

UCSF

UC San Francisco Electronic Theses and Dissertations

Title

Behavioral flexibility in the corticostriatal network

Permalink

<https://escholarship.org/uc/item/8z8729ks>

Author

Kondapavulur, Sravani

Publication Date

2021

Peer reviewed|Thesis/dissertation

Behavioral flexibility in the corticostriatal network

by

Sravani Kondapavulur

DISSERTATION

Submitted in partial satisfaction of the requirements for degree of

DOCTOR OF PHILOSOPHY

in

Bioengineering

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

AND

UNIVERSITY OF CALIFORNIA, BERKELEY

Approved:

DocuSigned by:

Karunesh Ganguly

537133AC2605407...

Karunesh Ganguly

Chair

DocuSigned by:

Philip Starr

DocuSigned by: 4C9...

Philip Starr

Tejal Desai

DocuSigned by: 4A0...

Tejal Desai

Michel Maharbiz

DocuSigned by: BE9F990F0B63467...

Michel Maharbiz

Committee Members

Copyright 2021
by
Sravani Kondapavulur

*To my parents, Mani and KT, my brother Satish, and the many communities
whose support made this endeavor possible*

Acknowledgements

There are many people who I would like to thank for supporting me throughout my PhD work. First, I would like to thank Karunesh Ganguly, my PhD advisor, for being a source of thought-provoking discussions, motivation, and support which have encouraged my scientific and professional curiosities. Coming into the lab as an engineer, with little prior neuroscientific experience, Karunesh provided endless knowledge of relevant papers to read, how to frame questions within the context of current scientific landscapes, and how to ground my scientific pursuits in clinically meaningful directions.

I would also like to thank my thesis committee members Philip Starr, Tejal Desai, and Michel Maharbiz. I truly appreciate the feedback and support received over the last three years. In particular, I would like to thank Philip Starr for providing guidance and support for my path towards becoming a neurosurgeon-scientist, within both scientific meetings and clinical preceptorships; Tejal Desai for providing mentorship on shaping my career as a bioengineer within the context of a physician-scientist path; and Michel Maharbiz for providing lighthearted insights into my research endeavors. I would also like to thank Sri Nagarajan and Rikky Muller for their contributions on my qualification exam committee.

I was incredibly fortunate to join the lab at a time when graduate student mentors were abundant. I would like to thank Stefan Lemke for being my introductory mentor into the field of systems neuroscience, and Ling Guo for teaching me how to apply this framework to study disease. Both of your patience, encouragement, and positivity taught me how to pursue specific scientific questions and take unsuccessful analyses in stride with the eventual successes. I would also like to thank Tess Veuthey and Kate Derosier, for providing experimental, analytical, and

professional mentorship along the way. When the pandemic disrupted standard research approaches, I was incredibly grateful for each of your willingness to take time remotely to provide guidance, both personally and professionally.

I would like to thank the graduate students who have joined the lab since: David Darevsky, Yasmin Graham, Sandon Griffin, and Francisco Aparicio. The enthusiasm with which each of you approach the science in our lab has been lovely and provided levity as we conducted science together during a tumultuous time. I would also like to thank April Hishinuma, Chelsey Rodriguez, and Sapeeda Barati for their support on many of our projects in the lab.

I would like to thank Preeya Khanna and Nikhilesh Natraj for providing domain specific expertise and support throughout my projects. Both of you were instrumental in teaching me how to ask and answer questions in a principled, methodical, and grounded way. Additionally, I would like to thank all the members of the Ganguly lab who have made the lab an enjoyable and fulfilling place to be a graduate student: Seok Joon Won, Adelyn Tu-Chan, Sergio Arroyo, Doug Totten, Daniel Silversmith, Jaekyung Kim, Reza Abiri, Nick Hardy, Harman Ghuman, Hoseok Choi, Kyungsoo Kim, and Sarah Seko. Much of my knowledge is due to your expertise, patience, and generosity in teaching and mentorship.

I would like to thank my prior research mentors at UC Berkeley, Luke Lee and Amy Herr, at Stanford, Lawrence Steinman, and at UCSF, Steven Hetts. Your time, teaching, and encouragement introduced me to critical thinking in research and the academic process. Additionally, I am grateful for your continued support, without which my path and opportunities at UCSF would not have been possible.

I would like to thank my additional neurosurgical mentors, Doris Wang, Shawn Hervey-Jumper, Ethan Winkler, and John Burke, for providing guidance on how to ground research

questions in clinically unmet needs, as well as encouraging my path towards surgeon-scientist training.

I would like to thank the UCSF-UC Berkeley Bioengineering Graduate Program, particularly SarahJane Taylor, Victoria Ross, and Christoph Schreiner for their help in navigating graduate school. The graduate students in this program have been welcoming and supportive, especially Hector Neira and John Kim who helped make the transition from medical to graduate school enjoyable.

I would like to thank Mark Anderson and Geri Ehle for their guidance and support over the past five years in the MSTP program. Additionally I would like to thank other MSTP faculty who have provided invaluable career mentoring including Cathy Lomen-Hoerth, Aimee Kao, Mercedes Paredes, Alexandra Nelson, and Aaron Tward. The longitudinal cohort of bioengineering MSTP students have been a great source of knowledge and friendship, and I would like to especially thank Leah Conant, Jeremy Bancroft-Brown, Adam Rao, and Karsyn Bailey for their guidance. Thank you to the entering MSTP class of 2016 for the companionship through trials and tribulations, especially Marissa Chou and Abrar Choudhury, who have spent countless hours in Zoom work sessions, sharing PhD memes, and enjoying sunshine in Golden Gate Park as we progress together. Finally, I would like to thank my friends from the entering MD class of 2016, including Zoya Qureshy, Martin Consunji, Kevin Huynh, Kevin Yu, Simon Ammanuel, Ogonna Nnamani Silva, and Max Feinstein, for the many runs, picnics, boardgame nights, and long chats.

This PhD would not be possible without a large community of social support along the way. Thank you to the aunts, uncles, and friends from Grand Oak Way and beyond who have taught me since childhood that humility, kindness, and friendship are the foundation of a healthy

future; in particular, I would like to thank the Kondas and Kongaras for being my second families. Thank you to undergraduate friends from the UC Berkeley Pre-engineering Program, in particular Ana Lua, Daniel Zezula, and Gio Gajudo, who have grounded and supported my aspirations since the start of adulthood; from Griffiths, Amit Goren, Stephanie Liu, and Jeff Dinakar, who have been cheerleaders and role models from the day we all met; and from research labs, Bridget Kilbride and Bradford Thorne, with whom science is the utmost combination of creativity and fun.

Thank you to John Pluvinae, for being a partner who has been supportive scientifically, clinically, and personally. Your unwavering faith in my abilities accompanied by comical scientifically ad-libbed songs have been invaluable. Additionally, thank you to the Pluvinae family for your support.

Finally, thank you to my family in India and the Bay Area, who have continuously asked about my research, supported me against cultural norms, and encourage me to pursue anything I have chosen. Thank you to my older brother Satish, for paving the way for me in many aspects of life. Thank you to my adopted puppy sibling Berkeley, for the endless enthusiasm and love this past year. I have no words to express the gratitude I feel for my parents Mani Kondapavulur and KT Venkateswara Rao. Your wisdom, generosity, and kindness have been endless over the years; thank you, I love you both.

This research was conducted with funding from the UCSF Medical Scientist Training Program and National Institute of Health, National Institute of Neurological Disorders and Stroke, Ruth L. Kirschstein National Research Service Award, Individual Predoctoral Fellowship (F31).

Behavioral flexibility in the corticostriatal network

Sravani Kondapavulur

Abstract

A critical question in skilled motor control is how brain circuits can flexibly apply learned motor behaviors to novel contexts. Plasticity in the corticostriatal network has been historically been implicated in motor learning across different species. In particular, the motor cortex and striatum are monosynaptically connected structures that play critical roles in early motor learning as well as consolidation of crystallized motor skills. However, when a well-learned skill must adapt to a new context, how does neural activity in these important motor regions support the behavioral transition?

This thesis begins by providing a perspective on how motor learning, adaptation, and habit formation have been classically studied, in addition to the current understanding of neural bases for each of these processes. I then outline the implications of motor task selection and behavioral intervention when making conclusions about underlying neural mechanisms. In Chapter 2, I present my work investigating how the corticostriatal network supports transfer learning of a previously acquired complex coordinated upper limb action, demonstrating that re-learning in a new context involves partial breakdown of previously crystallized neural activity. I conclude by presenting in Chapter 3 a model for network control of motor flexibility, proposing that cortico-striatal activity reflects dynamic optimization of neural patterns, adapting upstream cognitive contributions and driving existing downstream circuits for movement generation. Finally, I explore the implications of corticostriatal dysfunction in motor diseases.

Table of Contents

Chapter 1: Introduction	1
The motor behavior spectrum dilemma	1
What is a motor skill?	3
Network contributions to motor learning and related processes	5
Local field potentials vs. spiking, the debate	7
Analytical methods for measuring single- and cross-area communication	9
References	13
Chapter 2: Breakdown of corticostriatal patterns underlies motor transfer learning	21
Abstract	21
Introduction	22
Results	24
Discussion	52
Methods	56
References	68
Chapter 3: Two pieces of the puzzle – introducing a model for motor flexibility	74
Motor cortex – a node in multiple coexisting pathways	76
Implications of corticostriatal dysfunction	82
Summary	84
References	86

List of Figures

Chapter 1:

Figure 1.1: Similar figure skating jumps	1
Figure 1.2: Automated reach-to-grasp task for rodents	4
Figure 1.3: Tetzlaff, et. al. 2011 depiction of spiking and LFP detection	7
Figure 1.4: Gallego and Miller depiction of neural manifold	11

Chapter 2

Figure 2.1: Transfer learning of an automatic, skilled behavior	25
Figure 2.2: Breaking out of an automatic, skilled motor behavior is a multi-day process, with individuals exhibiting variability	26
Figure 2.3: In a rigid state, error-based updating occurs neither within trials nor across trials	28
Figure 2.4: In a rigid state, error-based updating occurs across days	29
Figure 2.5: Breakdown of reach-locked M1 and DLS unit spiking underlies transfer learning of reach-to-grasp skill	31
Figure 2.6: Both M1 and DLS units lose reach modulation of spiking activity with transfer learning	33
Figure 2.7: While unit modulation drops, M1 and DLS unit task-averaged activity does not change by session type	34
Figure 2.8: M1 temporal patterning breaks down and re-emerges with transfer learning	36
Figure 2.9: DLS temporal patterning also re-emerges following transfer learning	38

Figure 2.10: M1 and DLS single-trial population spiking activity loses temporal consistency with transfer learning	41
Figure 2.11: M1 and DLS single-trial reach-modulated population spiking breakdown underlies transfer learning	42
Figure 2.12: M1-DLS task population spiking drops correlation with preservation of modulation during transfer learning	44
Figure 2.13: M1-DLS reach-related cross-area dynamics drop during transfer learning	47
Figure 2.14: Motor behavior during the variable state has features of early and late learning	49
Figure 2.15: Unsupervised alignment of spiking activity suggests underlying trial structure in DLS	51
Chapter 3:	
Figure 3.1: Hypotheses for transfer learning of reach-to-grasp skill	75
Figure 3.2: Simplified interaction of network loops, from HH Yin and BJ Knowlton, 2006	77
Figure 3.3: Proposed model of parallel motor loops in early versus late motor learning	80
Figure 3.4: Proposed network model of behavioral flexibility in transfer learning	81

Chapter 1: Introduction

The motor behavior spectrum dilemma

When a child stumbles onto the ice for the first time, a global network process of motor learning begins. What are the leg motions that allow forward gliding motion while not tripping over the toe pick at the front of the boot? How does one maintain balance? Once the basics such as gliding and stopping are mastered, the child may decide to devote many hours of practice to the sport, to master complex spins, jumps, and ultimately competition routines set to music.

Now let us take one aspect of advanced skating, the mastery of two quite similar jumps, a flip and a lutz (**Figure 1.1**). The flip jump is generally learned first, and it involves leaning medially on the left foot and vaulting off the right toe pick to rotate in a counter-clockwise fashion in the air. Once this jump is mastered and becomes automatic, the lutz jump is learned; the lutz is identical in mechanics to the flip jump, except instead of leaning medially on the left foot, one leans laterally. Thus, learning a lutz involves transfer learning of the “vault off right toe into the air after gliding backward on the left leg” motor skill, with edge adaptation. However, due to initial learning of the flip to automaticity, some skaters become stuck in the medially leaning pattern – while initially leaning laterally, by the time they vault into the air the left foot is leaning medially, resulting in the motor error of a “flutz.”

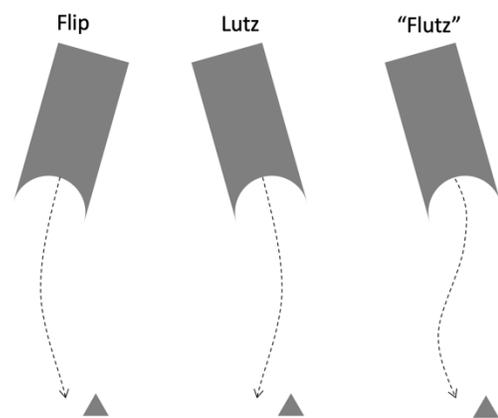


Figure 1.1. Similar figure skating jumps. L-R: flip, gliding backward on the left blade leaning medially, vaulting off the right toe pick (triangle); lutz, gliding backward leaning laterally, similarly vaulting off the right toe pick; flutz, erroneous jump with intended lateral lutz edge, but switching unintentionally to medial flip edge.

From this example, one can abstract many components of motor execution, demonstrating the “motor behavior spectrum dilemma.” The first axis is early learning to late learning: how does a flip jump become automatic? The second axis is between generalizable versus rigid motor actions: can learning the flip jump transfer skill to the lutz jump, or is learning the lutz jump a new motor skill? Finally, the third axis is between goal-directed versus habitual motor behavior – does a flutz occur due to intrusion of habitual medial edge on the intended jump? Most studies focus only on a single axis; however, there is overlap among the three in most motor behavior experiments. This chapter will define different components of motor behavior, and the current understanding of associated neural regions and activity patterns. The latter half will delve into merits and pitfalls of different readouts of neural activity and outline the current state of analytical methods for drawing conclusions about neural processing relevant to behavior.

What is a motor skill?

Learning a skill involves expansion of a physical repertoire, such that previously acquired movements are combined into a goal-directed action.¹ With practice, motor skills become faster and more accurate.² Subsequently, when speed and accuracy have reached an asymptote, behavior is considered automatic, as characterized by decreased cognitive effort independent of decreased dependence on reinforcement learning.^{3,4} Behavioral automaticity underlies habits, compulsions, and addictions. Of note, *both skills and habits can be automatic*, with decreased dependence on reward, and resilience to intrusion of competing motor actions.^{2,5}

The behavioral axis of goal-directed versus habitual performance of a single motor task has been extensively studied. Whereas habits are ingrained, and performed autonomously of outcome^{2,5-9}, goal-directed action involves active deliberation of action consequences and is therefore flexibly adaptive to changing environments⁵.

Classifications of motor behavior differ amongst animal models, as well as scientific fields, such as neuroscience versus psychology, further adding difficulties to across-study comparisons. For example, the definition of habit in humans and primates is quite broad – actions that are reflexive, performed without thought, stereotyped, and following a stimulus-response relationship¹⁰⁻¹³, reminiscent of the currently accepted definitions of automaticity³. In rodents however, the definition of habit is quite specific, characterized by behavioral automaticity that is insensitive to devaluation by satiety (i.e. pre-task delivery of reward) or paired aversion (i.e. pre-task pairing of reward with a nausogenic substance)¹⁴⁻¹⁶, reminiscent of human pathologies on the obsessive-compulsive spectrum¹⁷. In turn, random-ratio versus random-interval reward delivery

methods have promoted goal-directed and habitual performance respectively.^{16,18} Given the overlap in terminology, I will focus further behavior discussion and the neural correlates on early learning as compared to late, automatic skill execution.

Overlaid on the early to late learning axis is a second-order question – once a skill has been learned in the context of also learning task cues and environment, how can the skill be adapted to environmental modifications while keeping task structure? Particularly in rodents, this transfer of learning has been studied during auditory perception and discrimination, T-maze side preference, and water maze reversal learning¹⁹⁻²¹. Does establishing behavioral flexibility after learning, also known as transfer of learning, involve a minor tweaking of existing learned motor skill? Or rather does this process involve “re-learning”, so to speak?

These are the central behavioral questions that this thesis will delve into. In particular, in Chapter 2, I will explore the reach-to-grasp task in rodents (**Figure 1.2**), a correlate of gross proximal and dexterous distal movements that can also be seen in non-human primates and humans. Specifically, learning and execution of the reach-to-grasp task to a single location has been extensively studied in rodents²²⁻²⁴. However, it is unknown 1) whether this initially established behavior is rigid or generalizable; and 2) if the behavior is rigid, how a related reach-to-grasp skill can be learned.

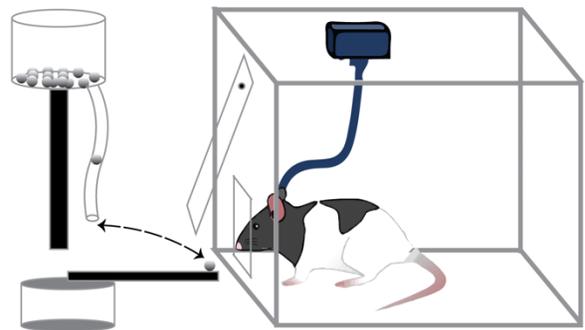


Figure 1.2: Automated reach-to-grasp task for rodents. For each trial, the pellet arm receives a pellet from the dispenser and is moved to the reach location. The trial begins with a cue tone with door opening, and the rodent reaches through a slit in the front wall to grasp and retrieve the sugar pellet reward.

Network contributions to motor learning and related processes

Generally, motor learning of a skill in a rodent is thought to involve a shift from associative cortex and striatum (i.e. secondary motor cortex (M2) and the dorsomedial striatum (DMS)) to the sensorimotor cortex and striatum (i.e. primary motor cortex (M1) and the dorsolateral striatum (DLS))^{5,25-28}. Additionally, thalamus has been demonstrated to provide feedback input to cortex and striatum for potentiation of optimal synaptic patterns²⁹.

Of note, the role of DLS in motor learning may differ based on the task itself. While complex sequences that involve increased dexterity, such as grasping, demonstrate an overall increase in DLS activation, the number of task-modulated units decreases with overlearning in simpler tasks, such as lever-pressing or navigational locomotion in a T-maze^{23,30,31}.

Specifically in the reach-to-grasp task, acquisition of skill is marked by the emergence of reach-related modulation of M2-M1 cross-area activity, with M2 local activity preceding M1 local activity³². When examining M1 and DLS during learning of the reach-to-grasp skill, 3-6 Hz theta LFP power and coherence increases, corresponding specifically with a decrease in kinematic variability of the proximal reaching motion²³. Additionally, there is evidence towards M1 being responsible for the grasping and DLS for reaching during this skill: inactivation of M1 preferentially disrupts grasping, while DLS inactivation preferentially decreases reach amplitude with intact grasping; dimensionality reduction of M1 population activity, rather than DLS population activity, via Gaussian process factor analysis can also separate out neural activity due to accurate versus inaccurate grasps. Finally, the role of M1 compared to thalamus has been recently explored. After the reach-to-grasp skill has been learned, temporary inactivation of M1

disrupts movement initiation but not execution, whereas temporary inactivation of thalamus disrupts movement execution itself²⁴. Thus, there is a growing body of evidence towards multi-site representation of complex learned movements in a well-learned state, demonstrating a need for more multi-region recordings throughout early learning, late learning, and task manipulations to further clarify the functional roles of different nodes in the motor network at different behavioral learning states.

Local field potentials vs. spiking, the debate

Analysis of neural activity often falls into two different categories: 1) study of local field potential (LFP) activity detected in a region within certain frequency bands, due to lower sampling frequency for recording and more stable chronic signal quality⁵²; 2) spiking activity, detecting depolarizations of single neurons extracellularly, providing a binary readout of unit firing. How do these two measures relate? LFP is often thought of as the aggregate electrical activity in a region, as demonstrated by Tetzlaff, et. al. 2011 (**Figure 1.3**)³³. What information can be gleaned from this signal, as it has lower resolution than individual unit spiking contributions to overall activity?

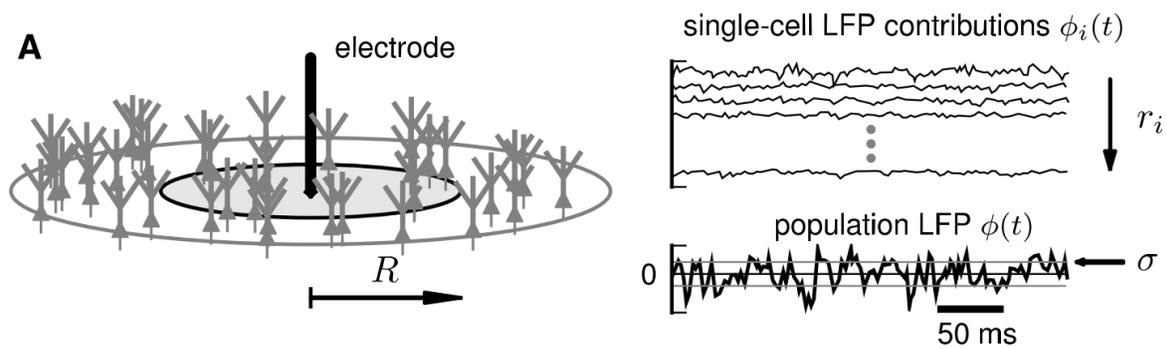


Figure 1.3. Tetzlaff, et.al. 2011 depiction of spiking and LFP detection. Electrodes inserted intracortically (left) can detect electrical activity due to single-neuron contributions (right, top) summing into a single “population” signal (right, bottom).

Classically, different LFP band oscillations have been linked to different brain states, such as delta (0.5-4Hz) with deep sleep, theta (4-8Hz) with movement, and alpha (8-12Hz) with sleep-wake cycles⁵³. However, newer studies have made it apparent that even within a given band, many more parameters can be decoded from the signal. For example, in primary motor cortex, one can find directional-related signals in delta, theta/alpha, as well as higher frequency

oscillations³⁵, further adding confusion as to how one should focus on interpreting aggregate LFP signals. One proposed method of reconciliation is the “communication through coherence” hypothesis³⁶, positing that matched-frequency neural oscillations between brain regions provide a communication pathway, through which spiking from an upstream region can more robustly promote spiking in a downstream region. Dual analyses of spiking and LFP in multiple regions during behavior will be necessary in the future to shed more light on how each represents network information transmission³⁷.

Analytical methods for measuring single- and cross-area communication

Given that spiking is considered a “ground truth” readout for neural activity, this section will outline some common methods utilized for interpretation of these data. Generally, there are two classes of spiking analyses, those that directly use time series of individual unit spiking, and those that use dimensionality reduction techniques on population spiking by leveraging relationships between sub-populations of unit firing.

One method of measuring spiking activity directly is population coupling, the measure of how likely an individual neuron is likely to spike relative to the local population spiking activity, giving an idea of whether neurons in a region are firing together or separately for a given behavior³⁸. Variability in neural response to a stimulus may be shared among neurons; thus another measure of correlating spiking activity between pairs of neurons across trials is via the noise correlation metric³⁹. With learning, spike-spike correlations may increase as neural patterns consolidate. Additional methods such as analyzing whether individual unit firing rates are tuned to a given stimulus or action direction have also been used, by analyzing peri-event time histograms (PETH) of firing around a given task or behavioral marker.

