal long-term representations. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1715), 20160161. https://doi.org/10.1098/rstb.2016.0161

Dayan, E., & Cohen, L. G. (2011). Neuroplasticity subserving motor skill learning. *Neuron*, *72*(3), 443–454. https://doi.org/10.1016/j.neuron.2011.10.008

Deitch, D., Rubin, A., & Ziv, Y. (2021). Representational drift in the mouse visual cortex. *Current Biology*, *31*(19), 4327-4339.e6. https://doi.org/10.1016/j.cub.2021.07.062

Dhawale, A. K., Miyamoto, Y. R., Smith, M. A., & Ölveczky, B. P. (2019). Adaptive Regulation of Motor Variability. *Current Biology*. https://doi.org/10.1016/j.cub.2019.08.052

Dhawale, A. K., Poddar, R., Wolff, S. B., Normand, V. A., Kopelowitz, E., & Ölveczky, B. P. (2017a). Automated long-term recording and analysis of neural activity in behaving animals. *ELife*, *6*, e27702. https://doi.org/10.7554/eLife.27702

Dhawale, A. K., Smith, M. A., & Ölveczky, B. P. (2017b). The Role of Variability in Motor Learning. *Annual Review of Neuroscience*, *40*(1), 479–498. https://doi.org/10.1146/annurev-neuro-072116-031548

Disse, G. D., Nandakumar, B., Pauzin, F. P., Blumenthal, G. H., Kong, Z., Ditterich, J., & Moxon, K. A. (2023). Neural ensemble dynamics in trunk and hindlimb sensorimotor cortex encode for the control of postural stability. *Cell Reports*, *42*(4), 112347. https://doi.org/10.1016/j.celrep.2023.112347

Dombeck, D. A., Graziano, M. S., & Tank, D. W. (2009). Functional Clustering of Neurons in Motor Cortex Determined by Cellular Resolution Imaging in Awake Behaving Mice. *Journal of Neuroscience*, *29*(44), 13751–13760. https://doi.org/10.1523/JNEUROSCI.2985-09.2009

Driscoll, L. N., Pettit, N. L., Minderer, M., Chettih, S. N., & Harvey, C. D. (2017). Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. *Cell*, *170*(5), 986-999.e16. https://doi.org/10.1016/j.cell.2017.07.021

Driscoll, L. N., Duncker, L., & Harvey, C. D. (2022). Representational drift: Emerging theories for continual learning and experimental future directions. *Current Opinion in Neurobiology*, *76*, 102609. https://doi.org/10.1016/j.conb.2022.102609

Economo, M. N., Viswanathan, S., Tasic, B., Bas, E., Winnubst, J., Menon, V., Graybuck, L. T., Nguyen, T. N., Smith, K. A., Yao, Z., Wang, L., Gerfen, C. R., Chandrashekar, J., Zeng,

H., Looger, L. L., & Svoboda, K. (2018). Distinct descending motor cortex pathways and their roles in movement. *Nature*, *563*(7729), 79–84. https://doi.org/10.1038/s41586-018-0642-9

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

Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, *9*(4), 292–303. https://doi.org/10.1038/nrn2258

Flint, R. D., Scheid, M. R., Wright, Z. A., Solla, S. A., & Slutzky, M. W. (2016). Long-Term Stability of Motor Cortical Activity: Implications for Brain Machine Interfaces and Optimal Feedback Control. *Journal of Neuroscience*, *36*(12), 3623–3632. https://doi.org/10.1523/JNEUROSCI.2339-15.2016

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

Galiñanes, G. L., Bonardi, C., & Huber, D. (2018). Directional Reaching for Water as a Cortex-Dependent Behavioral Framework for Mice. *Cell Reports*, *22*(10), 2767–2783. https://doi.org/10.1016/j.celrep.2018.02.042

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

Ganguly, K., & Carmena, J. M. (2009). Emergence of a Stable Cortical Map for Neuroprosthetic Control. *PLoS Biology*, *7*(7), e1000153. https://doi.org/10.1371/journal.pbio.1000153

