

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Bilateral population activity in the motor cortex and its role in upper limb control

Permalink

<https://escholarship.org/uc/item/0pt9j2jx>

Author

Dixon, Tanner Chas

Publication Date

2021

Peer reviewed|Thesis/dissertation

Bilateral population activity in the motor cortex and its role in upper limb control

by

Tanner Chas Dixon

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Joint Doctor of Philosophy
with the University of California, San Francisco

in

Bioengineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jose M. Carmena, Chair

Professor Richard B. Ivry

Professor Frederic E. Theunissen

Professor Karunesh Ganguly

Spring 2021

Bilateral population activity in the motor cortex and its role in upper limb control

Copyright 2021
by
Tanner Chas Dixon

Abstract

Bilateral population activity in the motor cortex and its role in upper limb control

by

Tanner Chas Dixon

Joint Doctor of Philosophy

with the University of California, San Francisco in Bioengineering

University of California, Berkeley

Professor Jose M. Carmena, Chair

A canonical understanding of how the brain controls arm movements assumes that each hemisphere delivers commands exclusively to the contralateral limb. This has been supported by a rich history of work in clinical neurology and experimental neuroscience. However, a growing body of research suggests that the ipsilateral hemisphere has some form of involvement as well. For example, neural activity recorded from a single hemisphere of the motor cortex may be used to accurately decode movements of the ipsilateral arm. The precise functions of these ipsilateral signals, and the bilateral nature of computations in the motor cortex more broadly, are poorly understood.

In this thesis, we investigate the computations underlying isolated control of a single arm and coordinated control of both arms. Single-unit activity was recorded from both hemispheres of the motor cortex in non-human primates performing a range of behavioral tasks designed to test specific hypotheses regarding functional laterality.

Moving our limbs requires a dynamic process of observing the environment we will interact with, selecting an appropriate action, preparing that action with the correct effector, and executing it. In Chapter 2, we ask: How does emerging population activity organize to form arm-specific signals as motor plans are prepared and executed? We find that the population signals are marked by two different components, which we refer to as “dedicated” and “distributed”. Dedicated signals were comprised of activity that was largely segregated for the two arms at the level of individual units, with activity most prominently located in the contralateral hemisphere. This component gradually emerged across preparation and movement and drove population activity for the two arms into divergent neural subspaces. In contrast, the distributed component represented a

fundamentally bilateral function, as it contained behaviorally-specific information for both arms but did not segregate the population signals for each arm into separate neural subspaces. These two components allow for both independence and interaction of bilateral arm signals.

Although nominally “motor” the motor cortex displays sensory responses as well. In Chapter 3, we ask: Is the postural state of both arms integrated into commands for a single arm? Using a task that manipulates the static posture of the stationary hand during unimanual reaching, we found that the relationship between neural activity and behavior was sensitive to the state of the stationary hand. Coding of the reach targets changed congruently with the posture of the stationary hand, i.e. the mapping between neural activity and reaching behavior shifted in the same direction that the posture of the stationary hand was moved. Our results offer mixed support for the intriguing hypothesis that ipsilateral activity reflects a parallel plan for how the stationary hand would move if it were selected for unimanual action. Alternatively, they may be interpreted in the context of bimanual coordination or reflect whole-body control that changes when the mechanical properties of the linked system are altered.

Finally, in Chapter 4, we ask: Do network dynamics spanning the two hemispheres change flexibly to meet the coordination requirements of different bimanual tasks? Daily usage of our hands typically requires the two limbs to work in concert with one another, adopting task-specific coordination patterns. Here, we introduce a novel bimanual task that differentiates not only bimanual control from unimanual control, but independent bimanual control from bimanual coordination. We find that the mapping between neural activity and reaching behavior is altered between isolated unimanual movements and coordinated bimanual movements. However, the dynamics governing the evolution of neural activity appear to be largely maintained across control conditions. One interpretation of these results is that activity in the motor cortex coordinates commands for the two limbs not by adopting new interhemispheric dynamics, but by occupying different regimes of a context-invariant dynamical landscape.

Our results regarding these questions offer new insights into the basic science of bilateral motor control and provide important foundations for future neurotechnologies like bimanual neuroprosthetics.

Dedicated to
Wallace and Obi-Wan

Contents

Chapter 1 Introduction	1
1.1 Descending and commissural motor pathways	1
1.1.1 Descending motor pathways	1
1.1.2 Interhemispheric communication	3
1.2 Motor planning and preparation.....	4
1.2.1 Preparatory processes and their neurophysiological implementation	4
1.2.2 Dynamical systems theory and “initial conditions”.....	6
1.3 Bilateral interactions in behavior and cortical neural activity.....	8
1.3.1 Spatiotemporal constraints on bimanual behavior	8
1.3.2 Insights on bimanual control from patient populations.....	10
1.3.3 Bilateral signals in the motor cortex.....	10
1.3.4 Implications for neurorehabilitation and assistive neurotechnologies	13
Chapter 2 Hybrid dedicated and distributed coding of bilateral arm signals	15
2.1 Introduction.....	15
2.2 Results	17
2.2.1 Behavior.....	17
2.2.2 Arm-dedicated units emerge across task phases while the overall distribution remains relatively arm-neutral.....	19
2.2.3 Modulation preferentially occurs within arm-dedicated units	22
2.2.4 The population signal is largely confined to arm-specific sub-populations ...	24
2.2.5 Neural subspaces for the two arms diverge across task phases.....	27
2.2.6 Subspace separation relies upon dedicated signals	30
2.2.7 An additional distributed signal contains target-specific information about both arms.....	32

2.2.8	The distributed signal is contained in a shared subspace for the two arms	.32
2.3	Discussion	35
2.3.1	Comparison to previous studies of bilateral arm signals in the motor cortex	35
2.3.2	Progressive segregation of arm-dedicated signals and its functional significance	36
2.3.3	Bilateral signals and their role in motor control	37
2.3.4	Interpretations from a dynamical systems perspective	38
2.4	Methods	40
2.4.1	Behavioral recordings and task	40
2.4.2	Surgical implantation	41
2.4.3	Electrophysiology	41
2.4.4	Modulation and Arm Preference metrics	42
2.4.5	Principal components analysis	43
2.4.6	Dimensionality estimation	44
2.4.7	Covariance alignment	45
2.4.8	PCA coefficient analysis	46
2.4.9	Linear discriminant analysis	47
2.4.10	Fine timescale analysis of population coding and subspace development (heatmaps)	47
2.4.11	Permutation testing procedures	47
2.5	Supplementary Figures	49
	Acknowledgements	54
Chapter 3 Interactions between bilateral limb posture and unilateral reaching signals		
		55
3.1	Introduction	55
3.2	Results	57
3.2.1	Task design	57
3.2.2	Single-unit responses	57

3.2.3	Target classification using population activity is sensitive to posture of the stationary hand.....	59
3.2.4	Target coding shifts congruently with shifts in the posture of the stationary hand.....	63
3.2.5	Target classifiers fail to generalize across the two hands	65
3.2.6	Analysis of the interactions between target and posture coding using dPCA	66
3.3	Discussion	69
3.3.1	Interpretation through the lens of bimanual coordination	69
3.3.2	Covert motor commands for the stationary arm.....	71
3.3.3	Additional considerations and future analysis ideas.....	72
3.4	Methods	73
3.4.1	Dataset details	73
3.4.2	Task design.....	74
3.4.3	Surgical implantation.....	75
3.4.4	Single-unit tuning	75
3.4.5	Comparison of tuning profiles across stationary arm postures	75
3.4.6	Analysis of similarity between ipsilateral and contralateral target responses	76
3.4.7	Linear discriminant analysis (LDA) generalization procedure.....	76
3.4.8	LDA exclusion criteria	77
3.4.5	Demixed principal components analysis (dPCA)	77
3.4.11	Permutation-based ANOVA	78
3.5	Supplementary Figures	79
	Acknowledgements	84
	Chapter 4 Interhemispheric population dynamics during bimanual coordination	85
4.1	Introduction	85
4.2	Results	86
4.2.1	Task design.....	86

4.2.2	Behavioral performance and task validation	90
4.2.3	Strength of movement representations in each hemisphere across task conditions.....	92
4.2.4	Population dynamics in the ipsilateral vs contralateral hemisphere.....	94
4.2.5	Comparison of dynamical mode properties across coordinated and independent control	97
4.2.6	Interhemispheric interactions.....	100
4.3	Discussion	102
4.3.1	Interhemispheric neural population dynamics.....	102
4.3.2	Context-invariant dynamics in PMd and M1.....	103
4.3.3	Inhibition versus coordination.....	104
4.4	Methods	107
4.4.1	Dataset details	107
4.4.2	Task design.....	107
4.4.3	Surgical implantation.....	108
4.4.4	Linear kinematic decoder	108
4.4.5	LDS fitting and testing	109
4.4.6	Characterization of dynamical modes	109
4.4.7	Procedure for “lesioning” dynamical modes	111
4.5	Supplementary Figures	112
	Acknowledgements	117
	Chapter 5 Conclusion.....	118
5.1	Summary of contributions.....	118
5.1.1	Independence and interaction of unimanual reaching signals.....	118
5.1.2	Integration of bilateral state for unilateral motor commands.....	119
5.1.3	Interhemispheric population dynamics involved in bimanual coordination	120
5.2	Future directions.....	121
5.3	Publications and presentations resulting from thesis	122
	References	124

Acknowledgements

It became clear to me early on that the relationships I would form during my time in Berkeley would be at the core of every good thing that came out of my PhD. I was surrounded by so many brilliant and compassionate people and they deserve every ounce of recognition that I can give. My mind would often wander during the more mundane daily tasks of my PhD (building grids, cleaning, etc.) and I would begin drafting this acknowledgements section in my mind. As I ruminated over the most recent person who I was feeling thankful for, I came to believe that this would be the most exciting portion of my thesis to write.

A fitting place to start is with the person at the helm of my PhD, my advisor, Jose M. Carmena. As I was preparing for grad school and consulting people at UW about the different labs I was interested in, every single person I talked to had incredibly positive things to say about Jose and his lab. My first in-person discussions with Jose after arriving in Berkeley are marked in my memory by two things that I still see as central to his personality. First, he was passionate. His excitement when talking about science and technology is infectious, and I knew that having an advisor like him would help me stay fueled for the long journey of a PhD. Second, his passion was most strongly displayed when talking about the people in his lab. I would hear things like, “My student Preeya is incredible, you should talk to her about this work she’s been doing” or, “Sam is brilliant, she’s been spearheading this awesome collaboration.” He takes so much pride in every person that comes through his lab and values us on a deep level. These are tremendous leadership qualities that I have been very encouraged by and hope to emulate. Thank you for your support, both professionally and personally, over the past years.

I was fortunate enough to have another mentor figure in Richard B. Ivry. I did not originally plan to do a rotation with Rich my first year, but I consider the decision to rotate in his lab my first term to be the greatest moment of serendipity during my PhD. On a practical level, it helped me begin to get involved with the “Ipsi” collaboration which would end up becoming my thesis work. I also learned a great deal about rigorous experimental design and behavioral analysis from my time with the Ivry lab. It is difficult to explain, but I quickly got the sense that Rich and the people in his lab were the kind of people I wanted to be around. It really began with Rich, Hyosub, and Darius. Whether I was getting humbled playing ping-pong or batting around adaptation modeling ideas, I just had fun with these guys. Each new person I met in the lab from then on was the same way. Brilliant scientists, but also people I simply enjoyed

being around. At the risk of sounding cliché, they felt a lot like family. When you find a group of people that bring you joy like that, you would be a fool to leave (and Alyssa explicitly told me as much). So, even though I joined Jose's lab, I kept coming to Ivry lab meetings and social events. No one ever kicked me out; maybe they didn't have the heart for it, or I just refused to listen. Thank you (the entire lab) for creating such a welcoming community, and also for partially transitioning the competition from the ping-pong table to the basketball court.

From the Ivry Lab, Christina Merrick deserves a particular shout-out. Christina was my partner-in-crime designing these experiments, as we were trying to orchestrate paralleled NHP and human studies, and she was leading the human end. Christina reminds me a lot of my sister, which is about as great of a compliment as I can give a person. A warm personality, good at all the things I'm bad at. I'm so glad that we will get to continue working together during our post-docs. Thank you for commiserating with me over publishing woes and being a constant source of positivity.

My first mentor from our lab (other than Jose) was Preeya Khanna, who I cannot say enough good things about. She is likely the most capable person I've ever met. She seems to know something about everything. Despite that fact, she is very humble. She's been one of my biggest role models during my PhD, and even to the final days was helping me with my research. On top of all that, she's been a tremendous friend. I've truly enjoyed our bike rides, getting me to try out ultimate frisbee, coming over for holidays and other fun events. Thank you, Preeya, for being such an incredible mentor and friend.

Samantha Santacruz was another pivotal mentor that I was fortunate to have. Sam was there actively supporting me during all the biggest moments of stress and celebration during my PhD. I'll never forget the first day that I recorded spikes in one of my animals. Sam had been helping me get the whole setup working and was on the audio monitor while I lowered electrodes. What a moment of relief and excitement when I heard those first few beautiful "pops." Our friendship really grew throughout our years of overlap in the lab as well, years during which I got to witness so many great things happening in her life: getting a faculty job at UT and beginning to grow her beautiful family. Thank you, Sam, for helping me make sure that my PhD did not fail before it began. I really look up to you.

David Piech and Albert You were both grad students who joined the lab the same year as me. I met both of them at other grad school interviews before ever coming to Berkeley. It's funny how the world works that we all ended up in the same lab

together. David, I love your excitement for all things science and technology. You have such a raw passion that is very inspiring. Thank you for your friendship and always interesting research feedback. Albert was like a brother to me. We spent multiple Christmas's together when we were both stuck in the Bay during the holiday season. He was the first person I would go to for help with any technical questions. He helped me build my first computer. He supported me with great compassion during some of the harder times of my PhD. Any grad student faces times where they look back and wonder if they made the right choice going to grad school, but realizing that I wouldn't have his friendship otherwise always helps affirm my decision. Truly, thank you for your friendship.

I have so many other members of the Carmena Lab, past and present, to thank. Vivek is the type of person that every lab needs as its cornerstone: active in discussion at lab meetings and always providing insightful questions and feedback. He is also a very bright personality and someone who I viewed as another major role model. Vivek, I feel robbed that you were in Portugal for my early years in the lab and I didn't get to meet you until later. Also, sorry that one of our first interactions involved me sending you a gif of Zac Efron from his High School Musical days. Suraj and Sid, I never overlapped with you in the lab, but thank you for inviting me to play basketball so that I could maintain my sanity. Nuria, I appreciate how real you are. Both in casual conversation and when talking research, I always find your thoughts to be very honest and original. Will, the interest that you took in my research was a real encouragement during my PhD. I enjoyed being able to slide my chair over and gab about any topic with you when I needed a distraction from whatever I was working on. Ellen, I've been so impressed by how quickly you came in and accomplished so much in the lab. You charted an incredibly challenging and ambitious research path. You clearly have that "X-factor" where I would bet on you being successful in any endeavor. I have tremendous respect for you and really value your friendship as well. Maki, how did I get anything done without you? Across the board, everything felt upgraded when you joined the lab – quality of the experiments, animal care, my own mental health. I deeply admired and appreciated the conscientiousness with which you approached animal work, always treating them with profound respect and care. I'll forever be grateful for you taking over my responsibilities in the lab when I needed to leave for a period of time. I learned a great deal from you about how to be a courteous lab member as well, which you set an excellent example of. Thank you for so very many things. Others – Tom, Ryan, Axel, Manu, Gabby – I am thankful for all of you as well. Allen, Cuong, and Maxwell, the undergraduates who helped me through the years, thank you.

I also leaned heavily on the Wallis lab during my PhD. In a very real sense, I could not have done my PhD without Joni. Thank you for lending your expertise to make my experiments and animal care the highest quality. I had full confidence in any aspect of the projects that you were involved in and any advice that you gave me. Having such confidence in you as another mentor figure had the trickle-down effect of helping me to develop my own confidence with regards to experimental methods and the myriad steps involved in actually collecting this type of data. Thank you. Thank you also to Zuzanna from Joni's lab, who began learning to run experiments around the same time as I did. She is a brilliant and resilient person who I have tremendous respect for. I enjoyed (?) commiserating over how difficult this work was and sharing any new methods we had discovered or successes we had come upon. I felt a strong camaraderie with Zuzanna as NHP grad students and have consistently been impressed with what she has accomplished. Thank you for your support and friendship.

I also owe a great deal of gratitude to my other dissertation committee members, Frederic E. Theunissen and Karunesh Ganguly. I sought out Frederic to be on my committee because he is known for having strong statistical methods and rigorous analytic approaches. These were qualities I wanted my own work to be known for, and I am very thankful that Frederic was able to serve on my committee and give me helpful feedback on these aspects of my project and others. Karunesh's work as a post-doc was what motivated my entire project, so there was something very special about having him on my committee as well. He is an incredibly thoughtful person, and I felt that there was always a great deal of traction in every conversation I had with him – meaning that we were battling around important considerations and I was coming out of it with some new insight or a new idea for analysis. Moreover, Karunesh is very clearly passionate both about the brain and helping people. It is inspiring to see the work that he's doing, and he is yet another person who has become a role model for me in these regards during my PhD. I really hope to chart a similar course of passionately exploring the intricacies of brain function while directing all those efforts towards things that will improve human lives.

Thank you also to Steve I. Perlmutter, my undergraduate faculty mentor. Steve, I learned so much from you during my undergrad and feel incredibly lucky to have had you as a mentor. You are kind and thoughtful, and truly the first person who I looked at and thought, "I can be happy becoming a scientist if I can follow that model." I remember running an experiment with you and Rob and dealing with noise issues in the recording, then being caught off-guard when you just stuck the ground connection straight into the wall socket. It was the first time I had seen that trick, but it went

against all the lessons my mom had taught me growing up. I believe that, officially, the first day I did the same thing to solve a noise issue was the first day I became a neurophysiologist. Thank you also to Bethany Kondiles, who was my graduate student mentor while in Steve's lab. Bethany, thank you for sharing the realities of graduate school with me so that I could be prepared for my own journey. Thank you for taking a real interest in me as a whole person and wanting to know how I'm "really doing." Some of the strongest encouragement I've gotten throughout my early research career has been you telling me that you're proud of me. Thank you also to Robert Robinson, who was another critical mentor during my time in Steve's lab. Rob, you probably taught me the most practical lessons that helped me get through grad school. You helped me understand hard work, and what that looks like in a research setting. I love that you never take anything too seriously and are always good for a laugh. I've been lucky to work with a lot of great people already, but truly you have been the easiest person for me to work with and have helped me discover the types of people and environments that I can be most effective in. I never had to wonder what you were thinking, and you were always very direct with me. You're also a tremendous friend, and I am thankful for you.

So many other friends and communities have supported me along this journey. I thank my best friend from back home, Daniel, whose friendship has really helped take the edge off the stresses of grad school. I thank my close friends Evan and Isabel who I met in stats classes at Berkeley – you have both become very special people in my life and I am grateful to have met you. Thank you to the friends in my department, who provided me my first sense of community as I adapted to life in Berkeley. Thank you to the community I gained through my gym, in particular MB, KB, and PB who created such a fun environment where I could step away from science and engineering. Thank you also to everyone from my church home group: DK, ES, DS, JS, CR, NR, BW, and all the others. You formed a pivotal community for Alyssa and me.

To my family, thank you for your love and support. Mom, thank you for being a shining example of strength in my life. Bailey, you have always been my best friend and the first person I call if I need to escape stress. Dad, from you I've learned to let passion fuel me, as I'm inspired by the excitement and tenacity you show in all your endeavors.

Finally, I must thank my wife, Alyssa, for putting up with this crazy amount of work and supporting me along the way. The cliché thing to say is that "I couldn't have done this without you," but I feel that's not quite the right perspective. When you have someone who you would always rather be spending time with instead of working, working becomes very hard. I think about the times when I would be in lab building my

grids for a recording by 7am, then you would text me in the evening asking when I would be home, and I would have to text back saying I still had a few more hours of cleanup to do. Or, during the pandemic when we were both home and you wanted to go to the beach, but I had to say no because I was busy writing my dissertation trying to meet some deadline. Your love and support were invaluable, but moments like those made my PhD more difficult in some ways. So, could I have done this without you? I don't think that's the right question, but I know that I wouldn't have wanted to.

Chapter 1

Introduction

The internal brain processes involved in arm movement are incredibly complex, and they are made only more complex by the fact that we have two of them often working in cooperation. Producing appropriate movements requires dynamic selection, planning, and control processes. Interactions within local brain networks, across hemispheres, and between cortical and subcortical structures all contribute to the necessary computations. In this introductory chapter, we provide a broad review of the literature surrounding neural control of unimanual and bimanual arm movements. We then close with discussion on the translational relevance of these topics.

1.1 Descending and commissural motor pathways

Discussion of lateralized function in the primate motor system must begin with a characterization of its anatomy. Here we review the primary descending pathways connecting the brain to interneurons and lower motor neurons in the spinal cord, with an emphasis on the presence or absence of midline crossing. Additionally, we discuss commissural connections that link the two sides of the nervous system. Attention will be focused on those pathways likely to be involved in control of the arm and hand.

1.1.1 Descending motor pathways

Direct connections between the cortex and spinal cord form the corticospinal tract. Corticospinal neurons that project to the cervical spinal cord can be found in both the primary motor cortex (M1) and several premotor areas of the brain, although they are most dense near the central sulcus [73]. Approximately 85% of corticospinal axons cross the midline at the medulla and enter the lateral funiculus, while the remainder descend ipsilateral to the hemisphere of origin and enter either the lateral or ventral funiculi [139, 183].

Neurons of the dorsolateral corticospinal tract, the fiber bundle in the lateral funiculus, are classically thought to terminate contralaterally in the dorsolateral intermediate zone and on motor neurons in the ventral horn [20]. However, these axons may also decussate again within the spinal cord or sprout bilateral branches [183]. This pathway influences both proximal and distal musculature in the arms [179]. Of particular interest are those cortical neurons that make monosynaptic connections to spinal motor neurons, called cortico-motoneuronal (CM) cells. These unique neurons are confined to the caudal portion of primary motor cortex, or “new” M1 [179] and seem to exclusively innervate the contralateral limb [202]. The direct and specific influence of CM cells is thought to provide increased fractionation of movement, like fine control of individual fingers [14, 146, 147].

The ventromedial corticospinal tract, which descends through the ventral funiculus, terminates bilaterally in the ventromedial intermediate zone [20]. This pathway primarily affects axial muscles of the trunk and proximal muscles of the shoulder girdle [20, 139].

The motor cortex may also exert its influence on movement via the reticulospinal tract. Several frontal motor areas in the cortex make connections with the reticular nuclei of the pons and medulla [139, 126]. These corticoreticular projections originate preferentially in areas rostral to M1, travel both ipsi- and contralaterally, and include corticospinal collaterals [36, 126, 146].

The reticulospinal tract descends within two primary bundles: the medial reticulospinal tract located in the ventral funiculus, and the lateral reticulospinal tract in the lateral funiculus [139]. The medial division originates in both the pons and medulla and descends ipsilaterally, while the lateral division originates in the medulla and descends bilaterally with a more dominant ipsilateral component. The reticulospinal tract is unique in its bilateral organization and is thought to be the main conduit for ipsilateral motor evoked potentials from cortical stimulation [6, 226]. The contralateral component of this pathway displays an extensor bias, while the ipsilateral component displays a flexor bias [11, 59, 60]. Reticulospinal output is thought to primarily affect axial and proximal muscles, although some evidence for grip control exists [11, 143]. It is also involved in regulating muscle spindle sensitivity via gamma motor neurons [83], potentially playing a role in preparation for movement [27, 168].

The vestibulospinal tract shares many features with the reticulospinal tract, and they are often grouped together as medially descending brainstem pathways. It originates in the vestibular complex of the lower pons and medulla, which receives afferent inputs through the vestibular nerve that conveys sensory information regarding orientation of the head [203]. Inputs from the cerebellum containing proprioceptive information may integrate head and body information in the vestibular nuclei [156, 174].

Fibers from the vestibulospinal tract descend within the ventral funiculus of the spinal cord with primarily ipsilateral termination points [139]. It likely plays only a peripheral role in control of the arm and seems primarily dedicated to orienting behavior and balance.

Connections between the red nucleus and the spinal cord constitute the rubrospinal tract and bear similarity to the features of the corticospinal tract. The motor cortex provides input to the red nucleus [138, 148], and functional overlap between the corticospinal and rubrospinal pathways has been illustrated by behavioral recovery [143] and physiological changes [15] following corticospinal tract lesions. The fibers of the rubrospinal tract descend contralaterally through the lateral funiculus, terminate in the lateral intermediate zone of the spinal cord, and make some direct connections to lower motor neurons [139, 147]. These connections provide more distal innervation than the other brainstem pathways, and, by virtue of their connections to short propriospinal interneurons and direct connections to lower motor neurons, are likely more specific in their influence on movement [139]. Work in cats has suggested that rubrospinal activation maintains a flexor bias [109, 139], yet the reverse appears to be true in bipedal primates (extensor preference) [97, 161]. The flexor-extensor relationship of rubrospinal innervation, however, may be flexible [15].

For thorough classical reviews of the descending motor pathways, see [139] and [147], which have been heavily cited throughout this introduction.

1.1.2 Interhemispheric communication

The largest white matter pathway connecting the cerebral hemispheres is the corpus callosum, comprised of approximately 200 million fibers in humans [78]. The anterior-posterior organization of fibers within the bundle map roughly onto the same anterior-posterior arrangement of their cortical targets; projections to frontal brain areas occupy the anterior two-thirds (rostrum, genu, and body), with supplementary and primary motor areas being located primarily in the anterior body [62, 144, 167]. The corpus callosum largely connects homologous regions of each hemisphere, which is most pronounced for the supplementary motor area (SMA) [157, 184]. Using retrograde labelling to compare transcallosal inputs to the hand areas of M1 and the SMA, Rouiller et al., 1994 found approximately ten times as many total transcallosal projections to the SMA [184]. Additionally, they found that a greater proportion of SMA inputs came from its contralateral homolog (57-63%), whereas the primary source of transcallosal M1 inputs come from the premotor cortex (PM: 46-52%, M1: 24-38%). This significant interconnection of homologous SMA regions across the hemispheres suggests a uniquely bilateral role in motor processing [172]. It has therefore been widely implicated in bimanual coordination, which will be further discussed in a later section.

The anterior commissure represents another major communication pathway between the two hemispheres of the cerebral cortex. However, it primarily connects the temporal lobes of the cortex [224], and there is no strongly agreed upon role in motor control [88]. While the anterior commissure may not be relevant to discussion in the motor domain for subjects with typical neurological development, it is worth noting that patients with congenital agenesis of the corpus callosum have enlarged anterior commissures relative to patients with acquired callosal damage [196]. This may suggest a capacity of the anterior commissure to compensate for lost communication along the corpus callosum, provided that this lost communication occurs early in development.

Interhemispheric integration also appears to happen at the level of the basal ganglia. Fibers from frontal motor areas make both ipsilateral and contralateral (proceeding through the corpus callosum) connections in the striatum [137, 158]. The most pronounced crossed projections originate in the SMA and M1 [158], although corticostriatal input coming from all cortical regions is primarily ipsilateral [137]. Similarly, pathways connecting the motor cortices to the cerebellum via the nucleus reticularis tegmenti pontis provide bilateral input [24, 25, 187], suggesting the cerebellum as another candidate location of interhemispheric integration. It is also important to note that both the basal ganglia [128], and cerebellum [127] form recurrent loops of connectivity with the cortex, and the three structures may function as an integrated circuit [19].

Lastly, spinal commissural interneurons provide communication across the midline within the spinal cord. While typically discussed in terms of lumbar circuitry for locomotion, commissural interneurons are also found in the cervical cord [155]. The rich circuitry within the spinal cord itself should therefore also be considered in discussion of laterality for upper limb control.

1.2 Motor planning and preparation

We turn our discussion now to the cognitive and neurophysiological bases of motor planning for movements of the arm and hand. While we will discuss these topics largely within an information processing framework that separates perceptual, cognitive, and motor steps, we note that these aspects of motor planning are likely to be highly integrated and not strictly serial in nature [50]. Additionally, a focus will continue to be placed on laterality of action whenever relevant.