With the advent of electrophysiology tools enabling thousands of neurons to be recorded simultaneously, studies have demonstrated that the amount of information required to explain most of the data variance is primarily located in a lower dimensional space than the full n -dimensional space, where n is the total number of neurons recorded. Thus, there exist various methods of leveraging inherent relationships between individual and sub-population neural firing to describe neural activity in a low-dimensional manner. Principal component analysis (PCA),

extracts a set of orthogonal principal components that are linear combinations of the original neural firing, capturing the greatest variance in the data⁴⁰⁻⁴³. An extension of this analysis, demixed PCA (dPCA) weights a given component such that one component can maximize representation of a task variable, such as auditory stimulus or behavioral task choice⁴⁴. Factor analysis (FA) has also been extensively used due to improved capability of separating changes in overall spiking rate from changes in spiking variability, by identifying orthogonal factors that preserve variance that is shared across neurons and discarding variance that is independent to each neuron^{41,43,45}. An extension of FA is Gaussian process factor analysis (GPFA); instead of smoothing the data first and then applying dimensionality reduction techniques, GPFA is a process that simultaneously smooths and reduces the dimensionality of the data, thus allowing for single-trial comparisons of neural trajectory activations^{43,46}. Finally, while each of these techniques mainly examines single-area dimensionality reduction, recently methods for examining cross-area low-dimensional activations have emerged. One method is canonical correlation analysis (CCA), a dimensionality reduction technique that defines axes of maximal correlation between two populations of activity^{32,47,48}. By optimizing for dual-area covariance, one can better identify shared area communication as compared to other dimensionality-reduction techniques, such as PCA and FA, that optimize for local variance.

How do these direct measures of spiking and dimensionality-reduced metrics relate? Pairwise and population metrics utilize the same spike count covariance matrices; while few studies report how pairwise and population metrics relate, those that do demonstrate similarities³². Changes in spike count correlation could be due to strength of shared variability (i.e. strength of common input), intra-area connectivity (i.e. intra-area co-fluctuation), or the dimensionality of the neural

activity in the area (i.e. if the area receives multiple competing inputs). Additionally, the standard deviation of pairwise spike count correlations can be due to changes in population co-variability, providing a complementary view to just the strength of correlated activity as an explanation of neural communication.

In turn, there has recently been a movement towards defining low-dimensional activity spaces as neural manifolds (**Figure 1.4**). A neural manifold reflects covariance patterns across population activity, with patterns of neural activity arising from structural network connectivity^{38,49,50}. How do spiking manifolds relate to learning? Motor adaptation, and adaptation of neural

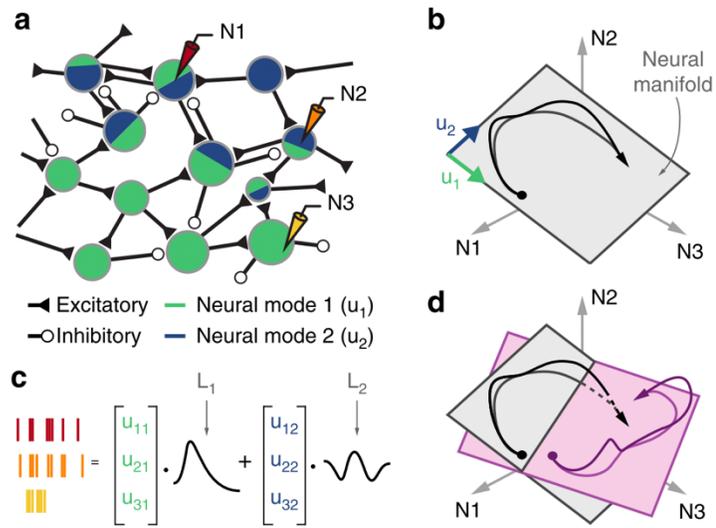


Figure 1.4. Gallego and Miller depiction of neural manifold. A) network of connected neurons of which we can experimentally sub-sample from. B) neural manifold derived directly from neural firing. C) depiction of dimensionality reduction of spiking into latent factors L1 and L2. D) Overlay of neural manifold from dimensionality reduction (pink) versus raw neural firing rates (grey); from Gallego...Miller, 2018.

patterns to external outputs (i.e., different reaching action, BCI cursor control) can occur quickly if the patterns of activity lie on the intrinsic manifold. Thus, within-manifold perturbations can be learned within a day. In contrast, off-manifold perturbations can be enforced with BCI learning paradigms. This learning of mapping between M1 activity and cursor control distinct from the original, intuitive mapping requires several days, but demonstrates a direct link between newly established activity and novel cursor control patterns^{45,51}. In the reach-to-grasp task in rodents, one can posit that a corticostriatal manifold develops with learning. In Chapter 2, I will

delve into whether transfer learning of the original skill is supported by the on- or off-manifold learning paradigm, as well as the implications of the result. In Chapter 3, I will then introduce a model for motor flexibility that will challenge the idea of M1 as a pattern-generating manifold, but rather one node of a larger motor network that can flexibly utilize on- or off-manifold patterns of activity to generate movement, along with the corresponding implications of corticostriatal network dysfunction.

References

1. Dayan, E., & Cohen, L. G. (2011, November 3). Neuroplasticity subserving motor skill learning. *Neuron*. Cell Press. <http://doi.org/10.1016/j.neuron.2011.10.008>
2. Graybiel, A. M., & Grafton, S. T. (2015). The striatum: Where skills and habits meet. *Cold Spring Harbor Perspectives in Biology*, 7(8), 1–14.
<http://doi.org/10.1101/cshperspect.a02169>
3. Ashby, F. G., Turner, B. O., & Horvitz, J. C. (2010). Cortical and basal ganglia contributions to habit learning and automaticity. *Trends in Cognitive Sciences*, 14(5), 208–215. <http://doi.org/10.1016/j.tics.2010.02.001>
4. Poldrack, R. A., Sabb, F. W., Foerde, K., Tom, S. M., Asarnow, R. F., Bookheimer, S. Y., & Knowlton, B. J. (2005). The neural correlates of motor skill automaticity. *Journal of Neuroscience*, 25(22), 5356–5364. <http://doi.org/10.1523/JNEUROSCI.3880-04.2005>
5. Lipton, D. M., Gonzales, B. J., & Citri, A. (2019). Dorsal striatal circuits for habits, compulsions and addictions. *Frontiers in Systems Neuroscience*, 13, 1–14.
<http://doi.org/10.3389/fnsys.2019.00028>
6. James, W. (1890). *The Principles of Psychology*.
7. Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Philosophical Transactions of the Royal Society B, Biological Sciences*.
<https://doi.org/10.1098/rstb.1985.0010>
8. Graybiel, A. M. (2008). Habits , Rituals , and the Evaluative Brain.
<http://doi.org/10.1146/annurev.neuro.29.051605.112851>
9. Robbins, T. W., & Costa, R. M. (2017). Habits. *Current Biology*, 27(22), R1200–R1206.
<http://doi.org/10.1016/j.cub.2017.09.060>

10. Lashley, K.S. Studies of Cerebral Function in Learning (VI). *Psychological Review*, 31(5), 369-375.
11. Desrochers, T. M., Amemori, K., & Graybiel, A. M. (2015). Habit Learning by Naive Macaques Is Marked by Response Sharpening of Striatal Neurons Representing the Cost and Outcome of Acquired Action Sequences. *Neuron*, 87(4), 853–868.
<http://doi.org/10.1016/j.neuron.2015.07.019>
12. Marien, H., Custers, R., & Aarts, H. (2019). Studying Human Habits in Societal Context: Examining Support for a Basic Stimulus–Response Mechanism. *Current Directions in Psychological Science*, 28(6), 614–618. <http://doi.org/10.1177/0963721419868211>
13. Sjoerds, Z., Dietrich, A., Deserno, L., de Wit, S., Villringer, A., Heinze, H. J., ... Horstmann, A. (2016). Slips of action and sequential decisions: A cross-validation study of tasks assessing habitual and goal-directed action control. *Frontiers in Behavioral Neuroscience*, 10,1–15. <http://doi.org/10.3389/fnbeh.2016.00234>
14. Smith, K. S., & Graybiel, A. M. (2013). A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron*, 79(2), 361–374.
<http://doi.org/10.1016/j.neuron.2013.05.038>
15. Smith, K. S., & Graybiel, A. M. (2016). Habit formation coincides with shifts in reinforcement representations in the sensorimotor striatum. *Journal of Neurophysiology*, 115(3), 1487–1498. <http://doi.org/10.1152/jn.00925.2015>
16. Gremel, C. M., & Costa, R. M. (2013). Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. *Nature Communications*, (May). <http://doi.org/10.1038/ncomms3264>

17. Vicente, A. M., Martins, G. J., & Costa, R. M. (2020). Cortico-basal ganglia circuits underlying dysfunctional control of motor behaviors in neuropsychiatric disorders. *Current Opinion in Genetics and Development*, 65, 151–159.
<http://doi.org/10.1016/j.gde.2020.05.042>
18. O'Hare, J. K., Ade, K. K., Sukharnikova, T., Van Hooser, S. D., Palmeri, M. L., Yin, H. H., & Calakos, N. (2016). Pathway-Specific Striatal Substrates for Habitual Behavior. *Neuron*, 89(3), 472–479. <http://doi.org/10.1016/j.neuron.2015.12.032>
19. Kurt, S., & Ehret, G. (2010). Auditory discrimination learning and knowledge transfer in mice depends on task difficulty, 2010, 21–24. <http://doi.org/10.1073/pnas.0912357107>
20. Andrade C, Alwarshetty M, Sudha S, Suresh Chandra J. (2001). Effect of innate direction bias on T-maze learning in rats: implications for research. *J Neurosci Methods*, 110(12):31-5. [http://doi.org/10.1016/s0165-0270\(01\)00415-0](http://doi.org/10.1016/s0165-0270(01)00415-0). PMID: 11564522.
21. Clark, E., Antoniak, K., Feniquito, A., & Dringenberg, H. C. (2017). Effects of the GluN2B-NMDA receptor antagonist Ro 25-6981 on two types of behavioral flexibility in rats. *Behavioural Brain Research*, 319, 225–233. <http://doi.org/10.1016/j.bbr.2016.11.032>
22. Wishaw, I. Q., & Tomie, J. A. (1989). Olfaction directs skilled forelimb reaching in the rat. *Behavioural Brain Research*, 32(1), 11–21.
[http://doi.org/10.1016/S01664328\(89\)80067-1](http://doi.org/10.1016/S01664328(89)80067-1)
23. Lemke, S. M., Ramanathan, D. S., Guo, L., Won, S. J., & Ganguly, K. (2019). Emergent modular neural control drives coordinated motor actions. *Nature Neuroscience*, 22(7), 1122–1131. <http://doi.org/10.1038/s41593-019-0407-2>

24. Sauerbrei, B. A., Guo, J. Z., Cohen, J. D., Mischiati, M., Guo, W., Kabra, M., ... Hantman, A. W. (2020). Cortical pattern generation during dexterous movement is input driven. *Nature*, 577(7790), 386–391. <http://doi.org/10.1038/s41586-019-1869-9>
25. Koralek, A. C., Long, J. D., Costa, R. M., & Carmena, J. M. (2010). Corticostriatal dynamics during learning and performance of a neuroprosthetic task. *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC'10*, 2682–2685. <http://doi.org/10.1109/IEMBS.2010.5626632>
26. Kawai, R., Markman, T., Poddar, R., Ko, R., Dhawale, A. K., Kampff, A.R., & Olveczky, B.P. (2015). Motor Cortex Is Required for Learning but Not for Executing a Motor Skill Article Motor Cortex Is Required for Learning but Not for Executing a Motor Skill, 800–812. <http://doi.org/10.1016/j.neuron.2015.03.024>
27. Kupferschmidt, D. A., Juczewski, K., Cui, G., Johnson, K. A., & Lovinger, D. M. (2017). Parallel, but Dissociable, Processing in Discrete Corticostriatal Inputs Encodes Skill Learning. *Neuron*, 96(2), 476–489.e5. <http://doi.org/10.1016/j.neuron.2017.09.040>
28. Hwang, E. J., Dahlen, J. E., Hu, Y. Y., Aguilar, K., & Yu, B. (2019). Disengagement of motor cortex from movement control during long-term learning, *Science Advances*, 5(10), 1–13. <http://doi.org/10.1126/sciadv.aay0001>
29. Dezfouli, A., & Balleine, B. W. (2012). Habits, action sequences, and reinforcement learning. *European Journal of Neuroscience*, 35(7), 1036–1051. <http://doi.org/10.1111/j.1460-9568.2012.08050.x.Habits>
30. Carelli, R. M., Wolske, M., & West, M. O. (1997). Loss of Lever Press-Related Firing of Rat Striatal Forelimb Neurons after Repeated Sessions in a Lever Pressing Task, *J Neurosci*, 17(5), 1804–1814.

31. Barnes, T. D., Kubota, Y., Hu, D., Jin, D. Z., & Graybiel, A. M. (2005). Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature*, *437*(7062), 1158–1161. <http://doi.org/10.1038/nature04053>
32. Veuthey, T. L., Derosier, K., Kondapavulur, S., & Ganguly, K. (n.d.). Single-trial cross-area neural population dynamics during long-term skill learning. *Nature Communications*, (2020), 1–15. <http://doi.org/10.1038/s41467-020-17902-1>
33. Tetzlaff, T., Potjans, T. C., Pettersen, K. H., & Gru, S. (2011). Modeling the Spatial Reach of the LFP. *Neuron*, *72*, 859–872. <http://doi.org/10.1016/j.neuron.2011.11.006>
34. Oliveira, S. C. De, Vaadia, E., Aertsen, A., Rotter, S., & Mehring, C. (2005). Encoding of Movement Direction in Different Frequency Ranges of Motor Cortical Local Field Potentials. *Journal of Neuroscience*, *25*(39), 8815–8824. <http://doi.org/10.1523/JNEUROSCI.0816-05.2005>
35. Oliveira, S. C. De, Vaadia, E., Aertsen, A., Rotter, S., & Mehring, C. (2005). Encoding of Movement Direction in Different Frequency Ranges of Motor Cortical Local Field Potentials, *25*(39), 8815–8824. <http://doi.org/10.1523/JNEUROSCI.0816-05.2005>
36. Akam, T. E., & Kullmann, D. M. (2012). Efficient “Communication through Coherence ” Requires Oscillations Structured to Minimize Interference between Signals. *PLOS Computational Biology*, *8*(11). <http://doi.org/10.1371/journal.pcbi.1002760>
37. Thorn, C. A., & Graybiel, A. M. (2014). Differential Entrainment and Learning-Related Dynamics of Spike and Local Field Potential Activity in the Sensorimotor and Associative Striatum, *34*(8), 2845–2859. <http://doi.org/10.1523/JNEUROSCI.1782-13.2014>

38. Okun, M., Steinmetz, N. A., Cossell, L., Iacaruso, M. F., Ko, H., Barthó, P., Moore, T., Hoyer, S.B., Mrsic-Flogel, T.D., Carandini, M., Harris, K. D. (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature*, *521*(7553), 511–515.
<http://doi.org/10.1038/nature14273>
39. Cohen, M. R., & Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nature Neuroscience*, *14*(7), 811–819. <http://doi.org/10.1038/nn.2842>
40. Mazor, O., & Laurent, G. (2005). Transient Dynamics versus Fixed Points in Odor Representations by Locust Antennal Lobe Projection Neurons. *Neuron*, *48*, 661–673.
<http://doi.org/10.1016/j.neuron.2005.09.032>
41. 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. <http://doi.org/10.1038/nature11129>
42. Ahrens, M. B., Li, J. M., Orger, M. B., Robson, D. N., Schier, A. F., Engert, F., & Portugues, R. (2012). Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature*, *485*(7399), 471–477. <http://doi.org/10.1038/nature11057>
43. Cunningham, J. P., & Yu, B. M. (2015). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, *17*(11), 1500–1509.
<http://doi.org/10.1038/nn.3776.Dimensionality>
44. Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C. E., Kepecs, A., Mainen, Z. F., Machens, C. K. (2016). Demixed principal component analysis of neural population data. *eLife*, *5*, 1–36. <http://doi.org/10.7554/eLife.10989>

45. Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P. (2014). Neural constraints on learning. *Nature*, *512*(7515), 423–426.
<http://doi.org/10.1038/nature13665>
46. Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, *102*(1), 614–635.
<http://doi.org/10.1152/jn.90941.2008>
47. Morcos, A. S., Raghu, M., & Bengio, S. (2018). Insights on representational similarity in neural networks with canonical correlation. *arXiv preprint*. arXiv:1806.05759.
48. Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M., & Kohn, A. (2019). Cortical Areas Interact through a Communication Subspace. *Neuron*, *102*(1), 249–259.e4.
<http://doi.org/10.1016/j.neuron.2019.01.026>
49. Gallego, J. A., Perich, M. G., Naufel, S. N., Ethier, C., Solla, S. A., & Miller, L. E. (2018). Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications*, *9*(1), 1–13. <http://doi.org/10.1038/s41467-018-06560-z>
50. Perich, M. G., Gallego, J. A., & Miller, L. E. (2018). A Neural Population Mechanism for Rapid Learning. *Neuron*, *100*(4), 964–976. <http://doi.org/10.1016/j.neuron.2018.09.030>
51. Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., ... Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(30), 15210–15215. <http://doi.org/10.1073/pnas.1820296116>

52. Jackson, A., & Hall, T. M. (2017). Decoding Local Field Potentials for Neural Interfaces. *IEEE transactions on neural systems and rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society*, 25(10), 1705–1714. <https://doi.org/10.1109/TNSRE.2016.2612001>
53. Destexhe, A., Contreras, D., & Steriade, M. (1999). Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *Journal of Neuroscience*, 19(11), 4595–4608. <https://doi.org/10.1523/JNEUROSCI.19-11-04595.1999>

Chapter 2: Breakdown of corticostriatal patterns underlies motor transfer learning

Abstract

Applying prior learned actions to novel contexts is critical for survival. Motor learning has been historically demonstrated to establish consistent neural patterns in primary motor cortex (M1), and more recently downstream structures as well, such as the dorsolateral striatum (DLS). However, when a single motor task has been learned in a single environmental context, it is unknown how the transfer of that learned motor task to a modified environment occurs, with competing possibilities of quick adaptation of existing neural patterns versus longer term re-learning involving aspects of an early learning process. Here, we overtrained rats to a single location in the reach-to-grasp task, and then studied how M1 and DLS activity patterns changed with reaching to a second pellet location. We found that early post-switch, reaching to the previously learned location persisted, accompanied by consistent patterns of activity in M1 and DLS. However, early in transfer learning, there was a state shift to variable patterns of activity supporting smooth reaching with variable endpoints for grasp between the old and new locations; subsequently, patterned activity returned in both regions. Together, these findings demonstrate that from an automatic motor state, establishing behavioral flexibility involves breakdown of corticostriatal spiking patterns, subsequently re-emerging with re-learning of the related skill.

Introduction

Motor cortex has been characterized as an engine for precise movement control through the generation of reliable neural activity patterns¹⁻⁴. Yet, in real world scenarios, organisms must also be capable of generating variable exploratory movements in order to adjust to changes in the environment. It is unclear how the primary motor cortex (M1) and downstream subcortical areas enable both exploratory variable movements as well as highly predictable movements in a well-known environment. For example, a prominent current hypothesis is that stable M1 population activity patterns -- for example sequential firing of neural activity^{2,5-7}-- are important for the production of reliable behaviors. Reliable sequencing of neural activity appears to emerge through the process of learning and consolidation^{3,6-8}; such activity is also associated with greater coupling between M1 and subcortical regions such as the dorsolateral striatum (DLS)^{4,7,9}. It remains unknown how established activity patterns across cortical and subcortical regions are altered in a setting where generated movements no longer result in successful outcomes.

How then might M1 and downstream structures allow for more flexible exploration when previously successful movement output must be adapted for reward? One possibility is that prior learning established a generalizable neural manifold in M1, such that modified movement patterns are largely similar to the previous ones, with minor differences in unit tuning. A second possibility is that while M1 reverts to an exploratory, variable state, downstream structures such as DLS maintain firing patterns to support robust learned movement. Finally, M1 and DLS could be functioning in tandem after initial learning of a skilled movement, re-establishing exploratory neural activity patterns together to enable behavioral flexibility.

In this study, we probed how behavioral flexibility is established by training rodents to automatic performance on the reach-to-grasp skill for a single location, and then switched the pellet reward location. We found that this transfer of learning was a multi-day process; we did not find evidence for rapid within-session adaptation. Establishment of initial behavioral flexibility during transfer learning corresponded with a return to an exploratory neural state in M1 and DLS, marked by decrease in consistency of task-modulated units and reaching spiking patterns, as well as a reduction of the coordination of cross-area spiking patterns. Subsequently, M1 and DLS exhibited coordinated re-establishment of task-aligned spiking structure, further adding evidence towards the third transfer learning hypothesis. Together, our results indicate that from an automatic motor state, establishing behavioral flexibility in a complex motor skill involves breakdown of M1-DLS spiking structure at single-unit, population, and shared communication levels, subsequently re-emerging with successful transfer learning.

Results

Transfer learning of an automatic skilled motor behavior is a multi-day process, with individuals exhibiting variability.

We recorded single-unit activity and local field potentials (LFPs) in M1 and DLS as rats ($n = 6$ animals) learned a reach-to-grasp skill. We then continued recording neural activity after the reach location was switched. Rats were trained for at least 1000 trials to Location A in automated behavioral boxes (**Figure 2.1A**)³⁶, after which the reach location was switched to Location B, 50-150 trials per day, for 3-5 days (**Figure 2.1B,C**). Decay curves of automatic reaches to A were fit for each animal (**Figure 2.1D, Figure 2.2A**), and the tau decay constant was correlated with how many trials to A rats had been exposed to (**Figure 2.2B**, $R^2 = 0.892$, $p = 0.00455$). Accuracy of each trial was dependent on both correct reach location, as well as dexterous grasping for reward (**Figure 2.2C**). For each session, we examined reach duration, defined as time from first reach to first grasp (**Figure 2.2D**), and reaction time, defined as time from trial start to first reach (**Figure 2.2E**). When examining sessions after switch of pellet location, there was a slight increase and subsequent decrease of reach to grasp duration, with a corresponding increase and decrease of normalized reaction time, from baseline to three and four sessions after switch, respectively (**Figure 2.2F-G**).