Georgopoulos, A. P., Ashe, J., Smyrnis, N., & Taira, M. (1992). The Motor Cortex and the Coding of Force. *Science*, *256*(5064), 1692–1695. https://doi.org/10.1126/science.256.5064.1692

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

Geva, N., Deitch, D., Rubin, A., & Ziv, Y. (2023). Time and experience differentially affect distinct aspects of hippocampal representational drift. *Neuron*, *111*(15), 2357-2366.e5. https://doi.org/10.1016/j.neuron.2023.05.005

Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B. L., Koay, S. A., Taxidis, J., Najafi, F., Gauthier, J. L., Zhou, P., Khakh, B. S., Tank, D. W., Chklovskii, D. B., & Pnevmatikakis, E. A. (2019). CaImAn an open source tool for scalable calcium imaging data analysis. *ELife*, *8*, e38173. https://doi.org/10.7554/eLife.38173

Gorko, B., Siwanowicz, I., Close, K., Christoforou, C., Hibbard, K. L., Kabra, M., Lee, A., Park, J.-Y., Li, S. Y., Chen, A. B., Namiki, S., Chen, C., Tuthill, J. C., Bock, D. D., Rouault, H., Branson, K., Ihrke, G., & Huston, S. J. (2024). Motor neurons generate pose-targeted movements via proprioceptive sculpting. *Nature*, *628*(8008), 596–603. https://doi.org/10.1038/s41586-024-07222-5

Guo, J.-Z., Graves, A. R., Guo, W. W., Zheng, J., Lee, A., Rodríguez-González, J., Li, N., Macklin, J. J., Phillips, J. W., Mensh, B. D., Branson, K., & Hantman, A. W. (2015). Cortex commands the performance of skilled movement. *ELife*, *4*. https://doi.org/10.7554/eLife.10774

Harrison, T. C., Ayling, O. G. S., & Murphy, T. H. (2012). Distinct Cortical Circuit Mechanisms for Complex Forelimb Movement and Motor Map Topography. *Neuron*, *74*(2), 397–409. https://doi.org/10.1016/j.neuron.2012.02.028

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

Hira, R., Ohkubo, F., Tanaka, Y. R., Masamizu, Y., Augustine, G. J., Kasai, H., & Matsuzaki, M. (2013). In vivo optogenetic tracing of functional corticocortical connections between motor forelimb areas. *Frontiers in Neural Circuits*, *7*, 55. https://doi.org/10.3389/fncir.2013.00055

Hocherman, S., & Wise, S. P. (1991). Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys. *Experimental Brain Research*, *83*(2). https://doi.org/10.1007/BF00231153

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

Hwang, E. J., Dahlen, J. E., Hu, Y. Y., Aguilar, K., Yu, B., Mukundan, M., Mitani, A., & Komiyama, T. (2019). Disengagement of motor cortex from movement control during

long-term learning. *Science Advances*, *5*(10), eaay0001. https://doi.org/10.1126/sciadv.aay0001

Jackson, A., Mavoori, J., & Fetz, E. E. (2007). Correlations Between the Same Motor Cortex Cells and Arm Muscles During a Trained Task, Free Behavior, and Natural Sleep in the Macaque Monkey. *Journal of Neurophysiology*, *97*(1), 360–374. https://doi.org/10.1152/jn.00710.2006

Jensen, K. T., Kadmon Harpaz, N., Dhawale, A. K., Wolff, S. B. E., & Ölveczky, B. P. (2022). Long-term stability of single neuron activity in the motor system. *Nature Neuroscience*, *25*(12), 1664–1674. https://doi.org/10.1038/s41593-022-01194-3

Katlowitz, K. A., Picardo, M. A., & Long, M. A. (2018). Stable Sequential Activity Underlying the Maintenance of a Precisely Executed Skilled Behavior. *Neuron*, *98*(6), 1133-1140.e3. https://doi.org/10.1016/j.neuron.2018.05.017