1.2.1 Preparatory processes and their neurophysiological implementation

In order to interact with our environment, we must perceive the relevant sensory information, digest and make judgements about that information, and prepare

appropriate motor commands to be deployed for action. These operations are, to some degree, performed in advance of movement and take time; if pre-cued with information (even partial information) about a required action, a subject can respond more quickly once instructed to move [40, 182]. We note, however, that movement initiation and movement preparation likely represent independent processes that both determine response times, depending on the behavioral context [49, 103].

A wide range of neurophysiological changes occurring prior to movement have been identified and associated with different preparatory processes. A rich literature using transcranial magnetic stimulation (TMS) protocols in human subjects has shown that corticomotor excitability is bilaterally modulated prior to cued movements of a single hand. Excitability in agonist muscles for the responding hand tend to increase during the reaction time following a go-cue, while excitability for the non-responding hand has been shown to decrease [149]. This selective excitation/suppression has been implicated in response selection. However, in behavioral paradigms where hand assignments on each trial are unpredictable and potential actions for each hand require the use of homologous muscles, increases in excitability have been reported in the non-responding hand [31, 159], suggesting that movements with both hands may be prepared in parallel to some extent. Studies using instructed-delay tasks have also provided evidence that bilateral decreases in excitability during delay periods reflect separate processes of selecting a response from candidate options (e.g. which hand to use) and withholding that response until the appropriate time of action [74, 75].

These observations of broad changes in corticomotor excitability have been complemented by higher resolution probes of preparatory activity using single-unit recordings in the cortex of non-human primates. While activity related to decision making, action selection, and motor preparation are widely distributed across the brain [50], we will focus our discussion here on the dorsal premotor cortex (PMd) and M1, which have been the most widely studied cortical regions in regards to motor aspects of preparation. Both PMd [40, 41, 45, 54, 110, 120, 181, 199, 219] and M1 [12, 45, 54, 121, 181, 209] neurons show instructed-delay activity related to upcoming movements. There is a general anterior-posterior gradient in the response properties across these adjacent areas: PMd shows stronger instructed-delay activity than M1, and its activity in general reflects more abstract features of behavior [53, 54]. Additionally, PMd activity during instructed-delay periods shows similar tuning for movements of both arms before becoming more effector-dependent during movement, while M1 activity is strongly effector-dependent throughout trial phases [45].

Early conceptions of how motor cortical activity relates to preparation tended to focus on a “rise-to-threshold” model. In this framework, each movement is associated with a particular threshold activation pattern, and inputs to the system describing parameters of the necessary action integrate over time to reach this threshold for

selecting an appropriate movement [48, 81, 101]. A simple form of this model would posit that preparatory activity is a sub-threshold version of what is observed during movement. This idea may find support in the TMS literature, where preparatory modulation of corticomotor excitability could bring the neural state closer to (or further from) this threshold.

However, tuning properties of motor cortical units often undergo changes between preparatory and movement phases of reaching tasks [43, 54, 120], suggesting a fundamental difference in their contributions to movement before and after initiation. Reaction times do not strictly correlate with the magnitude of preparatory neural responses either; more precisely, they seem to depend on proximity to a specific neural state [40, 42], such that higher firing rates may actually coincide with longer reaction times [40, 181]. Together, these observations have motivated a new perspective on preparatory activity, positing that the neural state prior to movement acts as the “initial condition” of a dynamical system instantiated in the cortical network [5, 43]. These initial conditions, while only abstractly related to the movement itself, seed the evolution of patterned activity that generates movement. This idea has quickly become the dominant working model of motor preparation in the field. We will discuss it further in the following section, along with the broader landscape of dynamical systems theory in motor control.

1.2.2 Dynamical systems theory and “initial conditions”

Many early studies of the motor cortex were principally interested in what features of movement neurons “represented.” Myriad motor parameters have been found to correlate with neural activity in the motor cortex during both preparation and movement, including reaction time [12, 40, 123, 181], speed of movement [41, 163], position, velocity, or direction of movement [41, 94, 95, 118, 163, 193, 218], and joint forces or muscle activity [39, 82, 118, 193]. In the prescient work of Fetz, 1992, it was suggested that chasing explicit representations of movement within the activation profiles of individual units might in fact be a red herring [86]. Rather, the contribution of each neuron may only be truly understood in the context of its interactions with the entire population. By virtue of existing within a densely interconnected network, units may display classical correlations with movement parameters, reflect complex or abstract features of the action, or seem entirely uninvolved when viewed in isolation, yet all contribute to the generation of population-level activity patterns that control movement.

In accord with this notion, the representational perspective has largely been replaced in recent years with one that views the cortical network as a dynamical system that generates time-varying patterns of activity for controlling movement. Convergent ideas have considered both the dynamics of behavior itself, e.g. feedback control [192, 212], as well as the dynamics inherent to the anatomy of the cortical network [44].

Neocortical circuits feature dense, highly recurrent connectivity [61]. This, in principle, results in two primary features that are of particular importance for the topic at hand [200]. First, activity will tend to be distributed across the entire network, yet likely be constrained to a lower-dimensional manifold governed by the functional connectivity and covariance of the population. Second, recurrence in the network will produce temporal dynamics that influence the evolution of the neural state. Under this theory, one must therefore consider neural activity not only in terms of its instantaneous contribution to behavior, but also in terms of its contribution to the evolution of the neural state.

The role of preparatory activity has been central to dynamical system theory of motor cortical function. As previously mentioned, preparatory activity is thought to act as an initial state for seeding the dynamics that are engaged during movement [5, 43]. The low-dimensional subspace in which this preparatory activity lies is distinct from the subspace containing movement activity, ostensibly making it null to the transformation that yields movement [80, 122]. Importantly, despite reports of the two activity spaces being orthogonal, preparatory activity is still strongly associated with subsequent movement responses [80], which is necessary for it to play a role in seeding movement dynamics. The neural state may evolve within the preparatory space towards an optimal initial position while awaiting some trigger for executing the movement. This initiation trigger likely represents a process independent of establishing the preparatory state [49, 103, 207]. Strong timing-related components of population responses in the motor cortex may reflect such a trigger [123]. Disruption of activity in premotor regions of the cortex late during preparation causes increased reaction times in monkey reaching [42] and ipsilateral response biases in rodent licking tasks [151]. Interestingly, unilateral inhibition early during planning has been shown to have very little effect on upcoming movements, yet if applied bilaterally, results in motor deficits [151]. This suggests that the networks involved in setting the preparatory state may be bilaterally distributed.

We end this section on a technical note. Neural network models are important tools that complement traditional neurophysiological recordings for investigating the dynamic behavior of neural signals [86]. By training recurrent networks to perform a behavioral task, one can identify common motifs in the network solution and compare these back with observed responses in their biological analogs [85, 207]. This has led to a more integrative and high-level pursuit of understanding activity in the motor cortex. For example, strong rotational components in population signals identified during movement [44] were recapitulated in a neural network model only when it was given incentive to find simple, i.e. highly regularized, solutions for producing similar output behavior [207]. This may suggest that rotational dynamics reflect the brain's pursuit of a simple solution. Thus, one may assign a theoretical "why" to otherwise puzzling experimental observations using such an approach. Importantly, this may operate in the

other direction as well: theory developed in network models may motivate experimentation. For more valuable discussion on neural population dynamics as they pertain to general computations performed in the brain, see [217].

1.3 Bilateral interactions in behavior and cortical neural activity

Most of the actions we perform with our arms require the engagement and coordination of both sides: an artist holding their sketchpad with one hand and drawing with the other, a pianist performing interleaved or simultaneous left and right hand key presses, a basketball player raising the ball with both hands in preparation to shoot, then removing their support hand just as they begin to flick the wrist on their shooting hand. Given the repertoire of normal behavior, it is no surprise that the nervous system has evolved to facilitate these kinds of bimanual interactions. Here, we first discuss behavioral observations of bimanual coordination and interference, as well as the insights obtained from comparing those features in certain patient populations. Furthermore, we survey the neurophysiology literature surrounding bilateral signals in the motor cortex and their relationship to unimanual and bimanual behavior. Finally, we discuss the translational relevance of these observations for informing neurorehabilitation approaches and developing assistive neurotechnologies.

1.3.1 Spatiotemporal constraints on bimanual behavior

Coordination of motor behavior happens at many levels, including agonist/antagonist muscle groups, synergistic movement of the joints along a single limb, and cooperative use of the two arms [68]. The unifying principle at each scale is that the overarching goal of behavior does not have any particular requirement on the individual components, but rather requires some correlational pattern across them. The general idea of divorcing action goals from their implementation is often referred to as “motor equivalence” [141] and has existed for nearly a century. Goal-directed actions may therefore be thought of in terms of variability that affects the intended outcome and variability that does not, termed the “controlled” and “uncontrolled” manifolds, respectively [188]. As one might expect from the name, variability tends to be greater along the uncontrolled manifold in well-coordinated actions. For example, joint angles tend to be variable across repeated reaching movements while the endpoint trajectories remain relatively stereotyped in comparison [102]. Stochastic optimal feedback control policies have been proposed that provide a formal computational and theoretic framework for interpreting these observations [212].

At present, we are chiefly concerned with coordination across the two arms or hands. Early investigations of bimanual behavior revealed a tendency to start and stop

movements of the two arms in synchrony, even when they move different distances [129, 130] or with different applied loads [195]. Tight synchronization of component sub-tasks is present in more complex bimanual actions, such as opening a drawer with one hand to retrieve an item with the other [124, 173]. Temporal constraints on bimanual behavior have also been revealed in repetitive movements, such as tapping the fingers at certain beats or tracing circles. Relative frequencies that are related by simple harmonics (e.g. 2Hz finger tapping with the left hand and 1Hz tapping with the right) are much easier to produce than those with more complicated structure [135]. Phase relationships are also pertinent. Subjects may produce stable anti-phase finger tapping or wrist flexion-extension movements at low frequencies, but as the frequency of movement increases, there reaches a point where a spontaneous transition to in-phase (co-activation of homologous muscles) movement occurs [131]. This interaction has been modeled using coupled dynamical systems, namely the seminal “Haken-Kelso-Bunz”, or “HKB”, model [104], which was among the first modelling endeavors to draw on the formalism of dynamical systems for understanding coordinated behavior. A “bimanual advantage” has also been observed during repetitive bimanual behavior, whereby the variability in cycle durations for a single hand is reduced when performing an action bimanually as compared to unimanually [107]. Clearly, the mechanisms that control timing of movement for the two arms are either shared or highly interactive.

In addition to these temporal constraints, movements of the two arms exhibit strong spatial interactions as well. For example, interaction of feed-forward motor commands has been demonstrated by placing an obstacle along the path of only one hand in a bimanual reaching task, which alters the trajectories for both arm movements [130]. Similarly, attempting to draw a circle with one hand while drawing a line with the other results in both drawings looking like something in-between a circle and a line [87]. Interactions in feedback control also exist and have revealed additional nuance. Perturbations randomly applied to a single arm elicit corrective responses in only the perturbed arm when a task entails two independent goals, yet identical perturbations cause bilateral corrective responses when the task requires integrated control from the two arms [66]. During integrated control, these corrective responses also adapt appropriately to changes in the required dynamics relating the two arms [67]. In a similar vein, certain spatial interactions have demonstrated sensitivity to the method of action cueing, i.e. symbolic or direct [63, 64], suggesting that they may involve higher level cognitive processes like response selection [116] or interactions in abstract perceptual space [160]. Collectively, these observations have led to an understanding that the two arms are controlled as a unified plant with somewhat flexible and context-dependent coupling.

Several excellent reviews of the cognitive and behavioral literature on bimanual coordination exist; for more reading, please see [68, 116, 208, 221].

1.3.2 Insights on bimanual control from patient populations

Psychophysics experiments have helped characterize the critical perceptual, cognitive, and motor aspects of bimanual behavior, allowing researchers to then look inward at the anatomy and physiology that implements them. One highly influential domain of research has been with patients that lack a corpus callosum, either from congenital agenesis or surgical transection. While spatial interactions in bimanual behavior are largely abolished in the absence of the corpus callosum [79, 88, 132], timing effects tend to be maintained, such as synchronization tendencies in discrete movements [114, 132, 196, 213] and the previously mentioned “bimanual advantage” [114]. Together, these results suggest that communication between cortical hemispheres subserves spatial interactions of the two arms, insofar as they can be dissociated from temporal aspects. Interestingly, certain well-practiced tasks that require spatial interaction of the two hands are relatively unaffected by callosotomy [90, 153, 176], yet there is a striking inability to learn new coordinative tasks [90, 175, 176]. These observations lend evidence that interhemispheric interactions in the cortex may be most critically involved in learning.

The cerebellum is another structure that has been broadly implicated in coordination, including the bimanual variety. There is evidence that the cerebellum supports temporal event structures for coordination [115, 116]. Patients with cerebellar lesions have shown desynchronization in bimanual behavior relative to healthy subjects [65, 194] yet maintain a “bimanual advantage” [89]. Cerebellar patients also show a deficit in adapting motor responses that require control of one arm to anticipate actions of the other [65], suggesting a role in learning/adaptation. Patients with Parkinson’s Disease (PD) reveal involvement of the basal ganglia as well. PD patients show decreased stability of anti-phase bimanual movement patterns compared to healthy subjects [32, 117], thus extending the classical role of the basal ganglia in timing to include temporal aspects of bimanual behavior.

1.3.3 Bilateral signals in the motor cortex

A classical understanding of the motor cortex assumes that one side of the brain controls the opposite side of the body. Indeed, modulation of single-unit firing rates tends to be stronger for the contralateral arm within the SMA [125, 133, 210], premotor [45, 71, 125], and primary motor cortices [8, 45, 71, 72, 108, 125, 133, 204, 210]. This is in broad agreement with the predominantly contralateral effects of cortical stroke [105], lesion [20], and stimulation [6, 162, 170, 171].

The first frontal motor area to be strongly implicated in bilateral control, and still the most prominent, was the SMA. The contralateral bias firmly established in M1 tends to be weaker, or even absent, in the SMA [21, 72, 125, 133, 210]. The notion that the SMA serves a uniquely bilateral function has been reinforced by observed deficits in

bimanual coordination following SMA lesions [23], extensive interconnection across the corpus callosum [184], and the presence of ipsilateral evoked responses following stimulation [162, 171]. As compared to M1, greater bilaterality in premotor cortex signals has also been noted [45, 210]. Our conversation that follows will focus primarily on PMd and M1.

Despite a largely lateralized organization that should not be overlooked, interest in the involvement of M1 during ipsilateral and bimanual movements has grown in recent years. The primary motor cortex exhibits responses to both ipsi- and contralateral arm movements during preparation [45, 71] and execution [8, 45, 71, 72, 92, 108, 125, 154, 204, 210]. Frequently, responses to both arms are observed even within the same neuron. Importantly, ipsilateral activity is weakly present during movements of the hand and fingers as well [69, 154, 210, 215]. The finding that such distal movements correlate with ipsilateral activity in the motor cortex is notable given that the known anatomical pathways connecting the motor cortex to motor neurons in the ipsilateral spinal cord involve primarily proximal musculature [20].

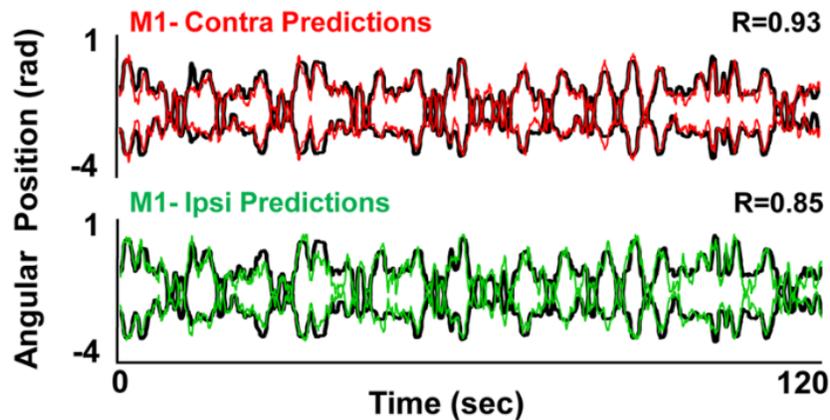


Figure 1.1 Contralateral and Ipsilateral decoding of arm movements. Kinematics during movements of a single arm (black lines) plotted with their decoded predictions using spiking activity in M1 of either the contralateral (red) or ipsilateral (green) hemisphere. Included with permission from [92].

Considering that the canonical role of the motor cortex is to control the contralateral body, what is the function of ipsilateral signals? The simplest hypothesis is that ipsilateral activity provides an independent control signal to operate alongside contralateral drive. Neurons that modulate exclusively during ipsilateral arm movements have been reported but are uncommon in both PMd and M1 [45, 71, 72, 108, 154]. Rather, these ipsilateral signals tend to be contained within neurons that are modulated for both arms. During preparation ipsilateral responses in PMd neurons show weaker magnitude but similar tuning to contralateral responses, while during movement tuning profiles become more distinct for the two arms [45]. Correlations between the tuning for ipsilateral and contralateral arm movements in M1 units have been reported but tend to be weak or absent [45, 108, 204]. However, strong correlations in bilateral activation patterns have been observed in more macro-level recordings such as functional magnetic resonance imaging (fMRI) [69], and electrocorticography (ECoG) [30]. Given the observations that ipsilateral activity tends to occur in units that are modulated for both arms, yet its relationship to behavior often differs from contralateral activation, a more intricate mechanism of signal independence has been recently proposed. Inspired by reports of orthogonal preparatory and movement subspaces [80], Ames and Churchland, 2019 and Heming et al., 2019 reported, in parallel, that bilateral arm signals could be separated at the level of the population using linear methods [8, 108]. Even if signals for the two arms are completely mixed within the same neurons, unique covariance structure specific to each limb may allow the population-level signals to orthogonalize. While this doesn't suggest any specific functional role for ipsilateral signals, it poses a new method of independence in the signal statistics that has functional implications.

Another hypothesis is that the brain prepares in parallel both left and right arm movements for accomplishing an action goal, despite a single side ultimately being selected for execution. Activity related to each potential action would still originate in the contralateral hemisphere. In this way, contralateral control would remain an apt descriptor of the system despite bilateral activity. Prior to action selection, activity patterns related to multiple potential actions have been simultaneously observed in cortical motor areas [47, 50, 111]. Similarly, signals may reflect effector-independent aspects of the computations necessary for producing movement. A related concept of "compositional coding" has recently been put forward by Willet et al., 2020 [223]. In this study, population-level signals in nominal hand areas of human premotor cortex were found to contain whole-body motor information. Furthermore, these signals could be separated into components representing effector-independent information about the movement and components representing which effector carried out the movement (e.g. left arm, right arm, left leg). However, the inconsistent correlations in tuning profiles between ipsi- and contralateral arm movements mentioned in the previous paragraph would argue against these hypotheses, at least in M1.

The presence of bilateral activity in the motor cortices may also be related to coordination of the two limbs. In an influential study by Donchin et al., 1998, it was shown that M1 and SMA neurons both displayed tuning to unique combinations of bimanual reaches [72]. Similar patterns in caudal premotor cortex have been reported in imaging studies with bimanual finger presses [69]. Activity supporting bimanual coordination may therefore be less constrained to SMA than once thought, and in fact be distributed across the frontal motor regions [72, 125, 204]. In the context of bimanual coordination, ipsilateral activity could provide state information or state-dependent modulation necessary for grooming contralateral drive and promoting cooperation between the limbs. Due to the largely proximal innervation patterns of ipsilateral and bilateral descending motor pathways [20], ipsilateral activity may also reflect postural stabilization relevant to reaches of both arms. This postural account would also predict unique interactions for simultaneous bimanual arm movements, as the stabilizing forces themselves would be unique.

As a final alternative, ipsilateral signals may reflect computations that span the hemispheres without being directly involved in movement production [8, 151]. Rodent work has shown that motor cortical neurons with intracortical projections exhibit reduced contralateral bias as compared to those with descending output [150, 201]. Furthermore, photoinhibition of not one, but both hemispheres during preparation was required to perturb subsequent responses (note that this was a directional licking task, yet it shares many aspects of laterality in arm movements) [151]. Together, these results implicate specific sub-populations of bilaterally distributed neurons in supporting preparatory network activity and dynamics. Such interpretations interface well with the dynamical systems theory presented earlier in this introduction.

1.3.4 Implications for neurorehabilitation and assistive neurotechnologies

Interest in characterizing bilateral motor cortical signals is largely rooted in its importance for the development of neurorehabilitation practices and assistive neurotechnologies for patients with motor impairments. The case of hemispheric stroke may be most relevant, where damage on one side of the brain affects output along corticofugal motor pathways and produces lateralized deficits in movement [105]. Interactions between the perilesional and contralesional cortex appear critical in determining functional recovery after stroke [35, 58, 100, 112, 164, 214]; however, a thorough mechanistic understanding of these interactions remains elusive. The compensatory capabilities of unaffected pathways originating in the contralesional hemisphere have also been studied [20, 51, 52, 58, 186]. Although their role in spontaneous recovery appears limited, more study is necessary to understand when these pathways might be beneficially involved in recovery and why. The ability of clinicians to

implement appropriate therapeutic approaches will therefore depend upon further investigation of the natural and pathological interactions of the two hemispheres.

Together with informing existing therapies, understanding the natural roles (and plastic capabilities) of bilateral signals will benefit the development of assistive technologies like brain-machine interfaces (BMI's). For example, if function is lost from one hemisphere following stroke, might the intact hemisphere be capable of producing two independent signals: one for continued control of the non-paretic limb, and the other to restore function of the paretic limb via a BMI? Biomimetic decoders seeded on ipsilateral movement activity have provided proof-of-principle that a single hemisphere contains signals for both arms that are rich enough for neuroprosthetic control [92]. Indeed, non-invasive recordings from the contralesional cortex in stroke patients (sometimes in tandem with the perilesional cortex) have been successfully implemented in BMI control [18, 26, 28, 29, 166]. These setups are often rehabilitative in nature and intended for temporary use, controlling some aspect of existing therapeutic approaches such as a robotic orthosis [91, 166] or functional electrical stimulation of paretic muscles (FES) [18]. One could also imagine a more traditional application of BMI, whereby chronic augmentation or replacement of function could be pursued in cases of extreme motor deficits [28]. The same concerns about independence and interaction of bilateral signals for these types of BMI would also be relevant for bimanual BMI's in tetraplegic patients. Currently, only one known study has demonstrated the use of a BMI with bimanual control in the published literature [113].

There is great promise in these neurorehabilitative approaches for restoring mobility and independence to large patient populations, yet it is clear that a stronger understanding of the basic neuroscience involved will be crucial for maximizing their potential. It is this human-centered concern that motivates the investigation of bilateral motor cortical signals described in this thesis.

Chapter 2

Hybrid dedicated and distributed coding of bilateral arm signals

2.1 Introduction

In the primate cortex, direct control of arm movement is primarily mediated by contralateral descending projections [20, 142, 202]. However, numerous studies have observed activity changes in the motor cortex during movements of the ipsilateral arm [8, 45, 72, 92, 108, 110, 154, 204] and hand [69, 210, 215]. The functional role of this ipsilateral activity has been the subject of considerable debate, with hypotheses ranging from a role in postural support, bimanual coordination, or an extrapyramidal control signal for unimanual movements.

Neurons in the primate dorsal premotor cortex (PMd) play a critical role in motor preparation [45, 110, 199, 210, 219]. Interestingly, their response properties and degree of laterality appear to change across the course of preparation. For example, within PMd, individual units exhibit a transition from effector-independent to effector-dependent encoding between preparatory and execution phases of reaching. In contrast, units in primary motor cortex (M1) mainly become active during movement itself and show a pronounced contralateral bias [45]. This suggests a transition from abstract planning to explicit specification of motor output parameters in the signals of individual neurons. A similar transition has been shown in the activation of different cell-types within rodent motor areas [150, 201]. These studies have found that neurons with intracortical projections show little lateral bias, particularly during pre-movement phases. In contrast, neurons with descending output display much stronger laterality, especially just before and during movement. This adds yet another level of granularity in the discussion of lateralized motor function. Collectively, these single-unit studies support a notion that there exist two distinct components within the motor cortex: one

that is bilateral and likely involved in abstract processing, and another that is dedicated to a single side of the body for execution.

The classical perspectives outlined above have been revisited in studies that focus on population-level analysis, considering instead how computations might be reflected in the way the network coordinates activity. Low-dimensional representations of large-scale neural recordings can be used to characterize these network patterns, revealing changes in covariance structure across behavioral settings that are not evident when looking at single neurons in isolation [55]. Ostensibly, these changes reflect reorganization of the population as it engages in different computational processes. Using these methods, pre-movement activity has been shown to evolve within an “output-null” subspace towards an optimal initial population state [43, 80, 122]. This initial state is advantageously positioned for engaging the internal dynamics of the network to produce patterned activity in an “output-potent” subspace for driving movement [44, 200, 207]. There is some evidence that bilateral activity may support these preparatory and dynamic properties [151]. Similar to the output-null and output-potent subspaces, arm-specific subspaces have also been observed in M1 during rhythmic movements [8] and in response to joint perturbations [108]. It remains unclear precisely what organizational principles produce these arm-specific subspaces, whether signals are fully separated at the level of the population, and how such properties develop across preparation and movement.

There are two fundamental and mutually non-exclusive ways that population signals may specify the selected arm across preparation and movement. (1) Signals may consolidate within dedicated sub-populations for each arm (i.e., within hemispheres, brain areas, or cell-types). (2) Signals for each arm may be distributed across the same units yet maintain unique covariance structure that separates them along arm-specific neural dimensions. Importantly, either of these architectures provides a way for downstream targets to discriminate signals and also yields the mathematical result of divergent subspaces. The first method is necessarily true: contralateral biases have been consistently observed during movement and, to a lesser extent, during preparation as well. Such lateral biases will trivially orthogonalize arm signals. The question is whether the second method is also true. Either signals that are mixed within units become separated (arm-specific) in the population readout, or they exist within a space where the same patterns of activity are involved in computations relevant to both arms. This is a vital distinction to make, as it constrains the possible roles that bilateral activity can play at each stage of processing and may point to an important heterogeneity in the population statistics.

In the present study, we recorded large populations of single-units in PMd and M1 bilaterally while monkeys performed an instructed-delay unimanual reaching task. As activity emerged during preparation, there was a tendency for units with stronger

arm preference to be more highly modulated, therefore representing a larger proportion of the population variance. As a result, the signals for each arm were largely segregated, primarily within contralateral PMd. During the transition to movement, M1 became more prominently involved and the signals for each arm became increasingly segregated. This unit-level segregation caused the subspaces corresponding to each arm to diverge across the trial. However, we did observe target-specific information that was mixed within single-units, indicating incomplete separation of arm signals. Importantly, the subspace containing this information did not segregate signals for the two arms at any stage. Taken together, the results point to two primary components in the population response: (1) A dedicated component that develops across preparation, reaches a maximum during movement, and mirrors the lateralized anatomy of corticospinal output with its contralateral bias. (2) A distributed component that represents far less variance, particularly during movement, and provides a space in which bilateral arm signals may easily interact.

2.2 Results

2.2.1 Behavior

Two macaque monkeys were trained to perform an instructed-delay reaching task in 3-D space (Figure 2.1A). Reaching movements were freely performed in an open area while kinematics were recorded using optical motion tracking. Visual feedback of endpoint position and task cues were provided through a virtual 3-D display. Each trial had three phases (Figure 2.1B). For the Rest phase, the monkey placed both hands in start targets positioned near the torso and remained still for 500 ms. For the Instruct phase, an instructional cue appeared at one of six target locations. The color of the cue specified the required hand for the forthcoming trial. The monkey was required to keep both hands in the rest positions while the cue remained visible for a variable interval (500 – 1500 ms). The Move phase was initiated when the start position marker for the reaching hand disappeared and the cue at the target location increased in size, which signaled the animal to reach. The monkey received a juice reward if it accurately reached the target and maintained the final position for 250 ms, while keeping the non-cued hand at its start position for the duration of the trial. 300ms representative windows from each phase were used in data analysis. Trials were blocked for each arm, with each block consisting of 2 trials per target in a randomized order (i.e. alternating 2 trials per target for the left arm, then 2 trials per target for the right; Figure 2.1C).