From a macro level, we observed that switching reaching to location B was a multi-day process (**Figure 2.1C,D**) that involved a transition from initiating trials to A to initiating trials to B, with little evidence of within-trial switching and or outcome-based switching from trial-to-trial . To assess for within-trial switching we analyzed whether extra reaches within a trial occurred (**Figure 2.3A, Figure 2.4A**). Subsequently, we characterized types of reach after the first reach,

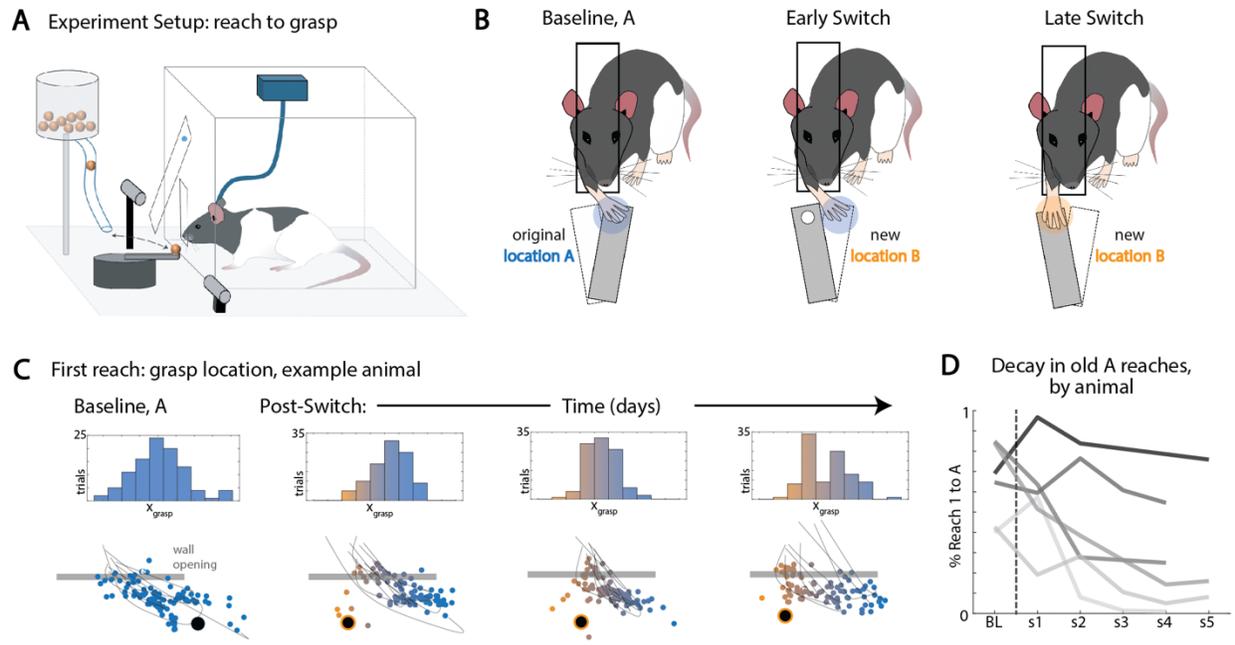


Figure 2.1. *Transfer learning of an automatic, skilled behavior.* (A) Reach-to-grasp experiment setup. A pellet is retrieved automatically from the dispenser and moved to reach position. A cue is then delivered, after which the door opens and rat reaches through a slit to retrieve the pellet. (B) Transfer learning paradigm. Rats are over-trained to reach to location A to retrieve the pellet reward (left). Then the pellet is moved to location B for retrieval (middle) over multiple days (right). (C) Top: histogram of x-position of grasp, across trials. Bottom: reach trajectories (lines) and grasp locations (small dots) relative to pellet location (large black circle). (D) Percentage of trials in a session with first reach to location A, as compared to low-amplitude reaches and/or reaches to location B.

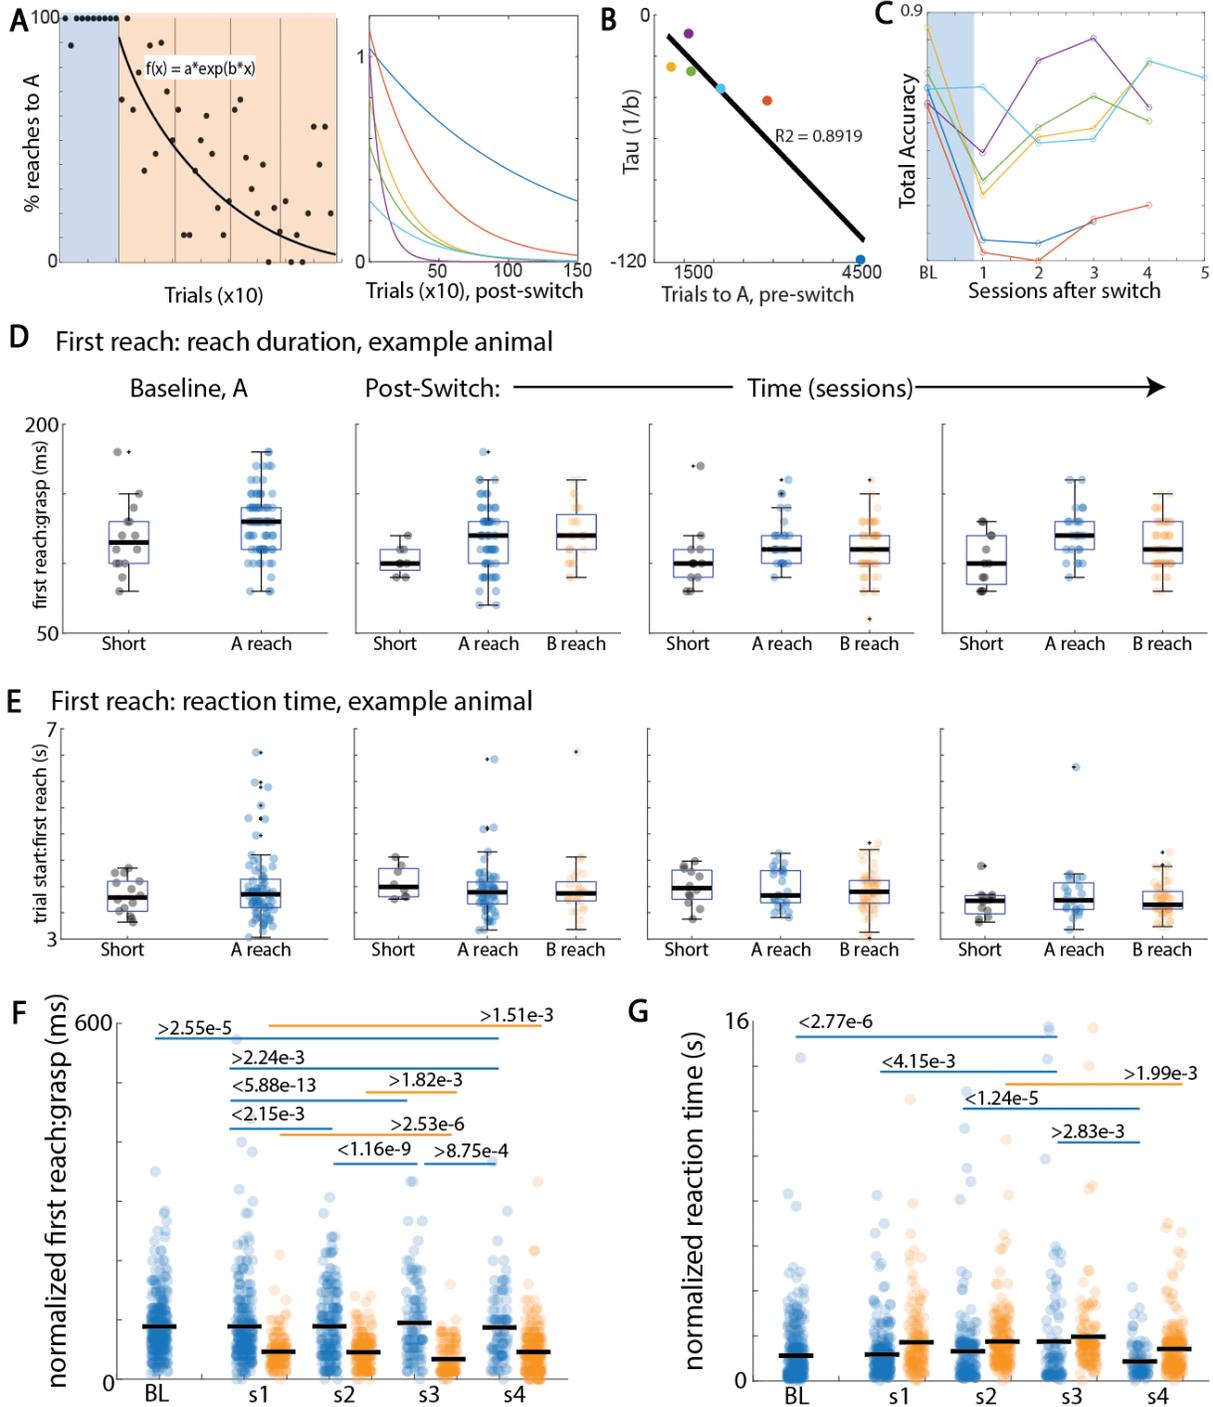


Figure 2.2. Breaking out of an automatic, skilled motor behavior is a multi-day process, with individuals exhibiting variability. (A) Percentage of reaches to A, per 10 trials, with exponential decay curve fit, for one animal (left), and across animals (right). (B) Slope of decay, tau, versus trials of training to A prior to switch. (C) Total accuracy across sessions, as defined by percent of trials in a session with successful pellet retrieval. (D) Reach duration, defined as first reach onset to first grasp time, across sessions for example animal. (E) Reaction time, as defined by reach onset time relative to trial start, example animal. (F) Normalized first reach to first grasp time across animals and sessions. (G) Normalized reaction time across animals and sessions.

within a trial (**Figure 2.3B, Figure 2.4B**). In the first session post-switch, rats perseverated in reaching to A as opposed to switching to B (**Figure 2.3C**, ANOVA, A vs. S, $p=4.3e-4$; A vs. B, $p=0.023$; **Figure 2.4C**). This perseveration did not happen on the few trials when reaches to B occurred (**Figure 2.4D**). Thus, when animals initiated a reach to A or B, they continued to reach to the same location within a trial. Finally, we probed whether early post-switch error-updating happened from trial to trial based on accurate retrieval of a pellet (**Figure 2.3D**). After reaching to B accurately, perseveration to A on the next trial's first reach continued for the first session post-switch (**Figure 2.3E**, ANOVA, A vs. S, $p=4.4e-3$; A vs. B, $p=0.057$). This was also true for inaccurate trials where the final reach was to A (**Figure 2.4E**), indicating that early on animals did not adjust their strategy based on recently received or missed rewards. Of note, in the low number of trials early post-switch where rats reached inaccurately to B, there were more reaches to A than to B on the subsequent trial, although not significant (**Figure 2.4F**). Overall, this indicated that early post-switch, successfully transfer learning to B did not happen within trials, or even across trials based on the previous trials' outcome, but rather was a multi-day process that required learning to consistently initiate trials to B.

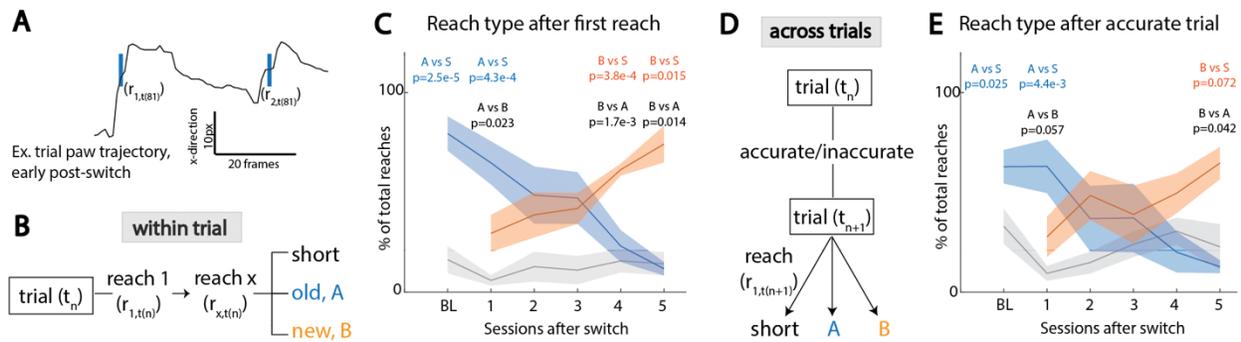


Figure 2.3. In a rigid state, error-based updating occurs neither within trials nor across trials. (A) Example x-trajectory of paw during a trial, with reach onset marked. (B) Within-trial updating for reaches after the first were classified into low-amplitude (short), old (A), or new (B) reaches. (C) For all reaches after the first reach in a trial, percentage of A, B, or short (S) reaches. (D) Across-trial updating of first reach was divided on whether the previous trial was accurate with successful pellet retrieval, or inaccurate. (E) For all reaches after accurate pellet retrieval, percentage of subsequent trial first reaches to A, B, or S.

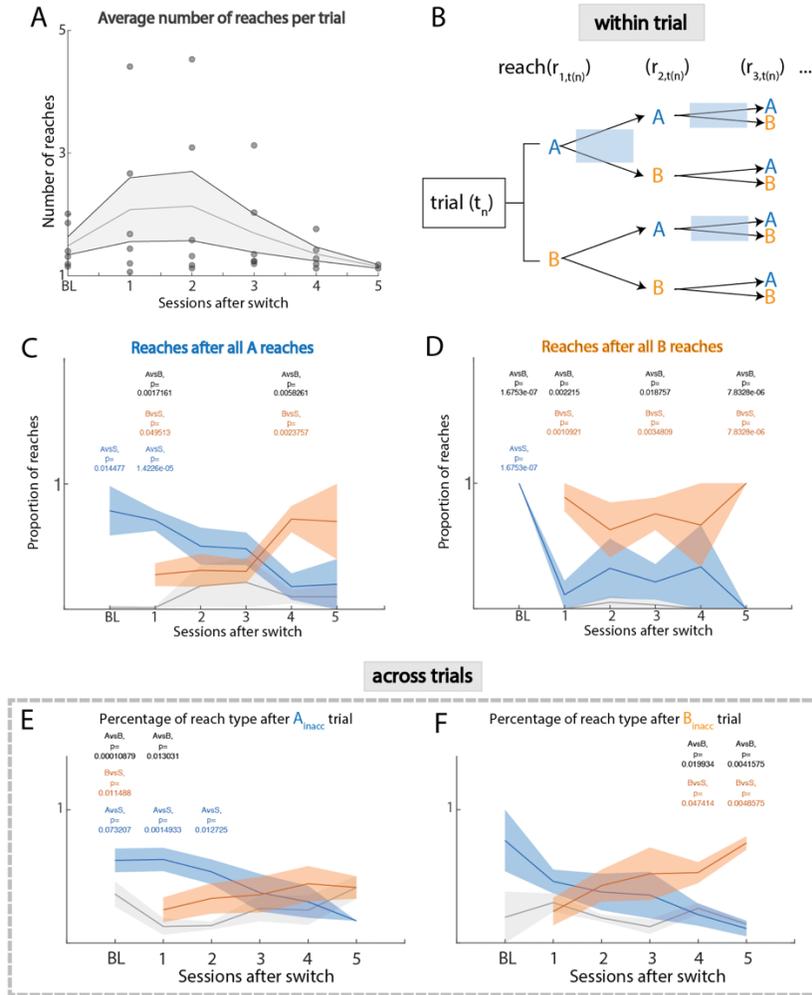


Figure 2.4. In a rigid state, error-based updating occurs across days. (A) Average number of reaches per trial, across sessions. (B) Within-trial updating for reaches after the first were classified into A, B, or low-amplitude (S, not shown). (C) Proportion of A, B, or S reaches for all within-trial reaches after a reach to A. (D) Proportion of A, B, or S reaches for all within-trial reaches after a reach to B. (E) Proportion of first A, B, or S reaches after inaccurate trial with final reach to A. (F) Proportion of first A, B, or S reaches after inaccurate trial with final reach to B.

Breakdown of reach-locked M1 and DLS unit spiking underlies transfer learning of reach-to-grasp skill

In a well-learned task, cortical and striatal units have predictable firing rate dynamics relative to the kinematics of behavior. However, it is unknown how this predictability of unit dynamics changes with transfer learning. When tracking M1 units over time, at sessions early to intermediate after the switch of pellet to B, we saw a breakdown of unit firing rate modulation aligned to reach onset (RO) across all trials in a session (**Figure 2.5A, Figure 2.7A**), despite rodents completing smooth, fast reaches (**Figure 2.2G**). While cortical dynamics of the reach-to-grasp task break down over learning a second similar skill, one might predict that striatal dynamics do not, given the similarity of gross movement and kinematics for reaches to A and B (**Figure 2.1C, Figure 2.2D,E**). However, DLS units also exhibited a similar pattern as those in M1, simultaneously dropping RO modulation on the same sessions during which M1 spiking modulation decreased (**Figure 2.5B, Figure 2.7B**). In turn, sessions were categorized into the following based on task parameters, behavior curves, and Fano factor (**Figure 2.5C**) of unit spiking activity (see Methods, Spiking analysis: Determination of session type): 1) Baseline (BL), pellet at location A, reaching to A, consistent reach-modulated spiking activity; 2) Automatic (auto), pellet at location B, reaching mostly to A, consistent reach-modulated spiking activity; 3) Variable (var), pellet at location B, reaching to A and B, local maximum of lowest 5 Fano factors per unit, aggregated, with significant deviation from baseline minimum Fano factors; 4) Relearned (rel), pellet at location B, reaching mostly to B, consistent reach-modulated spiking activity (**Figure 2.5D,E**).

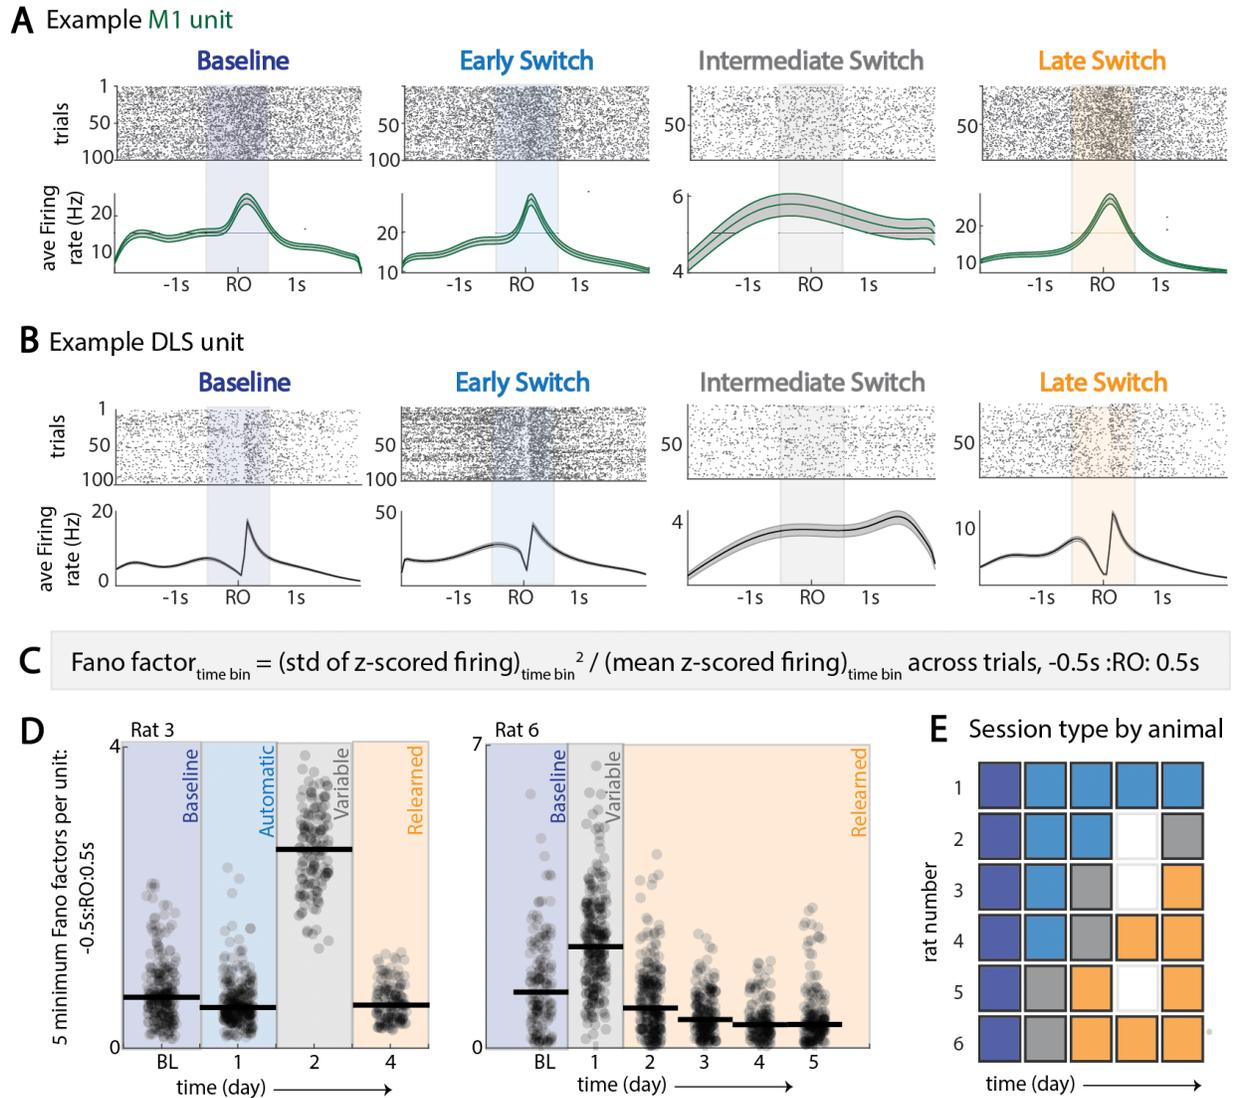


Figure 2.5. Breakdown of reach-locked M1 and DLS unit spiking underlies transfer learning of reach-to-grasp skill. (A) Top: trial rasters for a single example M1 unit, matched across sessions. Bottom: trial-averaged firing rate of the same unit (mean \pm SEM) for -2s to 2s around first reach onset. (B) Same as A, for an example DLS unit. (C) Equation for determining Fano factors for each 50ms time bin, with 80 total Fano factors per unit, per session. (D) Five minimum Fano factors per unit for -0.5s:RO:0.5s around first reach onset, across sessions, for two example animals. (E) Session type re-categorization based on Fano factors: indigo = baseline, blue = automatic (post-switch, before increase in Fano factors), gray = variable (post-switch, maximum increase in Fano factors), orange = relearned (post-switch, subsequent decrease in Fano factors).

Trial-averaged spiking modulation was calculated for each neuron in one session type per animal, with modulation defined as the z-scored absolute peak of firing rate between -0.5s and 0.5s around reach onset. In M1, there was a significant drop in trial-averaged unit modulation on the variable session for trials with first reach to A (**Figure 2.6A**, first reach to A (mean +/- sem), BL: 1.71 +/- 0.080; auto: 1.89 +/- 0.093; var: 0.801 +/- 0.057; rel: 1.76 +/- 0.095; linear mixed effects model p-values, BL vs var: 1.44e-14; auto vs. var: 2.66e-17; rel vs. var: 1.24e-15). Similarly, for trials with first reach to B, there was a drop in M1 reach-locked unit modulation on the variable session (**Figure 2.6A**, first reach to B (mean +/- sem), auto: 1.78 +/- 0.088; var: 1.02 +/- 0.064; rel: 1.75 +/- 0.11; linear mixed effects model p-values, auto vs. var: 4.48e-10; rel vs. var: 9.63e-10).

In DLS, the same pattern of unit spiking losing reach-modulation on the variable session was observed for both trials with first reach to A and trials with first reach to B (**Figure 2.6B**, first reach to A (mean +/- sem), BL: 1.32 +/- 0.11; auto: 1.78 +/- 0.099; var: 0.829 +/- 0.057; rel: 1.52 +/- 0.080; linear mixed effects model p-values, BL vs var: 3.46e-4; BL vs. rel: 3.60e-3; auto vs. var: 3.45e-12; rel vs. var: 3.19e-12; first reach to B (mean +/- sem), auto: 1.61 +/- 0.12; var: 0.811 +/- 0.046; rel: 1.62 +/- 0.086; linear mixed effects model p-values, auto vs. var: 2.29e-7; rel vs. var: 2.20e-14). Importantly, M1 and DLS average task firing rate does not change across session types (**Figure 2.7C,D**), and increased reach-related firing rate relative to non-reach periods per unit is preserved on the variable sessions (**Figure 2.7E**; M1 units, paired t-test, $p=1.17e-4$; DLS units, paired t-test, $p=8.98e-6$).