Kawai, R., Markman, T., Poddar, R., Ko, R., Fantana, A. L., Dhawale, A. K., Kampff, A. R., & Ölveczky, B. P. (2015). Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron*, *86*(3), 800–812. https://doi.org/10.1016/j.neuron.2015.03.024

Koh, N., Ma, Z., Sarup, A., Kristl, A. C., Agrios, M., Young, M., & Miri, A. (2024). Selective direct motor cortical influence during naturalistic climbing. bioRxiv. https://doi.org/10.1101/2023.06.18.545509

Lemon, R. N. (2008). Descending Pathways in Motor Control. *Annual Review of Neuroscience*, *31*(1), 195–218. https://doi.org/10.1146/annurev.neuro.31.060407.125547

Liberti, W. A., Markowitz, J. E., Perkins, L. N., Liberti, D. C., Leman, D. P., Guitchounts, G., Velho, T., Kotton, D. N., Lois, C., & Gardner, T. J. (2016). Unstable neurons underlie a stable learned behavior. *Nature Neuroscience*, *19*(12), 1665–1671. https://doi.org/10.1038/nn.4405

Liberti, W. A., Schmid, T. A., Forli, A., Snyder, M., & Yartsev, M. M. (2022). A stable hippocampal code in freely flying bats. *Nature*, *604*(7904), 98–103. https://doi.org/10.1038/s41586-022-04560-0

Lütcke, H., Margolis, D. J., & Helmchen, F. (2013). Steady or changing? Long-term monitoring of neuronal population activity. *Trends in Neurosciences*, *36*(7), 375–384. https://doi.org/10.1016/j.tins.2013.03.008

MacLean, J. N., Watson, B. O., Aaron, G. B., & Yuste, R. (2005). Internal Dynamics Determine the Cortical Response to Thalamic Stimulation. *Neuron*, *48*(5), 811–823. https://doi.org/10.1016/j.neuron.2005.09.035

Marks, T. D., & Goard, M. J. (2021). Stimulus-dependent representational drift in primary visual cortex. *Nature Communications*, *12*(1), 5169. https://doi.org/10.1038/s41467-021-25436-3

Masamizu, Y., Tanaka, Y. R., Tanaka, Y. H., Hira, R., Ohkubo, F., Kitamura, K., Isomura, Y., Okada, T., & Matsuzaki, M. (2014). Two distinct layer-specific dynamics of cortical ensembles during learning of a motor task. *Nature Neuroscience*, *17*(7), 987–994. https://doi.org/10.1038/nn.3739

Mason, C. R., Johnson, M. T. V., Fu, Q.-G., Gomez, J. E., & Ebner, T. J. (1998). Temporal profile of the directional tuning of the discharge of dorsal premotor cortical cells: *NeuroReport*, *9*(6), 989–995. https://doi.org/10.1097/00001756-199804200-00007

Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., & Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, *21*(9), 1281–1289. https://doi.org/10.1038/s41593-018-0209-y

Micou, C., & O'Leary, T. (2023). Representational drift as a window into neural and behavioural plasticity. *Current Opinion in Neurobiology*, *81*, 102746. https://doi.org/10.1016/j.conb.2023.102746

Micou, C., & O'Leary, T. (2024). Heavy-tailed statistics of cortical representational drift are advantageous for stabilised downstream readouts. bioRxiv. https://doi.org/10.1101/2024.09.26.614914

Mimica, B., Dunn, B. A., Tombaz, T., Bojja, V. P. T. N. C. S., & Whitlock, J. R. (2018). Efficient cortical coding of 3D posture in freely behaving rats. *Science*, *362*(6414), 584–589. https://doi.org/10.1126/science.aau2013

Miri, A., Warriner, C. L., Seely, J. S., Elsayed, G. F., Cunningham, J. P., Churchland, M. M., & Jessell, T. M. (2017). Behaviorally Selective Engagement of Short-Latency Effector Pathways by Motor Cortex. *Neuron*, *95*(3), 683-696.e11. https://doi.org/10.1016/j.neuron.2017.06.042