Average success rates were above 95% for both hands in both monkeys. Overall, reaction times averaged 308 ms for monkey O and 333 ms for monkey W. Distributions of reaction times for each hand/target combination are displayed in Figure 2.1D, which were fairly consistent across targets. Reach biomechanics varied across the workspace,

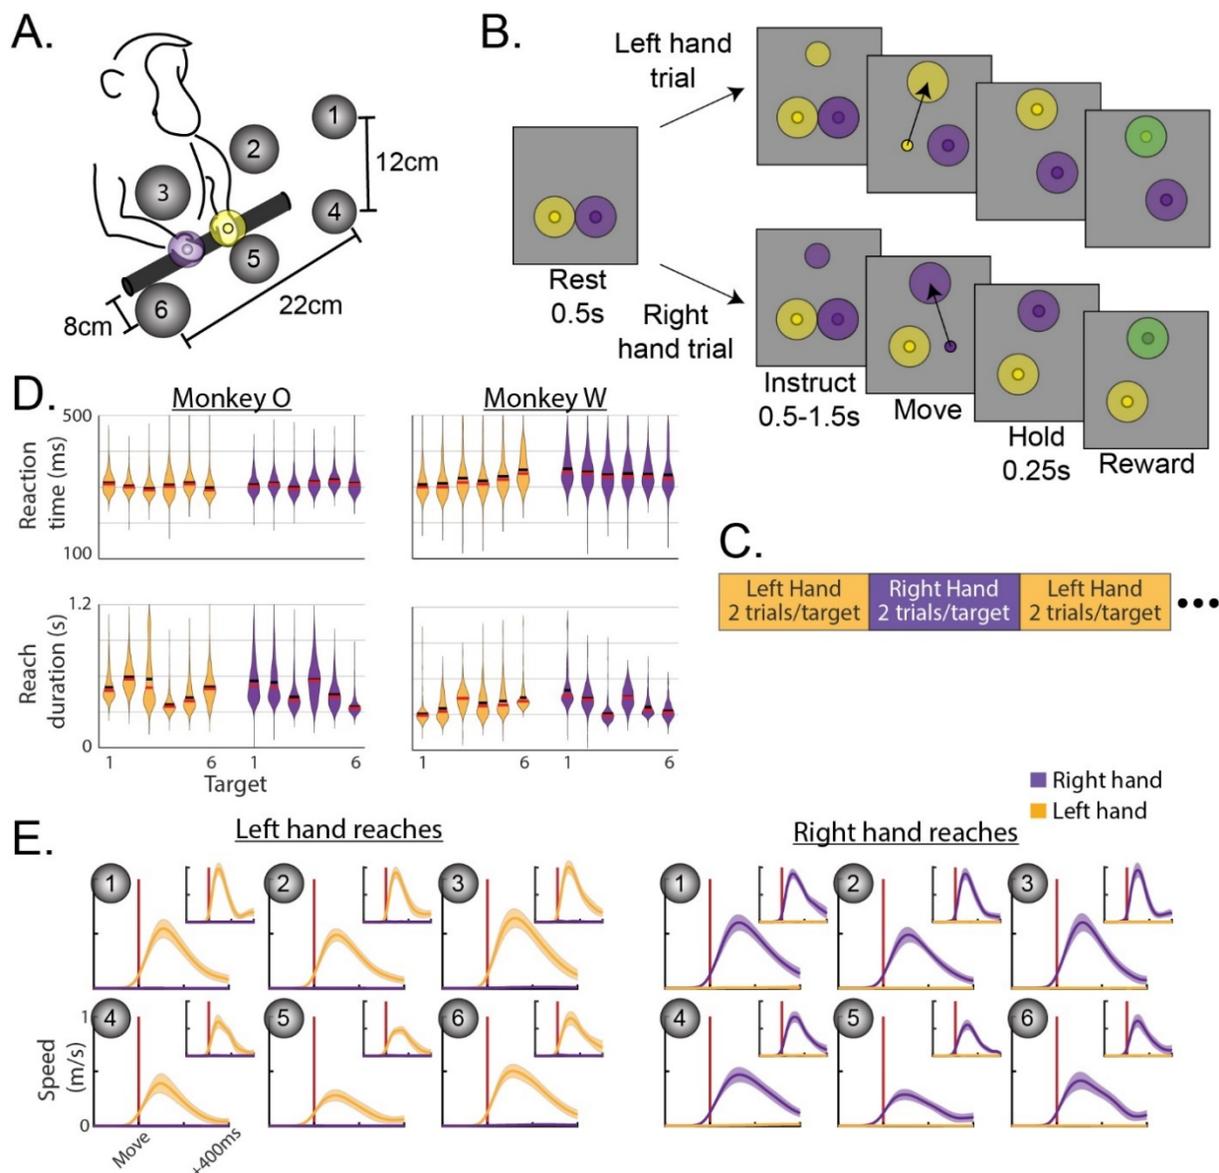


Figure 2.1. Behavior. (A) Monkeys reached to one of six virtual targets, indicated by grey spheres in the cartoon. During the task these would be invisible until one appeared to instruct the reach. (B) Trials consisted of 3 phases. Each trial was initiated by placing both hands in start targets and remaining still for 500ms (Rest phase). A small target then appeared at the location of the future reach in a color that indicated which hand to use. The monkey remained still during cue presentation for 0.5-1.5s (Instruct phase). The start target for the reaching hand then disappeared while the reach target enlarged to cue movement (Move phase). (C) Hand assignments followed a blocked schedule. (D) Distributions of reaction times (top row) and reach durations (bottom row) for each monkey, hand, and target. Left hand reaches in yellow, right in purple. Horizontal black bars show means, red bars show medians. (E) Speed profiles during left- or right-hand reaches. Both reaching and stationary hands are plotted in each, although stationary speeds are near 0 and hardly visible. Vertical red lines indicate threshold crossing to mark movement onset. Monkey O main, monkey W inset. Mean \pm standard deviation.

resulting in slightly different reach durations across targets (Figure 2.1D). In terms of kinematics, the initial feed-forward portions of reaches were smooth and stereotyped (Figure 2.1E). There was a very slight but significant increase in the speed of the non-reaching hand between Rest (mean – monkey O: 1.1 mm/s; monkey W: 2.9 mm/s) and Move (mean – monkey O: 3.6 mm/s; monkey W: 7.6 mm/s) phases of the task (permutation test – monkey O: $p=1.0e-4$; monkey W: $p=1.0e-4$). We note that the task was designed to mimic natural reaching without the use of physical restraints. As such, we assume the small movements in the non-reaching arm are part of the normal behavioral repertoire occurring during natural unimanual reaching. Nonetheless, we will address any reasonable impacts these small movements may have in our neural analyses.

2.2.2 Arm-dedicated units emerge across task phases while the overall distribution remains relatively arm-neutral

We recorded 433 and 113 single-units in the caudal aspect of dorsal premotor cortex (PMd) in monkeys O and W, respectively, and 331 and 289 single-units in primary motor cortex (M1) (Figure 2.2). Since both arms were used in the behavior, we can evaluate the ipsi- and contralateral responses in each unit. Units were pooled across hemispheres in the analysis, with contralateral summaries reflecting the collection of responses during trials performed with the contralateral arm, and vice-versa for trials performed with the ipsilateral arm. PMd and M1 units were analyzed separately. Firing rates were soft-normalized using the Rest phase mean and standard deviation, and modulation strength is expressed as

the mean squared value of these standard scores within the window of interest. This modulation metric is essentially variance, and may be thought of as variance for most purposes.

We first analyzed single units to determine the degree of modulation during the Instruct and Move phases of the task (Figure 2.3). Following instruction, many units in both PMd and M1 became significantly modulated for movements of one or both arms (Table S2.1). Units in PMd were, on average, more strongly modulated during the Instruct period than those in M1 (Figure 2.4A; permutation test – monkey O: $p=0.012$; monkey W: $p=3.2e-3$). This relationship reversed following movement, with average modulation in M1 becoming stronger than PMd (Figure 2.4A; permutation test – monkey O: $p=2.6e-3$; monkey W: $p=0.012$). These results are in line with the view that PMd plays a privileged role in motor preparation. The distributions of modulation values were heavy-tailed and contained some notably extreme values; however, we chose not to apply any outlier criteria. Controls are performed later in our population-level analyses to ensure that results are representative of trends across the entire population rather than a few extreme units.

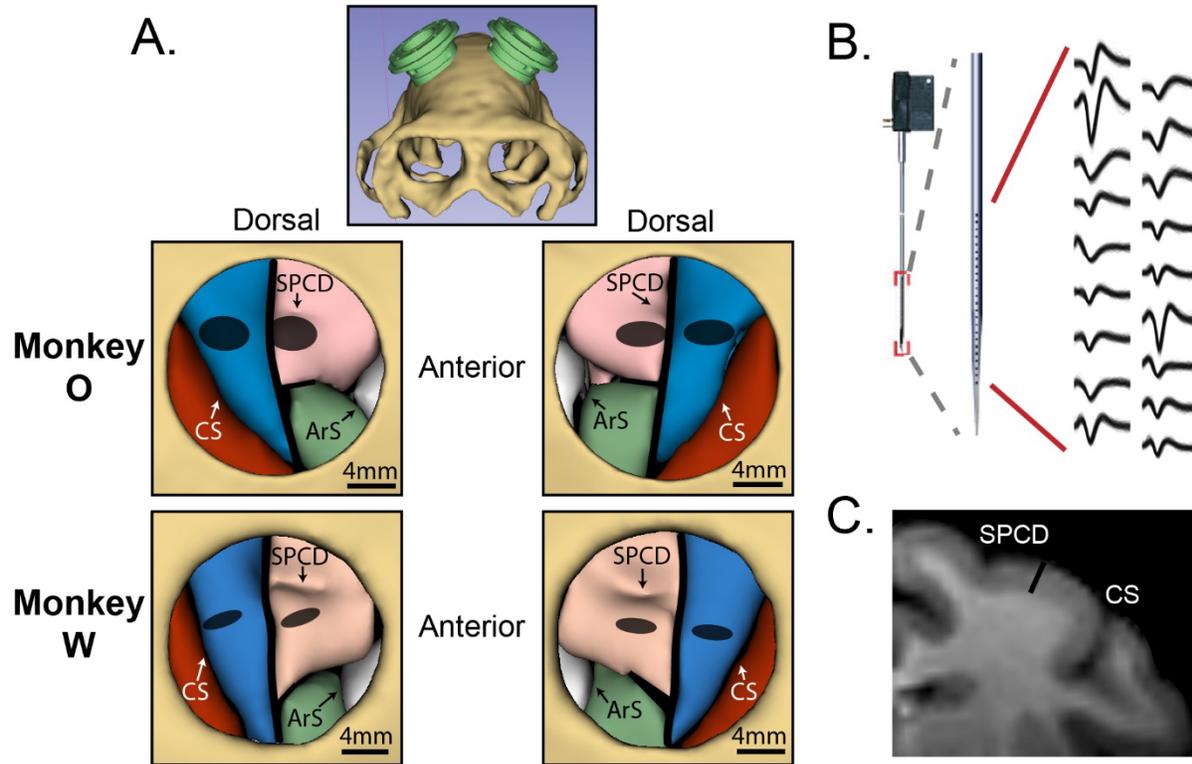


Figure 2.2. Neural recordings. (A) MRI-based renderings of the skull and target brain regions. Top panel shows the arrangement of the two chambers. Two bottom rows show segmented brain regions within the cranial window of each chamber, for each monkey. Region boundaries were assigned using Paxinos et al., 2000 [169]. Red - somatosensory cortex; blue - primary motor cortex (M1); pink - dorsal premotor cortex (PMd); green - ventral premotor cortex; white - frontal eye field. CS - central sulcus; SPCD - superior pre-central dimple; ArS - arcuate sulcus. Grey ellipses indicate regions sampled by recordings. (B) Interlaminar recordings were obtained using V- and S- probes (Plexon, Inc., Dallas, TX) with 24-32 electrodes aligned perpendicular to the cortical surface. Example waveforms were all simultaneously recorded from a single probe. (C) MRI coronal slice, monkey O. 3mm black bar is approximately equal to the distance spanned by electrodes on 32-channel probes. Same landmark labels as in (A).

We next considered the laterality of each unit by quantifying the relative modulation observed during ipsi- and contralateral trials. We expressed each unit's arm preference on a scale from -1 to 1, with 1 indicating exclusive contralateral modulation and -1 indicating exclusive ipsilateral modulation (Figure 2.4B). Although the cue for the forthcoming trial had yet to be presented during the Rest phase, arm selection could be implied from the blocked task structure (Figure 2.1C). However, except for a very small effect in PMd of monkey O (one-sample t-test – $\mu_{\text{Rest}}=0.06$, $p=9.7e-5$), there was no significant contralateral bias observed during the Rest phase in either brain area for both monkeys. Despite the lack of contralateral bias, both monkeys entered arm-specific population states during the Rest phase, which was more pronounced in PMd populations (mean difference between left and right arm firing rates – monkey O PMd: 1.85Hz, M1: 1.64Hz; monkey W PMd: 1.33Hz, M1: 0.98Hz; Figure 2.4C). For trials in which the same hand was repeated from the previous trial only, it was possible to classify the hand for the forthcoming movement from the population activity (Figure S2.1).

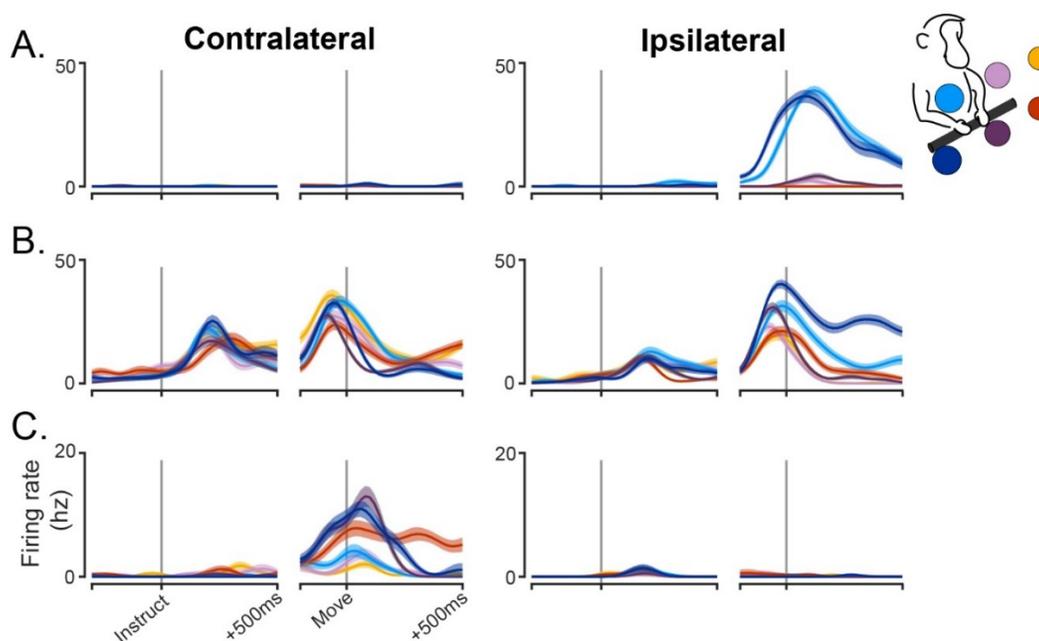


Figure 2.3. Firing rate traces of example single-units. Trial-averaged firing rates for 3 example single-units, all from the left hemisphere. Each color represents a different target according to the color-coding in the top right. Mean \pm SEM. (A) An M1 unit exclusively modulated during ipsilateral movements. (B) A PMd unit with both Instruct and Move phase modulation for both arms. (C) A PMd unit with modest contralateral modulation during the Instruct phase and strong contralateral modulation during movement, but no modulation on ipsilateral trials.

The emergence of laterality after the onset of the instruction cue mirrored the emergence of general unit modulation: A contralateral bias was present in PMd during the Instruct phase and then became present in both PMd and M1 during movement. Mean arm preference in PMd showed a modest but significant bias in the contralateral direction during the Instruct phase (one-sample t-test – monkey O: $\mu_{\text{Instruct}}=0.11$, $p=7.0e-8$; monkey W: $\mu_{\text{Instruct}}=0.16$, $p=1.8e-4$) and showed no significant change between Instruct and Move (paired-sample t-test – monkey O: $\mu_{\text{Move}}=0.15$, $p=0.11$; monkey W: $\mu_{\text{Move}}=0.13$, $p=0.65$). Mean arm preference in M1 did not show a significant contralateral bias until the Move phase (one-sample t-test – monkey O: $\mu_{\text{Instruct}}=0.03$, $p=0.13$; $\mu_{\text{Move}}=0.07$, $p=0.013$; monkey W: $\mu_{\text{Instruct}}=0.02$, $p=0.31$; $\mu_{\text{Move}}=0.20$, $p=5.1e-11$).

While shifts in the means were modest, changes in arm preference across phases were most evident in the tails of the distribution, corresponding to units that strongly preferred one arm or the other (Figure 2.4C). These arm-dedicated units typically preferred the contralateral arm, demonstrated by increased occupancy in the contralateral tails of the arm preference distributions; however, a small proportion of the population was exclusively modulated during ipsilateral trials as well (Figure 2.4B). Despite much of the population remaining arm-neutral (arm preference near 0) or preferring the ipsilateral arm, the emergence of strongly contra-dedicated units was sufficient to drive contralateral shifts in the population mean. In summary, despite much of the population remaining arm-neutral, an increasing number of highly arm-dedicated units emerged with each task phase, primarily favoring the contralateral arm.

2.2.3 Modulation preferentially occurs within arm-dedicated units

There are two primary means by which population signals can specify the selected arm at each phase. (1) The population may maintain unique covariance structure for each arm that separates signals along different neural dimensions, even if the constituent units are equally modulated for both arms. (2) Arm-dedicated units may dominate the population response, thereby representing the majority of population variance in dedicated sub-populations. The latter possibility is investigated over the following two sections. First, we consider whether modulation preferentially occurs in units that are strongly dedicated to one arm or the other.

We performed a regression analysis to quantify the relationship between strength of arm preference and modulation for the preferred arm. Importantly, arm preference and modulation were calculated from independent datasets to prevent artificial linkage between the two measures due to sampling noise. A slope of 1 corresponds to an order of magnitude increase in modulation, on average, when comparing fully arm-neutral units with fully arm-dedicated units. As seen in Figure 2.5A, the slopes are initially near zero

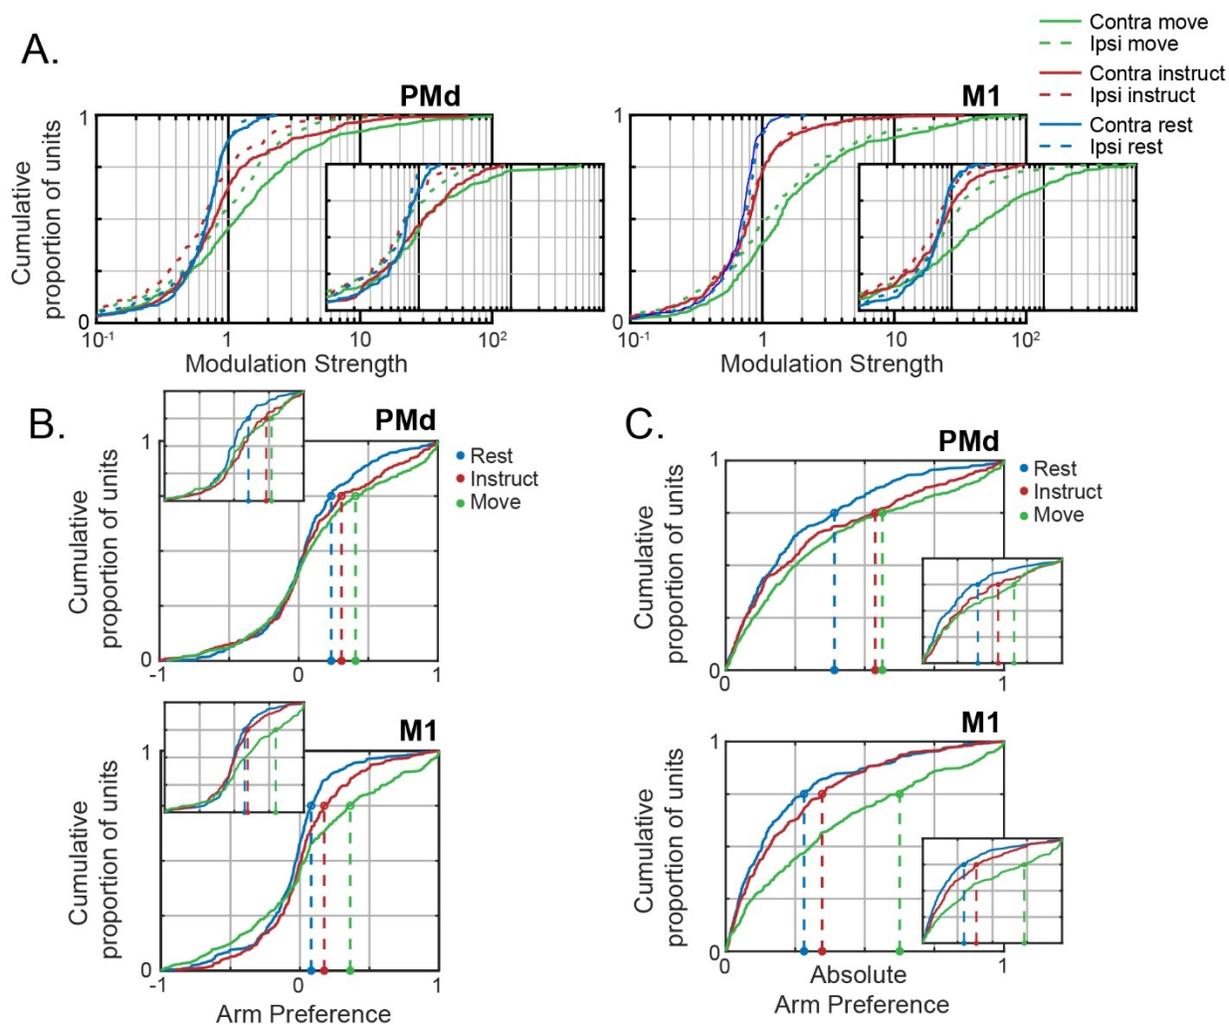


Figure 2.4. An increasing number of arm-dedicated units emerge with each task phase. (A) Cumulative distribution of single-unit modulation during each phase, arm. Left panel PMd, right panel M1. Large values cut off by plot: monkey O Contra Move [134(PMd), 133(PMd), 104(PMd)], Ipsi Move [234(M1), 181(M1), 130(M1)]; monkey W Contra Move [125(M1)]. (B) Cumulative distribution of arm preferences during each phase. Top panel PMd, bottom panel M1. Negative values are ipsi-preferring, positive values are contra-preferring. Circles and vertical dashed lines mark the upper quartile of each distribution (C) Same as (B), but using the absolute value of arm preference to indicate arm dedication, independent of hemisphere. For all plots: monkey O main, monkey W inset.

and then become positive over time. To quantify these changes, we used a multi-factorial permutation approach to test for effects of Area (PMd, M1), Phase (Rest, Instruct, Move), and Preferred Arm (Ipsi, Contra) on the population slopes.

We found a main effect of Phase in both animals (monkey O: $p=1.0e-4$, monkey W: $p=1.0e-4$): a positive correlation between arm preference and modulation strength emerged and strengthened across task phases (Figure 2.5A-B). By the Move phase, there was approximately a ten-fold increase in the modulation strength of units with an absolute arm preference of 1 (completed dedicated) relative to units with an arm preference near 0 (balanced modulation). Since PMd displayed greater modulation than M1 during preparation but not movement, we tested whether the two areas had differing slopes in each phase independently. We found a significant simple effect of Area during the Instruct phase (monkey O: $p=3.0e-4$; monkey W: $p=6.3e-3$) but not the Move phase (monkey O: $p=0.13$; monkey W: $p=0.91$). Thus, the relationship was more prominent within PMd prior to movement, while the two areas became roughly equivalent following movement initiation. This was confirmed with a test for 2x2 interaction (monkey O: $p=0.025$; monkey W: $p=9.9e-3$). Additionally, we analyzed the relationship between arm preference and modulation for the non-preferred arm to confirm that increased arm preference is associated selectively with increased modulation for the preferred arm (Figure S2.2). No significant positive relationships were observed in either monkey, either brain area, or any task phase; therefore, greater arm preference is associated with selective increases in modulation for a single arm.

Given the overall contralateral bias, we further tested whether this relationship held for both contra- and ipsi-preferring units. For the contra-preferring units, there was a significant simple effect of Phase (monkey O: $p=1.0e-4$; monkey W: $p=1.0e-4$). For the ipsi-preferring units, the Phase effect was significant for monkey O ($p=1.0e-4$), but only trended in this direction for monkey W ($p=0.087$), perhaps due to the lesser amount of ipsilateral modulation in monkey W. Slopes were generally steeper for contra-preferring units. The simple effect of Preferred Arm was significant during the Instruct phase for both monkeys (monkey O: $p=0.033$; monkey W: $p=1.0e-4$), and significant for Monkey W during the Move phase (monkey O: $p=0.53$, monkey W: $p=1.0e-4$). Given that there are also more contra-dedicated units than ipsi-dedicated units, these results suggest that a larger proportion of the contralateral signal exists within dedicated sub-populations compared to the ipsilateral signal. We directly test this conjecture in the following section where we consider population-level implications of these results.

2.2.4 The population signal is largely confined to arm-specific sub-populations

The preceding analyses establish that there is an increase across task phases in the proportion of units that are strongly dedicated to a single arm, and that those units

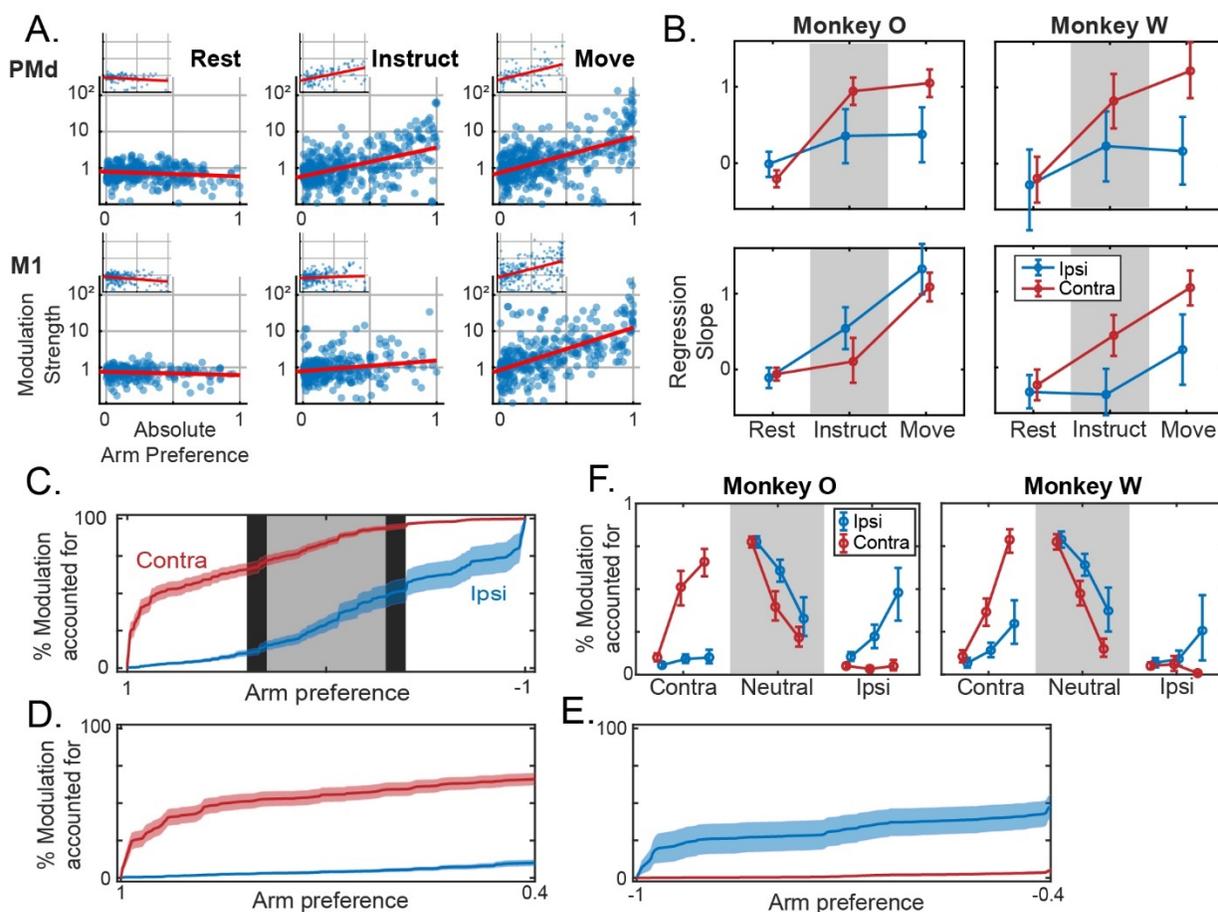


Figure 2.5. Neural activity is progressively consolidated within arm-specific subpopulations. (A) Modulation for the preferred arm plotted against arm preference, for all units in each brain area and task phase. Log-linear best fit lines are displayed in red. Inset figures belong to Monkey W. (B) Slopes of regression lines fit to data from (A), independently for ipsi- and contra-prefering sub-populations. Mean \pm bootstrapped 95% confidence interval. (C-E) For the Move phase in monkey O, cumulative modulation plotted against arm preference, i.e. each point indicates the proportion of modulation accounted for by all units with arm preference values to the left of the indexed position. Positive values on the x-axis indicate contra-prefering, and negative values indicate ipsi-prefering. Shaded error bars indicate bootstrapped standard error. (C) The full spectrum of arm preferences is shown. Shaded backgrounds indicate three partitions: Contra-dedicated [0.4, 1] and Ipsi-dedicated [-1, -0.4] in white, and Neutral [-0.3, 0.3] in grey. (D) Cumulative modulation within contra-dedicated regime. (E) Same as (D), but ipsi-dedicated. Note inverted axis. (F) The proportion of modulation within each partition from (C) during ipsi- or contralateral movements. Note that the total modulation is significantly lower for ipsilateral movements, particularly for Monkey W, and these data are only displayed as proportions. Mean \pm bootstrapped 95% confidence interval.

exhibit greater modulation in activity relative to those that are more neutral. This suggests that the population signal is progressively segregating at the level of individual units. To visualize this segregation, we ordered units based on arm preference and calculated the cumulative modulation at each value, i.e. the proportion of modulation across the entire population that is accounted for by units with arm preferences at or below a certain value (Figure 2.5C-E). Since PMd and M1 showed similar relationships in the previous analyses, we combined units from the two areas, analyzing them as a collective population. In the extreme case that population signals are entirely segregated, 100% of ipsilateral modulation would occur at an arm preference of -1, and 100% of contralateral modulation would occur at +1.