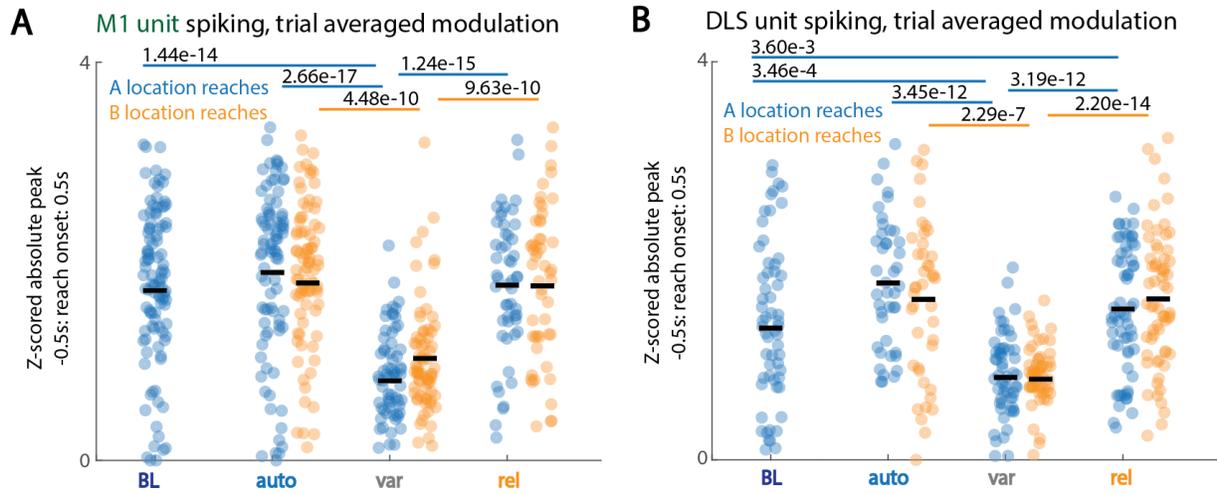


Figure 2.6. Both M1 and DLS units lose reach modulation of spiking activity with transfer learning. (A) Trial-averaged z-scored modulation peak per M1 unit across re-categorized session types, averaged separately for trials with first reach to A (blue dots) and trials with first reach to B (orange dots). (B) Same as A, for DLS units across session types.

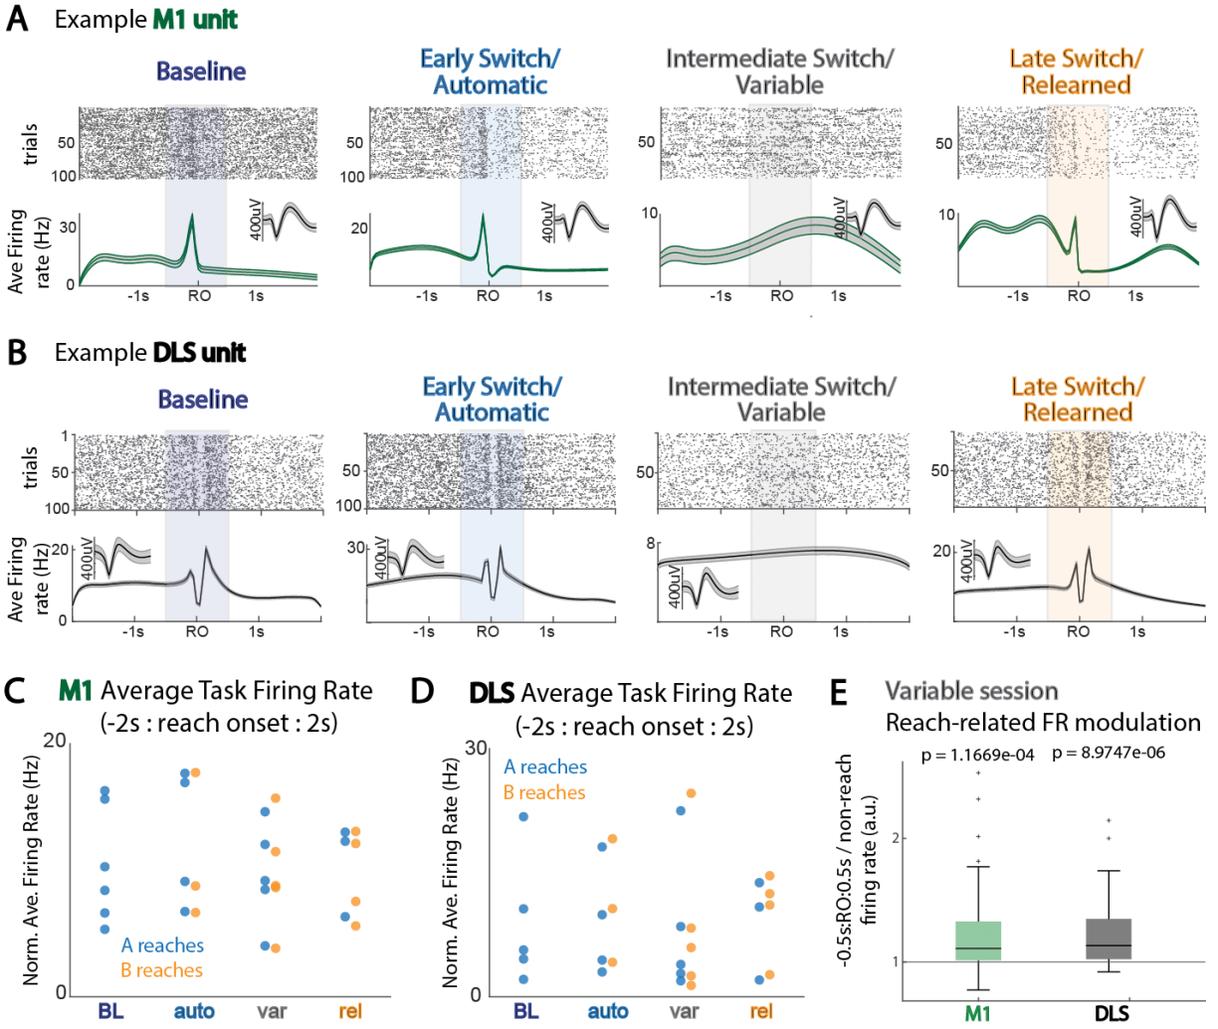


Figure 2.7. While unit modulation drops, M1 and DLS unit task-averaged activity does not change by session type. (A) Top: trial rasters for a single example M1 unit, matched across sessions. Bottom: trial-averaged firing rate of the same unit (mean \pm SEM) for -2s to 2s around first reach onset. (B) Same as A, for an example DLS unit. (C) Normalized average firing rate for all M1 units in an animal for A reaches (blue) and B reaches (orange) for -2s to 2s around reach onset. (D) Same as C, for DLS units. (E) Unit modulation during reach period, -0.5s to 0.5s around reach onset, relative to non-reach periods, for M1 (green) and DLS (black) units, paired t-test.

M1 and DLS temporal patterning re-emerges following variable session.

A hallmark of crystallized skills is a consistent temporal neural pattern of activation², which can be visualized in both a trial-averaged manner as well as in a single trial (**Figure 2.8A**).

Previously we have demonstrated a drop in single-unit modulation relative to reach onset (**Figure 2.5-2.7**). However, there could still be an underlying temporal structure in spiking activity from trial to trial. Consistency of pattern activation in a region can be assessed via correlation of a trial spiking sequence to the trial-averaged spiking sequence for a session (**Figure 2.8A**). A reach trials were compared to the template for all reaches to A; the same process was completed for B reach trials in a session, as described in Methods, Spiking analysis: Template matching. Strikingly, sequential activity similarly broke down on the variable session (**Figure 2.8B**), reminiscent of an early skill learning state^{4,7}. In M1, when comparing first reach to A spiking to the average neural template for first reaches to A, there was a drop in neural pattern consistency during the variable session as compared to the other sessions; interestingly, there was a slight increase in neural pattern consistency from baseline to the automatic and relearned sessions (**Figure 2.8C**, first reach to A (mean +/- sem), BL: 0.267 +/- 0.0096; auto: 0.342 +/- 0.0085; var: -0.0480 +/- 0.0094; rel: 0.356 +/- 0.014; linear mixed effects model p-values, BL vs auto: 1.07e-4; BL vs. var: 4.34e-78; BL vs. rel: 9.32e-7; auto vs. var: 1.20e-100; rel vs. var: 8.42e-69). For first reaches to B, M1 spiking activity was similarly more temporally consistent for the automatic and relearned sessions as compared to the variable sessions (**Figure 2.8C**, first reach to B (mean +/- sem), auto: 0.280 +/- 0.015; var: 0.0450 +/- 0.013; rel: 0.0322 +/- 0.0092; linear mixed effects model p-values, auto vs. var: 3.34e-8; rel vs. var: 3.31e-55).

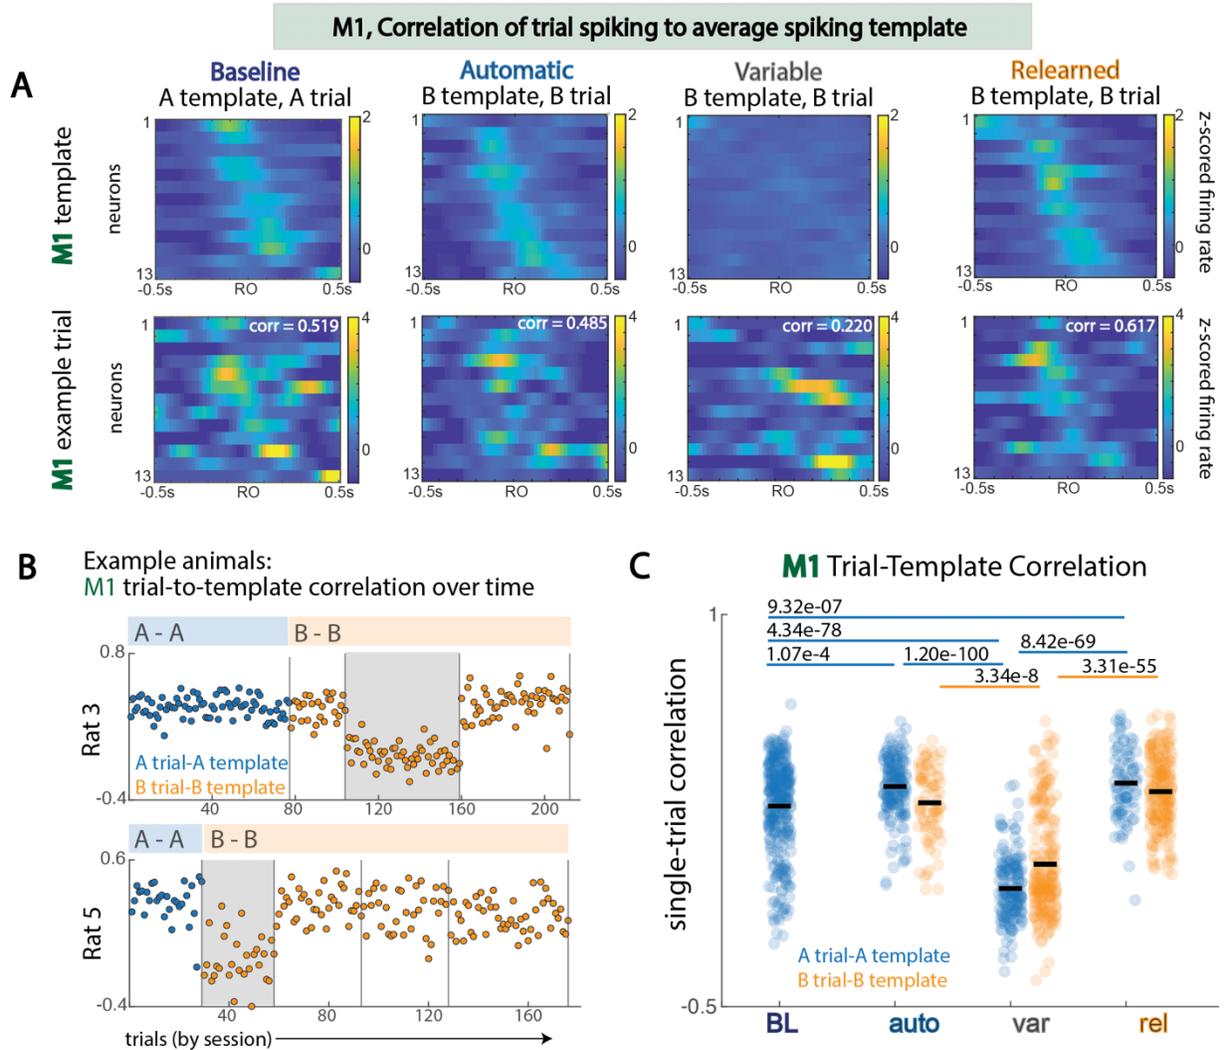


Figure 2.8. M1 temporal patterning breaks down and re-emerges with transfer learning. (A) Top: leave-one-out trial-averaged spiking pattern of M1 units for reach period, -0.5s to 0.5s around first reach onset, across session types. Bottom: left-out trial spiking correlation with corresponding session template. (B) Trial-to-template correlation across trials, with A trial-to-template correlations shown for the baseline session, and B trial-to-template correlations shown for subsequent sessions. Variable session is shaded in grey. (C) M1 trial-to-template correlations for all trials by session type, A trial-A template correlations in blue and B trial-B template correlations in orange.

This same pattern of changes in trial-to-template correlation across session types was observed in DLS (**Figure 2.9A**). Specifically, for first reaches to A and B, there was a decrease in spiking pattern consistency during the variable sessions, with a slight increase in pattern consistency for first reaches to A on the automatic and relearned sessions relative to baseline (**Figure 2.9B**, first reach to A (mean +/- sem), BL: 0.170 +/- 0.012; auto: 0.184 +/- 0.011; var: -0.0662 +/- 0.011; rel: 0.240 +/- 0.013; linear mixed effects model p-values, BL vs auto: 2.33e-3; BL vs. var: 3.23e-34; BL vs. rel: 9.54e-5; auto vs. var: 8.12e-20; rel vs. var: 4.04e-43; first reach to B (mean +/- sem), auto: 0.193 +/- 0.020; var: -0.0237 +/- 0.0088; rel: 0.2208 +/- 0.011; linear mixed effects model p-values, auto vs. var: 2.05e-25; rel vs. var: 3.35e-55). While the variability in pattern on the variable session can be inferred from the variability in spiking on this session (**Figure 2.5-2.7**), this analysis emphasizes that not only does spiking return to reach-locked modulation, but that there's a consistent temporal ordering of unit modulation soon after, indicating that consolidation of an optimal neural pattern follows the exploratory neural state in M1 and DLS.

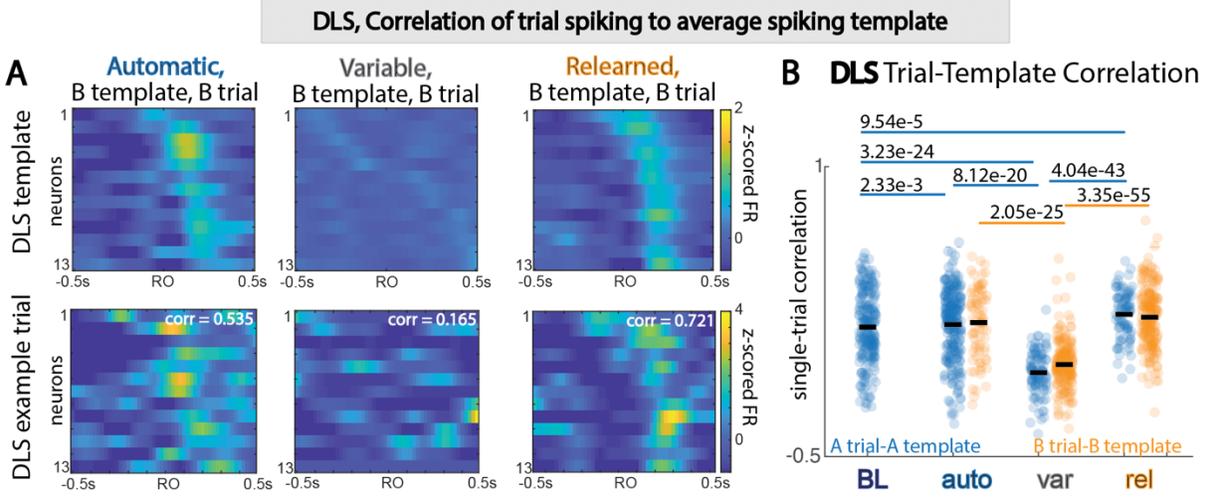


Figure 2.9. DLS temporal patterning also re-emerges following transfer learning. A) Top: leave-one-out trial-averaged spiking pattern of DLS units for reach period, -0.5s to 0.5s around first reach onset, across session types. Bottom: left-out trial spiking correlation with corresponding session template. (B) DLS trial-to-template correlations for all trials by session type.

M1 and DLS single-trial population modulation becomes less temporally consistent during transfer learning

While individual units lose reach-specific modulation of activity, there exist two possibilities: 1) the majority of spiking occurs with relation to trial start cues; 2) in an exploratory state, the majority of population spiking bursts are not aligned to the trial structure (i.e. trial start, reaching movement), despite overall increased firing in M1 and DLS during reach (**Figure 2.7E**). Thus, we examined M1 and DLS single-trial population spiking activity throughout the trial across the different session types (**Figure 2.10A,B**) demonstrating no consistent modulation relative to trial start across sessions.

On the variable session there was a drop in reach-related population spiking activity modulation, defined as z-scored peak of trial population activity -0.5s to 0.5s around first reach onset, in both M1 (**Figure 2.10C,D, Figure 2.11A,C**, first reach to A (mean +/- sem), BL: 2.12 +/- 0.0029; auto: 2.33 +/- 0.0026; var: 1.08 +/- 0.0060; rel: 2.24 +/- 0.0080; linear mixed effects model p-values, BL vs var: 6.87e-31; auto vs. var: 4.53e-33; rel vs. var: 4.72e-13; first reach to B (mean +/- sem), auto: 2.51 +/- 0.0062; var: 1.59 +/- 0.0045; rel: 1.98 +/- 0.0038; linear mixed effects model p-values, auto vs. var: 2.02e-5; rel vs. var: 4.27e-6) and DLS (**Figure 2.10C,D, 2.11B,D**, first reach to A (mean +/- sem), BL: 1.71 +/- 0.0033; auto: 1.99 +/- 0.0028; var: 1.28 +/- 0.0053; rel: 1.94 +/- 0.0078; linear mixed effects model p-values, BL vs. auto: 1.97e-5; BL vs var: 6.90e-3; auto vs. var: 3.14e-11; rel vs. var: 2.97e-9; first reach to B (mean +/- sem), auto: 2.09 +/- 0.0078; var: 1.34 +/- 0.0041; rel: 2.05 +/- 0.0034; linear mixed effects model p-values, auto vs. var: 2.97e-11; rel vs. var: 7.05e-19). However, when looking across the entire task period of -2s

to 2s around reach onset, there was little to no variation in peak firing rate in either M1 or DLS across session types (**Figure 2.11E,F**).

Is the drop in trial-averaged spiking modulation due to a shift in peak timing, as hinted by observation of trial firing (**Figure 2.10C,D**)? While peak firing rate throughout the task period remained relatively unchanged, for both A and B reaching, the time at which the peak of population firing occurred was significantly more spread out over the trial period during the variable session. This was true in both M1 (**Figure 2.10E**, first reach to A (mean time in seconds relative to reach onset +/- sem), BL: -0.0976 +/- 0.0023; auto: -0.00960 +/- 0.0022; var: -0.201 +/- 0.0072; rel: -0.202 +/- 0.0039; Bartlett's test p-values, BL vs. auto: 4.50e-3; BL vs var: 1.69e-19; auto vs. var: 1.05e-28; rel vs. var: 5.10e-11; **Figure 2.10E**, first reach to B (mean +/- sem), auto: -0.0250 +/- 0.055; var: -0.0330 +/- 0.049; rel: -0.328 +/- 0.0035; Bartlett's test p-values, auto vs. var: 6.50e-15; auto vs. rel: 2.73e-7; rel vs. var: 4.66e-6) and DLS (**Figure 2.10F**, first reach to A (mean +/- sem), BL: 0.0684 +/- 0.0032; auto: 0.0998 +/- 0.0035; var: -0.199 +/- 0.0070; rel: 0.0747 +/- 0.0090; Bartlett's test p-values, BL vs var: 5.97e-4; auto vs. var: 2.26e-5; rel vs. var: 1.38e-4; **Figure 2.10F**, first reach to B (mean +/- sem), auto: -0.0390 +/- 0.0092; var: -0.0637 +/- 0.0050; rel: -0.0671 +/- 0.0038; Bartlett's test p-values, auto vs. var: 4.67e-5; rel vs. var: 2.24e-6). Overall, this indicated that while M1 and DLS population spiking retained modulation, it became less RO-aligned, as examined on a trial-by-trial level, during the variable session early in transfer learning.

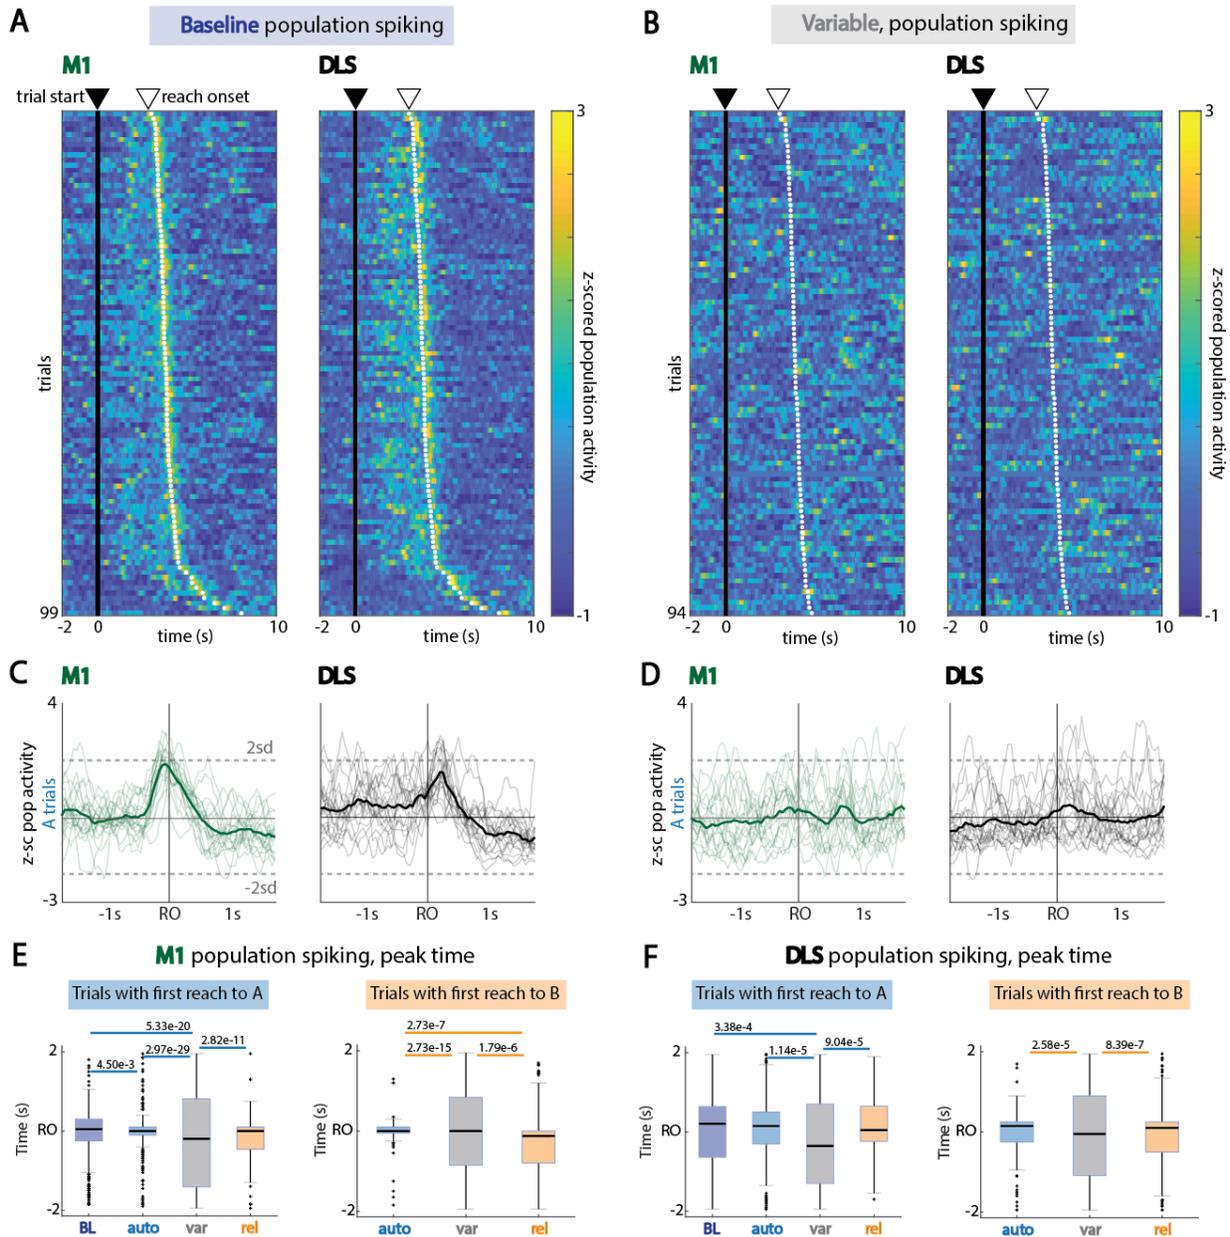


Figure 2.10. M1 and DLS single-trial population spiking activity loses temporal consistency with transfer learning. (A) M1 (left) and DLS (right) z-scored population spiking across entire trial, for all trials in a baseline session for an example animal. (B) Same as A, for example animal variable session. (C) Baseline session, example A reach trials, z-scored population spiking in M1 (left) and DLS (right) for shortened trial window, -2s to 2s around first reach onset. (D) Same as C, for example variable session trials with first reach to A. (E) Time of peak population spiking within -2s to 2s around first reach onset in M1 units for trials with first reach to A (left), and trials with first reach to B (right). (F) Same as E, for DLS units across session types.