Moran, D. W., & Schwartz, A. B. (1999). Motor Cortical Representation of Speed and Direction During Reaching. *Journal of Neurophysiology*, *82*(5), 2676–2692. https://doi.org/10.1152/jn.1999.82.5.2676

Mosberger, A. C., Sibener, L. J., Chen, T. X., Rodrigues, H. F. M., Hormigo, R., Ingram, J. N., Athalye, V. R., Tabachnik, T., Wolpert, D. M., Murray, J. M., & Costa, R. M. (2024). Exploration biases forelimb reaching strategies. *Cell Reports*, *43*(4), 113958. https://doi.org/10.1016/j.celrep.2024.113958

Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S., & Churchland, A. K. (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience*, *22*(10), 1677–1686. https://doi.org/10.1038/s41593-019-0502-4

Nath, T., Mathis, A., Chen, A. C., Patel, A., Bethge, M., & Mathis, M. W. (2019). Using DeepLabCut for 3D markerless pose estimation across species and behaviors. *Nature Protocols*, *14*(7), 2152–2176. https://doi.org/10.1038/s41596-019-0176-0

Nudo, R., & Frost. (2007). The Evolution of Motor Cortex and Motor Systems. In *The Evolution of Nervous Systems* (Vol. 3).

Omlor, W., Sipilä, P., Wahl, A.-S., Lütcke, H., Laurenczy, B., Chen, I.-W., Sumanovski, L. T., 't Hoff, M. van, Bethge, P., Voigt, F. F., Schwab, M. E., & Helmchen, F. (2019). Context-dependent limb movement encoding in neuronal populations of motor cortex. bioRxiv. https://doi.org/10.1101/588129

Oswald, M. J., Tantirigama, M. L. S., Sonntag, I., Hughes, S. M., & Empson, R. M. (2013). Diversity of layer 5 projection neurons in the mouse motor cortex. *Frontiers in Cellular Neuroscience*, *7*. https://doi.org/10.3389/fncel.2013.00174

Peters, A. J., Chen, S. X., & Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature*, *510*(7504), 263–267. https://doi.org/10.1038/nature13235

Peters, A. J., Lee, J., Hedrick, N. G., O'Neil, K., & Komiyama, T. (2017). Reorganization of corticospinal output during motor learning. *Nature Neuroscience*, *20*(8), 1133–1141. https://doi.org/10.1038/nn.4596

Renart, A., & Machens, C. K. (2014). Variability in neural activity and behavior. *Current Opinion in Neurobiology*, *25*, 211–220. https://doi.org/10.1016/j.conb.2014.02.013

Resendez, S. L., Jennings, J. H., Ung, R. L., Namboodiri, V. M. K., Zhou, Z. C., Otis, J. M., Nomura, H., McHenry, J. A., Kosyk, O., & Stuber, G. D. (2016). Visualization of cortical, subcortical and deep brain neural circuit dynamics during naturalistic mammalian behavior with head-mounted microscopes and chronically implanted lenses. *Nature Protocols*, *11*(3), 566–597. https://doi.org/10.1038/nprot.2016.021

Rokni, U., Richardson, A. G., Bizzi, E., & Seung, H. S. (2007). Motor Learning with Unstable Neural Representations. *Neuron*, *54*(4), 653–666. https://doi.org/10.1016/j.neuron.2007.04.030

Rule, M. E., Loback, A. R., Raman, D. V., Driscoll, L. N., Harvey, C. D., & O'Leary, T. (2020). Stable task information from an unstable neural population. *ELife*, *9*, e51121. https://doi.org/10.7554/eLife.51121

Rule, M. E., & O'Leary, T. (2022). Self-healing codes: How stable neural populations can track continually reconfiguring neural representations. *Proceedings of the National Academy of Sciences*, *119*(7), e2106692119. https://doi.org/10.1073/pnas.2106692119