We focused on two core questions. (1) Does the proportion of dedicated modulation increase across task phases, indicating a progression towards independent signals? (2) Does the amount of independent (or dedicated) modulation differ for ipsi- and contralateral activation? As expected, dedicated regimes of the arm preference distribution captured a large proportion of the modulation associated with movements of one arm and only a small proportion of the modulation associated with the other arm, primarily during execution (Figure 2.5C-F). For statistical testing, we split the arm preference domain into 3 equal width regimes, corresponding to contra-dedicated (arm preference > 0.4), ipsi-dedicated (arm preference < -0.4), and arm-neutral ($-0.3 < \text{arm preference} < 0.3$) units, and summarized the data by expressing the proportion of modulation contained within each regime (Figure 2.5F). We again used a multi-factorial permutation approach to test for effects of Phase (Rest, Instruct, Move), and Arm (Ipsi, Contra). We will refer to ipsilateral modulation in the ipsi-dedicated units simply as “ipsi-dedicated modulation” and vice-versa for contra-. We emphasize that these partitions were chosen to broadly isolate the extremes of the distribution; it is not intended that the precise boundaries map onto discrete cell-types or any similar interpretations.

For both animals, the effect of Phase was significant in the contralateral responses (monkey O: $p=1.0e-4$; monkey W: $p=1.0e-4$), with the proportion of contra-dedicated modulation increasing across phases (Figure 2.5F, red lines). Ipsi-dedicated modulation increased across task phases for both monkeys as well (Figure 2.5F, blue lines), although this effect was only significant for monkey O ($p=9.0e-4$; monkey W: $p=0.31$). There was a significant interaction between Arm and Phase for both monkeys (monkey O: $p=1.0e-4$; monkey W: $p=1.0e-4$), indicating the stronger emergence of contra-dedicated modulation as compared to ipsi-dedicated modulation. Both animals showed a simple effect of Hand during the Instruct phase (monkey O: $p=1.0e-4$; monkey W: $p=1.0e-4$), with more contra-dedicated modulation being observed than ipsi-. This effect was also significant during the Move phase for monkey W ($p=1.0e-4$) and approached significance for monkey O ($p=0.056$).

These results suggest that arm signals separate at the level of individual units throughout preparation. Moreover, contralateral signals are more independent than ipsilateral signals, in the sense that a larger proportion of the contralateral modulation was represented in dedicated regimes of the population. Since this characterization of the population response captures most of the modulation for each arm in mutually exclusive sub-populations, we will refer to it as the “dedicated” component.

It is possible, however, that the limited range of movement directions used in the task may influence the degree of dedicated modulation (the maximum angle between target vectors is 107°). For example, a unit that appears dedicated to one arm may only be unmodulated for the other arm over the range of movement directions being tested. That unit may in fact be modulated during different movements, which would cause it to appear neutral if sampled. As a post-hoc control for this possibility, we included data during the return movements following target acquisition, effectively doubling the range of sampled movement directions, and repeated the analyses of Figure 2.5. The relationship between arm preference and modulation strength remained for this Reach & Return data (Figure S2.3B-C). This relationship was significant for both monkeys, in both PMd and M1, and for both ipsi- and contra-preferring units, with one exception (permutation test of regression slopes – $p < 0.05$; monkey W, PMd, Ipsi-preferring units $p = 0.58$). The proportion of dedicated modulation was largely unchanged as well (Figure S2.3D), and replacing Move phase data with Reach & Return did not impact significance of the Phase effect (Contra-dedicated modulation – monkey O: $p = 1e-4$; monkey W: $p = 1e-4$; Ipsi-dedicated modulation – monkey O: $p = 1.8e-3$; monkey W: $p = 0.22$). Therefore, the dedicated signals that we observe persist even with a broad range of movement directions. Returning to the possibilities outlined at the beginning of the previous section, we therefore conclude that this dedicated component progressively becomes the dominant characterization of the population response – dominant in the sense that it represents the majority of modulation across the population.

2.2.5 Neural subspaces for the two arms diverge across task phases

We next sought to characterize the time course of changes in neural subspaces as movements were prepared and executed. We hypothesized that dedicated activation would drive population signals into diverging subspaces for the two arms. For these analyses, we pooled units from the left and right hemispheres. Using PCA on single-trial data, we first estimated the dimensionality of the neural subspace during each task phase using a cross-validated data reconstruction method (see Methods; Yu et al., 2009). This is an essential step to avoid drawing conclusions based on noise-dominated dimensions. Dimensionality was calculated separately for each session and arm. During Rest, the dimensionality was approximately 5-6, and decreased to approximately 4 during the Instruct and Move phases (Figure 2.6A). The high dimensionality during

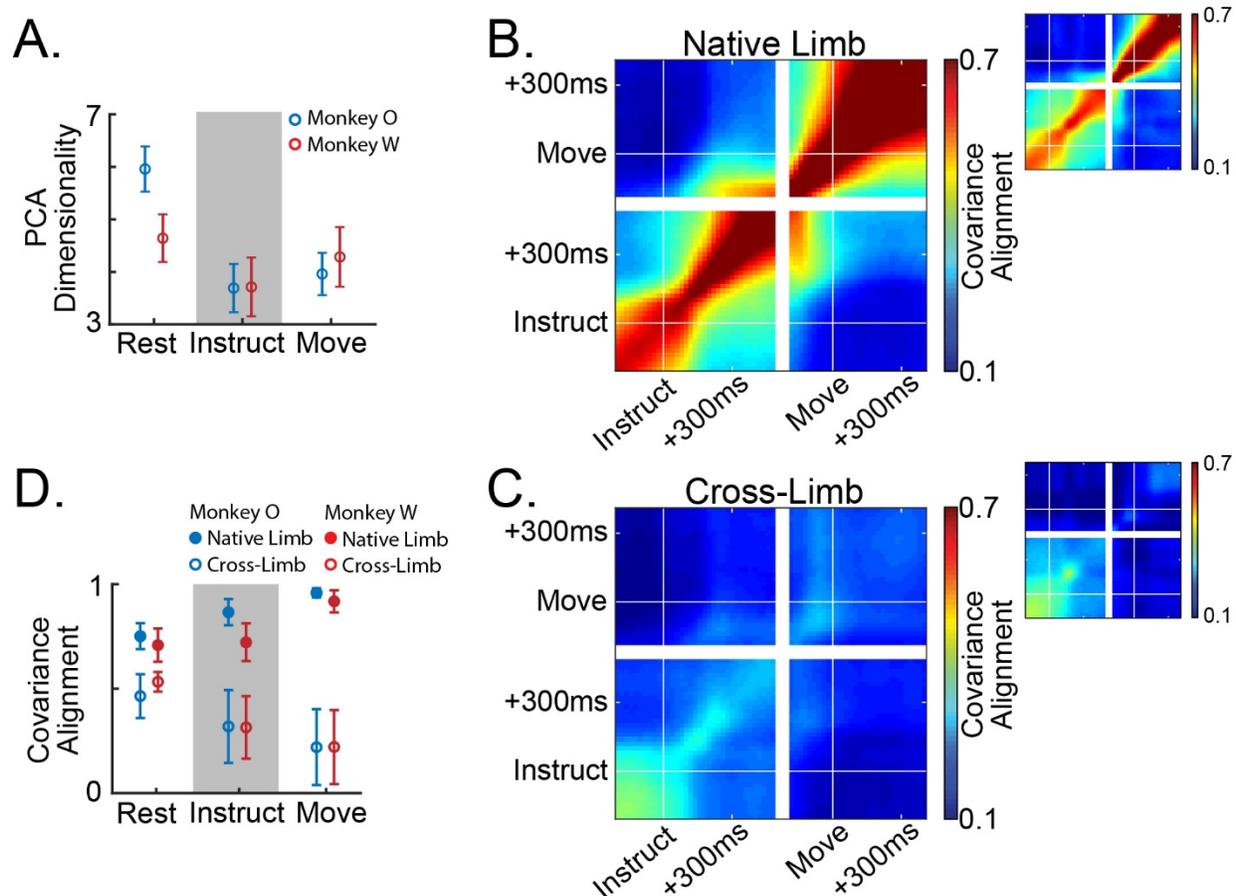


Figure 2.6. Population activity reorganizes and diverges for the two limbs throughout planning. (A) Dimensionality of the PCA subspace estimated as the number of components that minimizes the cross-validated reconstruction error of the full-dimensional neural data. Mean \pm standard error across datasets. (B,C) Heat maps indicate alignment of 4-dimensional PCA subspaces between all pairs of timepoints across the Instruct and Move phases of the task, averaged across sessions. (B) Compares subspaces across time for movements of the same arm. Three blocks forming along the diagonal indicate three distinct subspaces: a pre-instruction Rest space, a post-instruction Instruct space, and a peri-movement Move space. (C) Compares subspaces across time for movements of opposite arms. Prior to instruction there is a moderate alignment of the subspaces for each limb, however, the two subspaces diverge around 100ms post instruction. (D) Summary of the data in (B,C). Mean \pm standard deviation across datasets.

Rest aligns with recent reports [57], and values during Instruct and Move were comparable to those found in previous studies using similar methods [225]. We therefore chose to focus on only four components to represent the neural subspaces of each dataset.

We calculated the alignment between PCA subspaces associated with left or right arm movements using a metric that describes the proportion of low-dimensional variance for one dataset that is captured in the low-dimensional space of another (see Methods; [10]). If the network is organizing activity in the same way across datasets, then the covariance alignment is 1, regardless of signal magnitude. If activity is reorganized into orthogonal subspaces across datasets, then the covariance alignment is 0. Two types of alignment measurements were made: (1) Subspaces were fit to random partitions of trials for the same arm – what we will refer to as “native” alignment – giving us an estimate of natural variability in our subspace estimates when compared over the same time window, and describing the evolution of the motor plan when comparing across time windows. (2) Subspaces were fit separately using trials for either arm and compared with each other – what we will refer to as “cross” alignment – describing the divergence of the subspaces for the two arms at each task phase.

Using single-trial activity event-locked to the onset of instruction and movement, we were able to capture the fine-timescale evolution of any emerging or diverging subspaces (Figure 2.6B-C). When comparing the native alignment across task phases, we observed the emergence of distinct Instruct and Move period subspaces. Figure 2.6B shows these data displayed as a continuous heat map with block diagonal structure that coincides with the phase transitions. Within each phase native alignment was high, indicating consistent low-dimensional structure in the population activity that was specific to each stage (Figure 2.6B; Figure 2.6D filled circles).

As expected, subspaces for the two arms gradually diverged across task phases (Figure 2.6C; Figure 2.6D open circles). On the whole, subspaces for the two arms were significantly less aligned than the (cross-validated) comparisons within the same arm (Figure 2.6D open vs filled circles; two-way ANOVA, ME comparison type – monkey O: $p=3.3e-61$; monkey W: $p=5.2e-29$). Interestingly, subspace divergence was already apparent during the Rest phase (paired sample t-test, native-Rest vs cross-Rest – monkey O: $p=1.4e-11$; monkey W: $p=4.7e-7$). As mentioned in our analysis of single-unit arm preferences, this is likely due to predictable arm assignments from the blocked task structure (Figure 2.1C, Figure S2.1). Cross alignment decreased significantly as the trial unfolded, reaching a minimum during movement (one-way repeated measures ANOVA – monkey O: $p=8.0e-7$; monkey W: $p=2.4e-6$). These results map closely onto the progressive segregation of dedicated signals described in the previous section.

2.2.6 Subspace separation relies upon dedicated signals

Activity within mutually exclusive sub-populations naturally separates into distinct linear subspaces; as such, we can expect some level of subspace separation as a simple result of dedicated variance. However, it is possible that subspace separation could occur within a distributed representation as well [8, 108]. This question is especially important in considering units that show relatively balanced modulation for the two arms. Even though these units show similar levels of activity during contra- and ipsilateral movement, it is possible that their population-level contributions are different for each, and thus also contribute to subspace separation.

To investigate the extent to which subspace separation relied upon dedicated activation, we analyzed the structure of PCA subspaces via their coefficient weights. Since components of PCA models form an orthogonal basis set, each can be independently analyzed to determine its contribution to subspace divergence. We fit separate PCA models for each arm and task phase and calculated two statistics for each component: (1) To capture the contribution of a given component to subspace separation, we calculated the ratio of variance it captured for the two arms (right/left). (2) To capture the dependence of a given component on arm-dedicated units, we calculated a coefficient-weighted average of the arm preferences for all units (e.g., if non-zero weights were only given to right arm dedicated units, this value would be 1; if weights were evenly distributed across the spectrum of arm-preferences, this value would be 0). A strong relationship between these two metrics would suggest that subspace separation relies upon dedicated activation.

Indeed, this was the case during both the Instruct and Move phases. Figure 2.7A-C shows a single session example from the Move phase. The top principle components captured a large amount of the variance for the left arm while capturing little variance for the right arm. Components with a variance ratio strongly favoring the left arm almost exclusively weighted units that were themselves highly dedicated to the left arm. The lower components with more balanced variance ratios distributed weights more evenly across the arm preference spectrum. This pattern was evident in each phase throughout recordings from both monkeys. Figure 2.7D shows the relationship between right/left variance ratio and coefficient-weighted arm preference for the top five principal components of each dataset. Following the instruction cue, components that strongly discriminated between the two limbs (variance ratio far from 1) primarily weighted units that were themselves highly discriminating. This relationship remained strong as the range expanded during the Move phase. The analysis was also repeated using an alternative normalization method to mitigate the effect of highly modulated units. As expected, mitigating the effect of highly modulated units decreased the magnitude of subspace separation while maintaining the relationship between coefficient-weighted arm preference and variance ratio over the reduced range (Figure S2.4A). This

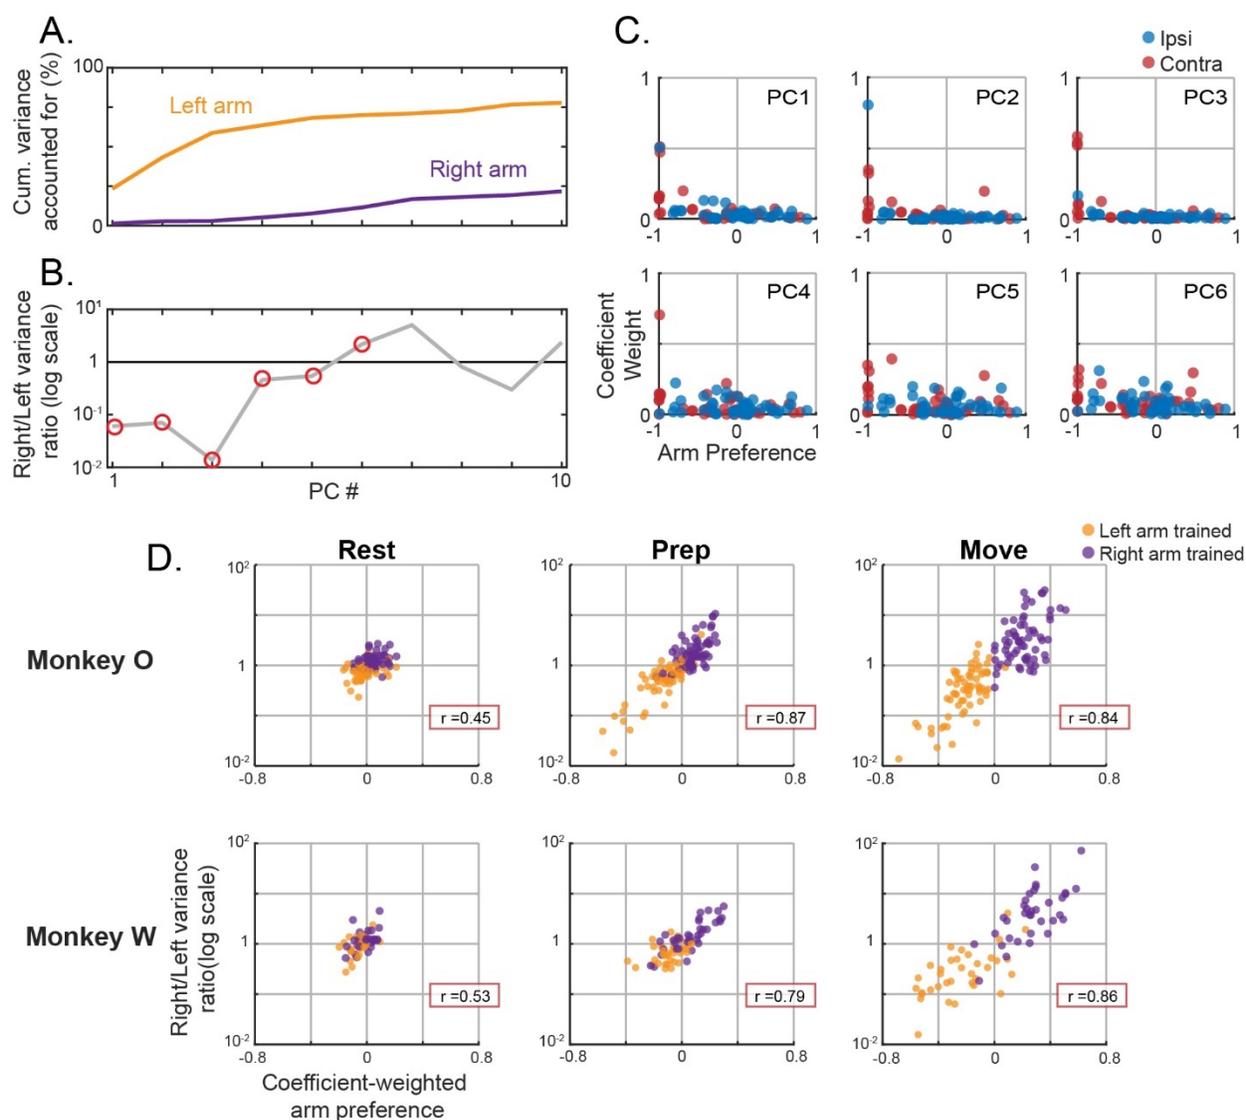


Figure 2.7. Separation of arm-specific subspaces relies upon unit-level segregation.

(A-C) Single session example of a PCA model trained to capture bi-hemispheric activity during left arm movements. Held-out testing data for 86 simultaneously recorded units were used. (A) Cumulative proportion of variance accounted for across the top 10 principal components. (B) For each component, the ratio of the explained variance between the two limbs. (C) Absolute values of the coefficient weights for each component plotted against the corresponding unit's arm preference. Top row represents components 1-3; bottom row represents components 4-6. Positive arm preference values indicate right arm preferring units. (D) The component variance ratio for the two arms plotted against a coefficient-weighted average of the arm preferences for each unit in that component. Datapoints represent the top 5 principal components of left or right arm trained models across all sessions. Separate models for each phase are plotted in each column. Pearson correlation coefficient for each dataset is displayed in the red box. Top row monkey O, bottom row monkey W.

further illustrates the dependence upon highly arm-dedicated, highly modulated units. In summary, these results suggest that the subspace separation described in the previous section relies upon signals that segregate at the level of individual units.

2.2.7 An additional distributed signal contains target-specific information about both arms

The preceding sections make clear that the population signal is dominated by a segregated organization. Nonetheless, it is likely that variance associated with the non-preferred arm of each unit also reflects a meaningful population component, albeit one that is much weaker in magnitude. Indeed, many of the units that we recorded in both PMd and M1 were significantly modulated for both arms throughout preparation and movement (Table S2.1). To assess the information content and strength of these secondary responses, we divided the entire population of units from both hemispheres and brain areas into two subgroups based on the preferred arm of each unit from a held-out dataset (Figure 2.8A). If the signals were entirely dedicated to one arm or the other, each subgroup would only contain information about its preferred arm (e.g., a left arm-preferring subgroup would be predictive of left but not right arm movements). If instead there is meaningful activation that is distributed across the same units, then each subgroup would contain both dedicated and distributed information about its preferred arm, but only distributed information about its non-preferred arm.

We first analyzed the time course of modulation for each subgroup during movements of the preferred and non-preferred arms. While modulation during preferred-arm trials was much stronger in the Instruct and Move phases, there was a small amount of modulation during trials of the non-preferred arm as well (Figure 2.8B). To determine whether this modulation carried target-specific information about the behavior, rather than non-specific changes related to task engagement or small movements of the non-selected arm, we trained linear discriminant analysis (LDA) classifiers to predict the target on each trial. Even though the units showed very little modulation when the non-preferred limb was used, prediction accuracy was well above chance (Figure 2.8C-E, paired sample t-test with Rest – monkey O: Instruct $p=1.5e-12$, Move $p=4.1e-21$; monkey W: Instruct $p=1.8e-3$, Move $p=1.1e-7$). This suggests that the population code is not entirely dedicated but contains a meaningful distributed component as well. We refer to this as “distributed” in the sense that the contributing units carry information about both arms.

2.2.8 The distributed signal is contained in a shared subspace for the two arms

We next asked whether subspace separation exists specifically within the distributed portion of population activity. To isolate distributed signals, we again partitioned the population based on preferred arm and fit 4-D PCA models to neural

activity during only trials of the non-preferred arm. This is a conservative approach for fitting only the distributed activity, since dedicated activity will be absent during reaches of the non-preferred arm. This approach can also be interpreted as directly isolating the effect of covariance differences by removing the effect of magnitude differences. We will refer to the subspace spanned by these models as the “distributed” subspace.

If population activity for each arm separates along orthogonal neural dimensions, even in the absence of dedicated variance, then the distributed subspace would preferentially capture variance for the non-preferred arm, since that is what it was fit with. Despite having greater magnitude in the ambient space, preferred arm activity would exist largely in the null space of this projection and none of its variance would exist in the distributed subspace. Alternatively, if signals for the non-preferred arm exist within a shared subspace for the two arms, then the patterns of activity for either arm would be preserved through the projection, and we would expect as much or more variance captured for the preferred arm.

Across all task phases and for both animals, more variance was observed in the distributed subspace during preferred arm trials than during non-preferred arm trials (Wilcoxon signed rank – $p < 0.05$ for all six comparisons). The ratios of variance captured for each arm were expressed as non-preferred over preferred and were computed using the raw variance, not the proportion of total variance. Variance ratios were below 1 for nearly every individual dataset and became even lower with each subsequent phase (Figure 2.8F). Additionally, mean coefficient weights were not significantly different across PMd and M1 for either the Instruct or Move phase, except for the Move phase data in monkey O where M1 weights were slightly larger (permutation test – $p > 0.05$ for monkeys O and W Instruct phase, monkey W Move phase; monkey O Move phase $p = 1.4e-3$, $|\overline{w_{M1}}| = 0.12$, $|\overline{w_{PMd}}| = 0.092$). This indicates that the two areas were similarly contributing to the effect. Together these results suggest that across the entire process of preparation and execution of movements, arm signals that are mixed at the level of individual units occupy a shared subspace and are not differentiated through linear population readouts. We again used an alternative firing rate normalization method to confirm that this result was not dependent on overrepresentation of units with the strongest modulation, and the same results were observed (Figure S2.4B). In summary, the subspace capturing distributed activity is not unique to the arm it was fit to, but rather represents a shared subspace for population activity associated with either arm.

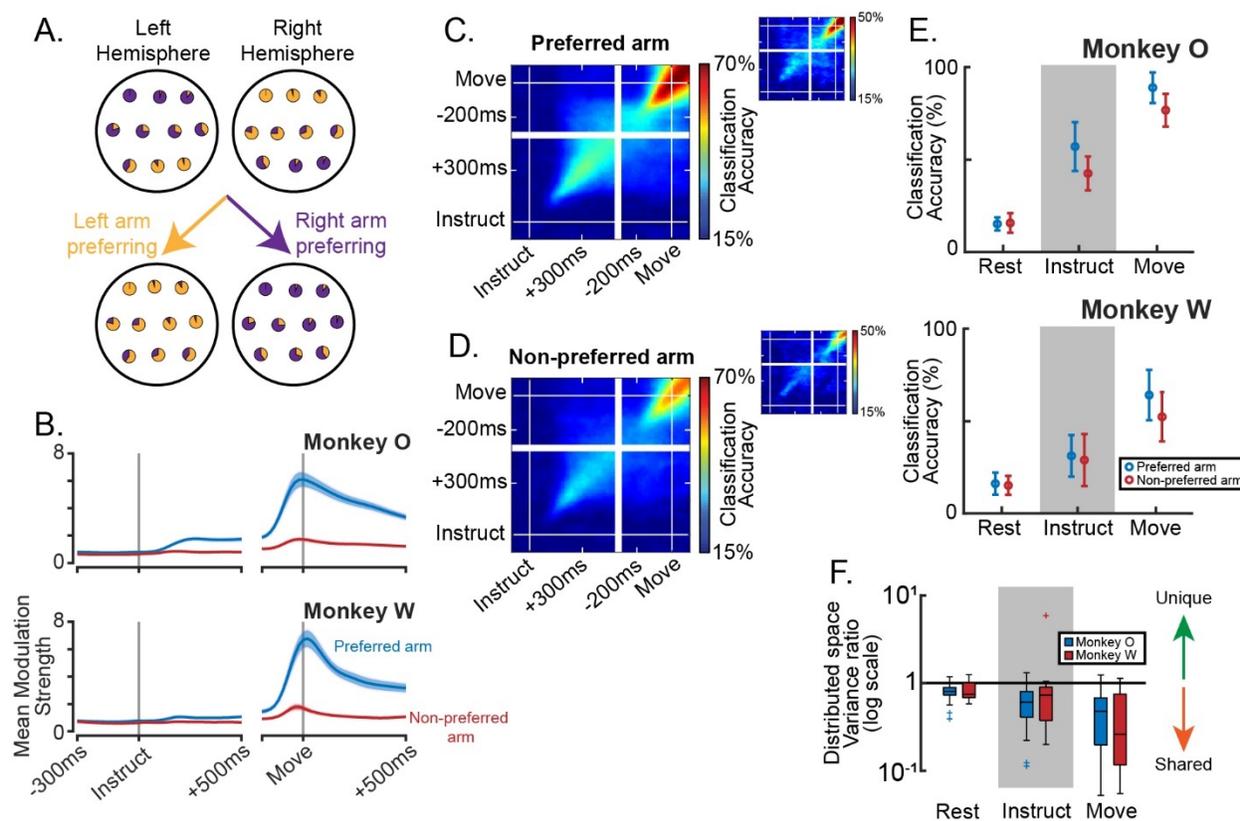


Figure 2.8. Behaviorally specific information exists within a subspace that captures bilateral activity. (A) Illustration of the population partitioning approach. Each unit is represented as a pie-chart displaying the relative modulation during left and right arm trials. Most units in the left hemisphere are more strongly modulated during right arm movements (mostly purple pie-charts), yet some prefer left arm movements (mostly yellow pie-charts). Regardless as to which hemisphere each unit is in, the population may be subdivided into left and right arm preferring sub-populations. On the extreme that all information about each arm is contained within dedicated sub-populations, this simple division will fully segregate the signals such that movements of the non-preferred arm cannot be classified. (B) Modulation as a function of time, taken as the mean over all units during trials of their preferred or non-preferred arm, +/- standard error. (C) Target classification accuracy using LDA for movements of the preferred arm. Models are trained on each time point and tested on each time point to provide high temporal resolution and inform cross-phase generalization of the classifier. Plots are averaged over all sessions (13 Monkey O, large plots; 7 Monkey W, small plots) and both sub-populations (left-preferring, right-preferring). (D) Same as (C), but for non-preferred arm movements. (E) Summary data of (C,D) for monkey O, top panel, and monkey W, bottom panel. Mean +/- standard deviation across datasets. (F) Ratio of the variance captured in the distributed subspace for the two limbs.