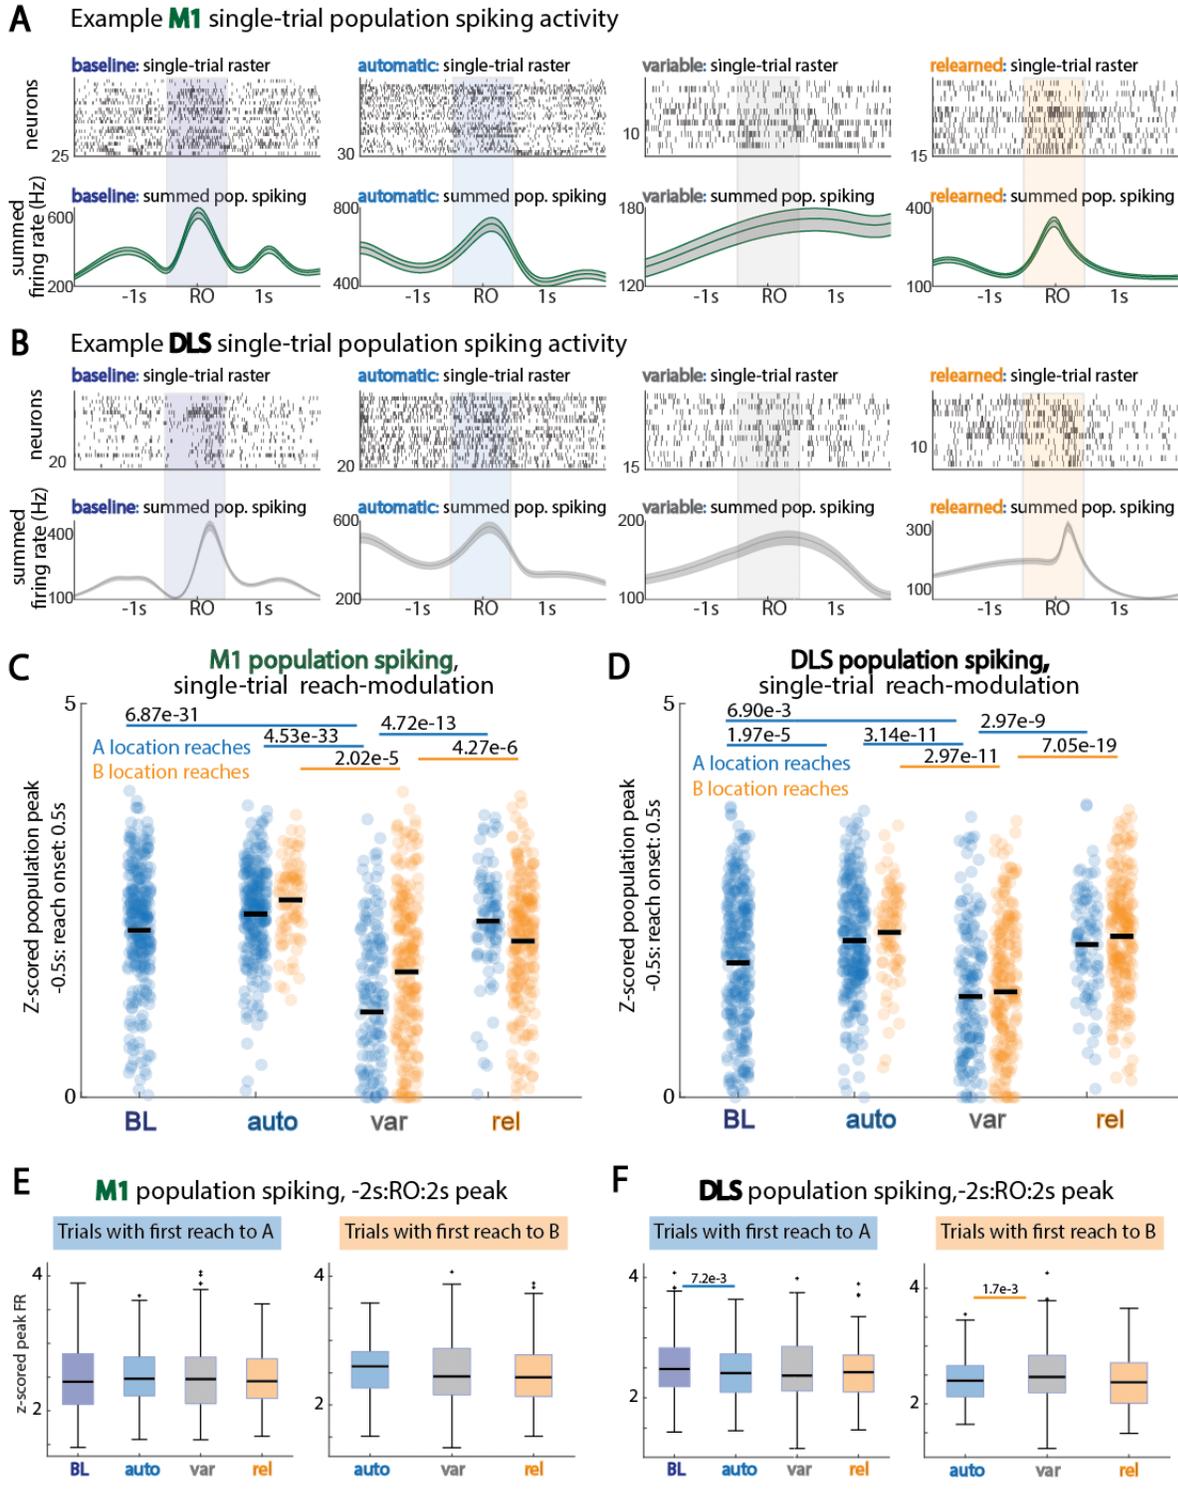


Figure 2.11. M1 and DLS single-trial reach-modulated population spiking breakdown underlies transfer learning. (A) Top: M1 unit rasters for a single example trial, across session types. Bottom: Summed M1 population spiking for the corresponding trial. (B) Same as A, for DLS population and example trials across session types. (C) M1 single-trial population spiking modulation during reach period, -0.5s to 0.5s around first reach onset, across session types for trials with first reach to A (blue) and trials with first reach to B (orange). (D) Same as C, for DLS single-trial population spiking across session types. (E) M1 single-trial population spiking modulation over expanded trial window, -2s to 2s around first reach onset, across session types. (F) Same as E, for DLS single-trial population spiking modulation over expanded trial window.

M1-DLS task population spiking drops correlation with preservation of modulation during transfer learning

Given that M1 and DLS are monosynaptically connected structures implicated in motor skill formation and consolidation, which independently demonstrated breakdown and re-emergence of task activity, we then examined how correlation of M1 and DLS population spiking activity evolved over session types. During the variable session, M1 and DLS spiking correlation drops during the task period of -2s to 2s around reach onset (**Figure 2.12A-C**; A reach trials, BL<auto, $p=1.16e-10$, auto>var, $p=3.19e-11$, var<rel, $p=4.00e-4$; B reach trials, auto>var, $p=2.28e-5$, var<rel, $p=5.22e-7$). However, the modulation of M1 and DLS peaks themselves, from trial to trial, still occurred in a coordinated manner during the variable session (**Figure 2.12D,E**; A reach trials, $R^2 = 0.301$, $p=2.40e-14$; B reach trials, $R^2 = 0.529$, $p=1.83e-38$).

M1-DLS reach-related cross-area dynamics drop during learning of switched contingency

We then explored how coordinated communication of M1 and DLS evolved during the breakdown of neural patterning seen with transfer learning using canonical correlation analysis (CCA). As opposed to directly comparing correlations of single-trial M1 and DLS population spiking activity (**Figure 2.12**), CCA finds maximally correlated combinations of simultaneous M1 and DLS activity, measuring cross-area dynamics¹⁰⁻¹². Through this method, axes of maximal correlation are identified for M1 and DLS (**Figure 2.13A**), with subsequent projection of high-dimensional neural activity on these axes to examine shared signals in a lower dimension (**Figure 2.13B**). To establish that CCA models of M1-DLS cross-area activity were behaviorally significant, we compared the R^2 of CCA models fit on actual data to the R^2 of CCA models fit on trial-shuffled data. From this generated distribution of shuffled R^2 values, a CV from the true

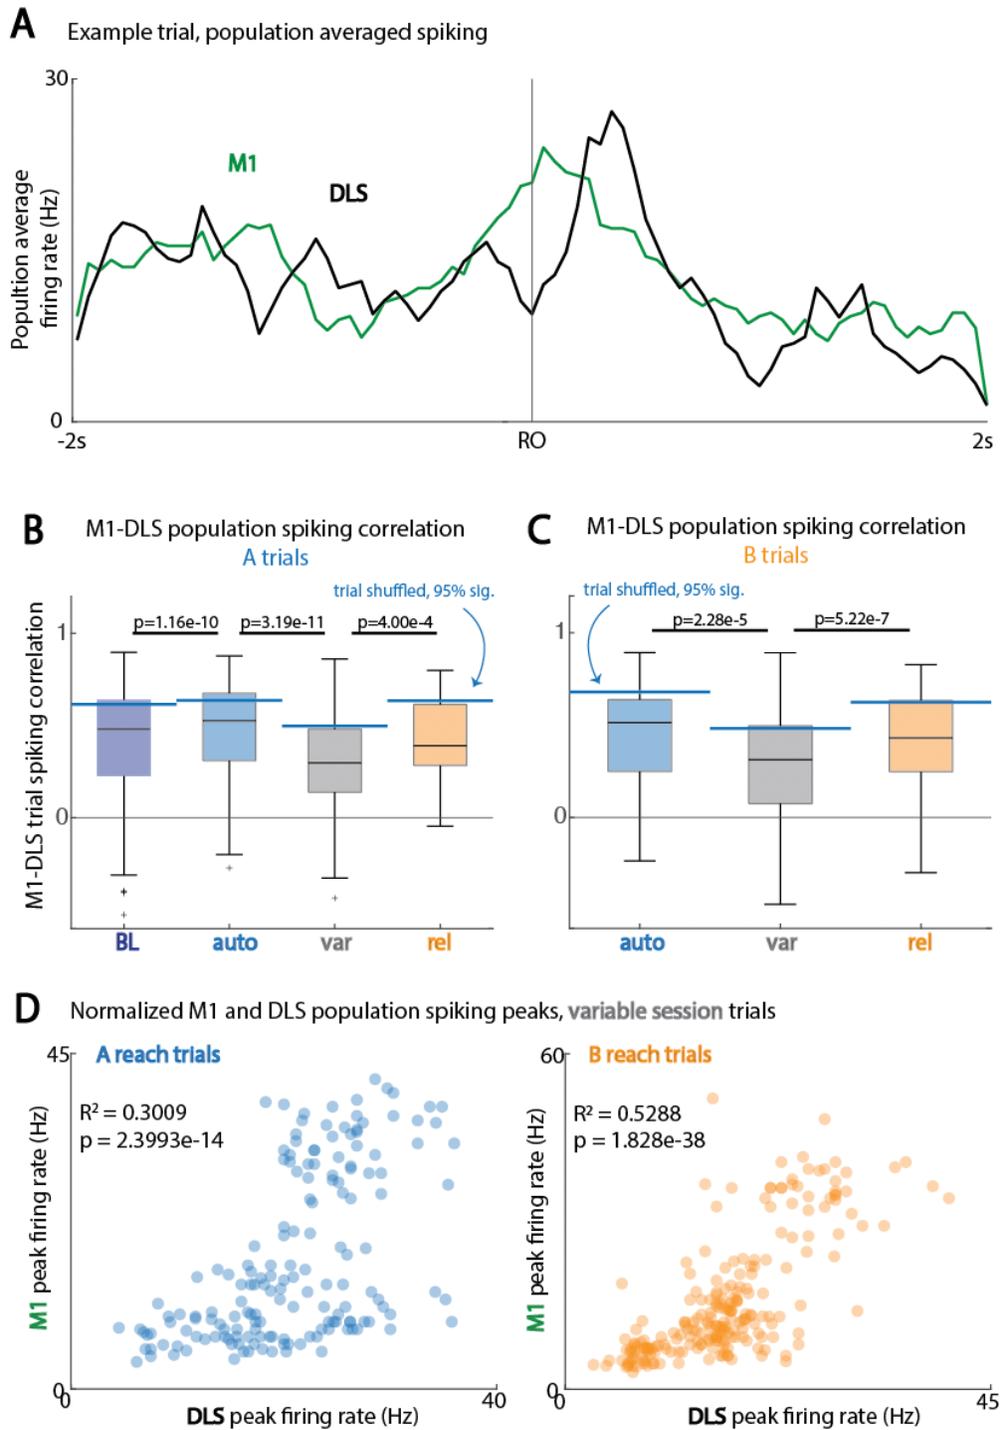


Figure 2.12. M1-DLS task population spiking drops correlation with preservation of modulation during transfer learning. (A) Average population spiking for M1 units (green) and DLS units (black) over a single trial, relative to first reach onset. (B) Correlation of M1 and DLS average population firing traces across session types for trials with first reach to A. (C) Same as B, for trials with first reach to B. (D) Variable session, single-trial M1 and DLS peak firing rate within -2s to 2s around first reach onset window, for trials with first reach to A (left, blue) and trials with first reach to B (right, orange).

dataset was considered significant if its R^2 value exceeded the 95th percentile of the reference distribution (**Figure 2.13C**). Most datasets had 1-3 significant CVs, demonstrating that CCA could identify shared low-dimensional activity across M1 and DLS (**Figure 2.13D**); datasets that had no significant CVs were excluded from analyses.

To examine whether M1-DLS shared communication is task-relevant or reach-relevant, we compared M1 cross-area activity vs. DLS cross-area activity prior to reach onset and during reach for the different behavioral states (**Figure 2.13E**). We then calculated the relative modulation index (RMI) of CCA activation per trial, to compare how reach-related cross-area activity was modulated over the process of unlearning an automatic action (**Figure 2.13F**). For trials with first reach to A, M1 cross-area activity was higher at the automatic state, as compared to the baseline and variable states (**Figure 2.13G**, reach modulation index, linear mixed effects model p-values, BL vs auto: $2.65e-12$; auto vs. var: $5.19e-10$). Cross-area activity in DLS exhibited a similar pattern, with reach-related cross-area activity higher at the automatic state, as compared to the variable state (**Figure 2.13G**, reach modulation index, linear mixed effects model p-value, auto vs. var: $6.01e-4$).

Even more importantly, this drop of reach-related cross-area activity on the variable sessions was pronounced for the new reach being learned to B, with moderate re-emergence during relearning. M1 cross-area activity was higher at the automatic state as compared to the variable state, with increase in cross-area activity from the variable to relearned state (**Figure 2.13H**, reach modulation index, linear mixed effects model p-values, auto vs. var: $1.34e-31$; var vs. rel: $1.31e-9$). Cross-area activity in DLS followed the same pattern, with decrease in activity from

automatic to variable state and increase in activity from variable to relearned state (**Figure 2.13H**, reach modulation index, linear mixed effects model p-values, auto vs. var: $7.75e-16$; var vs. rel: $7.44e-8$). Thus, with increased sequence variability in M1 and DLS (**Figure 2.8**, **Figure 2.9**), there was a corresponding decrease in M1-DLS shared cross-area activity during reach (**Figure 2.13G,H**), indicating that while M1-DLS spiking activity was coordinated enough to identify a shared subspace, in an exploratory state there is minimal modulation of cross-area communication during the reach-relevant period of the task on the variable sessions.

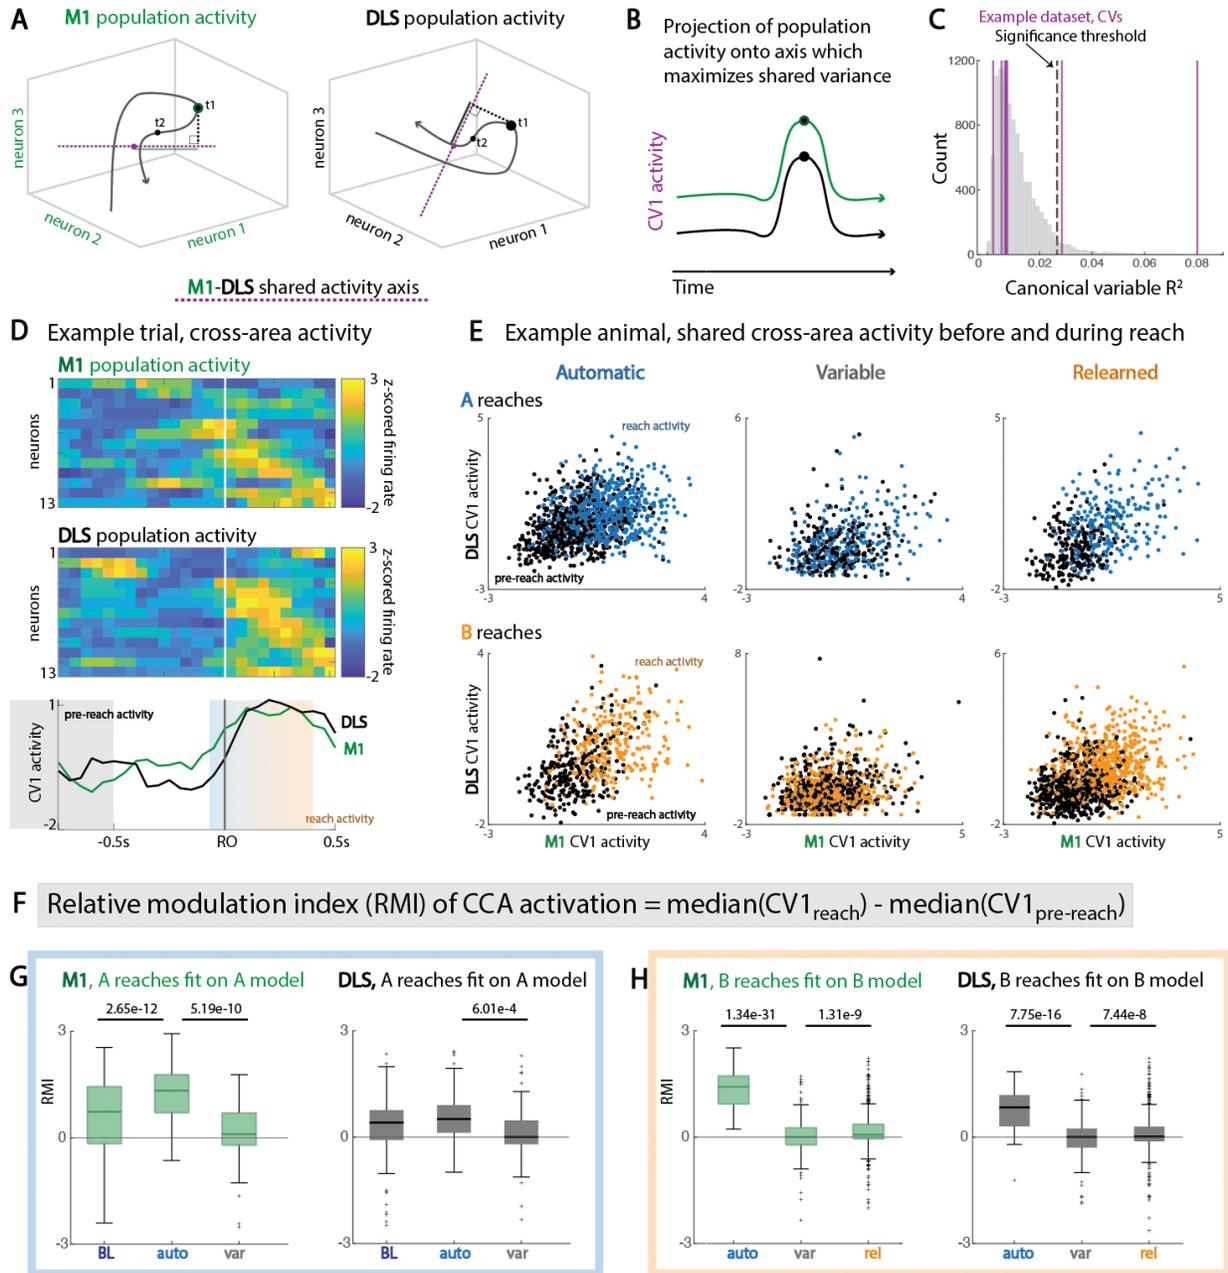


Figure 2.13. M1-DLS reach-related cross-area dynamics drop during transfer learning. (A) Canonical correlation analysis (CCA): example identification of axis with maximal shared variance between M1 and DLS activity (purple dotted line), with example value of projected population activity at time t_1 (solid line on shared axis) for M1 population activity (left, green) and DLS population activity (right, black). (B) Projection of population activity onto axis which maximizes shared variance, over time. (C) Example bootstrap shuffled distribution of R² values for identification of significant canonical variable (CV) axes. (D) Top: single-trial M1 spiking activity. Middle: single-trial DLS spiking activity. Bottom: corresponding single-trial M1 (green) and DLS (black) activations along CV1. Pre-reach period (grey) is -1 to -0.5s before first reach onset. Reach period (blue-orange) is -0.1s to 0.4s around first reach onset. (E) M1 and DLS CV1 activity for pre-reach period (black dots) vs. CV1 activity for reach period for A reaches (blue, top) and B reaches (orange, bottom), across session types. (F) Equation for relative modulation index (RMI) of CCA activation, comparing reach and pre-reach periods. (G) RMI for trials with first reach to A, for M1 activations (left, green) and DLS activations (right, green), across session types. (H) Same as G, for trials with first reach to B, across session types.

Behavior and spiking activity on the variable session have features of both early and late learning states

How might one explain how robust behavior is observed during the variable sessions without consistent M1 and DLS neural patterns? First, with re-categorization of sessions neurally, we observe an increase in normalized reaction time from the automatic to variable session when combining both first reach to A and first reach to B trials (**Figure 2.14A**, linear mixed-effects model, $p=5.31e-15$). Surprisingly, normalized reach-to-grasp time decreases with transfer learning (**Figure 2.14B**, BL vs. var: $7.42e-11$; BL vs. rel: $7.42e-13$; auto vs. var: $2.73e-11$; var vs. rel: $1.69e-6$), indicating that there may be greater cognitive exploration during this session accompanied by execution of the previously learned reach-to-grasp skill.