Rule, M. E., O'Leary, T., & Harvey, C. D. (2019). Causes and consequences of representational drift. *Current Opinion in Neurobiology*, *58*, 141–147. https://doi.org/10.1016/j.conb.2019.08.005

Russo, A. A., Bittner, S. R., Perkins, S. M., Seely, J. S., London, B. M., Lara, A. H., Miri, A., Marshall, N. J., Kohn, A., Jessell, T. M., Abbott, L. F., Cunningham, J. P., & Churchland, M. M. (2018). Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron*, *97*(4), 953-966.e8. https://doi.org/10.1016/j.neuron.2018.01.004

Sacrey, L.-A. R., Alaverdashvili, M., & Whishaw, I. Q. (2009). Similar hand shaping in reaching-for-food (skilled reaching) in rats and humans provides evidence of homology in release, collection, and manipulation movements. *Behavioural Brain Research*, *204*(1), 153–161. https://doi.org/10.1016/j.bbr.2009.05.035

Sadeh, S., & Clopath, C. (2022). Contribution of behavioural variability to representational drift. *ELife*, *11*, e77907. https://doi.org/10.7554/eLife.77907

Saiki-Ishikawa, A., Agrios, M., Savya, S., Forrest, A., Sroussi, H., Hsu, S., Basrai, D., Xu, F., & Miri, A. (2024). Hierarchy between forelimb premotor and primary motor cortices and its manifestation in their firing patterns. bioRxiv. https://doi.org/10.1101/2023.09.23.559136

Saleh, M., Takahashi, K., Amit, Y., & Hatsopoulos, N. G. (2010). Encoding of Coordinated Grasp Trajectories in Primary Motor Cortex. *The Journal of Neuroscience*, *30*(50), 17079–17090. https://doi.org/10.1523/JNEUROSCI.2558-10.2010

Saleh, M., Takahashi, K., & Hatsopoulos, N. G. (2012). Encoding of Coordinated Reach and Grasp Trajectories in Primary Motor Cortex. *The Journal of Neuroscience*, *32*(4), 1220–1232. https://doi.org/10.1523/JNEUROSCI.2438-11.2012

Sauerbrei, B. A., Guo, J.-Z., Cohen, J. D., Mischiati, M., Guo, W., Kabra, M., Verma, N., Mensh, B., Branson, K., & Hantman, A. W. (2020). Cortical pattern generation during dexterous movement is input-driven. *Nature*, *577*(7790), 386–391. https://doi.org/10.1038/s41586-019-1869-9

Schwartz, A., Kettner, R., & Georgopoulos, A. (1988). Primate motor cortex and free arm movements to visual targets in three- dimensional space. I. Relations between single cell discharge and direction of movement. *The Journal of Neuroscience*, *8*(8), 2913–2927. https://doi.org/10.1523/JNEUROSCI.08-08-02913.1988

Sergio, L. E., Hamel-Pâquet, C., & Kalaska, J. F. (2005). Motor Cortex Neural Correlates of Output Kinematics and Kinetics During Isometric-Force and Arm-Reaching Tasks. *Journal of Neurophysiology*, *94*(4), 2353–2378. https://doi.org/10.1152/jn.00989.2004

Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., & Ziv, Y. (2017). Tracking the Same Neurons across Multiple Days in Ca2+ Imaging Data. *Cell Reports*, *21*(4), 1102–1115. https://doi.org/10.1016/j.celrep.2017.10.013

Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience*, *36*(1), 337–359. https://doi.org/10.1146/annurev-neuro-062111-150509

Shmuelof, L., Krakauer, J. W., & Mazzoni, P. (2012). How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. *Journal of Neurophysiology*, *108*(2), 578–594. https://doi.org/10.1152/jn.00856.2011

Sternad, D. (2018). It's not (only) the mean that matters: variability, noise and exploration in skill learning. *Current Opinion in Behavioral Sciences*, *20*, 183–195. https://doi.org/10.1016/j.cobeha.2018.01.004