2.3 Discussion

We have shown that the combined population response spanning PMd and M1 across hemispheres contains two primary components with regards to laterality.

The first is characterized by signals that are segregated at the level of individual units, which we call the “dedicated” component. Activity that emerged following the instructional cue was most prominent in PMd and showed a tendency for stronger modulation in units with greater arm preference. This caused the signals for each arm to begin segregating into mutually exclusive sub-populations and occupy divergent low-dimensional subspaces. During the transition to movement, M1 became strongly engaged and segregation of arm signals became even more pronounced.

The second component leveraged signals that were mixed within units, which we call the “distributed” component. We showed that arm signals were not completely segregated by splitting the population in two based on each unit’s preferred arm and analyzing the responses during non-preferred arm trials. Despite being very small in magnitude, these signals contained target-specific information. In contrast to the natural separability of the dedicated component, however, subspaces fit to this activity captured at least as much variance for the other arm during each phase of the task, suggesting a shared subspace for the two arms that persists across preparation and movement.

2.3.1 Comparison to previous studies of bilateral arm signals in the motor cortex

Our study adds to a growing body of existing work reporting activity related to both arms in the same motor cortical units during either preparation [45, 110, 210] or movement [8, 45, 72, 108, 125, 204, 210]. Two recent studies have addressed the puzzling presence of bilateral activity in M1 during unimanual behavior [8, 108]. Despite many units being active for either arm (similar to our own observations in Figure 2.4B, Table S2.1), both studies reported separation of population-level arm signals into distinct neural subspaces. Furthermore, they attributed this separation to covariance changes between the two modes that would cause even signals that are mixed within units to contribute to the effect. In the present study, we build upon these observations and reveal an underlying organizational structure that suggests arm signals are not as mixed within units as they would appear based on distributions of single-unit arm preferences. We show that dedicated signals contribute more to the overall population variance and solely account for the presence of arm-specific subspaces. Signals that are mixed within units reflect a different feature of the population. Our example in Figure 2.5C-E demonstrates that segregation of arm signals into distinct neural subspaces likely arises from activation of exclusive sub-populations – similar to the top principal components from Ames and Churchland, 2019 and Heming et al., 2019 (as well as those in our Figure 2.7A-B example), dedicated regimes of the population in our study

captured large amounts of variance for one arm while capturing very little for the other. Analysis of model structure confirmed that dedicated variance drove PCA results (Figure 2.7, 2.8). Furthermore, we mapped the development of this signal feature across preparation and movement, showing that it begins to emerge even during preparation (Figure 2.5, 2.6) and involves both PMd and M1 (Figure 2.5).

Not only does this clarify the mechanism of signal separation in motor cortical activity, but it acknowledges a critical heterogeneity in the population response. Dedicated signals represent most, but not all, of the population variance. Signals that are mixed within units (distributed signals) reveal a portion of the population activity that is not independent for the two arms. Effector-independent coding has long been appreciated in PMd [45, 210]. Recently, Willet et al., 2020 identified separate “limb-coding” and “movement-coding” dimensions in neural activity from the hand knob area of human premotor cortex [223]. The “movement-coding” component represented movements for different effectors in the same fashion. This bears similarity to our “distributed” component, although we only show that activity for each arm resides in the same space, not that its relationship to behavior is invariant. Notably, PMd units were not more heavily weighted in our distributed models. This suggests that even in M1, signals for the two arms are not fully independent.

2.3.2 Progressive segregation of arm-dedicated signals and its functional significance

To our knowledge, this study is the first to compare low-dimensional population structure during preparation of left vs right arm reaching in neurologically-intact subjects (see [223] for preparation of attempted movements in a human participant with C4 spinal cord injury). It has been proposed that neural subspaces reorganize between preparation and execution of reaching movements [80], which we observe in our own data (Figure 2.6B). A principle interest of this study was to determine how the emergence of arm-specific signals maps onto this reorganization process. Since previous work has shown that the transition from preparation to movement coincides with an increased proportion of lateralized units [45, 150, 201], we expected activity to progressively segregate at the level of individual units, represented primarily in the contralateral hemispheres, as the population reorganizes between task phases. This was indeed the case, and began even during the instruction phase (Figure 2.5). Careful inspection of model structure revealed that segregation of arm signals at the individual unit level drove separation of arm-specific neural subspaces gradually throughout the trial (Figure 2.6C-D; Figure 2.7). This segregation reached its maximum during movement, thus reducing any concern that small movements of the non-selected arm had an impact on our results or conclusions.

Importantly, contralateral signals were more independent than ipsilateral ones; a larger proportion of contralateral modulation occurred in contra-dedicated units than the reverse case for ipsi- (Figure 2.5). This was not a surprising result, as contralateral bias in the functional organization of motor cortex has been clearly revealed by effects stroke [105], lesion studies [20], and cortical stimulation [6, 162, 170]. One candidate hypothesis for the presence of ipsilateral activity is that it supplies an independent control signal. There is some evidence that ipsilateral cortex plays an increased role in movement following hemispheric damage [20, 58, 112, 222], though not necessarily a beneficial or compensatory one. The magnitude of ipsilateral encoding increases with the degree of movement complexity [215] and may involve spatially distinct neural populations [38, 226]. However, the corticospinal tract (CST) is almost entirely contralateral, and the effectiveness of the ipsilateral component has been debated [12, 140, 183, 202]. Ipsilateral cortex may also exert its influence via connections made in the reticular formation [6, 12, 222], which projects bilaterally to the spinal cord. Our results showed a small amount of independent ipsilateral activity (monkey O more so than monkey W), with more of the ipsilateral signal coming from non-dedicated units (Figure 2.5). Thus, if the ipsilateral hemisphere provides any independent control signal, it is much weaker than the contralateral signal. Rather, our results suggest that ipsilateral signals are involved in some form of bilateral control, which we now discuss.

2.3.3 Bilateral signals and their role in motor control

Correlated activity for movements of the two arms has been widely reported in the literature, primarily using macro-scale neurophysiological approaches. Increases in excitability of homologous effectors during transcranial magnetic stimulation (TMS) [159] and symmetric activation patterns in functional magnetic resonance imaging (fMRI) [69, 216] suggest that bilateral motor cortical circuits are organized with mirrored properties. Similar correlated structure has also been reported in human ECoG [30] and premotor spiking activity [223]. Mirror activation and other forms of interhemispheric communication have been proposed to support intermanual skill transfer [69] or shaping of contralateral activity patterns during complex behavior [215]. However, correlations between the tuning for ipsilateral and contralateral arm movements in M1 units tend to be weak or absent [45, 108, 204]. In the present study we have not directly compared directional tuning, yet we did observe that the distributed component of bilateral signals existed within a shared subspace for the two arms (Figure 2.8). Mirror activity would necessarily reside in the same neural subspace for each arm, provided that subspace is linear, as all linear subspaces are invariant with respect to reflection. Our results are therefore consistent with functional hypotheses of ipsilateral cortex involving mirror symmetric activation, and more generally for hypotheses that predict linear correlations between activation patterns for the two arms.

We note, however, that while consistent linear correlations in the tuning properties of individual neurons would deterministically result in shared neural subspaces, a lack of linear correlation does not mean that neural subspaces will be orthogonal.

The distributed component that we have characterized may also play a role in bimanual coordination. Distinct bimanual activity patterns have been observed in caudal premotor regions using human fMRI [69] and in M1 using single-unit recordings in monkeys [72, 204]. Surgical transection of the corpus callosum, the primary direct connection between hemispheres [93], disrupts typical spatial coupling and continuous synchronization of arm movements as well [88, 132], suggesting a cortical locus for these forms of bilateral control. These studies suggest that bilaterally distributed networks involving PMd/M1 may facilitate bimanual coordination, a function historically attributed to the supplementary motor area [22]. Our task involved unimanual movements, containing no component of coordination. However, the result that target-specific information existed within a shared subspace (Figure 2.8) is consistent with a role in coordination. We make limited claims on this hypothesis due to our simplified behavior, and stress that implicating a role in bimanual coordination does not simply mean revealing a shared substrate for signals of both limbs. Nonetheless, a bi-hemispheric network structure may underly computations for controlling the two arms as a unified plant [220]. M1 has been implicated in multi-joint integration for voluntary movement and feedback control [177, 191]. Bimanual behaviors have a similar task of overcoming redundant degrees of freedom [17]; many patterns of behavior for each arm independently may help one achieve an action goal so long as cooperation of the two remains intact (“motor equivalence” [141]). This lower-dimensional behavioral coordination space, sometimes called “the uncontrolled manifold” [188], would likely have a similar neural manifold in which bilateral arm signals interact (for related discussion and review, see [68, 208, 221]). The distributed space that we report may reflect such a manifold.

2.3.4 Interpretations from a dynamical systems perspective

One unified explanation for the two components identified in this study is that they represent the computational (or “hidden”) layers and the output layer of cortical processing. In this framework, the distributed signal would reflect a bilateral network that plays a supportive role in motor processing rather than direct output. The idea that bilaterally distributed networks contribute to computations that do not directly represent the output has been previously proposed by Ames and Churchland, 2019 [8]. Preparatory activity in motor areas reflects abstract features of action and may lack a strong contralateral bias [45, 110]. The distinctive lack of laterality in the distributed signal we observed is consistent with other reports of abstract preparatory responses. It played a relatively stronger role during preparation as well, since the dedicated

component did not fully develop until movement. This aligns with reports that behaviorally specific features become more apparent in motor cortical signals during active behavior, including laterality [45, 199].

From a dynamical systems perspective, distributed signals could serve to enforce internal dynamics of the overall population. Preparatory signals in pre- and primary motor cortex are thought to converge on an ideal population state, or initial condition, such that internal circuit dynamics will guide appropriate patterns of activity for the upcoming movement [40, 151, 200]. Rodent studies have shown that preparatory activity in motor cortical neurons projecting to other cortical areas lacks strong laterality, while neurons with descending output exhibit pronounced contralateral bias and became active closer to movement onset (left vs right directional licking task [150]; left vs right arm pedal pressing task [201]). Furthermore, these bilaterally distributed networks provide robustness to unilateral perturbation during preparation, and it has been hypothesized that the two hemispheres operate together to maintain the network state (left vs right directional licking task [151]). The two components that we have identified generally align with this form of network structure. In addition to setting the initial state, persistence of the distributed component during movement may reflect the ongoing dynamics of pattern generation [200, 207].

Within this interpretation, progressive segregation of arm signals may reflect emergence of descending output from the network that mirrors the well-established laterality of anatomical pathways [20, 202]. Alternatively, it may reflect a timing signal for triggering action or transitioning the network from preparation to movement [123, 207] while simultaneously specifying the selected effector. Like the dedicated signals we observed, signals that reflect the timing of movements, but not their direction, have been shown to capture the most variance in PMd/M1 population responses [123]. Premotor activity has also been shown to contain “limb-coding” dimensions that specify a movement effector independently of the movement type [223]. The large dedicated signals that we observe bear similarity to both of these previously identified response features, and all three could reflect the same underlying computational process.

In summary, we present a statistical description of arm signals spanning M1 and PMd throughout reach preparation, characterizing in detail both lateralized and non-lateralized features of the population response. The two components that we have identified will be crucial for contextualizing current theory on bilateral motor cortical processing as well as designing future experiments that investigate the independence and interaction of signals across the hemispheres.

2.4 Methods

2.4.1 Behavioral recordings and task

Kinematic data were collected using LED-based motion tracking of several points along each arm (Phasespace Inc, San Leandro, CA). 3D positions of each LED were sampled at 240Hz. Prior to offline analysis, these positions were smoothed using a cubic spline and smoothing parameter 0.005 (*cspaps* function – MATLAB). The most distal LED, located on the back side of each hand just below the wrist, was used for online endpoint feedback and all offline analysis.

Monkeys were trained to perform a variant of an instructed-delay reaching task (Figure 2.1B). Endpoint feedback of each arm and all visual stimuli were presented to the animal using a custom-built virtual reality 3D display. This display consisted of two mirrors that projected shifted images independently to each eye to produce stereopsis. Cursors, indicating effector endpoint position, were color coded for the left (yellow) and right (purple) hands, as were all associated stimuli.

Each trial began with the appearance of the start positions for each hand (spherical targets, radius 4cm, with centroids separated by 8cm), located near the body on top of a physical bar that the monkey rested its hands on (Figure 2.1A). In a self-initiated manner, the monkey would assume the start position by placing both cursors in their appropriate starting positions and maintaining that position for 500 ms (“Rest” phase). Our threshold for detecting movement online was 9cm/s; breaking this threshold would abort the trial.

Marking the beginning of the “Instruct” phase, a cue (spherical target, radius 3cm) would appear at one of six locations within a fronto-parallel plane 8cm in front of the start positions (Figure 2.1A). The color of the cue indicated the required arm, and position of the cue was the target location for the forthcoming reach. The instruction cue remained visible through the delay period, a duration that was sampled uniformly on the interval 500-1500ms. Movement beyond the speed threshold with either hand would abort the trial.

At the end of this period, two simultaneous changes signaled the monkey to move and marked the start of the “Move” phase. First, the sphere defining the start position for the cued arm disappeared. Second, the cue at the target location enlarged (3cm to 4cm radius). The monkey then reached toward the target and once at the terminal location, had to maintain that position for 250ms. To earn a juice reward, the animal had to initiate the reach within 500ms of the onset of the imperative, terminate the movement within the target’s circumference, and keep the non-reaching hand stationary for the duration of the trial. To further emphasize that the trial was successful, the target turned green.

300ms windows were used to represent each phase in data analysis. For the Rest phase, we used the final 300ms before the onset of the instruction cue. For the Instruct phase, we used data in the interval between 200ms to 500ms post-cue. For the Move phase, we used the first 300ms following the onset of movement, defined as when speed of the reaching hand exceeded 10cm/s. We used a late window for the Rest phase to avoid any residual activity associated with moving to the start positions. The steady state neural response was used to position the Instruct phase window; this was reached approximately 200ms after the onset of the instruction cue (see Figure 2.7B). The Move window was selected to capture peak neural activity associated with movement while including only the feed-forward portion, which typically lasted 250-300ms (Figure 2.1C, bottom row). Reach durations were calculated as the time between movement onset and the first point where (1) movement speed dropped below 20cm/s, and (2) velocity in the depth direction reached 0.

2.4.2 Surgical implantation

All procedures were conducted in compliance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals and were approved by the University of California at Berkeley Institutional Animal Care and Use Committee under protocol ID AUP-2014-09-6720-1. This protocol approval and all surgical methods apply to each chapter of this thesis. Two adult male rhesus monkeys (*Macaca mulatta*) were implanted bilaterally with custom acute recording chambers (Grey Matter Research LLC, Bozeman, MT). Partial craniotomies within the chambers allowed access to the arm regions of dorsal premotor (PMd) and primary motor (M1) cortices in both hemispheres. Localization of target areas was performed using stereotactically aligned structural MRI collected just prior to implantation, alongside a neuroanatomical atlas of the rhesus brain [169].

2.4.3 Electrophysiology

Unit activity was collected using 24-32 channel multi-site probes (V-probe - Plexon Inc, Dallas, TX), with contacts separated by 100um and positioned axially along a single shank. Probes were lowered deep enough to cover roughly the full laminar structure of cortex (Figure 2.2B-C). The depth of insertion was determined by (1) measurements of the dural surface prior to recording, and (2) presence of spiking activity across all channels. 2 probes were typically inserted in each hemisphere daily and removed at the end of the session, one in PMd and one in M1. A total of 12 insertion points across PMd and M1 of each hemisphere were used across 13 recording sessions in Monkey O, and 6 insertion points across 7 sessions for Monkey W (Figure 2.2A).

Neural data were recorded using the OmniPlex Neural Recording Data Acquisition System (Plexon Inc, Dallas, TX). Spike sorting was performed offline (Offline Sorter – Plexon Inc, Dallas, TX). Single-unit waveforms were isolated in multi-dimensional feature space (including principal components, non-linear energy, waveform amplitudes) and rejected if either (1) the waveform clusters were not stable over the course of the session, or (2) $>0.4\%$ of inter-spike-intervals were below 1ms. For population level analyses (PCA, LDA), a small number of multi-units were included. A multi-unit was defined by waveform clusters that separated from the noise cluster and were stable over time, but did not quite meet the inter-spike-interval criteria or contained what might be multiple unit clusters that could not be easily separated. For monkey O, the average proportion of multi-units in each single session population sample was 17%, ranging 12-25%. For monkey W, average 20%, ranging 12-32%.

Spiking data were binned in 20ms non-overlapping bins, square-root transformed to stabilize variance, and smoothed with a 50ms gaussian kernel for all analyses [225].

2.4.4 Modulation and Arm Preference metrics

As a time-varying value, modulation was calculated as:

$$M_t = \left(\frac{x_t - \mu_{Rest}}{\sigma_{Rest} + 1} \right)^2 = z_t^2,$$

where

x_t : instantaneous firing rate at time t

μ_{Rest} : mean firing rate during Rest

σ_{Rest} : standard deviation during Rest

This unitless metric reflects the deviation from baseline activity, normalized by baseline fluctuations. The constant 1 was added to the denominator for soft-normalization to ensure that units which were silent during rest did not have exploding values and were not overly emphasized in the dataset. Because some units had slightly different activity on left and right arm trials even before instruction, the standard deviation during Rest was calculated separately for each arm and σ_{rest} was calculated as the mean of the two.

Single values of modulation representing each phase were obtained using the same 300ms phase windows that were used in behavioral analysis (see Methods – Behavioral recordings and task). Phase-specific modulation data were concatenated across trials into a $(16m \times 1)$ vector, where m is the number of trials and 16 is the number of samples within a 300ms phase window. The mean over this vector provided our scalar estimate of modulation. Note that this is simply a normalized form of variance, which makes the comparison of single-unit and population-level results more

straightforward. To make the relationship with variance explicit, modulation can be rewritten as:

$$M_{phase} = \frac{1}{(\sigma_{Rest} + 1)^2} E[(X_{phase} - \mu_{Rest})^2]$$

Where the expectation on the right is essentially variance using the Rest mean.

Arm Preference was calculated independently for each phase of the task using the formula:

$$AP_{phase} = \frac{M_{contra,phase} - M_{ipsi,phase}}{M_{contra,phase} + M_{ipsi,phase}}$$

An arm preference of 1 corresponds to a unit that is exclusively modulated during contralateral trials, while an arm preference of -1 is the same for ipsilateral trials. In Figure 2.7, Figure S2.4, and the accompanying analyses, the convention of left arm and right arm was used in place of ipsi and contra. In analyses that used arm preference along with other features, independent datasets were used to calculate each to avoid any artificial coupling due to sampling noise, e.g. modulation and arm preference. Note also that the scaling factor used in the modulation calculation cancels out of the arm preference calculation, making it invariant to the choice of normalization.

2.4.5 Principal components analysis

Principal components analysis (PCA) was used to identify low-dimensional representations of population activity with the *pca* function in MATLAB. PCA computes an orthogonal basis set that reflects the principal axes of variation in the data. Individual components do not strictly correspond to observed activity patterns, and one should be wary of interpreting them as such, yet the low-dimensional space spanned by the top few components has been frequently used in systems neuroscience as a helpful descriptor of coordinated ensemble activity [55]. PCA was selected over other dimensionality reduction techniques for its widespread use and relative lack of assumptions. Additionally, PCA was used in two recent papers covering similar topics to this one [8, 108]. Therefore, using PCA over other alternatives was also intended to improve generalization of our results to the existing literature.

Prior to fitting the models, firing rate data were soft-normalized using the same method as in the modulation strength calculation:

$$z_t = \frac{x_t - \mu_{Rest}}{\sigma_{Rest} + 1}$$

An alternative normalization factor was used to create Figure S2.4, replacing the denominator by the full firing rate range + 5Hz [8, 80, 108].

Since Rest phase mean activity was already subtracted from individual units, we did not de-mean again prior to computing PCA models. Measures of variance accounted for were not inflated by capturing means because they were computed using the variance of the component scores (Figure 2.7, Figure 2.8F):

$$V = \text{Tr}(\text{Cov}(XP))$$

Where X is a $(16m \times n)$ data matrix of concatenated trials and P is an $(n \times p)$ projection matrix, with m trials, 16 samples in the phase window of each trial, n units, and p principal component dimensions.

Cross-validation approaches were used for all analyses and figures to address overfitting. This provided accurate and generalizable estimates of variance capturing metrics that could also be appropriately compared across datasets (i.e. across time or arms).

2.4.6 Dimensionality estimation

Dimensionality of the PCA subspace was estimated by optimizing the cross-validated reconstruction of full-dimensional neural data from component scores. Given n units, m trials, and 16 samples within the phase window for each trial, the following procedure was used:

1. Leave out the i^{th} trial (16 samples) from the data matrix, yielding training data, $X^{(-i)} \in \mathbb{R}^{16(m-1) \times n}$, and testing data, $X^{(i)} \in \mathbb{R}^{16 \times n}$.
2. Train PCA model of dimension $p < n$ on $X^{(-i)}$, using singular value decomposition (SVD) to compute the projection matrix, $P^{(-i)} \in \mathbb{R}^{n \times p}$
3. Leave out the j^{th} unit from the testing data and projection matrix by removing the j^{th} column and row from each, respectively, yielding $X_{-j}^{(i)} \in \mathbb{R}^{16m \times (n-1)}$ and $P_{-j}^{(-i)} \in \mathbb{R}^{(n-1) \times p}$. This is the current unit that will be reconstructed.
4. Using the Moore-Penrose pseudoinverse, find a new projection matrix with the j^{th} unit removed, whose transpose is $(P_{-j}^{(-i)})^+ \in \mathbb{R}^{p \times (n-1)}$. This matrix projects the $(n-1)$ dimensional neural activity into the original p dimensional PC space, therefore computing component scores in the absence of unit j .
5. Calculate the component score for the i^{th} trial using the remaining units and the new projection matrix, then estimate the j^{th} unit from that component score by projecting back into the ambient space. As a single step, this calculation is:

$$\hat{X}_j^{(i)} = \left[P^{(-i)} \left(P_{-j}^{(-i)} \right)^+ \left(X_{-j}^{(i)} \right)^T \right]_j$$

7. Repeat for trials $i=1, \dots, m$
8. Repeat for units $j=1, \dots, n$
9. Repeat for component numbers $p=1, \dots, 10$
10. Take the number of components that minimizes the predicted residual error sum of squares (PRESS) statistic:

$$PRESS = \sum_{i=1}^m \sum_{j=1}^n \left(X_j^{(i)} - \hat{X}_j^{(i)} \right)^2$$

This method reconstructs the full-dimensional neural data, independent of the training set, by identifying consistent population structure. There are no mathematical constraints favoring increased dimensionality, i.e. it is robust to overfitting. As such, the number of components that minimizes the reconstruction error provides a conservative estimate of the dimensions that meaningfully reflect population structure. Similar methods have been used previously for assessing dimensionality reduction techniques for neural data [225]. Using heuristics, such as the number of components to explain 90% variance, would be inappropriate for our analyses. They are prone to overfitting, which would include meaningless components and impair analysis of model structure via coefficient weights.

2.4.7 Covariance alignment

We computed a measure of similarity between pairs of subspaces that we call Covariance Alignment. Our method is essentially the same as that previously used for comparing low-dimensional spaces via factor analysis [10]. In short, this measure computes the proportion of low-dimensional variance from one dataset that is also captured in the low-dimensional space of another dataset.

Given data matrices $X_A, X_B \in \mathbb{R}^{16m \times n}$, the following procedure was used:

1. Train PCA models of dimension $p < n$ on X_A and X_B , using SVD to compute the projection matrices $P_A, P_B \in \mathbb{R}^{n \times p}$

2. Project X_A into its own p -dimensional space and compute the variance as:

$$V_A = Tr(Cov(X_A P_A)) = Tr(Cov(T_A))$$

3. Project the p -dimensional representation of X_A , which is T_A , into the p -dimensional space identified using X_B and compute the variance as:

$$V_{A_{in}_B} = Tr(Cov(X_A P_A P_A^T P_B)) = Tr(Cov(T_A P_A^T P_B))$$

4. Return the proportion of p -dimensional variance from dataset A that is also captured in dataset B 's subspace using the ratio:

$$CA = \frac{V_{A \text{ in } B}}{V_A} = \frac{\text{Tr}(\text{Cov}(X_A P_A P_A^T P_B))}{\text{Tr}(\text{Cov}(X_A P_A))} = \frac{\text{Tr}(\text{Cov}(T_A P_A^T P_B))}{\text{Tr}(\text{Cov}(T_A))}$$

This metric is subtly different from the alignment indices used in [80] and [108]. The key difference here is the double projection in the numerator, which means that we are specifically capturing the proportion of low-dimensional variance from one dataset that is captured in the low-dimensional space of another, rather than the ratio of overall variance captured in two different subspaces.

2.4.8 PCA coefficient analysis

Since components of PCA models form an orthogonal basis set, each was independently analyzed to determine its contribution to subspace divergence. Two statistics were calculated for each component using held-out datasets.

First, we projected activity during trials of each arm onto a single component, calculated the variance of the projections for each arm, and expressed them as a ratio. This captured each component's contribution to discrimination between the arms. For component C , this calculation is:

$$V_{C,R/L} = \frac{\text{Var}(X_R P_C)}{\text{Var}(X_L P_C)}$$

Where $X_R, X_L \in \mathbb{R}^{16m \times n}$ are data matrices for the right and left arms, respectively, and $P_C \in \mathbb{R}^{n \times 1}$ is the projection matrix for component C . The log of this ratio will be far from 0 if there is much more variance for one arm than the other along the axis defined by P_C .

Second, we calculated a coefficient-weighted average of the arm preferences for all units. If non-zero weights were only given to right arm dedicated units, this value would be 1; if weights were evenly distributed across the spectrum of arm-preferences, this value would be 0. Therefore, this measure captured the dependence of a given component on arm-dedicated units. The coefficient-weighted arm preference, CAP , for component C was calculated as

$$CAP_C = \frac{A|P_C|}{\sum_{i=1}^n |P_{C,i}|}$$

Where $A \in \mathbb{R}^{1 \times n}$ is the vector of arm preferences for each unit.

2.4.9 Linear discriminant analysis

Population coding of movement was analyzed using Linear Discriminant Analysis (LDA) with the *fitdiscr* function in MATLAB. LDA assumes that each class (target x limb combination) is associated with a multivariate normal distribution over the predictor variables (spiking activity of multiple units) having identical covariance but different means.