Earlier analyses demonstrated that the timing of M1 and DLS population spiking peaks on the variable session were distributed away from reach onset (**Figure 2.10C-F**), and recent work has demonstrated that single-trial neural dynamics can be due to variable movements independent of expert movement execution¹³. What behavior, if any, is occurring during these events of peak firing within the trial? We found that the majority of behavior during M1 peak activity involved movements, such as preparatory paw activity, movement onset before pausing, reaching, and grasping (**Figure 2.14C**). Spiking activity for each neuron during the reach-related periods of the variable session still remained higher than non-reach periods of spiking, in both M1 and DLS (**Figure 2.15A**, pairwise t-test p-values, M1: $1.17e-4$; DLS: $8.97e-6$), indicating de-coupling of population spiking bursts and underlying reach-supporting neural activity.

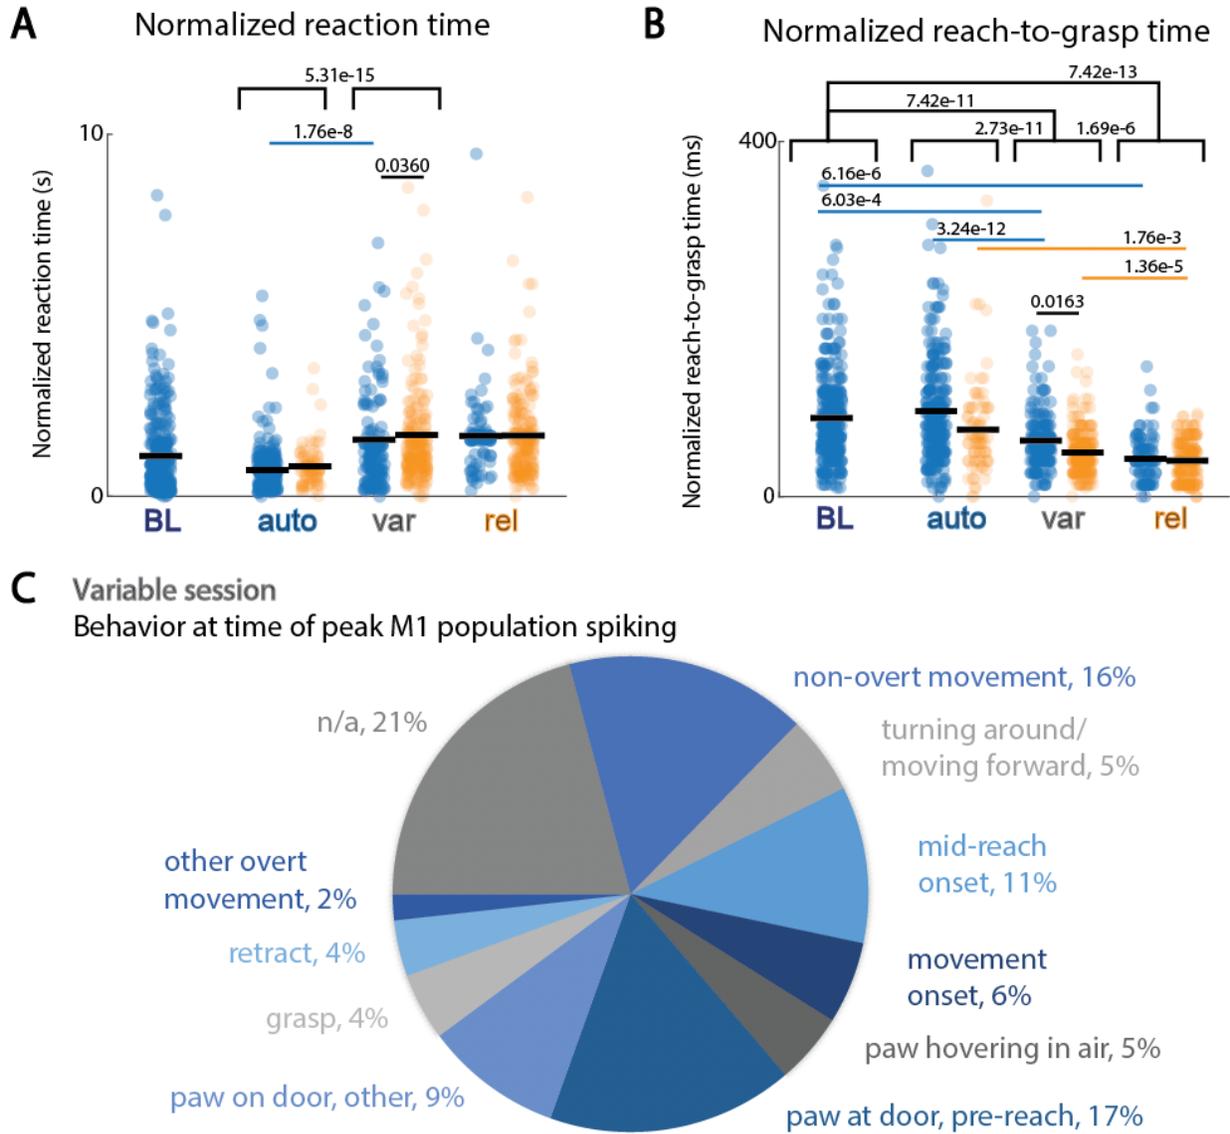


Figure 2.14. Motor behavior during the variable state has features of early and late learning. (A) Normalized reaction time, defined as time from trial start to first reach adjusted by minimum reaction time per animal, across re-categorized session types. (B) Normalized reach-to-grasp time, defined as time from first reach to first grasp adjusted by minimum reach-to-grasp time per animal, across re-categorized session types. (C) Behavior at peak of M1 population spiking in the -2s to 2s around first reach onset window.

If peak timing varies, is there still a consistent neural sequence of preparatory movement activity that is not locked to reach onset that co-exists with an overlying population activity burst?

SeqNMF is a newly introduced algorithm that utilizes convolutional non-negative matrix factorization to identify patterns of neural activity in an unsupervised fashion (**Figure 2.15B**)¹⁴. In turn, with peaks of an optimal sequence identified, one can measure the robustness of this neural activity pattern using trial-to-template correlation as completed previously (**Figure 2.8**, **Figure 2.9**), with alignment instead to the seqNMF-identified peak of onset (**Figure 2.15C**). In M1, there was a decrease in seqNMF peak-aligned trial-to-template correlation from baseline to variable state (**Figure 2.15D**, top, linear mixed effects model, $p=9.03e-5$), and trend towards increase from variable to relearned states (**Figure 2.15D**, top, linear mixed effects model, $p=0.0501$). In DLS, there was a drop in seqNMF peak-aligned trial-to-template correlation from the baseline to automatic state (**Figure 2.15D**, bottom, linear mixed effects model, $p=6.01e-3$) but not to the variable state with Bonferroni correction ($\alpha = 8.33e-3$), hinting towards preservation of an underlying sequence of activity in DLS with some neural pattern exploration in M1 during the variable session. Additionally, the overall linear mixed-effects model was nonsignificant in both M1 and DLS for differences in trial-to-template correlation, indicating that there may be preservation of weak M1 and DLS spiking patterns for the reach-to-grasp skill across session types that cannot be found when aligning neural activity to observed behavior markers, thus allowing for downstream generation of robust movements.

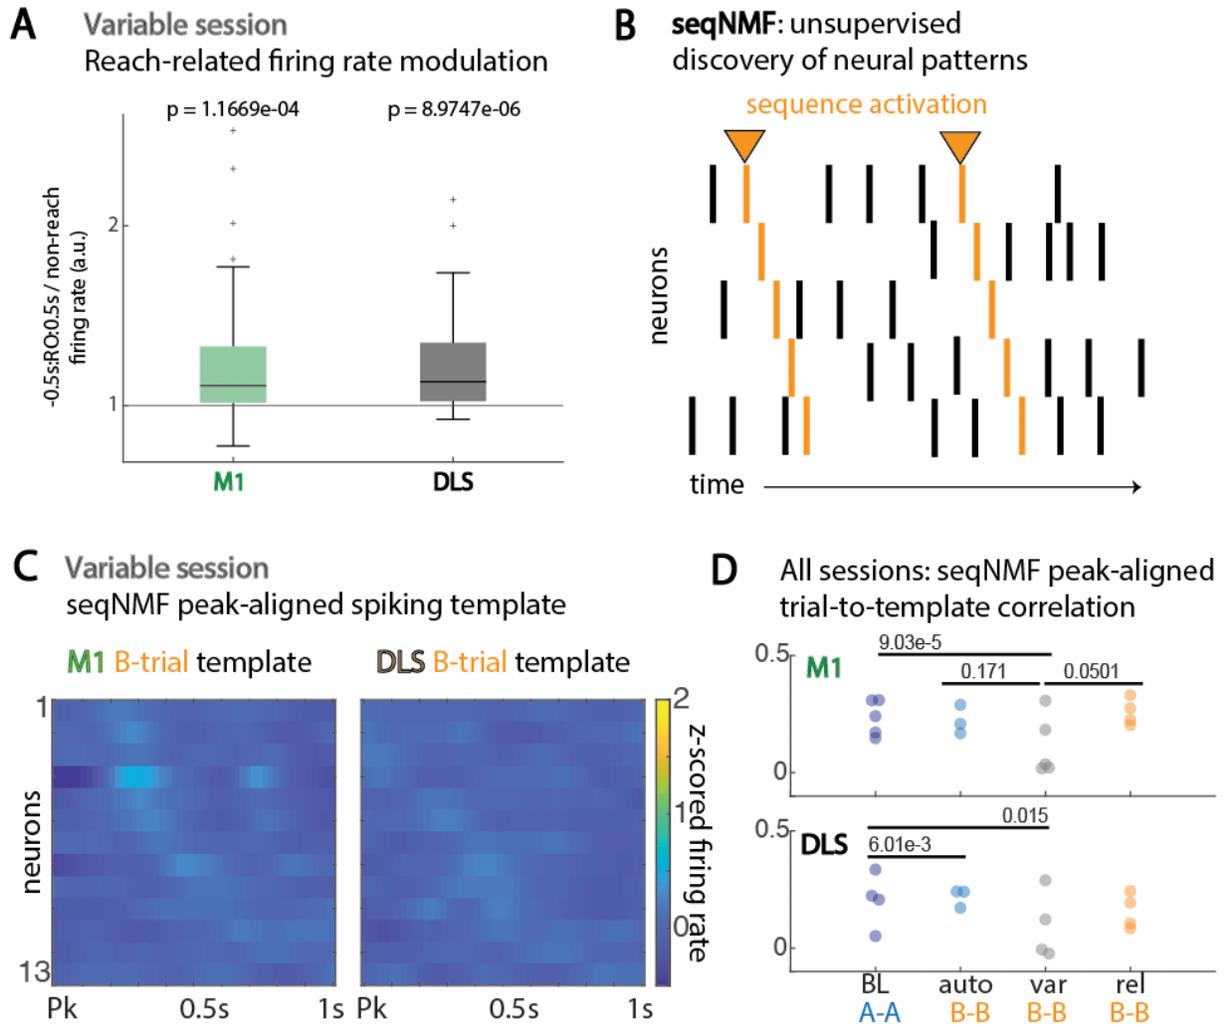


Figure 2.15. Unsupervised alignment of spiking activity suggests underlying trial structure in DLS. (A) Variable session, modulation of M1 (green) and DLS (black) unit firing rate during reach period, -0.5s to 0.5s around reach onset, relative to non-reach periods. (B) Unsupervised discovery of neural patterns with sequential non-negative matrix factorization (seqNMF). (C) M1 (left) and DLS (right) B trial-average spiking templates on the variable session for an example animal, aligned to peak of seqNMF activation in window of -2s to 2s around first reach onset. (D) Trial-to-template correlations with seqNMF peak activation alignment for M1 activity (top) and DLS activity (bottom), across session types.

Discussion

In this study, we find that automaticity in the reach-to-grasp task entails establishment of a rigid, inflexible motor skill. Learning to slightly modify behavior to adjusted task parameters from this rigid state is a multi-day process. In both M1 and DLS, neural pattern breakdown underlies this shift from rigid to flexible behavior. At a single neuron level, the reach-specific modulation across trials decreased during this return to behavioral exploration. Similarly, at the population level, the variable and exploratory behavioral state was characterized by equivalent modulation of activity from trial to trial, with increased temporal spread of when this population spike bursting occurs. After this variable state, there was subsequent establishment of consistent spiking patterning in M1 and DLS for the newly learned movement. Finally, we examined the communication subspace between M1 and DLS spiking. In the well-learned and automatic states, M1-DLS communication during reach was high, as compared to pre-reach. However during the exploratory state, there was little reach-modulated communication, consistent with the idea that M1-DLS patterns of activity are more consolidated when a motor skill is well-learned. In turn, we demonstrate that traversing from an automatic, rigid state to a more behaviorally flexible state involves a breakdown of M1 and DLS spiking patterns, with subsequent re-establishment of the cross-area manifold with re-learning.

A major question underlying this work is whether this rigid behavioral state is general to automatic behaviors, or whether this could be a sign of habit formation. Classically, establishment of habit has been promoted via random-interval delivery of reward, in contrast with random-ratio delivery of reward for goal-directed movement^{15,16}. This is difficult to decouple in the reach-to-grasp task, due to the reward being part of the grasping component of the

skill. In turn, more nuanced versions of this task, such as a reach-grasp-pull of a joystick for example, would enable more specific determination of the neural basis of habit formation and breakdown in a complex skilled forelimb task.

The prolonged nature of skill re-learning seen in this task paradigm is perhaps more akin to reinforcement learning than within-session adaptation¹⁷⁻¹⁹. Reinforcement learning has classically been studied with regard to reward-informing dopaminergic neurons in the striatum determining optimal patterning for synaptic strengthening²⁰⁻²³. In this study, we are unable to classify subtypes of neurons being recorded, or verify that recorded neurons are the same day-to-day. With emerging technologies, such as the combination of optogenetic tagging with calcium imaging^{24,25}, more detailed questions, such as whether dopaminergic neuron activity specifically increases during the variable session, can be addressed. Additionally, we are unable to verify that the same neurons are being recorded over time, as a majority of this data is a combination of single unit and multi-unit activity. Thus, we cannot determine whether the corticostriatal manifold is the same from a baseline to automatic state, and whether the automatic A reach manifold is the same as that when accidental reaches are made to A following relearning.

While this study examines M1 and DLS activity across a re-learning paradigm, it is unknown whether M1 neural variability is driving DLS neural variability during the exploratory state. While M1 and DLS are monosynaptically connected structures^{4,7}, and well-learned motor actions may follow a feedforward neural pattern from M1², recent studies have introduced some exceptions. In a forelimb joystick task, it has been shown that M1 exhibits variability in patterning over longer-term motor learning²⁶. Additionally, lesion of M1 after well-learned

simple motor skill has been demonstrated to have no effect on action execution^{27,28}, hinting that downstream structures such as DLS may be driving gross forelimb movement, such as reaching, rather than M1. Thus, a future direction would be to inactivate M1 prior to the introduction of the second reach location, to study whether DLS exhibits neural variability on its own, let alone being able to support reaching to a new location itself. Alternatively, a third structure, such as the secondary motor cortex¹⁰ or prefrontal cortex²⁹, both of which have projections to M1 and DLS, may be coordinating this simultaneous establishment of neural variability. Thus, simultaneous recording in these four regions during early motor skill learning, automatic performance, and introduction of a second motor skill, would more definitively address directionality of M1-DLS coordinated activity during this task.

This study could have clinical implications for the disruption of extreme patterns of motor behavior, such as those in obsessive compulsive disorder³⁰⁻³². One could posit that the mechanism for breaking down automatic behavior not only requires consistent behavioral monitoring to cognitively decide to change the behavior, but also an internal breakdown of established patterns of corticostriatal dynamics. Some forms of electrical stimulation (intra-cortical microstimulation³³, transcranial magnetic stimulation^{34,35}) have been demonstrated to disrupt manifold activity in a single region. There is an open question as to how selective disruption of one manifold can influence downstream dynamics. By extension, however, it is still unclear how to manipulate one set of multi-area dynamics for a pathological behavior, while still preserving multi-area dynamics for the broad distribution of functionally beneficial motor behaviors. This study demonstrates that multi-area manifold activity dynamically re-organizes with new learning in the context of a previously learned task. Future work elucidating the mechanism of this

process will inform more nuanced, specific neuromodulation therapies for establishing behavioral flexibility.

Methods

Animal Care

All procedures were in accordance with protocols approved by the Institutional Animal Care and Use Committee at the San Francisco Veterans Affairs Medical Center. Adult male Long-Evans rats between 3 and 6 months old (n=6, 300-500g; Charles River Laboratories) were housed in a controlled temperature room with a 12-h light/12-h dark cycle. All experiments were conducted during the light cycle.

Surgery

All surgeries were performed using sterile surgical technique under 2% isoflurane (5% at induction). Surgery involved exposure and cleaning of the skull, preparation of the skull surface (using cyanoacrylate) and implantation of skull screws on the perimeter for headstage stability. Reference screws were implanted posterior to lambda, contralateral to neural recordings. Ground screws were implanted posterior to lambda, ipsilateral to neural recordings. Craniotomy and durotomy were performed, followed by implantation of neural probes. Neural probes (32- or 64-channel 33um polyimide-coated tungsten microwire electrode arrays (Tucker-Davis Technologies)) were implanted in the forelimb area of M1 (centered at 3.5mm lateral and 0.5mm anterior to bregma; layer V at a depth of 1.5mm) and DLS (centered at 4mm lateral and 0.5mm anterior to bregma; at a depth of 4.5mm). The final location of the electrodes was confirmed by electrolytic lesion. Post-operative recovery regimen included the administration of 0.02 mg/kg buprenorphine for 2 days, and 0.2 mg/kg meloxicam, 0.5 mg/kg dexamethasone and 15 mg/kg trimethoprim-sulfadiazine for 5 days. All animals were allowed to recover for 1 week prior to further behavioral training.

Histology

Rats were anesthetized with isoflurane and transcardially perfused with 0.9% sodium chloride, followed by 4% paraformaldehyde. The harvested brains were then postfixed in 4% paraformaldehyde for 24 hours and immersed in 20% sucrose for 2 days prior to drop freezing. Coronal cryostat sections (40um thickness) were mounted with permount solution (Fisher Scientific) on super-frosted coated slides (Fisher Scientific). Microscope images of the whole section were taken by a Zeiss microscope.

Behavioral training

Rats naive to motor training were first assessed for forelimb preference: approximately 10-20 pellets were placed in front of the animal, with preference defined by the limb which reached to the pellets the most. 4 out of 6 animals then underwent surgery for electrode implantation followed by a recovery period. Rats were then trained within an automated behavior box to perform dexterous reach-to-grasp movements³⁶ to a single location A for at least 1000 trials. 2 out of 6 animals were trained similarly, with electrode implantation one week after training the reach-to-grasp motor skill to A. Following surgical recovery, these 2 animals were re-trained to baseline level of reach to A. Overall, this initial behavioral training required minimal user intervention, as the automated reach-box was controlled by custom MATLAB scripts and an Arduino microcontroller. Each trial consisted of a pellet dispensed on the end of an arm with pellet holder groove, movement of the pellet arm to the pre-programmed location, beep alerting the animal the trial was beginning, and then door opening. Animals needed to then reach through a slit, grasp, and retrieve the pellet within 15 seconds. An IR sensor centered over the pellet was used to detect when there was no longer a pellet in the groove, indicating the trial was over; the

door then closed. Each animal was trained to plateau performance of reaching to A (~100-150 trials per day), prior to pellet location being switched to a second location, B, that was equidistant from the center of wall slit as pellet location A and still reachable with a rat's preferred forelimb. Animals were then trained on the second pellet location B for 4-6 days (~100 trials per day).

Behavioral analysis

For 2/6 animals, rat behavior was video recorded using a side-view camera. For the remaining 4/6 animals, rat behavior was recorded using both top-down and side-view cameras. Three types of cameras were used: Microsoft LifeCams, which captured videos at 30Hz; Basler cameras, which captured videos at 75Hz; and Point Grey/FLIR cameras, which captured videos at 75-100Hz. Reach trajectories were captured from video using DeepLabCut³⁷ to track the center of the rat's paw. Reach trajectories consisted of paw trajectory from each reach onset to subsequent grasping motion that occurred beyond the slit. Reach videos and trajectories were viewed and scored to obtain trial success, reach type (low amplitude, endpoint at old location A, endpoint at new location B), and timepoints for reach onset, pellet touch, grasp onset, and retract onset.

To characterize motor performance, we quantified reaction time, reach duration, pellet retrieval success for each trial, and location of reach endpoints both within and across trials. Reaction time was defined as the time taken from when the door opened for the start of the trial to when the rat began to reach, combining both attentional and cue-related motion behaviors (Figure 2.2G, Figure 2.14A). Reach duration was defined as the time from first reach onset to first grasp onset (Figure 2.2F, Figure 2.14B). Percent reach success was defined as the percent of trials on

which a pellet was successfully retrieved out of total trials with a full amplitude reach within a session (Figure 2.2C). Low amplitude reaches were those in which the center of the paw reached past the slit but digits did not reach the vertical plane where the pellet was located. Reaches to the new location B were defined as those where at least half of the paw covered the pellet on grasp. All other full amplitude reaches were classified as reaches to old location A.

For denoting behavior at time of maximal task M1 population firing (Figure 2.14C), 25 frames from peak of firing onwards were viewed, and description of rodent behavior at first frame was qualitatively described. Movements were then grouped by similarity of behavior relative to the trial (i.e. reach, grasp, turn).

Electrophysiology data collection

We recorded extracellular neural activity, including units and local field potentials (LFP), using an RZ2 system (Tucker-Davis Technologies). Spike data was sampled at 24,414Hz; LFP data was sampled at either 1,017Hz (n=2) or 24,414Hz (n=4). Snippets of data that crossed a high signal-to-noise threshold (at least 4 standard deviations away from the mean) were deemed spiking events; time stamps and peak-aligned waveforms were stored for any event that crossed the threshold. Spike sorting was then performed using Offline Sorter v.4.3.0 (Plexon) with a k-means-based clustering method followed by manual inspection. Spikes were sorted separately for each session. We accepted units based on waveform shape, clear cluster bounding in principal component space, and 99.0% of detected events with an ISI > 2ms. Clusters interpreted to be single units or multi-units were kept for analysis; those determined to be noise were discarded.

Neural data analysis

Analyses were conducted using a combination of custom-written scripts and function in MATLAB R2018B (MathWorks), along with functions from the EEGLAB (<http://sccn.ucsd.edu/eeglab/>) and the Chronux (<http://chronux.org/>) toolboxes.

LFP analysis

For the four animals with LFP recorded at 24,414Hz, raw LFP signals were decimated channel by channel with an 8th order Chebyshev Type I low pass filter to a tenth of the original signal (2,414Hz). LFP for all animals was then pre-processed with the following steps: z-scoring the entire recording session, channel by channel; artifact rejection (manually removing noisy/broken channels, identifying trials with motion artifact on the majority of channels); common-mode referencing using the median signal (the median signal across all non-noisy channels in a region was calculated at every time point and subtracted from each channel to decrease common noise and minimize volume conduction). Common-mode referencing was performed independently for channels in each region, M1 and DLS.

We filtered the LFP signals to isolate and display the low-frequency (3-6Hz) component of the signal. Filtering was performed using the EEGLAB function 'eegfilt.' To examine power across multiple frequency bands, we calculated movement-related LFP spectrograms and power spectra within each region using wavelets with the EEGLab function 'newtimef.' This function was also used to calculate the inter-trial coherence (ITC) of LFP signals.