Stevenson, I. H. (2018). Omitted Variable Bias in GLMs of Neural Spiking Activity. *Neural Computation*, *30*(12), 3227–3258. https://doi.org/10.1162/neco_a_01138

Stevenson, I. H., Cherian, A., London, B. M., Sachs, N. A., Lindberg, E., Reimer, J., Slutzky, M. W., Hatsopoulos, N. G., Miller, L. E., & Kording, K. P. (2011). Statistical assessment of the stability of neural movement representations. *Journal of Neurophysiology*, *106*(2), 764–774. https://doi.org/10.1152/jn.00626.2010

Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., & Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, *364*(6437), eaav7893. https://doi.org/10.1126/science.aav7893

Suresh, A. K., Goodman, J. M., Okorokova, E. V., Kaufman, M., Hatsopoulos, N. G., & Bensmaia, S. J. (2020). Neural population dynamics in motor cortex are different for reach and grasp. *ELife*, *9*, e58848. https://doi.org/10.7554/eLife.58848

Tennant, K. A., Adkins, D. L., Donlan, N. A., Asay, A. L., Thomas, N., Kleim, J. A., & Jones, T. A. (2011). The Organization of the Forelimb Representation of the C57BL/6 Mouse Motor Cortex as Defined by Intracortical Microstimulation and Cytoarchitecture. *Cerebral Cortex*, *21*(4), 865–876. https://doi.org/10.1093/cercor/bhq159

Ueno, M., Nakamura, Y., Li, J., Gu, Z., Niehaus, J., Maezawa, M., Crone, S. A., Goulding, M., Baccei, M. L., & Yoshida, Y. (2018). Corticospinal Circuits from the Sensory and Motor Cortices Differentially Regulate Skilled Movements through Distinct Spinal

Interneurons. *Cell Reports*, *23*(5), 1286-1300.e7. https://doi.org/10.1016/j.celrep.2018.03.137

Ueta, Y., Otsuka, T., Morishima, M., Ushimaru, M., & Kawaguchi, Y. (2014). Multiple Layer 5 Pyramidal Cell Subtypes Relay Cortical Feedback from Secondary to Primary Motor Areas in Rats. *Cerebral Cortex*, *24*(9), 2362–2376. https://doi.org/10.1093/cercor/bht088

Vargas-Irwin, C. E., Shakhnarovich, G., Yadollahpour, P., Mislow, J. M. K., Black, M. J., & Donoghue, J. P. (2010). Decoding Complete Reach and Grasp Actions from Local Primary Motor Cortex Populations. *Journal of Neuroscience*, *30*(29), 9659–9669. https://doi.org/10.1523/JNEUROSCI.5443-09.2010

Wagner, M. J., Kim, T. H., Kadmon, J., Nguyen, N. D., Ganguli, S., Schnitzer, M. J., & Luo, L. (2019). Shared Cortex-Cerebellum Dynamics in the Execution and Learning of a Motor Task. *Cell*, *177*(3), 669-682.e24. https://doi.org/10.1016/j.cell.2019.02.019

Wang, S., Laittre, E. A. de, MacLean, J., & Palmer, S. E. (2024). Quantifying stimulus-relevant representational drift using cross-modality contrastive learning (No. arXiv:2305.11953). arXiv. https://doi.org/10.48550/arXiv.2305.11953

Wang, X., Liu, Y., Li, X., Zhang, Z., Yang, H., Zhang, Y., Williams, P. R., Alwahab, N. S. A., Kapur, K., Yu, B., Zhang, Y., Chen, M., Ding, H., Gerfen, C. R., Wang, K. H., & He, Z. (2017). Deconstruction of Corticospinal Circuits for Goal-Directed Motor Skills. *Cell*, *171*(2), 440-455.e14. https://doi.org/10.1016/j.cell.2017.08.014