The feature matrix $X_{LDA} \in \mathbb{R}^{m \times n}$ consisted of a single sample per trial for each of the n units. For fine timescale analysis, this was the instantaneous firing rate. For models representing an entire phase, this was the mean firing rate during the 300ms phase window. Uniform priors were enforced for all models. As it was expected that the covariance may change across use of the two arms during reaching, LDA models were trained separately for each limb to allow fitting of arm-specific covariance matrices. LDA was chosen for its robustness to violations of the given assumptions and its history of success with neural data [69, 180].

2.4.10 Fine timescale analysis of population coding and subspace development (heatmaps)

The same basic method was used for displaying fine timescale changes in population coding of movements (via LDA) and covariance structure (via PCA, Covariance Alignment). The method is depicted schematically in Figure S2.5. Neural data were organized as 3D tensors (units, time windows, trials). Comparisons were made between all possible pairs of time windows, using fully independent trial sets to prevent overfitting. For LDA models, this consisted of leave-one-out cross-validation; for Covariance Alignment, random partitioning into two datasets of equal trial numbers. Averages of the cross-validated results provided the 2D matrices visualized using heatmaps in Figure 2.6B-C and Figure 2.8C-D. A single row or column therefore reflects the similarity of population coding or covariance between a single timepoint and all other timepoints across the trial. Block diagonal structure in the heatmaps reveals locally consistent structure within task phases.

2.4.11 Permutation testing procedures

Permutation tests were used for both single and multi-factorial hypothesis testing when parametric tests were inappropriate. Null distributions were constructed by constraining permutations to only data that were exchangeable under the null hypothesis whenever possible [9]. For example, we maintained the crossed structure of Phase (Rest, Instruct, Move), by only permuting Phase labels within units. 10,000 permutations were used for all analyses, and p-values were estimated as the proportion

of permutations resulting in test statistics that were at least as extreme as what was observed. In cases where the observed test statistic was more extreme than any permutations, we assigned a p-value of $1/\text{number of permutations} = 1.0e-4$.

2.5 Supplementary Figures

	Number SU's	Significantly modulated					
		Instruct			Move		
		Ipsi	Bi	Contra	Ipsi	Bi	Contra
PMd	433	68(16%)	154(36%)	105(24%)	52(12%)	245(57%)	94(22%)
	113	11(10%)	31(27%)	26(23%)	19(17%)	42(37%)	21(19%)
M1	331	57(17%)	95(29%)	61(18%)	47(14%)	194(59%)	60(18%)
	289	39(13%)	35(12%)	59(20%)	22(8%)	110(38%)	87(30%)

Table S2.1. Proportions of significantly modulated single-units across task phases.

For well isolated single-units in each brain area, the proportions of the total population that were significantly modulated when compared with the Rest phase (two-sample t-test, $p < 0.05$) are displayed in each cell of the table. For each phase, single-units were classified as uniquely ipsi, contra, or bilaterally modulated. Top row in each pair of rows represents Monkey O, bottom row Monkey W.

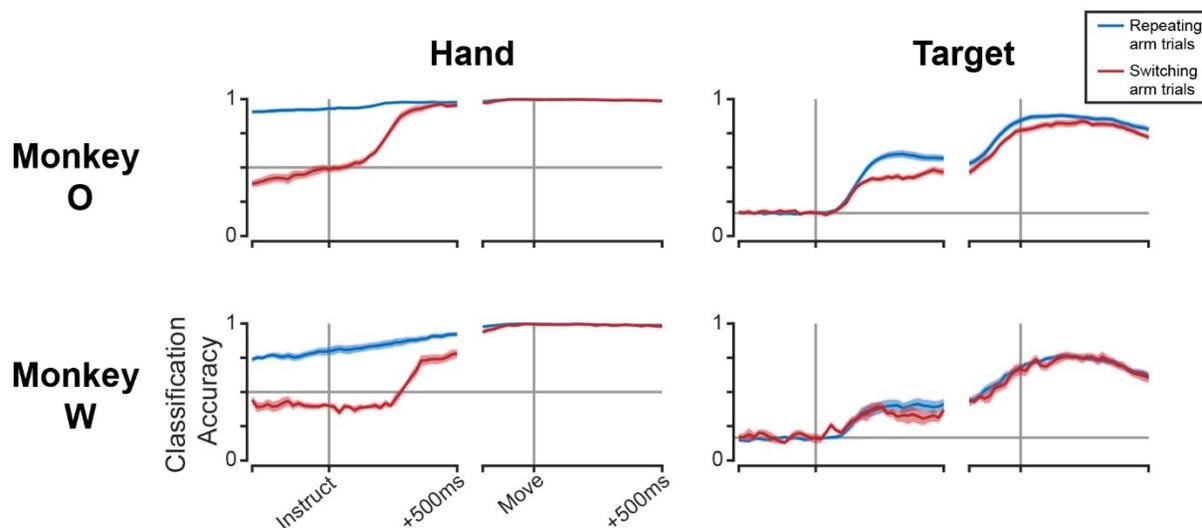


Figure S2.1. Arm-specific neural patterns exist during Rest on predictable trials.

Cross-validated classification accuracy for hand (left column) and target (right column) assignments. LDA models were trained on only trials that required use of the same arm as the previous trial, then tested on either held-out repeating arm trials (blue lines) or switching arm trials (red lines). Separate models were used for each timepoint. Horizontal grey lines indicate chance level. 13 Sessions for monkey O (top row); 7 sessions for monkey W (bottom row). Mean \pm standard error across sessions.

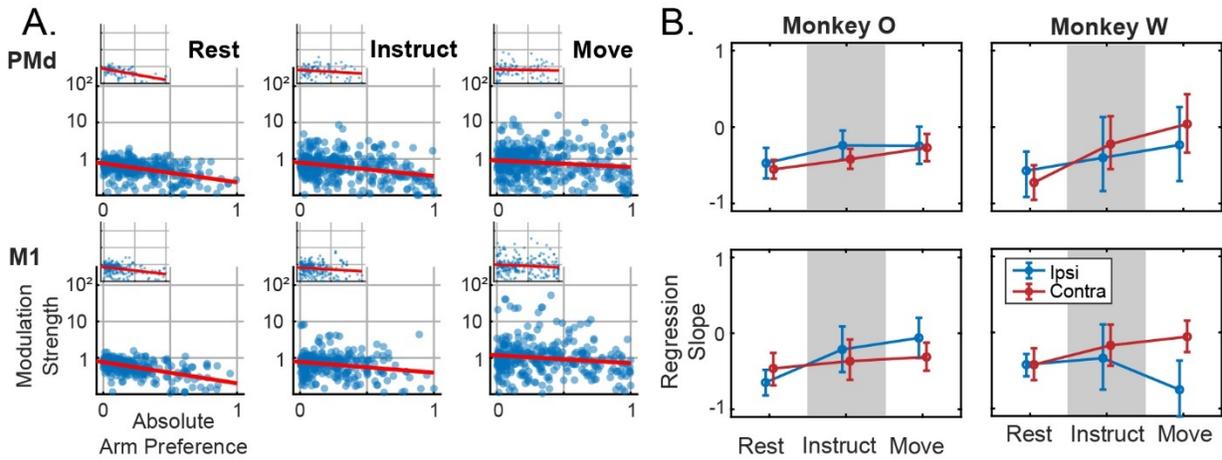


Figure S2.2. Modulation for the non-preferred arm does not increase with greater arm preference. Companion figure for Figure 5A-B. **(A)** Modulation for the non-preferred arm plotted against arm preference, for all units in each brain area and task phase. Log-linear best fit lines are displayed in red. Inset figures belong to Monkey W. **(B)** Slopes of regression lines fit to data from (A), independently for ipsi- and contra-prefering sub-populations. Mean \pm bootstrapped 95% confidence interval. Note the different y-axis from Figure 4B. The slope was not significantly greater than 0 in any condition, meaning that increased arm preference is associated uniquely with greater modulation in the preferred arm, as opposed to an increase for both arms that is just larger for the preferred arm.

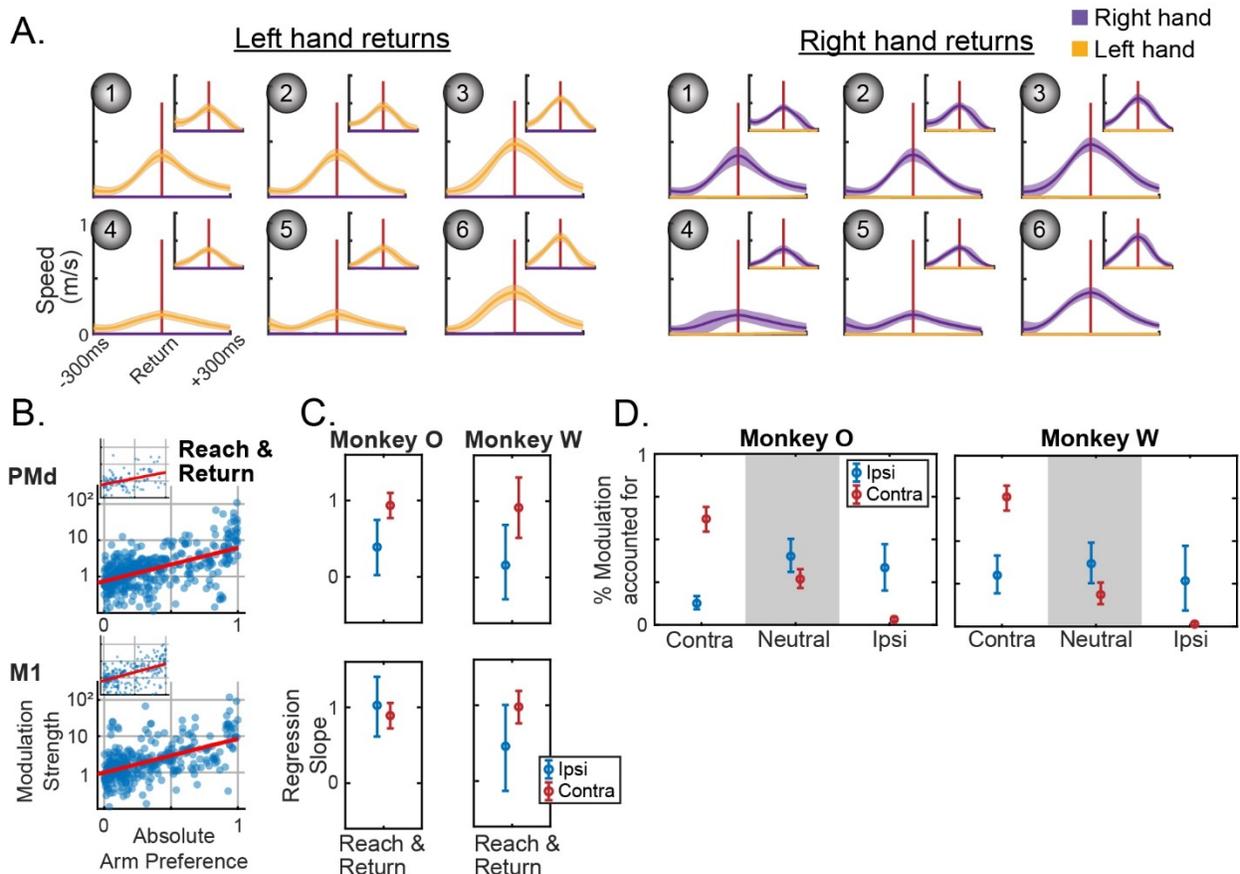


Figure S2.3 Dedicated signals persist with increased kinematic range. To determine whether having a limited range of reach directions was responsible for the observation of arm-dedicated signals, select analyses were performed again on data that included return movements. By including these movements, which were opposite the direction of the outward reaches used in the primary analyses, the range of sampled behavior was greatly increased. (A) Speed profiles for return movements following target acquisition during left- or right-hand trials. Individual trials were aligned to peak return speed, indicated by the vertical red line. Both reaching and stationary hands are plotted in each. Despite being unconstrained by the task, the non-selected hand remained still during the return. Monkey O main, monkey W inset. Mean \pm standard deviation. (B-D) Analyses from Figure 5 repeated using Move phase data concatenated with 300ms of data beginning 200ms before the point of peak return speed, i.e. reach and return. (B) Compare to Figure 5A. Modulation for the preferred arm plotted against arm preference, for all units in each brain area. Log-linear best fit lines are displayed in red. Inset figures belong to Monkey W. (C) Compare to Figure 5B. Slopes of regression lines fit to data from (B), independently for ipsi- and contra-prefering sub-populations. Mean \pm bootstrapped 95% confidence interval. (D) Compare to Figure 5F. The proportion of modulation within each partition from (C) during ipsi- or contralateral movements. Note that the total modulation is significantly lower for ipsilateral movements, particularly for Monkey W, and these data are only displayed as proportions. Mean \pm bootstrapped 95% confidence interval.

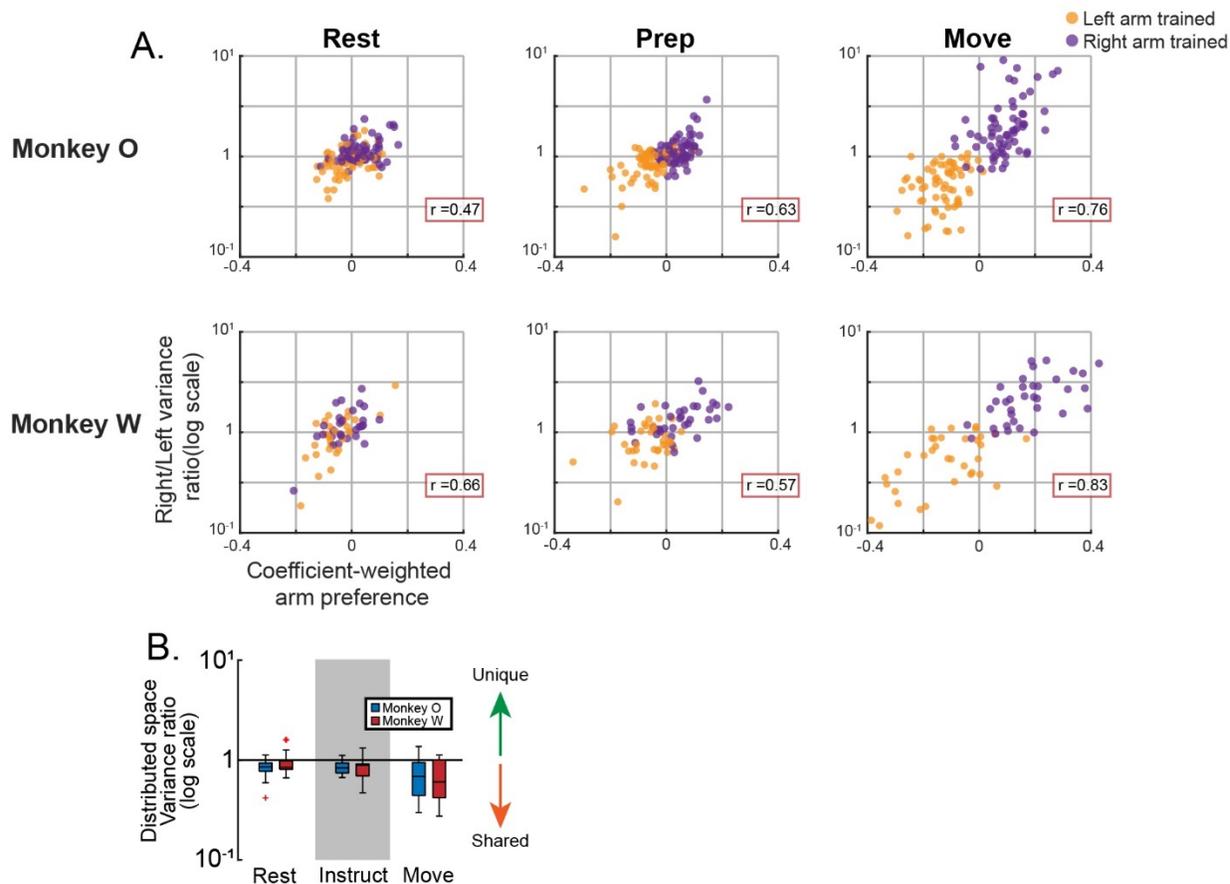


Figure S2.4 Subspace analysis using alternative firing rate normalization. Prior to performing PCA, an alternative method of normalizing firing rates was used for these plots. Rather than dividing by the standard deviation at Rest, each unit's firing rate trace was divided by the full firing rate range + 5Hz [8, 80, 108]. This will mitigate the effect of highly modulated units, which PCA will preferentially represent otherwise. (A) Repetition of Figure 7D. (B) Repetition of Figure 8F.

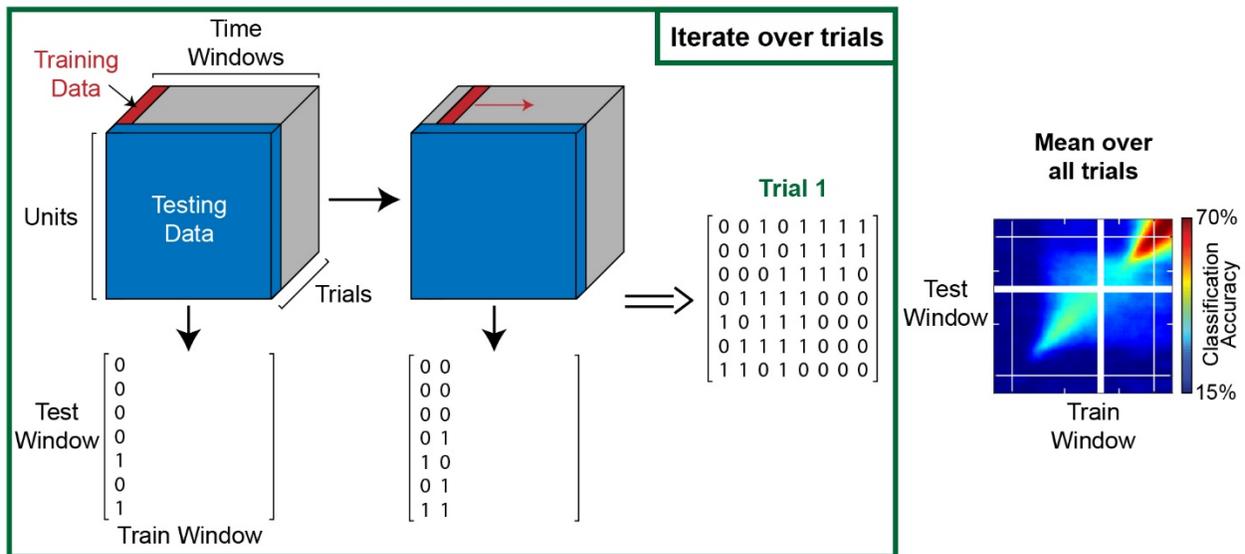


Figure S2.5 Method for fine timescale analysis of population coding and subspace separation. This schematic outlines the process for fine timescale analysis of population coding using LDA and leave-one-out cross-validation. Neural data were organized as 3D tensors (units, time windows, trials). Models were trained to predict targets using a single time window and all but one trial. Those models were then used to predict the target on the held-out trial, making separate predictions based on neural data from each time window. The process was then repeated using the next time window as training data until all possible pairs of time windows had been used as training and testing data. This constituted a 2D matrix of “hit” booleans (number time windows x number time windows) for the predictions of a single trial. After iterating over all trials to be used as held-out test data, the mean was taken across trials to construct a single 2D matrix of classification accuracy. The same basic process was used for visualizing the development of subspace separation, but instead of leave-one-out cross-validation trial sets were repeatedly divided into two random halves of equal size. Covariance alignment was then computed between all possible pairs of timepoints for the two disjoint trial sets.

Acknowledgements

This work has been posted as a preprint to bioRxiv [71] and is currently under review, but it has not yet been accepted to a peer-reviewed journal at the time of this thesis submission.

We thank M. Kitano for help in NHP care and handling, and A. You, E. Formento, W. Liberti and Z. Balewski for helpful discussions. This work was supported by the National Defense Science and Engineering Graduate Fellowship to T.C.D, and a grant from the National Institute of Health (NS097480) to J.M.C.

Chapter 3

Interactions between bilateral limb posture and unilateral reaching signals

3.1 Introduction

The previous chapter established fundamental properties of the signals relating to movements of the two hands. In particular, it characterized the level of independence between the two arm signals during isolated movements and identified a meaningful overlap between them as well. The present chapter now seeks to understand bilateral signal interactions from a functional perspective, focusing largely on how posture of the stationary hand impacts unimanual reaching signals.

Despite being nominally “motor”, passive movement of the arms influences activity in both the pre- and primary motor cortices, suggesting that sensory information also exists in these areas [7, 84, 145]. The presence of this information may allow contextually appropriate feedback responses [177] or reflect the forecasting of sensory consequences of movement [106, 205]. The concept of forecasting or mentally rehearsing a non-executed movement is particularly interesting when considering the presence of ipsilateral activity in the motor cortex. Evidence for mental rehearsal of actions has been shown in the activity of PMd [46, 211] and M1 neurons [77, 211]. An intriguing idea is that activity in these areas during unimanual movements of the ipsilateral arm reflects a covert motor command simulating what the non-selected (contralateral) arm would do if it were performing the action.

During delay periods prior to movement when requirements for the upcoming action are left ambiguous, activity associated with multiple potential actions has been

observed in motor cortical activity [47, 50, 111]. Even after action selection, activity related to this parallel action may continue covertly throughout movement and yield spurious “ipsilateral” modulation that is truly related to the contralateral limb. Notably, however, preparatory activity appears to be largely effector-independent, reflecting the action and its visual consequences rather than actual implementation [45, 165]. Nonetheless, effector-independent preparatory signals could seed the evolution of multiple parallel movements, with effector selection happening independently. If this were the case, one would expect to see correlated activity profiles for movements of the two arms. This has been observed to some extent within PMd neurons [45], and is more commonly reported using macroscopic methods like functional magnetic resonance imaging (fMRI) [69] and electrocorticography (ECoG) [30]. However, correlations between ipsilateral and contralateral tuning in M1 neurons have typically been reported as weak or absent [45, 108, 204].

Sensory modulation of motor cortical activity may also be important for bimanual coordination. Neurons in M1 have been shown to integrate multi-joint sensory information to produce appropriate feedback responses following joint-torque perturbations [177]. In much the same way, activity in a single hemisphere that ultimately drives movement of the contralateral limb may integrate state information regarding both limbs in order to coordinate bilateral interaction. Indeed, signals in primary motor and caudal premotor cortex during bimanual behavior have been shown to reflect non-linear interactions of unimanual signals [69, 72, 133]. The state of one limb therefore alters the relationship between neural activity and behavior for the other limb, a feature critical for coordination.

In the current study, we devised an experiment to test the sensitivity of unimanual reaching signals to posture of the stationary hand. Two macaque monkeys were trained in a variant of an instructed-delay reaching task where the two hands began each trial in one of three starting configurations: left hand eccentric with right hand centered, both hands centered, or left hand centered with right hand eccentric. We found that neural activity associated with reaching movements was sensitive to the current posture of the stationary hand, despite the fact that it was not involved in the movement itself. Classifiers trained to predict reach targets incurred a generalization cost when tested on trials where the stationary hand was held in a different posture than the trials used for model training. Errors showed a tendency to happen in the same direction as the change in stationary hand posture, suggesting that target coding shifted congruently with the postural state. Furthermore, we isolated components of the population response that captured the interaction between posture and target coding. Separation of target traces along these components largely reflected the horizontal geometry of physical target locations and revealed a preferential impact on reaches to the contralateral space. These results are in keeping with a covert motor command that

quietly simulates actions of the non-selected hand, yet some other results were inconsistent with this hypothesis as well. However, the findings may also be interpreted within the context of bimanual coordination, where bilateral postural state is integrated into the motor commands for a single limb, or whole-body control, where mechanical properties of the body that are altered by postural changes must be accounted for in the reach command.

3.2 Results

3.2.1 Task design

The data reported in this chapter were collected during the same recording sessions as those of the previous chapter (see Methods 3.4.1 and Chapter 2 for dataset details). We introduced a manipulation to the basic instructed-delay reaching task for the purpose of investigating the interaction between passive posture of the stationary arm and population coding of reaching signals. At the start of each trial, monkeys placed their hands in one of 3 different starting configurations: left hand eccentric with right hand centered, both hands centered, or left hand centered with right hand eccentric (Figure 3.1B, corresponding panels ordered left to right).

This modified posture creates at least 3 changes that are potentially relevant to the computations being performed in the motor cortex during unimanual reaching. First, this postural change may simply modulate activity to represent the current state of each limb. Another possibility that expands upon the previous one is that the relative state of both limbs is pertinent to the population reaching signals, potentially serving a role in bimanual coordination or whole-body control. Lastly, the lateral shift from centered to eccentric starting position greatly alters each of the target vectors for the hand that occupies the new starting position (Figure 3.1B, boxed schematic). Consider, for example, target vectors for the top-right target (Figure 3.1B, red lines). Shifting to the eccentric posture induces a 79° change in target vectors, and the vector connecting the eccentric position to the top-right target becomes identical to the vector connecting the centered position to the top-middle target. If some degree of activity arises in association with the necessary movements for both arms to reach towards a single target, then the postural shift would therefore alter that activity as well.

3.2.2 Single-unit responses

We first broadly characterized the tuning properties of single units. Units were categorized as being tuned to movements of the ipsilateral arm, contralateral arm, or both arms in each phase. Note that tuning is not the same as modulation as described in the previous chapter; tuning requires that the firing rates for at least one target be significantly different from the rest. The proportions of units in each tuning category are

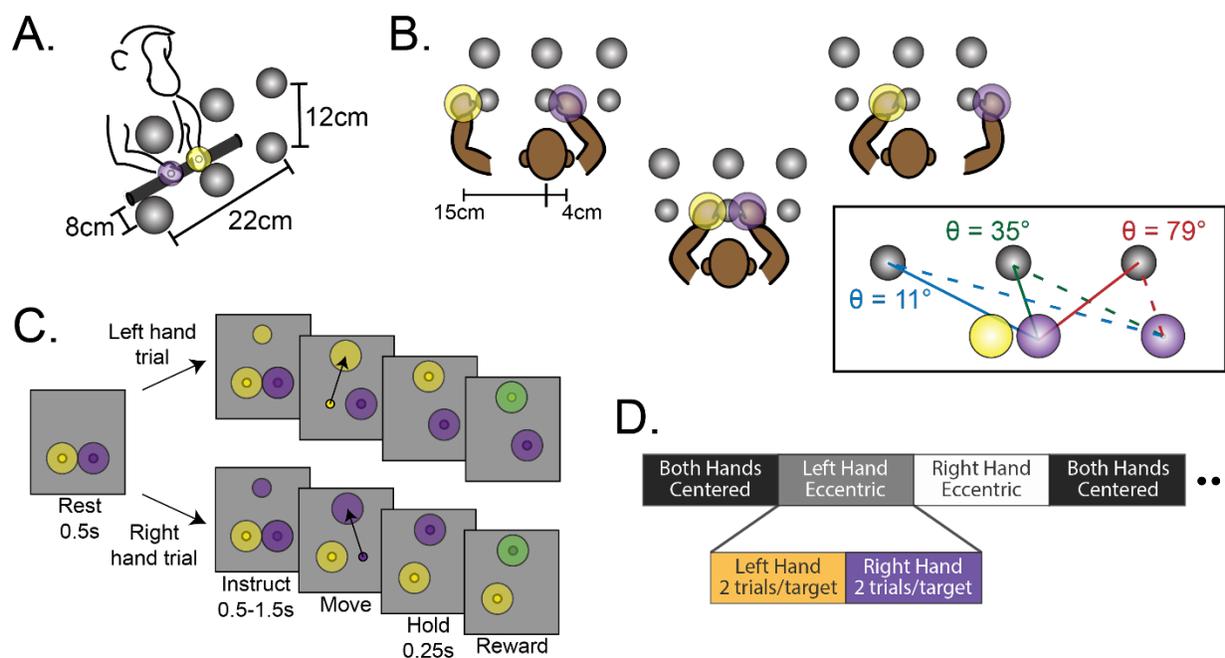


Figure 3.1. Task design. (A) Monkeys reached to one of six targets arranged in a 2x3 fronto-parallel grid. (B) At the start of each trial, monkeys placed their hands in one of three starting configurations: left hand eccentric with right hand centered (left panel), both hands centered (middle panel), or left hand centered with right hand eccentric (right panel). Note that when the stationary hand was in the centered position only 5 targets were included in the possible target set; the bottom target in the contralateral space was removed to prevent the two hands from coming in close proximity with one another. The boxed image indicates the angular difference between target vectors originating from the centered or eccentric starting positions for each of the targets in the top row. (C) Each trial consisted of 3 phases (repeated from Figure 2.1B). Trials were initiated by placing both hands in their color-coded starting positions for 500ms (Rest). A small target then appeared at the location of the future reach in a color that indicated which hand to use. The monkey remained still during cue presentation for 0.5-1.5s (Instruct). The start target for the reaching hand then disappeared while the reach target enlarged to cue movement (Move). (D) Starting configurations and hand assignments were organized in a nested block structure. Using a single starting configuration, monkeys would perform 2 trials per target for the left hand then 2 trials per target for the right hand. The starting configuration would then change, and the blocks of left- and right-hand reaches would repeat. The order of reach target locations was randomized within each block.

summarized in Table S3.1. As expected, Instruct phase tuning was most prominent in PMd units. 65%/61% of PMd units were tuned to movements of at least one arm during the Instruct phase compared to 56%/25% of units in M1 for monkeys O and W, respectively. These proportions became approximately equal during the Move phase. Importantly, many units were tuned to movements of both arms. The proportion of bilaterally tuned units tended to be slightly higher in PMd, particularly during the Instruct phase. 26%/19% of units in PMd were bilaterally tuned during the Instruct phase, compared to 22%/2% in M1. Both animals exhibited an increase in the number of bilaterally tuned units within each region during the Move phase (PMd: 63%/40%, M1: 63%/33%). These results are in broad agreement with the existing literature [8, 45, 108].