Spiking analysis

Unit modulation

All spiking analyses were aligned for first reach onset (RO). To determine unit modulation (Figure 2.5A,B; Figure 2.6A,B; Figure 2.7A-D), peri-event time histograms (PETHs) were generated by averaging spiking activity for each neuron across trials in a session, locked to first reach onset, and binned at 50ms. PETHs were then fit with a smoothing spline using a custom MATLAB function. To determine task-related unit modulation, we z-scored each unit's average firing rate for the session, found all peak prominences and times of activity from 2s before reach onset to 2s after reach onset with the MATLAB function 'findpeaks', and identified the maximum peak of z-scored activity from -0.5s to 0.5s around first reach onset (Figure 2.6A,B). This process was conducted separately for trials with first reach to pellet location A and trials with first reach to pellet location B across sessions.

Determination of session type

Sessions were categorized into the following based on task parameters, behavior curves, and Fano factors of unit spiking activity: 1) Baseline (BL), pellet at location A, reaching to A, consistent reach-modulated spiking activity; 2) Automatic (auto), pellet at location B, reaching mostly to A, consistent reach-modulated spiking activity; 3) Variable (var), pellet at location B, reaching to A and B, local maximum of lowest 5 Fano factors per unit, aggregated, with significant deviation from baseline minimum Fano factors; 4) Relearned (rel), pellet at location B, reaching mostly to B, consistent reach-modulated spiking activity (Figure 2.5D,E).

Fano factors describe how variable spiking activity is at a given time point across trials; with increased spike count consistency within a given time bin, there is a decrease in the Fano factor value for that time bin. To calculate Fano factors for each unit the following process was followed (Figure 2.5C). First, for each animal, the minimum trial number across sessions was identified for sub-sampling. For each unit and each trial in a session, spike counts from -2s to 2s around RO were binned at 50ms, z-scored, and smoothed with a 5-bin moving average. Trials in a session were then sampled to the minimum trial number, and the Fano factor of each time bin was calculated, where $\text{Fano_timebin} = \frac{\text{standard deviation of spike counts in that time bin across trials}^2}{\text{mean spike count across trials}}$. For each unit in a given session, this process was repeated with 1000 total trial sub-samples, and the final Fano factor per time bin was the median value of the 1000 sub-sampled Fano factors per bin, with each unit having 80 total Fano factors spanning -2s:RO:2s per session.

To examine consistency of unit spiking activity, we identified the 5 Fano factors per unit per session with lowest value, thus allowing for timing-agnostic sampling of minimum variation in unit firing (Figure 2.5D). To determine which sessions had a significantly different distribution of Fano factors from baseline, we used the MATLAB function 'kstest2'; if multiple sessions were significantly different from baseline, the session with the greatest Fano factor median was determined to be the 'variable' session. Subsequently, the 'automatic' session was the session recorded closest in time before the variable session, and the 'relearned' session the closest in time after for each animal. One animal did not have a variable or relearned session, one animal did not have a relearned session, and two animals did not have automatic sessions (Figure 2.5E).

Average task firing rate

To calculate average task firing rate, the minimum number of neurons in each region for an animal across sessions was identified. For A or B trials in a session, average firing rate for -2s to 2s around first reaches in the session across the sub-sampled neurons was calculated 1000 total times, with re-sampling of neurons each time to account for unit-to-unit firing rate variability. The median value of average firing rate was taken for each session per rat.

Reach-related firing rate unit modulation

To calculate reach-related spiking activity versus that at non-reach periods for the variable session, average trial firing rate from -0.5s to +0.5s around first reach onset was compared to the average session firing rate for that unit outside of the reaching time periods (Figure 2.7E, Figure 2.15A).

Template Matching

To assess how temporally consistent single trial spiking activity was across session types, for each animal on each session we separated out trials with first reach to A and first reach to B. If there were at least 5 trials of one type for a given rat and session, each trial's spiking activity for a region (e.g. M1 units) was compared to the average template spiking from the remaining trials of that type. Specifically, regional spiking activity from 500ms before first reach RO to 500ms after first RO was binned at 20ms, smoothed with a 60ms Gaussian kernel standard deviation, and concatenated across units for a given trial. Given the variable number of units in a session, the minimum number of units an animal had in a region, across sessions, was determined to be the number of units to subsample for this analysis; a minimum of 5 units was necessary for a

session to be included. Each trial spiking activity, with sub-sampled units, was correlated to the average spiking activity from trials with the same sub-sampled units (Figure 2.8A, Figure 2.9A). This process was repeated 1000 times per set of trials, keeping the mean correlation across trials for each iteration; the median of these means is reported as the session trial-template correlation (Figure 2.8C, Figure 2.9B).

Population modulation

To characterize population spiking activity modulation (Figure 2.11A,B), z-scored unit activity from 2s before first RO and 2s after first RO for each trial was smoothed using a 5-point moving average, summed and then divided by the number of neurons for normalization. Trial activity for each session was then grouped into trials with first reach to A or trials with first reach to B. Normalized population single-trial spiking activity was then averaged across trials to A or B within a session for both M1 and DLS units separately. All peak prominences and times of activity from 2s before first RO to 2s after first RO were detected with the MATLAB function 'findpeaks' with no minimum peak prominence. Maximum modulation prominence out of all local maxima in period from -0.5s to 0.5s around first reach onset was identified as peak of population reach-modulation (Figure 2.11C,D). This same data was used to identify population spiking peak firing rate in the larger trial window (Figure 2.11E,F) and times of peak firing rate relative to reach onset (Figure 2.10C,D). Correlation of population peak activity (Figure 2.12D) was computed by comparing the normalized M1 and DLS peak prominences from each trial of the variable session, across animals.

Cross-area spiking correlation

Correlation of M1-DLS population spiking activity (Figure 2.12A-C) was computed by binning spikes at 1ms, from -2s to +2s around first reach onset, averaging population activity across neurons, and applying the MATLAB function 'xcorr.' M1 activity was trial-shuffled relative to DLS activity 1000 times across animals per session to compute 95% significance level,

Cross-area neural subspace

Shared cross-area subspaces between M1 and DLS were identified using canonical correlation analysis (CCA), which defines axes that maximally correlates activity between the two areas¹⁰⁻¹². Neural data in M1 and DLS were binned at 50ms, and data from -2s to +0.5s surrounding first reach onset was concatenated across trials with first reach to A and trials with first reach to B separately; mean activity in each group was subtracted. CCA models were then fit using the MATLAB function 'canoncorr'. The number of canonical variables (CVs) output by CCA is the minimum number of neurons in either M1 or DLS for that session. The R² value for each CV was computed using tenfold cross-validation, with 95% significance determined by comparison to a bootstrap distribution of top CV R² created from trial-shuffled data (10⁴ shuffles), as described previously¹⁰ (Figure 2.13A-C). Only sessions with minimum 5 units in both areas were included for analysis; sessions with no significant CVs were subsequently removed from analyses. For evaluating cross-area signals (Figure 2.13D-H), only the top CV was used for consistency across datasets.

Cross-area task representation

To calculate the difference in cross-area activity before first reach versus during first reach, we defined a pre-reach period as -1 to -0.5s before first reach onset, and reach period as -0.1 to +0.4s around first reach onset (Figure 2.13D,E). Cross-area median activity within a trial was calculated for each time period, and compared across session types (Figure 2.13F-H).

Unsupervised sequence detection

To detect neural sequences independent of imposed behavioral marking alignment, we applied a convolutional non-negative matrix factorization method, seqNMF¹⁴. First, spike times from -2s to +0.5s around first reach onset were binned at 20ms and concatenated across trials per unit. The factorization parameter L (length in time bins of each factor) was set to 50, indicating search for a 1s maximum length sequence. The regularization parameter lambda was set to 0.001. Next, a 5-fold cross-validation with up to 5 possible sequences was run on each session, to identify a local minimum of identifiable sequences. Across all sessions, the local minimum was one sequence. Subsequently, we used the function ‘seqNMF’ in MATLAB with K=1 (number of sequences to identify), L=50 (maximum 1s sequence), lambda = 0.001 (regularization parameter) to identify the top sequence and strength of activation over the concatenated time series. This process was repeated ten times and the average signal was used for peak detection.

To detect time of peak sequence strength in a trial, maximum activation during each trial period (-2s to +0.5s around first reach onset) was found using the MATLAB function ‘findpeaks.’ This marked the start of the 1s sequence (Figure 2.15B). To measure consistency of the sequence detected across trials, trial-to-template correlation was performed as detailed in the ‘Template

Matching' section, with spiking times now from 0 to +1s relative to maximum sequence peak detection time (Figure 2.15C,D).

Statistics

Linear mixed-effects models were used to test the significance of differences across both behavioral and neural measures when comparing differences in group means. These models account for units or trials coming from the same animal, which are more correlated than those from different animals, thus providing a stricter computation of statistical significance. For comparison of distribution broadness (Figure 2.10C,D), the Bartlett's test was used to determine whether samples came from populations with equal variances.

References

1. Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical control of arm movements: A dynamical systems perspective. *Annual Review of Neuroscience*, *36*, 337–359. <http://doi.org/10.1146/annurev-neuro-062111-150509>
2. Peters, A. J., Chen, S. X., & Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature*.
<http://doi.org/10.1038/nature13235>
3. Ramanathan, D. S., Gulati, T., & Ganguly, K. (2015). Sleep-Dependent Reactivation of Ensembles in Motor Cortex Promotes Skill Consolidation. *PLoS Biology*, *13*(9), 1–25.
<http://doi.org/10.1371/journal.pbio.1002263>
4. Peters, A., Fabre, J., Steinmetz, N., Harris, K., & Carandini, M. (2021). Striatal activity topographically reflects cortical activity. *Nature*, (July 2019).
<http://doi.org/10.1038/s41586-020-03166-8>
5. Cunningham, J. P., & Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, *17*(11), 1500–1509. <http://doi.org/10.1038/nn.3776>
6. Runyan, C. A., Piasini, E., Panzeri, S., & Harvey, C. D. (2018). Distinct timescales of population coding across cortex. *Nature*, *548*(7665), 92–96.
<http://doi.org/10.1038/nature23020>.Distinct
7. Lemke, S. M., Ramanathan, D. S., Guo, L., Won, S. J., & Ganguly, K. (2019). Emergent modular neural control drives coordinated motor actions. *Nature Neuroscience*, *22*(7), 1122–1131. <http://doi.org/10.1038/s41593-019-0407-2>

8. Veuthey, T. L., Derosier, K., Kondapavulur, S., & Ganguly, K. (2020). Single-trial cross-area neural population dynamics during long-term skill learning. *Nature Communications*, (2020), 1–15. <http://doi.org/10.1038/s41467-020-17902-1>
9. Ashby, F. G., Turner, B. O., & Horvitz, J. C. (2010). Cortical and basal ganglia contributions to habit learning and automaticity. *Trends in Cognitive Sciences*, 14(5), 208–215. <http://doi.org/10.1016/j.tics.2010.02.001>
10. Veuthey, T. L., Derosier, K., Kondapavulur, S., & Ganguly, K. (n.d.). Single-trial cross-area neural population dynamics during long-term skill learning. *Nature Communications*, (2020), 1–15. <http://doi.org/10.1038/s41467-020-17902-1>
11. Morcos, A. S., Raghu, M., & Bengio, S. (2018). Insights on representational similarity in neural networks with canonical correlation. *arXiv preprint*. arXiv:1806.05759.
12. Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M., & Kohn, A. (2019). Cortical Areas Interact through a Communication Subspace. *Neuron*, 102(1), 249–259.e4. <http://doi.org/10.1016/j.neuron.2019.01.026>
13. 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. <http://doi.org/10.1038/s41593-019-0502-4>
14. Mackevicius, E. L., Bahle, A. H., Williams, A. H., Gu, S., Denisenko, N. I., Goldman, M. S., & Fee, M. S. (2019). Unsupervised discovery of temporal sequences in high-dimensional datasets , with applications to neuroscience, 1–42.
15. Gremel, C. M., & Costa, R. M. (2013). Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. *Nature Communications*, (May). <http://doi.org/10.1038/ncomms3264>

16. O'Hare, J. K., Ade, K. K., Sukharnikova, T., Van Hooser, S. D., Palmeri, M. L., Yin, H. H., & Calakos, N. (2016). Pathway-Specific Striatal Substrates for Habitual Behavior. *Neuron*, 89(3), 472–479. <http://doi.org/10.1016/j.neuron.2015.12.032>
17. Shadmehr, R., & Mussa-ivaldi, F. A. (1994). Adaptive Task of Dynamics during Learning of a Motor. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 14(5), 3208–3224.
18. Izawa, J., & Shadmehr, R. (2011). Learning from sensory and reward prediction errors during motor adaptation. *PLoS Computational Biology*, 7(3), 1–11. <http://doi.org/10.1371/journal.pcbi.1002012>
19. Dezfouli, A., & Balleine, B. W. (2012). Habits, action sequences, and reinforcement learning. *European Journal of Neuroscience*, 35(7), 1036–1051. <http://doi.org/10.1111/j.1460-9568.2012.08050.x.Habits>
20. Mink, J.W. (1996) Basal ganglia: focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, 50(4), 381-425. [https://doi.org/10.1016/S0301-0082\(96\)00042-1](https://doi.org/10.1016/S0301-0082(96)00042-1)
21. Lee, D., Seo, H., & Jung, M. W. (2012). Neural basis of reinforcement learning and decision making. *Annual Review of Neuroscience*, 35, 287–308. <http://doi.org/10.1146/annurev-neuro-062111-150512>
22. Rubin, J. E., Vich, C., Clapp, M., Noneman, K., & Verstynen, T. (2020). The credit assignment problem in cortico-basal ganglia-thalamic networks: A review, a problem and a possible solution. *European Journal of Neuroscience*, 1–20. <http://doi.org/10.1111/ejn.14745>

23. Gerfen, C.R., & Surmeier, D.J. (2012). Modulation of striatal projection systems by dopamine. *Annual Reviews Neuroscience*, 2(34), 441–466.
<http://doi.org/10.1146/annurev-neuro-061010-113641>.Modulation
24. Tolö, J., Gordus, A., Orger, M. B., Severi, K. E., Macklin, J. J., Patel, R., ... Looger, L. L. (2013). Genetically encoded calcium indicators for multi-color neural activity imaging and combination with optogenetics, 6(March), 1–29.
<http://doi.org/10.3389/fnmol.2013.00002>
25. Kim, C. K., Adhikari, A., & Deisseroth, K. (n.d.). Integration of optogenetics with complementary methodologies in systems neuroscience. *Nature Publishing Group*.
<http://doi.org/10.1038/nrn.2017.15>
26. 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), 1–13.
<http://doi.org/10.1126/sciadv.aay0001>
27. Kawai, R., Markman, T., Kawai, R., Markman, T., Poddar, R., Ko, R., ... Olveczky, B.P. (2015). Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron*, 86, 800–812. <http://doi.org/10.1016/j.neuron.2015.03.024>
28. Dhawale, A. K., Wolff, S. B. E., Ko, R., & Ölveczky, B. P. (2019). The basal ganglia can control learned motor sequences independently of motor cortex. *bioRxiv*.
<http://doi.org/10.1101/827261>
29. Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., ... Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6), 860–868. <http://doi.org/10.1038/s41593-018-0147-8>

30. Ahmari, S. E., Spellman, T., Douglass, N. L., Kheirbek, M. A., Simpson, H. B., Deisseroth, K., ... Hen, R. (2013). Repeated cortico-striatal stimulation generates persistent OCD-like behavior. *Science*, *340*(6137), 1234–1239.
<http://doi.org/10.1126/science.1234733>
31. Lipton, D. M., Gonzales, B. J., & Citri, A. (2019). Dorsal striatal circuits for habits, compulsions and addictions. *Frontiers in Systems Neuroscience*, *13*, 1–14.
<http://doi.org/10.3389/fnsys.2019.00028>
32. Vicente, A. M., Martins, G. J., & Costa, R. M. (2020). Cortico-basal ganglia circuits underlying dysfunctional control of motor behaviors in neuropsychiatric disorders. *Current Opinion in Genetics and Development*, *65*, 151–159.
<http://doi.org/10.1016/j.gde.2020.05.042>
33. Vyas, S., O’Shea, D.J., Ryu, S.I., & Shenoy K.V. (2020). Causal role of motor preparation during error-driven learning. *Neuron*, *106*(2), 329-339.
<https://doi.org/10.1016/j.neuron.2020.01.019>
34. Pasley, B. N., Allen, E. A., & Freeman, R. D. (2009). Article State-Dependent Variability of Neuronal Responses to Transcranial Magnetic Stimulation of the Visual Cortex. *Neuron*, *62*(2), 291–303. <http://doi.org/10.1016/j.neuron.2009.03.012>
35. Jenkinson, N., & Miall, R. C. (2010). Disruption of Saccadic Adaptation with Repetitive Transcranial Magnetic Stimulation of the Posterior Cerebellum in Humans. *Cerebellum*, *9*, 548–555. <http://doi.org/10.1007/s12311-010-0193-6>
36. Wong, C. C., Ramanathan, D. S., Gulati, T., Joon, S., & Ganguly, K. (2015). An automated behavioral box to assess forelimb function in rats. *Journal of Neuroscience Methods*, *246*, 30–37. <http://doi.org/10.1016/j.jneumeth.2015.03.008>

37. Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., & Bethge, M. (2018). user-defined body parts with deep learning. *Nature Neuroscience*, *21*(September). <http://doi.org/10.1038/s41593-018-0209-y>

Chapter 3: Two pieces of the puzzle – introducing a model for motor flexibility

As demonstrated in Chapter 2, transfer learning of a complex motor skill is a multi-day process, involving a breakdown of consistent spiking activity in both M1 and DLS, in contrast with the adaptation hypothesis that prior learning could generalize for within-session updating of spiking patterns (**Figure 3.1**). Most prior studies of the reach-to-grasp skill have focused on the neural evolution of M1 and connected motor structures during learning¹⁻⁴, or during execution of a well-learned single-location skill⁵⁻⁷. However, studies of motor learning conflate reach-to-grasp specific learning with learning of trial structure, cues, reward identification, and more. Similarly studies of motor execution are limited by reach-to-grasp skill performance to a single location, such that it is unknown whether the skill and corresponding neural state is rigid or generalizable.

Here, I demonstrate that overtraining the reach-to-grasp skill to a single location promotes a rigid state, and the behavioral manipulation of environmental modification, pellet location for reward, can probe motor learning specifically, independent of task structure learning. From the variable state, it is apparent that other regions in the larger brain network can support the learned fast reach-to-grasp movement without robust patterned input from M1 and DLS. The following section will delve into a possible model for how behavioral flexibility could be enabled in the context of skilled motor execution.

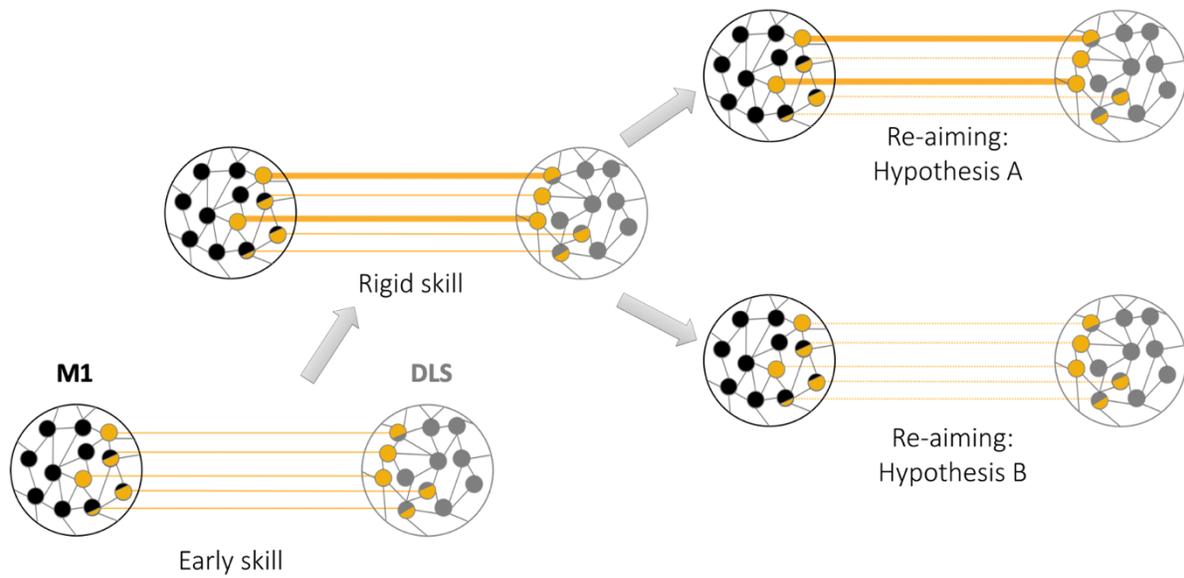


Figure 3.1. Hypotheses for transfer learning of reach-to-grasp skill. From early learning to establishment of a rigid skill (left), corticostriatal communication has been demonstrated to increase, particularly during motor execution. How then does corticostriatal patterning support re-aiming the initial reach for transfer learning? Hypothesis A (top, right) predicts an adaptive response, with communication maintained and minor adjustment of secondary patterns of communication, such as neural tuning. Hypothesis B (bottom, right) predicts some amount of preserved communication, with some features of early learning as the previously rewarded neural pattern breaks down to allow for exploration.

Motor cortex – a node in multiple coexisting pathways

Primary motor cortex has been recently described as a pattern generator for downstream structures, serving as the primary driver for a motor dynamical system⁷⁻¹¹. While this could be true for execution of well-learned movements, it is apparent from neural variability in early motor learning and now transfer learning that other nodes in the larger brain network can support movement execution outside of the well-learned optimized skill state. How do parallel processing loops (i.e. limbic, associative, sensorimotor) interact with each other? Knowlton and Yin, 2006¹² illustrated this beautifully in a simplified schematic, showcasing how cortical, striatal, and thalamic nodes can interact across both shorter timescales (**Figure 3.2a**) and proposing how parallel loops interact on longer timescales for increased automaticity (**Figure 3.2b**). However, what are specific regions other than M1 and DLS directly implicated in movement pattern generation, from cortical regions through peripheral neurons controlling muscle activation?

The prefrontal cortex (PFC) along with dorsomedial striatum (DMS) have been heavily implicated in associative learning across species¹³⁻¹⁷. Thus, with transfer learning there could be a shift in pattern generation from M1-DLS to PFC-DMS for modification of directionality in reaching. Of note, due to recurrent networks involving each of these node pairs, a parallel hypothesis is that specific patterns seen in a well-learned M1-DLS state are a byproduct of strong recurrent network activation. In turn, with exploratory behavior, there could be less consistent neural patterning in PFC-DMS as well, but increased firing rates overall are sufficient to drive downstream motor patterns with constrained variability.