Warriner, C. L., Fageiry, S. K., Carmona, L. M., & Miri, A. (2020). Towards Cell and Subtype Resolved Functional Organization: Mouse as a Model for the Cortical Control of Movement. *Neuroscience*, *450*, 151–160. https://doi.org/10.1016/j.neuroscience.2020.07.054

Warriner, C. L., Fageiry, S., Saxena, S., Costa, R. M., & Miri, A. (2022). Motor cortical influence relies on task-specific activity covariation. *Cell Reports*, *40*(13), 111427. https://doi.org/10.1016/j.celrep.2022.111427

Weiler, N., Wood, L., Yu, J., Solla, S. A., & Shepherd, G. M. G. (2008). Top-down laminar organization of the excitatory network in motor cortex. *Nature Neuroscience*, *11*(3), 360–366. https://doi.org/10.1038/nn2049

Whishaw, I. Q., Dringenberg, H. C., & Pellis, S. M. (1992). Spontaneous forelimb grasping in free feeding by rats: motor cortex aids limb and digit positioning. *Behavioural Brain Research*, *48*(2), 113–125. https://doi.org/10.1016/S0166-4328(05)80147-0

Whishaw, I. Q., Faraji, J., Kuntz, J., Mirza Agha, B., Patel, M., Metz, G. A. S., & Mohajerani, M. H. (2017). Organization of the reach and grasp in head-fixed vs freely-moving mice provides support for multiple motor channel theory of neocortical organization.

Experimental Brain Research, *235*(6), 1919–1932. https://doi.org/10.1007/s00221-017-4925-4

Whishaw, I. Q., Ghasroddashti, A., Mirza Agha, B., & Mohajerani, M. H. (2020). The temporal choreography of the yo-yo movement of getting spaghetti into the mouth by the head-fixed mouse. *Behavioural Brain Research*, *381*, 112241. https://doi.org/10.1016/j.bbr.2019.112241

Whishaw, I. Q., Pellis, S. M., & Gorny, B. P. (1992). Skilled reaching in rats and humans: evidence for parallel development or homology. *Behavioural Brain Research*, *47*(1), 59–70. https://doi.org/10.1016/S0166-4328(05)80252-9

Whishaw, I. Q., Travis, S. G., Koppe, S. W., Sacrey, L.-A., Gholamrezaei, G., & Gorny, B. (2010). Hand shaping in the rat: Conserved release and collection vs. flexible manipulation in overground walking, ladder rung walking, cylinder exploration, and skilled reaching. *Behavioural Brain Research*, *206*(1), 21–31. https://doi.org/10.1016/j.bbr.2009.08.030

Xia, J., Marks, T. D., Goard, M. J., & Wessel, R. (2021). Stable representation of a naturalistic movie emerges from episodic activity with gain variability. *Nature Communications*, *12*(1), 5170. https://doi.org/10.1038/s41467-021-25437-2

Yamawaki, N., Raineri Tapies, M. G., Stults, A., Smith, G. A., & Shepherd, G. M. (2021). Circuit organization of the excitatory sensorimotor loop through hand/forelimb S1 and M1. *ELife*, *10*, e66836. https://doi.org/10.7554/eLife.66836

Zhou, P., Resendez, S. L., Rodriguez-Romaguera, J., Jimenez, J. C., Neufeld, S. Q., Giovannucci, A., Friedrich, J., Pnevmatikakis, E. A., Stuber, G. D., Hen, R., Kheirbek, M. A., Sabatini, B. L., Kass, R. E., & Paninski, L. (2018). Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *ELife*, *7*, e28728. https://doi.org/10.7554/eLife.28728

Zhu, F., Grier, H. A., Tandon, R., Cai, C., Agarwal, A., Giovannucci, A., Kaufman, M. T., & Pandarinath, C. (2022). A deep learning framework for inference of single-trial neural population dynamics from calcium imaging with subframe temporal resolution. *Nature Neuroscience*, *25*(12), 1724–1734. https://doi.org/10.1038/s41593-022-01189-0