For all tuning profiles that were significant, we also tested whether they were significantly different across the two postures of the stationary arm (Figure S3.1). Despite the stationary hand being uninvolved in the movement, many tuning profiles differed across postures. We refrain from comparing the proportion of significantly altered tuning profiles across phases or brain regions, as there are various reasons why differences in proportions could exist, each with very different interpretations. For example, a larger proportion may indicate larger magnitude shifts in tuning, or simply reflect more precise tuning profiles in each posture that are better statistically powered for determining a significant difference. The core observation is simply that these tuning profiles do shift, a phenomenon that we will investigate more thoroughly at the population level.

Finally, we performed an analysis of the similarity between tuning profiles for the ipsi- and contralateral arms in all units that were bilaterally tuned. Target activations were modestly correlated across the two arms (Figure S3.2). We assessed correlated tuning in both extrinsic and intrinsic reference frames by reflecting target labels across the midline, but no clear and consistent preference for either reference frame was observed (Figure S3.2.E). In general, the correlation did not differ strongly across task phases or brain regions either, with Spearman’s correlation coefficients typically falling in the range of 0.3-0.5.

3.2.3 Target classification using population activity is sensitive to posture of the stationary hand

We next analyzed the influence of stationary arm posture on population-level reaching signals. It is expected that the population state may be impacted by the passive posture of the limbs, yet it is unclear whether that impact will be relevant to readouts of the active behavior. To test this, we trained linear discriminant analysis (LDA) classifiers to predict which targets monkeys reached to on individual trials (Figure 3.2). These models identify axes in neural state space that separate the different

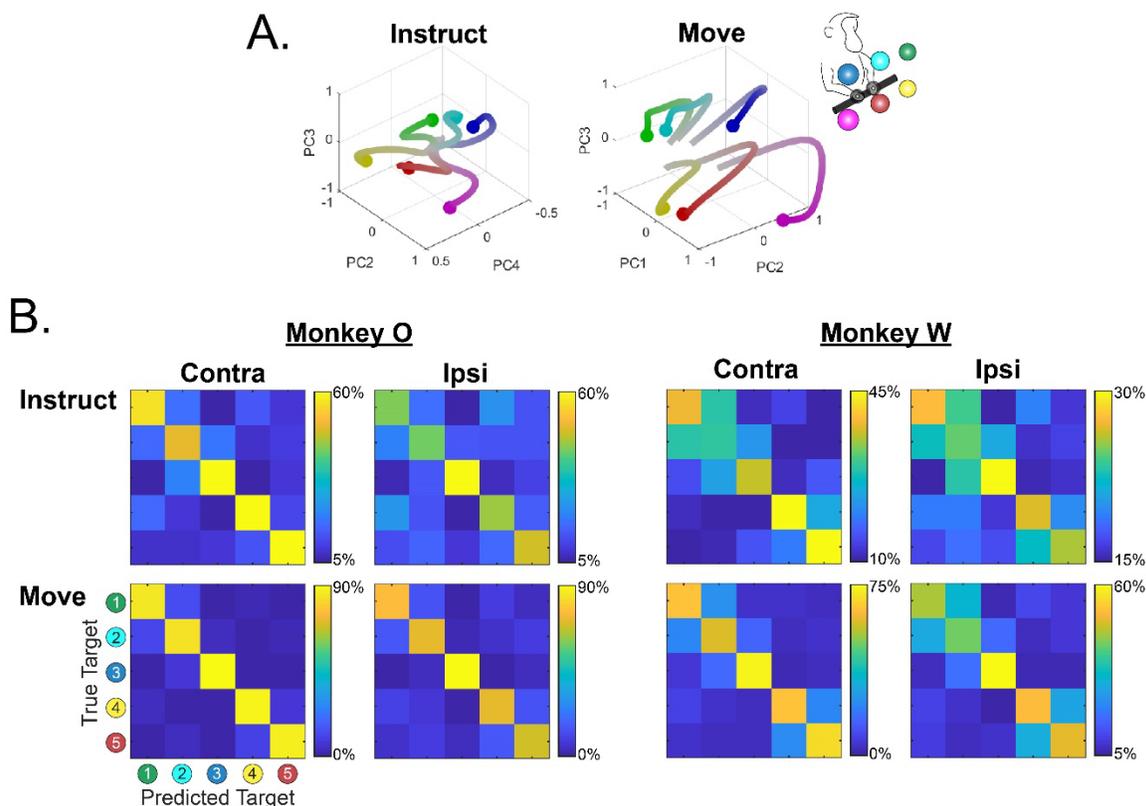


Figure 3.2. Neural population activity reflects the geometry of reaching targets during instructed-delay and movement. (A) PCA projections of trial-averaged neural activity from M1 for each target during the Instruct (left panel) and Move phase (right panel). The Instruct phase plot shows data from 100ms before onset of the instructional cue to 500ms after. The Move phase plot shows data from 200ms before movement onset to 300ms after. The traces in both plots fade from grey to the corresponding target color to indicate the advancement of time. The 3D perspectives were manually rotated to optimize the view of target separation for comparison with the physical geometry of the reach targets (displayed in the cartoon on the right with corresponding target colors). PC's 2-4 were displayed for the Instruct example, as PC1 appeared largely target invariant. These examples are from the right hemisphere of Monkey O during left arm trials and include both M1 and PMd activity. (B) Confusion matrices for cross-validated target classification using LDA. The color of cells within a subplot indicates the row-normalized frequency of a predicted target for a given true target. The target labels in the bottom left subplot are color coded to the corresponding targets in the cartoon of panel A. Classifiers were only created for trials where the reaching hand began in the centered start position, and separate classifiers were trained for each posture of the stationary hand. Right hand target labels were mirrored to combine with left hand trials. Both M1 and PMd activity was included in each model. Note that target 6 was omitted, as this target was unreachable when the stationary hand was held in a centered position. Note also that each subplot uses a different colorbar scale.

reach targets, and therefore reflect dimensions of the signal that distinguish movement types and may ignore changes in neural activity that are common to all movements like stationary arm posture. To determine whether these classifier models were sensitive to posture of the stationary arm, we implemented a model generalization testing approach. The same testing set of neural data was used to evaluate models trained using either 1. held-out trials with the same posture of stationary hand (“Within” training condition), or 2. trials with the stationary hand in the other posture (“Across” training condition). A drop in performance between these two training conditions would indicate an interaction between posture of the stationary arm and target coding.

Models were trained using populations from a single hemisphere so that ipsi- and contralateral sensitivity could be analyzed separately and compared. Models were also trained separately for each task phase. Units from PMd and M1 were combined. Accuracy of target predictions for each of these conditions is summarized in Figure 3.3A, with each pair of connected datapoints indicating the same testing set data evaluated in each of the two training conditions. We found that training models using a different posture of the stationary hand did in fact incur a generalization cost (Figure 3.3B). We analyzed the effects on model accuracy using a permutation-based repeated measures ANOVA for the following factors: Hand (2 levels – Ipsi, Contra), Phase (2 levels – Instruct, Move), and Training Condition (2 levels – Within, Across). The three-way interaction was significant in both monkeys (monkey O: $p=0.042$, monkey W: $p=0.025$), so follow-up two-way tests were performed.

For ipsilateral classifiers, there was a significant two-way interaction between Phase and Training Condition (monkey O: $p_{\text{Ipsi}}=8.2e-3$, monkey W: $p_{\text{Ipsi}}=2.9e-3$). Across-posture classifiers performed worse than Within-posture classifiers during both phases (one-way test – monkey O: $p_{\text{Ipsi,Instruct}}=1.0e-5$, $p_{\text{Ipsi,Move}}=1.0e-5$; monkey W: $p_{\text{Ipsi,Instruct}}=0.037$, $p_{\text{Ipsi,Move}}=2.0e-4$), but the generalization cost was greater during the Move phase. On average, the performance drop was -5.9%/-4.5% during the Instruct phase and -10.5%/-11.9% during the Move phase. For contralateral classifiers, there was no interaction between Phase and Training Condition for either animal (monkey O: $p=0.90$, monkey W: $p=0.98$), and a performance drop across arm postures was only significant for monkey O (monkey O: $p_{\text{Contra}}=2.1e-4$; monkey W: $p_{\text{Contra}}=0.24$). However, this performance drop for contralateral classifiers in monkey O was less than the drop for ipsilateral classifiers in both phases, averaging -3.0% during Instruct and -3.2% during Move. The interaction effect representing this differential sensitivity was significant for the Move phase, but just below the threshold for significance in the Instruct phase (monkey O, Hand x Training Condition interaction – $p_{\text{Instruct}}=0.057$, $p_{\text{Move}}=4.4e-4$).

In summary, changes in posture of the stationary hand counterintuitively affected classification of reach targets on individual trials. Ipsilateral signals were most sensitive

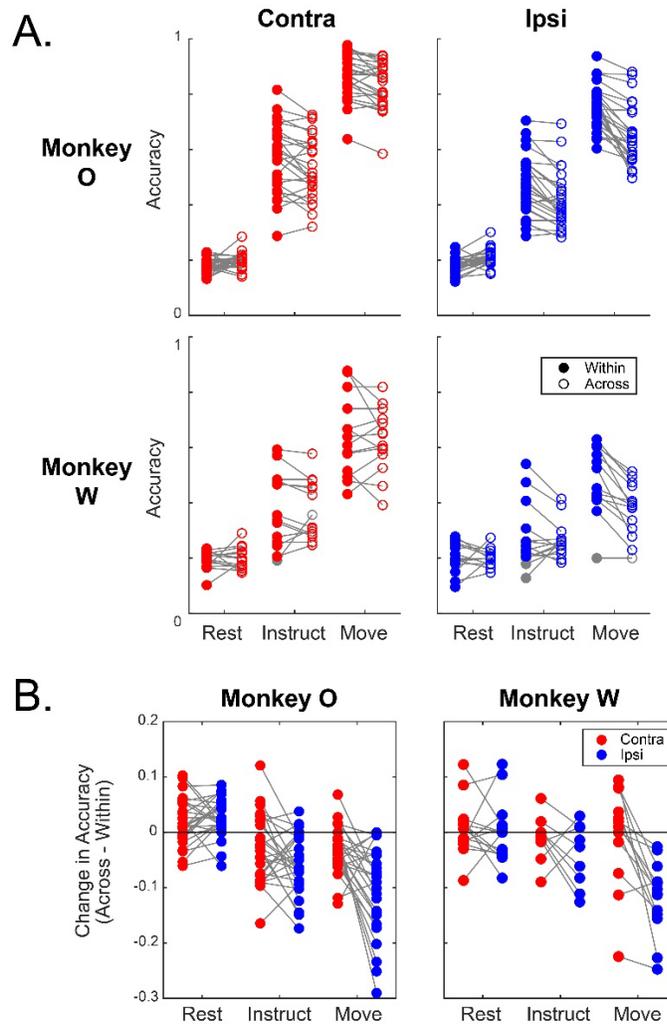


Figure 3.3. Posture of the stationary hand impacts the ability to classify reach targets. (A) Performance of LDA target classifiers across postures of the stationary hand. Models were trained to predict reach targets using neural data from both M1 and PMd in a single hemisphere. Classifiers were trained on either trials with the stationary hand in the same posture as the test set (“Within” closed circles) or on trials of the other posture (“Across” open circles) using fully cross-validated methods to make the two measures comparable. Each pair of datapoints connected by grey lines corresponds to data from a single recording session. Only the 5 targets that were in the target set for both stationary hand postures were used in the models, making chance level 0.2. Because classifiers at or below this value would not contain meaningful information about model generalization, any datapoints meeting those criteria in the “Within” condition were removed from statistical analysis but remain plotted in grey. (B) Generalization cost across stationary hand postures. The drop in accuracy between the two model training conditions (Across - Within) was computed for each set of datapoints to measure the generalization cost associated with switching postures of the stationary hand. Pairs of datapoints connected by grey lines represent the same set of units recorded in a single recording session from a single hemisphere evaluated during either ipsi- (blue dots) or contralateral (red dots) arm trials.

to this perturbation, incurring a greater generalization cost when models were trained and tested on different postures. The effect was also most prominent during the movement itself, as compared to the instructed-delay period.

3.2.4 Target coding shifts congruently with shifts in the posture of the stationary hand

In order to investigate the specific ways in which classifier performance was being affected, we analyzed the confusion matrices associated with these models. We computed the relative change in prediction frequencies for each of the targets between the two training conditions and plotted the results in Figure S3.3. For example, if neural activity during trials for target 2 was more frequently predicted as target 1 when the classifier was trained using a different stationary hand posture, then we would observe an increase in the matrix cell corresponding to True Target 2 and Predicted Target 1.

We were specifically interested in the direction of shifts in these erroneous predictions, and therefore summarized the results by computing changes in frequency for target misclassifications in each direction: high, low, same direction as the postural shift, or opposite direction of the postural shift. These frequencies represented predictions of the closest target in the given direction from the true target. Consider, for example, target predictions for left arm movements when the stationary right arm is held in the centered position (Figure 3.4A). The Across-posture condition would represent predictions for a model trained with the stationary right hand shifted rightward to the eccentric position. Therefore, if the top-middle target were misclassified as the top-right target, it would be considered a miss in the same direction as the postural shift. The resulting changes in directional miss frequencies that we observed are displayed in Figure 3.4B. We were particularly interested in whether predictions shifted horizontally in the same direction as the postural shift, as this is what would be predicted if a component of the population signal reflected the movement that the stationary arm would perform if it were selected for action. Additionally, this feature could point to coding of relative arm states or bilateral state-dependent modulation that could be important for bimanual coordination.

As in the previous section, we performed a permutation-based repeated measures ANOVA for the following factors: Hand (2 levels – Ipsi, Contra), Phase (2 levels – Instruct, Move), and Miss Direction (2 levels – Same, Opposite). No significant interactions were found involving the Miss Direction factor; however, a main effect of Miss Direction was observed for both monkeys (monkey O: $p=3.0e-5$; monkey W: $p=1.0e-5$). The change in prediction frequency was higher for misses in the same direction as the postural shift than it was for the opposite direction. In fact, while misses in the same direction as the postural shift increased on average (monkey O: $\mu=3.4\%$; monkey W: $\mu=2.6\%$; mean taken across hands and phases), across-posture classifiers

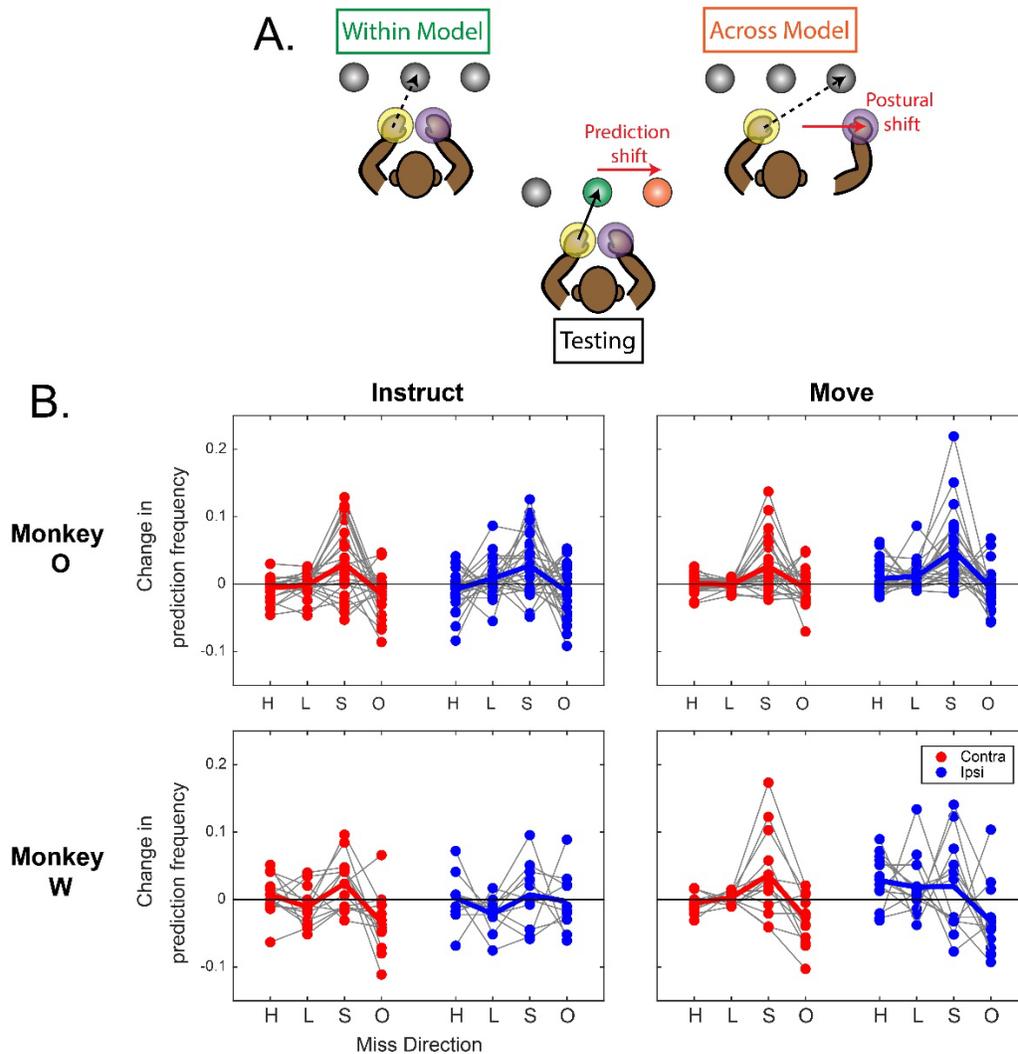


Figure 3.4. Target coding shifts congruently with changes in stationary arm posture. (A) Schematic illustrating misclassification directionality. Consider classifiers of left arm reaches tested on trials with the stationary right hand in the centered posture. Classifiers are trained using either trials of the same posture (“Within Model”) or with the stationary hand shifted rightward to the eccentric position (“Across Model”). In this example, the two trained classifiers have learned to associate the same neural state to 2 different reach targets, indicated by the dotted lines in the top panels. When tested, the Within Model predicts the correct target (green) while the Across Model predicts the target to the right (orange). Since the postural and prediction shifts are in the same direction, this Across Model prediction is labeled as a same direction miss. (B) Each plot displays the change in frequency for misclassification in each of 4 different directions: high (H), Low (L), same direction as the postural shift (S), and opposite direction of the postural shift (O). To construct these plots, the number of times that a classifier predicted the target adjacent to the true target in the specified direction was computed and divided by the total number of trials. The plotted values represent the difference in these proportions between the Within-posture and Across-posture classifier conditions. A positive value indicates an increase in the directional miss frequency in the Across-posture case.

actually became less likely to miss in the opposite direction of the shift (monkey O: $\mu=-0.9\%$; monkey W: $\mu=-1.9\%$). These results suggest that across both hands and phases, target coding at the population level shifts congruently with changes in the posture of the stationary hand.

3.2.5 Target classifiers fail to generalize across the two hands

Since target coding displayed a tendency to shift congruently with the posture of the stationary hand, it is possible that a component of the population signal reflects what that stationary hand would do if it had been selected for execution of the action. This component of the signal would be present during movements of either arm, albeit likely greater in magnitude when it represents the selected movement. Therefore, an additional prediction of this hypothesis is that target coding will show some level of generalization across the two arms.

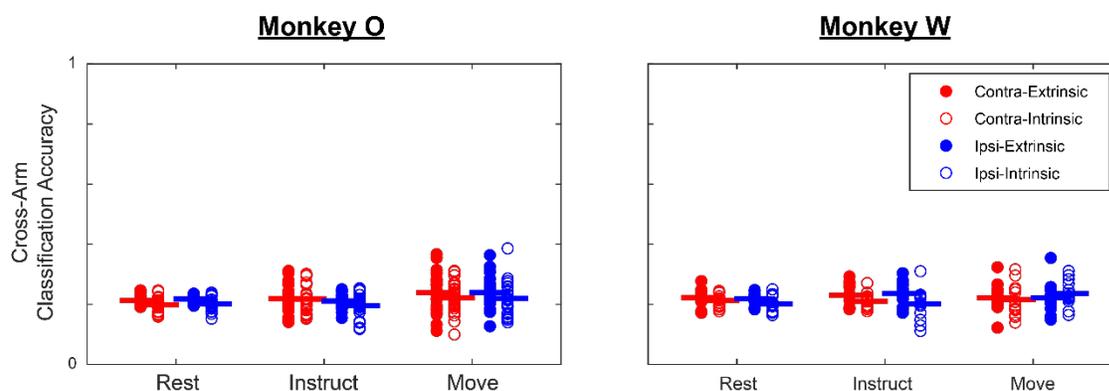


Figure 3.5. Target classifiers fail to generalize across the hands. Classifiers were trained to predict movements of one hand using activity from one hemisphere, then tested on trials of the other hand. Plot colors indicate which hand the classifiers were tested on (red: contra, blue: ipsi). Generalization was tested in both extrinsic and intrinsic reference frames, where target labels were mirrored across the midline to obtain predictions in the intrinsic space. Filled circles (left side of each set) represent the extrinsic reference frame, and open circles (right side of each set) represent the intrinsic reference frame. Each datapoint represents a single recording, and thick horizontal lines indicate the mean across recordings.

We tested this prediction by training LDA models to predict movements of one hand and testing them on the other. Models were evaluated in both directions: trained on ipsilateral movements, tested on contra, and vice-versa. We also evaluated generalization in both extrinsic and intrinsic (mirrored target labels) reference frames. These classifiers failed to generalize across the arms in all circumstances (Figure 3.5). Target-specific modulation in both the Instruct and Move phases caused cross-arm classification accuracy to become more variable simply because there was more variance in the signals. This is indicated by the increase in spread during the Instruct and Move phases in Figure 3.5. However, the relationship between this modulation and the behavior was not strongly preserved across the two arms and thus resulted in mean accuracy being similar to classifiers trained using neural data prior to target presentation (Rest phase). These results suggest that the associations between neural activity and behavior that are modeled by these LDA classifiers are not strongly maintained across the two arms.

3.2.6 Analysis of the interactions between target and posture coding using dPCA

To further analyze the interactions between target and posture coding, we used demixed principal components analysis (dPCA) [136]. dPCA attempts to decompose the population signal into components that each exclusively represent variance associated with a single task parameter or set of parameter interaction terms. We applied the method using 3 task parameters: Time (condition-independent), Target, and Posture (of the stationary arm). The marginalization procedure yields both main effect and interaction terms, which we grouped into 4 categories: Time (only Time term), Target (Target, Time x Target), Static Posture (only Posture term), and Posture Interaction (all interactions involving Posture). Our primary interest was in the Posture Interaction group, as these components would be those responsible for the generalization cost observed during LDA classification.

Unit data were trial-averaged, normalized, and combined across recording sessions for this method. Since we were no longer limited by the number of units simultaneously recorded during a single session (which was the case using LDA), our unit counts were effectively larger and allowed us to separate PMd and M1 as well as data from each hemisphere in order to analyze the areas independently. Each model was therefore trained using data from a single monkey, a single area (PMd/M1), a single hemisphere, and a single hand (ipsi/contra). The number of components was selected independently for each model as the minimum number that provided at least one component in each term category and a total of 8 condition-specific (non-Time) components [123].

Results from two example datasets are displayed in Figure 3.6, using PMd activity in either hemisphere of monkey O during ipsilateral reaches. The remaining datasets are displayed in the supplement (Figure S3.3-4). As designed, the model identified components that primarily captured variance for specific marginalizations of the data. The bars in Figure 3.6A-B display the proportion of total variance accounted for by each component. These bars are actually composed of horizontally stacked segments that represent the variance from each term source, yet each bar is primarily a single color as a result of the demixing. The models ranged in size from 12-18 components and captured an average of 88% of the total variance. As in previous reports in forelimb motor tasks (e.g. [123]), the majority of variance was related to the condition-independent Time term (average 60% of total; Figure 3.6A-B, Figure S3.3-4, blue bars). On average, target components explained the next most variance (average 20% of total; Figure 3.6A-B, Figure S3.3-4, orange bars).

The amount of variance captured along Static Posture dimensions was greater for ipsilateral arm movements but did not differ substantially across PMd and M1. These components accounted for an average of 10% in Ipsi PMd, 1% in Contra PMd, 11% in Ipsi M1, 2% in Contra M1 (averages include 4 samples: 2 hemispheres from each monkey). The Posture Interaction dimensions followed this same trend (Ipsi>Contra) and captured the least variance, averaging 3% in Ipsi PMd, 1% in Contra PMd, 2% in Ipsi M1, 1% in Contra M1. To consider the relative impact on movement-specific activity patterns, we also calculated the ratio of variance captured across Target dimensions and Posture Interaction dimensions. This ratio, expressed as Posture Interaction divided by Target, averaged 0.18 in Ipsi PMd, 0.04 in Contra PMd, 0.13 in Ipsi M1, 0.03 in Contra M1. These results complement our previous observation that LDA target classifiers were more sensitive to posture when classifying ipsilateral reaches.