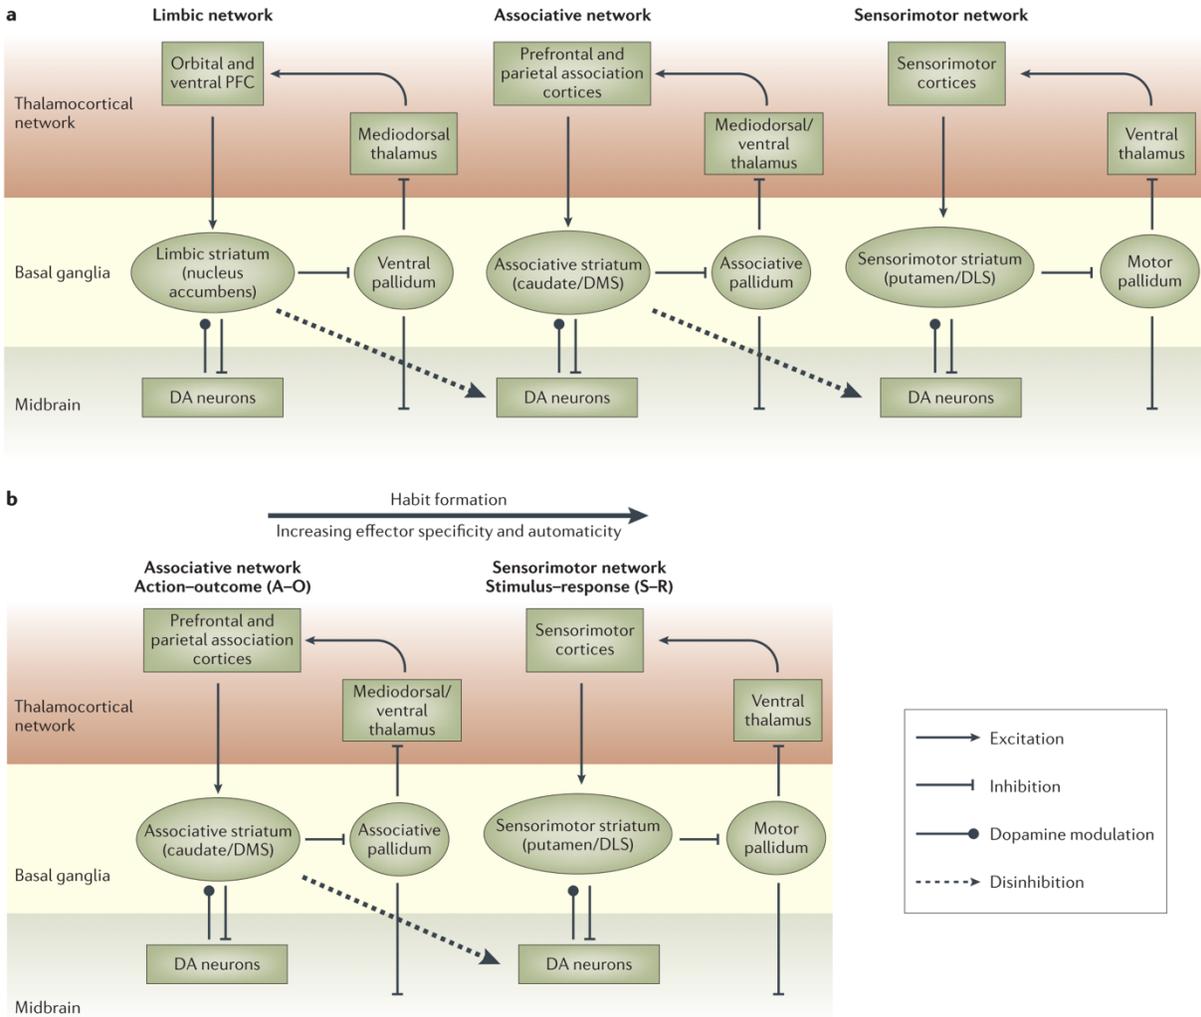


Figure 3.2. Simplified interaction of network loops, from HH Yin and BJ Knowlton, 2006. (A) Theory of cortical, striatal, and brainstem nodes may interact across the limbic, associative, and sensorimotor network loops. (B) Hypothesis as to how increased automaticity and habit formation shifts neural representation of motor action from associative network to sensorimotor network.

What are downstream nodes of M1 that could support the reach-to-grasp behavior after initial learning? There are two parallel pyramidal tract (PT) projection systems from M1, with both projecting to basal ganglia and brainstem: 1) the first type forms a recurrent loop with thalamic regions back to motor cortex; 2) the second type bypasses thalamus and projects directly to motor nuclei in the medulla¹⁸. The former has been implicated in early preparatory activity, and the latter in later preparatory activity and motor execution. In turn, the variability in temporal firing seen in M1 and DLS on the variable session could be due to disruption of consistency in the motor preparatory pathway, with some preservation of consistency in the movement execution pathways from M1 which cannot be discretely separated from higher order electrophysiology techniques used in the study from Chapter 2. A third type of projection system exists in M1, the intratelencephalic (IT) type neuron pathway, which only projects to cortical and basal ganglia regions, but not to downstream structures such as thalamus and brainstem¹⁹. Similarly, the IT pathway alongside cortico-thalamic PT pathway could be responsible for overall variability in sequential firing seen in M1 and DLS during transfer learning. As recurrent loops are disrupted to enable behavioral variability, movement generating PT pathway activity is likely preserved to enable fast and constrained reaching. Evidence towards this hypothesis includes validation that population bursting in M1 is linked to movements, but not necessarily the same movement trial-to-trial, hinting towards a smaller subset of neural activity underlying the actual desired task movement²⁰.

In turn, I propose the following simplified model of neural support for flexible movement generation. First, with learning of a motor skill, co-activation of associative and sensorimotor loops establishes an optimal neural patterning, or manifold, over time, driven predominantly by

recurrent loops (**Figure 3.3**). If a variable skill is learned during this time (i.e. reaching to multiple locations), the manifold can support generalizable, adaptive skill that is seen in non-human primates and humans performing center-out reaching and variable reach-to-grasp skills²¹⁻²⁵. However, if the skill is too specific, and a broadening of the learned skill needs to occur via transfer learning instead of adaptation, the intrinsic motor manifold must alter in order to support these new movements^{26,27}. Specifically in the case of transfer learning of a fast, automatic skill, downstream motor pathways, such as motor brainstem nuclei, and connected spinal outputs have likely consolidated an optimal muscle output pattern. In turn, variability in spiking patterns at the cortico-striatal level could reflect a pause in strong recurrent activity, adding a constrained variability to the generated motor output, which still is likely influenced by non-recurrent PT projections from M1 (**Figure 3.4**).

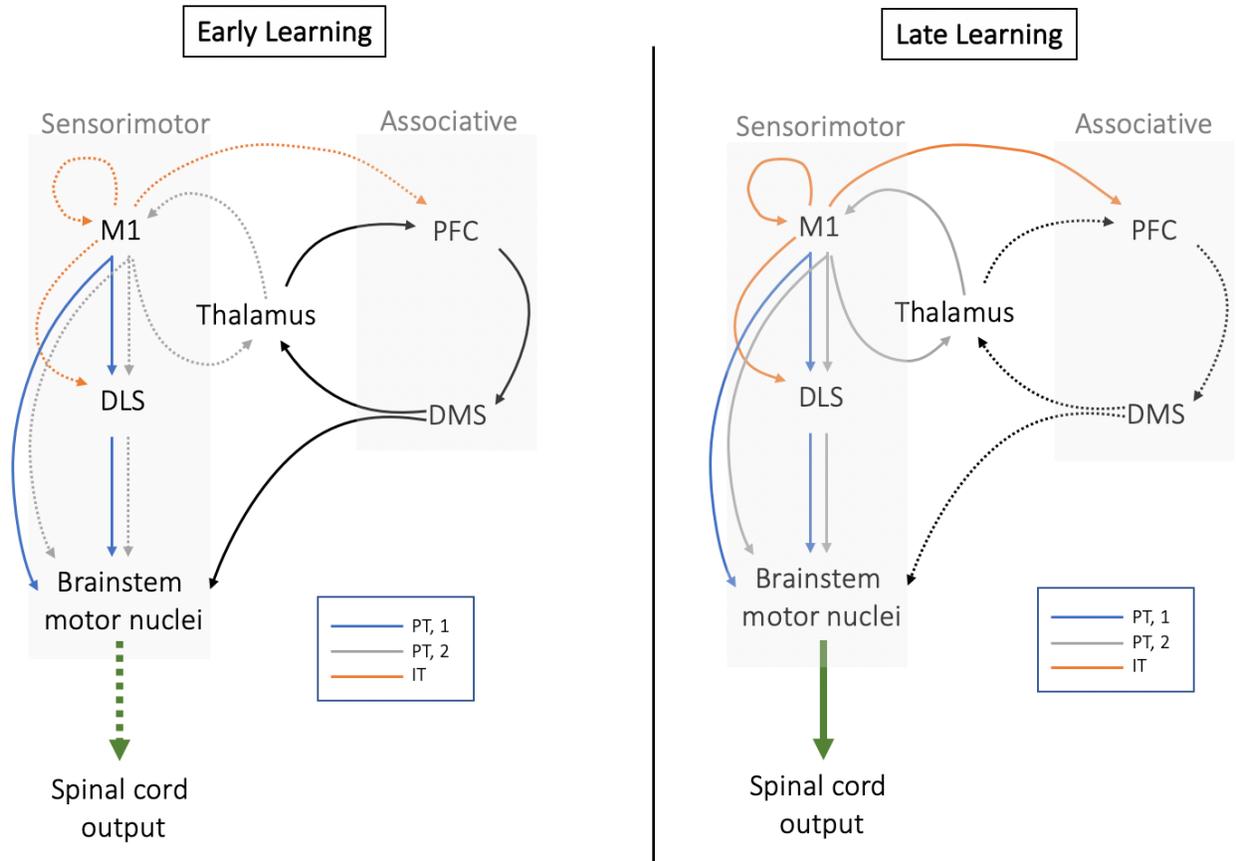


Figure 3.3. Proposed model of parallel motor loops in early versus late learning. In early learning (left) recurrent PT and IT pathways have not yet been consolidated; rather direct PT motor pathways and tandem activation of associative pathways support exploratory movement. In late learning (right), recurrent PT and IT pathways are more prominent, providing an optimal short-circuit for skilled execution.

Implications of corticostriatal dysfunction

Many neuropsychiatric diseases with motor symptoms exhibit disrupted corticostriatal functioning, such as obsessive-compulsive disorder (OCD) and autism spectrum disorder (ASD)²⁸. With well-learned healthy motor movement, there is co-engagement of sensorimotor cortical and striatal patterning that disengages following completion of the motor skill, demonstrating innate flexibility between movement execution modes^{2,15}. In Chapter 2, it becomes apparent that transfer learning, a subset of behavioral flexibility, also appears to rely upon corticostriatal patterning flexibility to break out of a rigid skill state and expand to include a second learned skill. How then might behavioral dysfunction arise with disruption of normal corticostriatal function?

One intriguing idea is that corticostriatal patterning may become rigid, with the inability to flexibly adapt back to a neutral state or to novel situations. The motor patterns in ASD and OCD can be quite stereotyped, demonstrating minimal variance, or perseverative, occurring repeatedly despite any goal-directed intention, respectively^{28,29}. What is the neural basis of these symptoms? While there are many competing hypotheses as to what drives compulsions in humans^{30,31}, mouse models of OCD recapitulate similar symptomatology and demonstrate altered corticostriatal synaptic function, specifically reduction of glutamate-mediated synaptic transmission in *Sapap3*-ko mice³², increased orbitofrontal cortex (OFC) activity and disrupted striatal morphology in *Slitrk5* mice³³, and reduction of striatal synapses in *Hoxb8* mice³⁴, among others. Similarly in ASD, mice models demonstrate perseverative behavioral features (i.e. increased grooming, impaired reversal learning), with accompanying striatal dysfunction (i.e. reduced striatal A2a receptor function, decreased dendritic morphology)³⁵⁻³⁷.

Even more intriguingly, repeated stimulation of OFC-ventromedial striatum (VMS) connections, in contrast to acute stimulation, promoted persistent grooming behaviors that were reversed by administration of fluoxetine, an OCD treatment³⁸. Most of these studies have been focused on two associative nodes of the cortico-striatal network, and almost none have examined the role of motor generating networks (i.e. M1-DLS) in the execution of these repetitive movement patterns. In particular, few electrophysiology studies examine the neural substrates of compulsive motor behavior in animal models of OCD or ASD. Thus, future studies examining how associative networks, such as those influenced by OFC, and motor networks interact in neuropsychiatric disorders with motor symptoms would likely shed light on: 1) neural firing patterns in motor network nodes (i.e. M1, DLS) that underlie these repetitive behaviors; 2) whether M1-DLS connections in disease animal models are as flexible as demonstrated in wildtype rodents; 3) how associative networks influence motor network firing patterns. By examining neuropsychiatric diseases from a higher-level network perspective, there can be further symptom-specific target discovery for either pharmacological or electrical intervention, enabling design of better therapeutic interventions.

Summary

In Chapter 1, I explored in depth the dilemmas faced when studying motor behavior, due to multiple axes being reflected in tasks currently studied in neuroscience (i.e. early versus late learning, goal-directed versus habitual action, and rigid versus generalizable skill). In turn, comparisons of the neural bases of behaviors across task types, both in terms of which skill and when each skill is studied, makes it difficult to abstract higher order conclusions about how motor regions interact during learning and execution of a variety of motor skills. However, more careful characterization of behavioral states, as well as newer methods of analyzing higher-dimensional neural data, are setting the scene for current neuroscientific endeavors to more carefully parse out network contributions to nuanced motor behavior.

I was particularly interested in exploring how motor network activity could support the transfer of a learned motor skill to a slightly modified environmental context, as detailed in Chapter 2. Surprisingly, I found that constrained learning of the reach-to-grasp skill promotes a rigid neural state, such that re-learning how to successfully retrieve the pellet from a new location was a multi-day process. Even more striking was the discovery that both M1 and DLS neural activity returned to an exploratory state early in transfer learning, while executed movements were still smooth and fast in contrast to true early learning of the skill. Thus, it is apparent that specific patterns of motor network activity may be more of a reflection of optimized recurrence from learning, that downstream motor structure may play a role in constraining learned behavior, and that parallel associative processing may enable the motor network to explore novel environmental parameters for subsequent re-consolidation of optimal recurrent patterns.

A model reconciling these hypotheses is introduced in Chapter 3, with supporting evidence from the literature regarding establishment of manifolds in the motor network during learning. In particular, differential M1 pathways, such as the recurrent pyramidal tract (PT) and intratelencephalic (IT) tract may provide the pattern generation seen in well-learned movements, whereas the non-recurrent PT pathway may support the actual motor output signal provided to downstream structures for movement execution. Subsequently I review the implications of disrupted corticostriatal flexibility in human disease, with future research directions.

References

1. Whishaw, I. Q., & Tomie, J. A. (1989). Olfaction directs skilled forelimb reaching in the rat. *Behavioural Brain Research*, *32*(1), 11–21.
[http://doi.org/10.1016/S01664328\(89\)80067-1](http://doi.org/10.1016/S01664328(89)80067-1)
2. Lemke, S. M., Ramanathan, D. S., Guo, L., Won, S. J., & Ganguly, K. (2019). Emergent modular neural control drives coordinated motor actions. *Nature Neuroscience*, *22*(7), 1122–1131. <http://doi.org/10.1038/s41593-019-0407-2>
3. Veuthey, T. L., Derosier, K., Kondapavulur, S., & Ganguly, K. (n.d.). Single-trial cross-area neural population dynamics during long-term skill learning. *Nature Communications*, (2020), 1–15. <http://doi.org/10.1038/s41467-020-17902-1>
4. Ramanathan, D. S., Gulati, T., & Ganguly, K. (2015). Sleep-Dependent Reactivation of Ensembles in Motor Cortex Promotes Skill Consolidation. *PLoS Biology*, *13*(9), 1–25. <http://doi.org/10.1371/journal.pbio.1002263>
5. Dolbakyan, E., Hernandez-Mesa, N., & Bures, J. (1977). Skilled forelimb movements and unit activity in motor cortex and caudate nucleus in rats. *Neuroscience*, *2*, 73–80. [https://doi.org/10.1016/0306-4522\(77\)90069-0](https://doi.org/10.1016/0306-4522(77)90069-0)
6. Storozhuk, V. M., & Bracha, V. (1984). 317 unit activity of motor cortex during acoustically signalled reaching in rats. *Behavioral Brain Research*, *12*, 317–326.
7. Sauerbrei, B. A., Guo, J. Z., Cohen, J. D., Mischiati, M., Guo, W., Kabra, M., ... Hantman, A. W. (2020). Cortical pattern generation during dexterous movement is input driven. *Nature*, *577*(7790), 386–391. <http://doi.org/10.1038/s41586-019-1869-9>
8. Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of

- neural population activity. *Journal of Neurophysiology*, *102*(1), 614–635.
<http://doi.org/10.1152/jn.90941.2008>
9. 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>
 10. Cunningham, J. P., & Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, *17*(11), 1500–1509. <http://doi.org/10.1038/nn.3776>
 11. Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation through Neural Population Dynamics. *Annual Review of Neuroscience*, *43*, 249–275.
<http://doi.org/10.1146/annurev-neuro-092619-094115>
 12. Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation, *7*(June), 464–476. <http://doi.org/10.1038/nrn1919>
 13. Killcross, S., & Coutureau, E. (2003). Coordination of actions and habits in the medial prefrontal cortex of rats. *Cerebral Cortex*, *13*(4), 400–408.
<http://doi.org/10.1093/cercor/13.4.400>
 14. de Wit, S., Watson, P., Harsay, H. A., Cohen, M. X., van de Vijver, I., & Ridderinkhof, K. R. (2012). Corticostriatal connectivity underlies individual differences in the balance between habitual and goal-directed action control. *Journal of Neuroscience*, *32*(35), 12066–12075. <http://doi.org/10.1523/JNEUROSCI.1088-12.2012>
 15. Kupferschmidt, D. A., Juczewski, K., Cui, G., Johnson, K. A., & Lovinger, D. M. (2017). Parallel, but Dissociable, Processing in Discrete Corticostriatal Inputs Encodes Skill Learning. *Neuron*, *96*(2), 476–489.e5. <http://doi.org/10.1016/j.neuron.2017.09.040>

16. Makino, H., Hwang, E. J., Hedrick, N. G., & Komiyama, T. (2016). Circuit Mechanisms of Sensorimotor Learning. *Neuron*, *92*(4), 705–721.
<http://doi.org/10.1016/j.neuron.2016.10.029>
17. Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., ... Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, *21*(6), 860–868. <http://doi.org/10.1038/s41593-018-0147-8>
18. Economo, M. N., Viswanathan, S., Tasic, B., Bas, E., Winnubst, J., Menon, V., ... Chandrashekar, J. (2018). Distinct descending motor cortex pathways and their roles in movement. *Nature*. <http://doi.org/10.1038/s41586-018-0642-9>
19. Shepherd, G. M. G. (2014). Corticostriatal connectivity and its role in disease. *Nature Reviews Neuroscience*, *14*(4), 278–291. <http://doi.org/10.1038/nrn3469>.Corticostriatal
20. 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. <http://doi.org/10.1038/s41593-019-0502-4>
21. Rajan, K., Harvey, C. D., Tank, D. W., Rajan, K., Harvey, C. D., & Tank, D. W. (2016). Recurrent Network Models of Sequence Generation and Memory Recurrent Network Models. *Neuron*, *90*(1), 128–142. <http://doi.org/10.1016/j.neuron.2016.02.009>
22. Athalye, V. R., Ganguly, K., Costa, R. M., & Carmena, J. M. (2017). Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Article Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control. *Neuron*, *93*(4), 955–970.e5. <http://doi.org/10.1016/j.neuron.2017.01.016>

23. Perich, M. G., Gallego, J. A., & Miller, L. E. (2018). A Neural Population Mechanism for Rapid Learning. *Neuron*, *100*(4), 964–976.e7.
<http://doi.org/10.1016/j.neuron.2018.09.030>
24. Gallego, J. A., Perich, M. G., Naufel, S. N., Ethier, C., Solla, S. A., & Miller, L. E. (2018). Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications*, *9*(1), 1–13. <http://doi.org/10.1038/s41467018-06560-z>
25. Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A., & Miller, L. E. (2020). Long-term stability of cortical population dynamics underlying consistent behavior. *Nature Neuroscience*, *23*(2), 260–270. <http://doi.org/10.1038/s41593-019-0555-4>
26. Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., ... Batista, A. P. (2014). Neural constraints on learning. *Nature*, *512*.
<http://doi.org/10.1038/nature13665>
27. Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., ... Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(30), 15210–15215. <http://doi.org/10.1073/pnas.1820296116>
28. Vicente, A. M., Martins, G. J., & Costa, R. M. (2020). Cortico-basal ganglia circuits underlying dysfunctional control of motor behaviors in neuropsychiatric disorders. *Current Opinion in Genetics and Development*, *65*, 151–159.
<http://doi.org/10.1016/j.gde.2020.05.042>

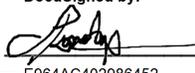
29. Lewis, M., & Kim, S. (2009). The pathophysiology of restricted repetitive behavior, *Journal of Neurodevelopmental Disorders, 1*, 114–132. <http://doi.org/10.1007/s11689-009-9019-6>
30. Robbins, T.W., Vaghi, M.M., Banca, P. (2019). Obsessive-compulsive disorder: puzzles and prospects. *Neuron, 102*, 27-47. doi: 10.1016/j.neuron.2019.01.046
31. Geramita, M.A., Yttri, E.A., Ahmari, S.E. (2020). The two-step task, avoidance, and OCD. *Journal of Neuroscience Research, 98*, 1007-1019. <https://doi.org/10.1002/jnr.24594>
32. Welch, J.M., Lu, J., Rodriguiz, R.M., Trotta, N.C., Peca, J., Ding, J-D., Feliciano, C., Chen, M., Adams, J.P., Luo, J., ... Feng, G. (2007). Cortico-striatal synaptic defects and OCD-like behaviours in *Sapap3*-mutant mice. *Nature, 448*, 894-900. <https://doi.org/10.1016/j.biopsych.2013.01.008>
33. Shmelkov, S. V, Hormigo, A., Jing, D., Proenca, C. C., Bath, K. G., Milde, T., ... Rafii, S. (2010). Slitrk5 deficiency impairs corticostriatal circuitry and leads to obsessive-compulsive – like behaviors in mice. *Nature Medicine, 16*(5), 598–603. <http://doi.org/10.1038/nm.2125>
34. Attwells, S., Setiawan, E., Wilson, A. A., Rusjan, P. M., Mizrahi, R., Miler, L., Xu, C., Richter, M. A., Kahn, A., Kish, S. J., Houle, S., Ravindran, L., & Meyer, J. H. (2017). Inflammation in the Neurocircuitry of Obsessive-Compulsive Disorder. *JAMA psychiatry, 74*(8), 833–840. <https://doi.org/10.1001/jamapsychiatry.2017.1567>
35. Silverman, J.L., Yang, M., Lord, C., Crawley, J.N. (2010). Behavioural phenotyping assays for mouse models of autism. *Nature Reviews Neuroscience, 11*, 490-502. doi:10.1038/nrn2851

36. Amodeo, D. A., Cuevas, L., Dunn, J. T., Sweeney, J. A., & Ragozzino, M. E. (2018). The Adenosine A_{2A} Receptor Agonist ,CGS 21680, Attenuates a Probabilistic Reversal Learning Deficit and Elevated Grooming Behavior in BTBR Mice, *Autism Research*, *11*, 223–233. <http://doi.org/10.1002/aur.1901>
37. Bechar, A., & Lewis, M. (2012). Modeling Restricted Repetitive Behavior in Animals. *Autism – Open Access, SI*. <http://doi.org/10.4172/2165-7890.S1-006>
38. Ahmari, S. E., Spellman, T., Douglass, N. L., Kheirbek, M. A., Simpson, H. B., Deisseroth, K., ... Hen, R. (2013). Repeated cortico-striatal stimulation generates persistent OCD-like behavior. *Science*, *340*(6137), 1234–1239. <http://doi.org/10.1126/science.1234733>

Publishing Agreement

It is the policy of the University to encourage open access and broad distribution of all theses, dissertations, and manuscripts. The Graduate Division will facilitate the distribution of UCSF theses, dissertations, and manuscripts to the UCSF Library for open access and distribution. UCSF will make such theses, dissertations, and manuscripts accessible to the public and will take reasonable steps to preserve these works in perpetuity.

I hereby grant the non-exclusive, perpetual right to The Regents of the University of California to reproduce, publicly display, distribute, preserve, and publish copies of my thesis, dissertation, or manuscript in any form or media, now existing or later derived, including access online for teaching, research, and public service purposes.

DocuSigned by:

E964AC402986452... Author Signature

4/7/2021
Date