The method by which we selected the number of components ensured that there would be Posture Interaction dimensions in our models, but what we were most interested in was precisely how the relationship between target signals changed as a function of stationary hand posture, and how this relates to the physical geometries of the target configuration and postural states. A general trend across the datasets is illustrated in the bottom plots of Figure 3.6C-D where projections onto the top Posture Interaction components from two example models are displayed. During Rest, the two postural states were shifted with a simple DC offset from each other (left of vertical red dotted line). Following instruction (right of vertical red dotted lines), the neural states along this dimension deviated in opposite directions for the two postures, either converging on each other or even becoming mirrored with respect to their Rest phase relationship. Leading up to and during movement (vertical blue dotted lines), target traces began to separate from each other in a way that reflected their horizontal geometry. This resulted in movements toward the contralateral space being most

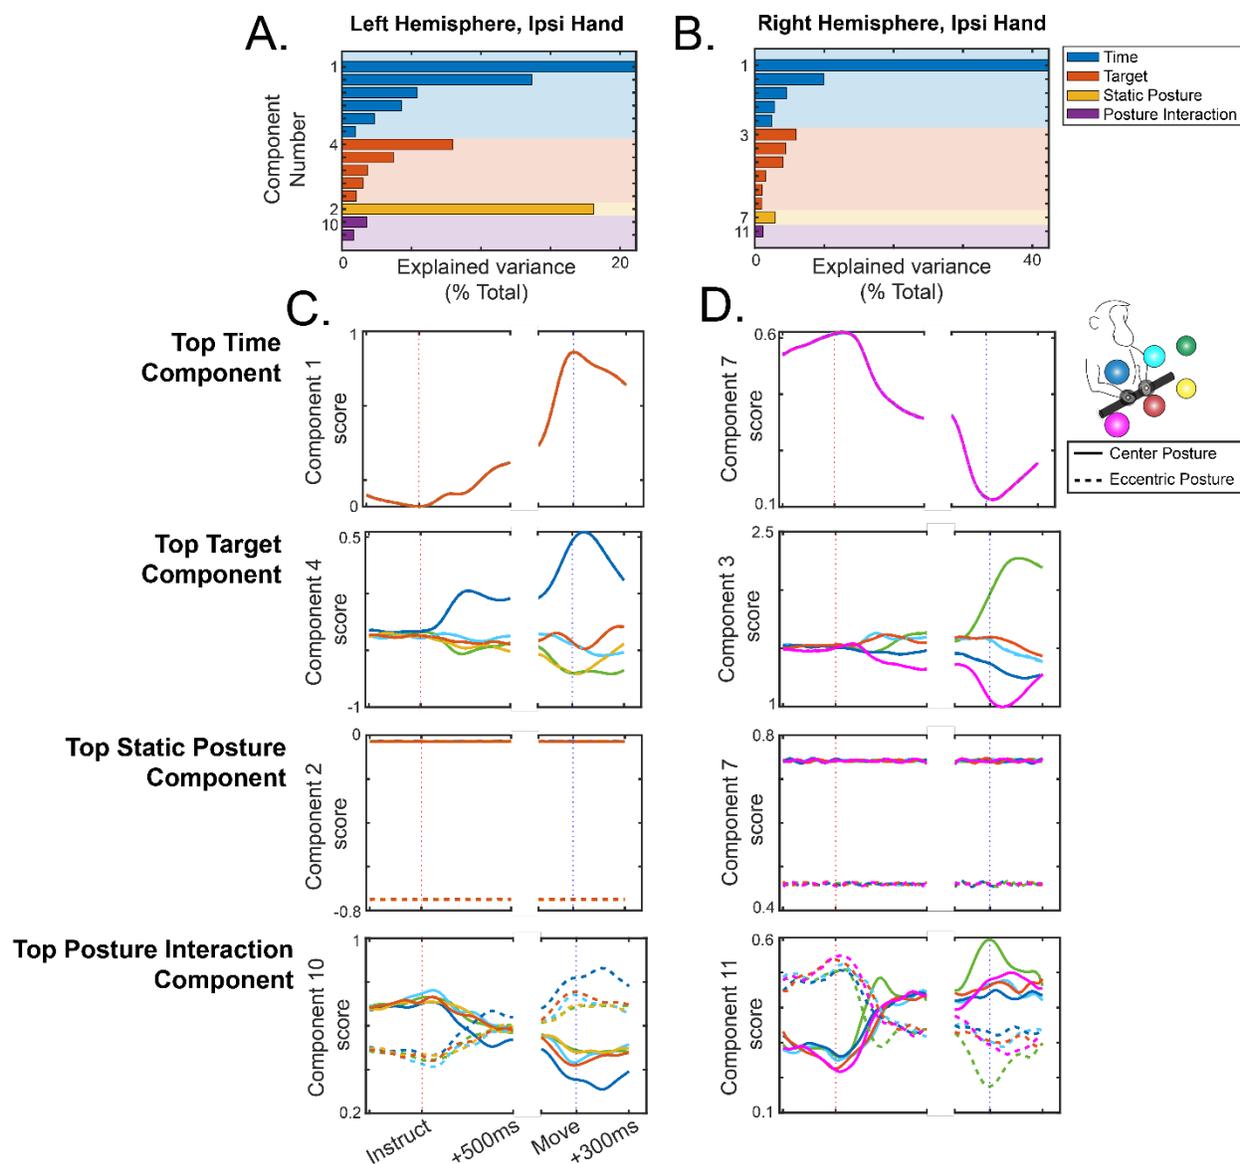


Figure 3.6. Analysis of the interactions between posture and movement coding using dPCA. Example dPCA results from PMd of monkey O during ipsilateral reaching trials. The left column uses the left hemisphere (ipsi=left hand) and the right column uses the right hemisphere (ipsi=right hand). (A-B) The proportion of total variance explained by each component for the left hemisphere (A) and right hemisphere (B) models. Each color represents a certain set of marginalization terms. Component numbers reflect the ranked order based on variance accounted for, and only the top component from each grouping is labeled. The bars are actually composed of horizontally stacked segments that represent the variance from each term source, yet each bar is primarily a single color as a result of the demixing. (C-D) Projections onto the top component for each term set for the left hemisphere (C) and right hemisphere (D) models. Each plot shows all 10 signal traces (2 stationary arm postures x 5 targets). Solid lines represent the centered stationary arm posture and dashed lines represent the eccentric posture. Target traces are color coded to the cartoon in the top right.

affected (Figure 3.6C bottom plot blue lines, 3.6D bottom plot green lines), which is indicated in the plots by the deviation between centered and eccentric postures for corresponding target traces. We note that these reaches towards the contralateral space correspond to the targets that experience the greatest angular change in the vectors connecting the stationary hand position to target position in physical space (Figure 3.1B, boxed schematic).

Target separation along the top Posture Interaction components was rarely (and weakly) observed during the Instruct period, as in the examples of Figure 3.6. Quantitative comparisons between ipsi- and contralateral movements or PMd and M1 have not yet been performed, although the same trends appear to be present across most of the models. A notable exception was PMd activity for monkey W, where Posture Interaction Components appeared noisier and only very weakly showed the mirroring properties, if at all (Figure S3.3, right column). This may be due to the lower cell counts entered into those models (left hemisphere: 66 units, right hemisphere: 73 units), which should be taken into consideration during interpretation.

3.3 Discussion

In the present study we have demonstrated that reaching signals at the single-unit and population levels are sensitive to bilateral arm posture. Classifiers trained to predict reaching targets incurred a generalization cost when the posture of the stationary hand was different between training and testing sets. This sensitivity in the mapping between neural state and behavior was more prominent for activity ipsilateral to the reaching hand and contralateral to the stationary hand that experienced the postural shift. The direction of classification errors shifted congruently with the changes in the postural state. Components of the population response capturing the interaction between posture and target coding primarily affected the horizontal dimension with preferential impact on reaches to the contralateral space.

3.3.1 Interpretation through the lens of bimanual coordination

The motor cortex secures its name by being the cortical region with the most projections to spinal motor neurons (including pre- and primary sub-regions [73]) and requires the lowest amplitude stimulation for evoking movements [170]. However, it has long been known that the motor cortex exhibits significant sensory responses as well [84, 106, 198]. There is evidence that motor cortical function is modulated by limb state [99, 189, 190]. Sensory signals in the motor cortex have been discussed in terms of long-latency reflex responses and feedback control capable of integrating information across the joints to provide appropriate responses [98, 177]. One might similarly ask whether

the motor cortex integrates state information across the two arms to provide appropriate motor output during bimanual coordination.

We report that the posture of the stationary limb influences the mapping between neural activity and reaching movements in the motor cortex, as captured by LDA classifier models. During active movement of both limbs, previous studies have identified distinct bimanual activity patterns in the motor cortex that are not simple combinations of what is observed during unimanual behavior [69, 72, 133]. While evidence of these bimanual patterns has been shown in both M1 and PMd, it has been recently argued that neural populations closer to the central sulcus maintain stronger independence between signals for the two arms [55]. Rather than looking at interactions between two active movement signals, our study instead investigates the interaction between active reaching signals of a single arm and the passive state of the other arm. It is important to note that we cannot definitively say whether the different postural states represent a truly passive manipulation, as we did not record EMG and cannot rule out differences in muscle activity for active maintenance of the two postures. However, we may speak more generally in terms of the “state” of the limb, which may include efference as well as afference that do not change over time. A recent imaging study has shown that although generic suppressive responses in the ipsilateral motor cortex are observed during both active and passive finger movements, only active movement produces activity specific to the movement type (i.e. finger-specific [16]). Our results nicely complement this work, suggesting that although passive changes in ipsilateral limb state may not induce movement-specific activity patterns, these changes may still impact the functional organization of the circuit. A potential future analysis for drawing further parallel with this imaging study would be to test performance in classifying postural states during the Rest period using ipsilateral or contralateral activity. If the results for finger movements translated to arm postures, then we would expect weak or non-existent performance from the ipsilateral hemisphere.

We now discuss the two complementary results that: (1) classification errors when generalizing across postures tend to occur congruently with the postural shift (Figure 3.5), and (2) neural dimensions that capture the interaction between target and posture coding reflect the horizontal geometry of reaches, and preferentially capture signal changes during reaches to the contralateral space (Figure 3.6). These results may have a few different interpretations within the lens of bimanual coordination. First, it stands to reason that signal interactions between the two arms would be most pertinent when they are close together in physical space. For two effectors to interact during behavior, they must be close together. The biased effect during reaches to the contralateral space may reflect this concept: reaches away from the stationary hand are relatively unaffected by its postural state, since they will not be interacting, but reaches toward the stationary hand experience a greater modulation due to posture, since they

are entering into the same workspace. The signal interactions that we observe could also speak to a form of whole-body control. Since the two arms are physically linked through the torso, changing the posture of the stationary hand would impact mechanical properties of the system as a whole. These changes in the mechanical properties of the linked body may require modification of the reaching command, since the rotational forces induced by reaching will be experienced differently after the postural adjustment. These final ideas, while reasonable, should be considered speculative until more specifically tested.

3.3.2 Covert motor commands for the stationary arm

Covert motor commands, which we operationally define as any activity patterns that mirror true movement while the limb remains stationary, have been reported in several forms. “Mirror” neurons display an intriguing behavior whereby similar patterns of activity are produced during both action and action-observation. These mirror neurons were first identified in premotor cortex [70], but have also been reported in M1 [211]. The activity of these cells may also be related to the concept of mental rehearsal, where a non-executed action is simulated in neural circuits. Indeed, it is possible to train decoders for brain-machine interfaces (BMI) control using activity during passive observation of a task [206], suggesting that those activity patterns are similar to what the subject may naturally attempt to produce when controlling the BMI.

Might it be possible that circuits in the motor cortex could produce covert motor commands while simultaneously executing another? This could provide an account of motor signals ipsilateral to the active hand as being weakened simulations of what the contralateral (stationary) hand would do should it be selected for action. Our result that unimanual reaching signals are sensitive to posture of the stationary arm would align with this sort of account. Changing the position of the stationary arm would alter the implied reach trajectory to each target (Figure 3.1B) and would therefore modify a putative simulation of the neural activity required to execute that trajectory. We also observed that target classification errors tended to shift congruently with the posture of the stationary hand (Figure 3.4B) and the same targets that would be expected to see the greatest change in implied trajectory were the ones that were most distinguished along neural dimensions that captured posture and target coding interactions (Figure 3.6).

This hypothesis would rely upon the same neurons being involved in movements of either arm, having the same response properties for both arms as well. If some component of the activity during ipsilateral arm movements were a weakened simulation of what the contralateral arm would do, then one would expect to see similar patterns of activity regardless of which arm was selected. We did observe many bilaterally tuned

units (Table S3.1) but tuning across those sub-populations was only modestly correlated for the two arms and didn't show a preference for a particular reference frame (Figure S3.2). Additionally, LDA classifiers failed to generalize across the two arms (Figure 3.5). However, it is important to note that changes in the magnitude of the signal may cause these classifiers to fail, rather than the patterns per se, i.e. similar patterns underneath a global suppression of activity might show similar results. Simultaneous representations of multiple different action choices have been previously reported in PMd neurons [47, 50], but only during periods where the upcoming action requirement is left ambiguous to the subject. In line with what one might expect, after a cue arrives to indicate which action to choose, activity associated with the non-selected action fades away. Our results therefore offer some support for a covert motor command hypothesis, but do not produce other key predictions that, together with the existing literature, make it dubious to assign such a specific meaning to the presence of ipsilateral activity.

3.3.3 Additional considerations and future analysis ideas

In the current work, dPCA results have been presented in a largely qualitative manner. The trends can be seen visually but should be captured within a quantitative measure that can be further compared across brain areas (PMd/M1), laterality (ipsi/contra), and task phase (Instruct/Move). A first step may be to compute the root mean square (RMS) distance between target traces for the two postures in the Posture Interaction components to measure the relative impact that postural state has on each target along these neural dimensions. Additionally, consider the fact that components from different term categories are designed to be uncorrelated, but not necessarily orthogonal. Measuring the actual alignment between these dimensions could yield additional insight as to how the postural state is affecting target readouts from LDA. For example, if posture and target components are not orthogonal, then their additive effect could impact target coding by translating the neural state of all targets in one posture into a region of state space that is closest to the neural state associated with a single target in the other posture. Additionally, Laterality could be added as a parameter to the dPCA marginalization to examine which aspects of the task are limb-dependent [55, 223]. In general, it may be possible to further tap the dPCA results for information regarding the geometric changes in target coding that result from postural changes of the stationary arm.

In a similar vein, our directional miss results using LDA classifiers (Figure 3.4) could be nicely complemented by regressing the reach vectors in continuous space on the population neural activity. This may allow us to investigate the specific geometry of the effects induced by changing the stationary arm posture without being constrained to target representations.

An additional consideration to make is whether the change in postural state induces target-specific micromovements of the stationary hand. We have found that the stationary hand does move very slightly during the Move phase (Figure 2.1E, Results 2.2.1). EMG activity was not recorded, but these movements did not seem target-specific in nature. However, this has not been exhaustively tested. The most plausible scenario in which stationary arm movements could be influencing our results can be described as follows: Movements to the contralateral space come in closer proximity to the stationary hand than those to the ipsilateral space and induce some resulting movement in the stationary hand that spuriously codes for the horizontal target position. When the stationary hand is moved to the eccentric position, it lies farther from the reaching path of the active hand. Those induced movements of the stationary hand are therefore mitigated or otherwise change in nature, resulting in an interaction between the spurious target code and stationary hand posture. This consideration was made during the experimental design, which is the reason why the bottom target to the contralateral space was removed from the set of possible reach targets. Visually, this seemed a conservative way of ensuring that the two hands would not exhibit any form of movement interaction. A more generic form of this concept would be that the change in postural state affects the rotational forces on the body during movement, which begins to mesh with hypotheses surrounding bimanual coordination as we have described earlier. In the future, we may perform control analyses that attempt to decode reach target and/or postural state from the (small) velocities of the stationary hand. This would not necessarily preclude the more generic form of this idea, where target- and posture-specific changes in motor output would reside in the axial musculature of the trunk. Performing new experiments with EMG recordings could provide testing of that idea by decoding movements from the axial musculature.

3.4 Methods

3.4.1 Dataset details

The datasets presented here overlap with those of the previous chapter. The same recording sessions are analyzed but include additional behavioral manipulations. In the caudal aspect of the dorsal premotor cortex (PMd) we recorded 433 (242 left hemisphere) and 113 (57 left hemisphere) single-units in monkeys O and W, respectively. In the primary motor cortex (M1) we recorded 331 (196 left hemisphere) and 289 (195 left hemisphere) single-units. For population level analyses (LDA, dPCA), a small number of multi-units were included. A multi-unit was defined by waveform clusters that separated from the noise cluster and were stable over time, yet did not quite meet the criteria for confident single-unit isolation (see Methods 2.4.3). For monkey O, the average proportion of multi-units in each single session population sample was 17%, ranging 12-25%. For monkey W, average 20%, ranging 12-32%. These

data were collected across 13 sessions for monkey O and 7 sessions for monkey W. In each session, monkey O performed between 606 and 1000 successful trials, and monkey W performed between 307 and 525 successful trials.

3.4.2 Task design

The basic form of the instructed-delay reaching task is described in the previous chapter. The key manipulation that was included in the present experiments was the starting configuration of the hands (Figure 3.1B). To initiate a trial, monkeys were required to place both hands in designated start positions (spherical targets, 4cm radius) that were arranged in one of three configurations: left hand eccentric with right hand centered, both hands centered, or left hand centered with right hand eccentric (Figure 3.1B, corresponding panels ordered left to right). Centered starting positions were 4cm from the midline, and eccentric starting positions were 15cm from the midline. Reach targets were arranged in a 2x3 fronto-parallel grid at a depth of 8cm from the starting positions. The three columns of targets were located at the midline and +/-11cm on the horizontal axis (Figure 3.1A). The vectors connecting the start position to each target were therefore dramatically altered between configurations (Figure 3.1B, boxed image). This manipulation was designed to test if and how the mapping between neural activity and reaching behavior was sensitive to the position of the stationary hand.

The trial timeline consisted of three phases (Figure 3.1C). These phases are described in more detail in the previous chapter and summarized here. A “Rest” phase began each trial, during which both hands remained stationary in the designated starting configuration. An “Instruct” phase then followed, when the monkeys were given instructions about the upcoming movement but were required to remain still. Finally, monkeys reached towards the target to obtain a reward during the “Move” phase.

Trials were organized in nested blocks, with the starting configuration being the outer block and the hand assignment being the inner block (Figure 3.1D). Using a single starting configuration, monkeys would perform two trials per target reaching with the left hand, then two trials per target with the right. The starting configuration would then switch, and the blocks of left- and right-hand reaches would repeat with new randomized target orders. Note that in certain configurations one of the targets was removed from the set of possible targets presented to the animal. When the stationary hand was in the centered position, we did not present the bottom target in the contralateral space (e.g. left hand reaches to the bottom-right target), as the reach path to that target put the two arms in close proximity to one another and we wanted to prevent the risk of the two arms physically interacting at all (e.g. bumping into each other).

3.4.3 Surgical implantation

All procedures were conducted in compliance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals and were approved by the University of California at Berkeley Institutional Animal Care and Use Committee under protocol ID AUP-2014-09-6720-1. This protocol approval and all surgical methods apply to each chapter of this thesis. Two adult male rhesus monkeys (*Macaca mulatta*) were implanted bilaterally with custom acute recording chambers (Grey Matter Research LLC, Bozeman, MT). Partial craniotomies within the chambers allowed access to the arm regions of dorsal premotor (PMd) and primary motor (M1) cortices in both hemispheres. Localization of target areas was performed using stereotactically aligned structural MRI collected just prior to implantation, alongside a neuroanatomical atlas of the rhesus brain [169].

3.4.4 Single-unit tuning

To determine whether single-units were tuned to reaches of a given arm, a one-way ANOVA was performed to determine whether firing rates were different for at least one of the reach targets. Tuning was determined using trials where the stationary hand was in the eccentric position so that all 6 targets were reachable by the active hand. The mean firing rate within phase windows was computed for each trial. For the Instruct phase, this window was 200ms to 500ms post-instruction. For the Move phase, this window was the first 300ms following movements onset. The ANOVA therefore operated on 6 vectors for a single phase, where each vector contained the mean firing rates on each trial performed for that target. An alpha level of 0.05 was used to determine significance.

3.4.5 Comparison of tuning profiles across stationary arm postures

For testing whether tuning profiles changed significantly across postures of the stationary arm, the average firing rate across trials was computed for each reach target (i.e. the firing rate on each trial was first averaged within the phase window, then averaged across trials for each target). Only the 5 targets that were reachable with both postures of the stationary arm were included in these analyses, thus the tuning profile was represented as a vector $X \in \mathbb{R}^5$. The observed angle between postures a and b could then be computed as:

$$\theta_{ab} = \arccos\left(\frac{X_a \cdot X_b}{\|X_a\| \|X_b\|}\right)$$

A reduced 2-target example of this method is illustrated in Figure S3.1. A permutation testing approach was used to determine whether observed angular differences were significant. Posture labels were shuffled across trials within individual targets (i.e. All trials for target 1 from both postures were combined, then reassigned randomly to 2 new groups of the same size as the original posture groups. This would then be repeated for each target.). The angular difference in tuning was then computed between these permuted datasets, and a p-value was assigned as the proportion of permutations that resulted in an angular difference at least as large as what was observed.

3.4.6 Analysis of similarity between ipsilateral and contralateral target responses

The similarity between ipsilateral and contralateral tuning was assessed by analyzing the correlations between the mean firing rates associated with reaches to each target. Only those units which were considered tuned to movements of both arms were included in analysis. The starting configuration with both hands in the center was used, and only the 4 targets that were reachable from that configuration for both hands (top row, bottom-middle target) were included. Two vectors were constructed by concatenating the mean firing rates for these 4 targets across all units, yielding $X_{Contra}, X_{Ipsi} \in \mathbb{R}^{4n}$, where n is the number of units. Spearman's correlation coefficients were then computed between these two vectors. Correlations were assessed in both an extrinsic reference frame, where target labels were maintained for the two hands, and an intrinsic reference frame, where target labels were reflected across the midline so that cross-body movements were considered equivalent.

3.4.7 Linear discriminant analysis (LDA) generalization procedure

Linear discriminant analysis (LDA) was used in this study to model associations between neural population activity and behavior, and then test whether those associations depended upon experimental manipulations. LDA was applied using the same preparation as in the previous chapter (see 2.4.9 Linear discriminant analysis). Mean firing rates on each trial were taken within the same phase windows as those used for single-unit analyses. These mean firing rates were then used as predictors in the LDA models.

To assess how sensitive classifiers were to posture of the stationary hand, two different training conditions were used to predict the same target reaches. Leave-one-out

cross-validation was used for the “Within” posture evaluation. For the “Across” posture evaluation, the entire trial set from the other posture was used to train the classifier model. Cross-validation therefore allowed comparison of the two conditions without the risk of overfitting effects. Each value of accuracy represented the mean of both postures, i.e. Each “Within” accuracy datapoint was the mean of center-trained, center-tested and eccentric-trained, eccentric-tested performance. Each “Across” accuracy datapoint was the mean of center-trained, eccentric-tested and eccentric-trained, center-tested performance.

Generalization across reaches of the two arms was also tested. It was expected that changing the reaching arm would negatively impact classifier performance, and the question was instead whether generalizing across the two arms would provide classification above chance level at all. Models were therefore trained for one hand and tested on the other in each phase. Separate models were trained and tested separately for each of the 3 starting arm configurations, and their predictions were all combined for computing a single accuracy value. Only the targets reachable by both hands were included in analysis, and that target set was different depending on the configuration. The chance level was therefore not simply $1/\text{number of targets}$. The resulting accuracies for the Instruct and Move phase were compared against the Rest phase values, since the Rest phase came prior to target presentation and therefore should not contain any information about the forthcoming reach.

3.4.8 LDA exclusion criteria

The analyses in this study generally followed a fully within-subject design: neural data from a single hemisphere was evaluated for each hand, phase, and model. For assessing generalization of LDA classifiers across postures of the stationary hand, it was necessary that the within-posture classifiers perform above chance level, since performance drops would not be meaningful otherwise. For each statistical test, any classifier that performed at or below the chance level ($1/\text{number of targets}$) was therefore removed. To maintain the within-subject design, all data associated with that same dataset was also removed.

3.4.5 Demixed principal components analysis (dPCA)

Interactions between target and posture coding were analyzed using demixed principal components analysis (dPCA) [136]. dPCA attempts to decompose the population signal into components that each exclusively represent variance associated with a single task parameter or set of parameter interaction terms. In this way, it marginalizes and decomposes the data like an ANOVA. We applied the method using 3

task parameters: Time (condition-independent), Target, and Posture (of the stationary arm). The marginalization procedure yields both main effect and interaction terms, which we grouped into 4 categories: Time (only Time term), Target (Target, Time x Target), Static Posture (only Posture term), and Posture Interaction (all interactions involving Posture). Components within each term group are designed to be orthogonal or near-orthogonal. However, components across term groups are only designed to be uncorrelated, and may not necessarily be orthogonal.

Unit firing rates were first binned in 20ms non-overlapping bins, then smoothed using a gaussian kernel with standard deviation 50ms. These firing rate estimates were then trial-averaged and soft normalized using the full firing rate range +5Hz [8, 80, 108, 123]. Units were then combined across recording sessions. The number of components to use in each model was determined independently, using the minimum number of components necessary to obtain at least one component for each term category and 8 total condition-specific (i.e. not Time) components. Using 8 condition-specific components has been used previously under the rationale that 6-8 components typically capture much of the condition-specific variance in similar tasks [123].

3.4.11 Permutation-based ANOVA

Permutation tests were used for both single and multi-factorial hypothesis testing when parametric tests were inappropriate. Multi-factorial repeated-measures ANOVA tests constrained permutations to exchangeable units where possible, resulting in approximate interaction testing and exact main effects testing [9]. p-values were estimated as the proportion of permutations resulting in test statistics that were at least as extreme as what was observed. In cases where the observed test statistic was more extreme than any permutations, we assigned a p-value of $1/\text{number of permutations} = 1.0e-4$.

3.5 Supplementary Figures

	Number SU's	Significantly tuned					
		Instruct			Move		
		Ipsi	Bi	Contra	Ipsi	Bi	Contra
PMd	433	61(14%)	112(26%)	108(25%)	37(8%)	278(63%)	75(17%)
	113	7(6%)	21(19%)	41(36%)	11(10%)	45(40%)	22(19%)
M1	331	57(17%)	72(22%)	56(17%)	30(9%)	209(63%)	56(17%)
	289	16(6%)	7(2%)	50(17%)	36(12%)	95(33%)	68(24%)

Table S3.1. Proportions of significantly tuned single-units across task phases. For well isolated single-units in each brain area, the proportions of units that were significantly tuned to different target reaches (one-way ANOVA comparing firing rates across targets, $p < 0.05$) are displayed in each cell of the table. Note that this is different from the significant modulation data in Table 2.1 – tuning refers to activity for at least one of the targets being different from the rest, while modulation referred to significant deviation from baseline activity in general and could be identical across targets. For each phase, single-units were classified as uniquely ipsi, contra, or bilaterally tuned. Top row in each pair of rows represents Monkey O, bottom row Monkey W.

	Significantly different tuning between postures			
	Instruct		Move	
	Ipsi	Contra	Ipsi	Contra
PMd	29/173 (17%)	22/220 (10%)	84/315 (27%)	72/353 (20%)
	6/28 (21%)	3/62 (5%)	7/56 (13%)	14/67 (21%)
M1	16/129 (12%)	25/128 (20%)	49/239 (21%)	65/265 (25%)
	6/23 (26%)	3/57 (5%)	23/131 (18%)	18/163 (11%)

Table S3.2. Proportions of tuning profiles that were sensitive to posture of the stationary hand. Each cell of the table displays the proportion of significant tuning profiles that were also significantly different between the two postures of the stationary hand (permutation test, $p < 0.05$). Top row in each pair of rows represents Monkey O, bottom row Monkey W.

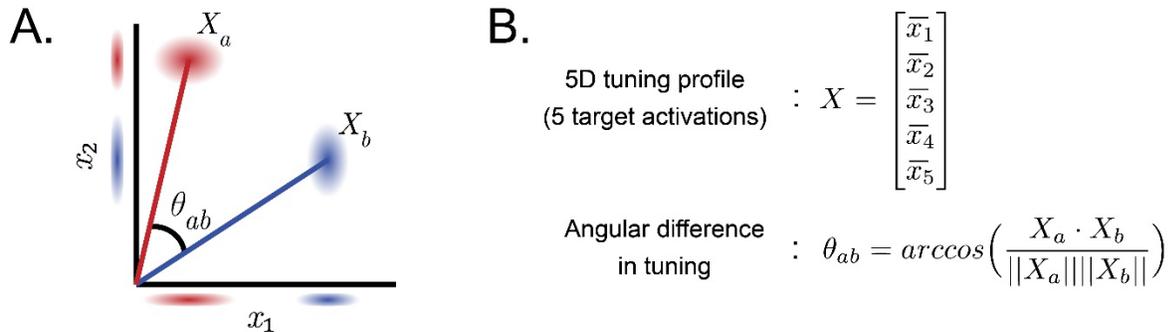


Figure S3.1. Quantifying differences in target tuning. (A) Schematic of a reduced example in 2-dimensional space. Consider tuning profiles defined by firing rates to target 1 (x_1) and target 2 (x_2). Datasets a and b each have their own distributions of firing rates for trials to these two targets, and the means of those distributions can be taken to represent each dataset's tuning profile as a vector in 2-dimensional space (X_a, X_b). The angle between those resulting vectors quantifies the difference in target tuning across the two datasets. This provides a non-parametric method of assessing tuning in replacement of parametric methods like cosin-tuning, which requires a polar coordinate frame. (B) Mathematical formulation in the 5-target case.

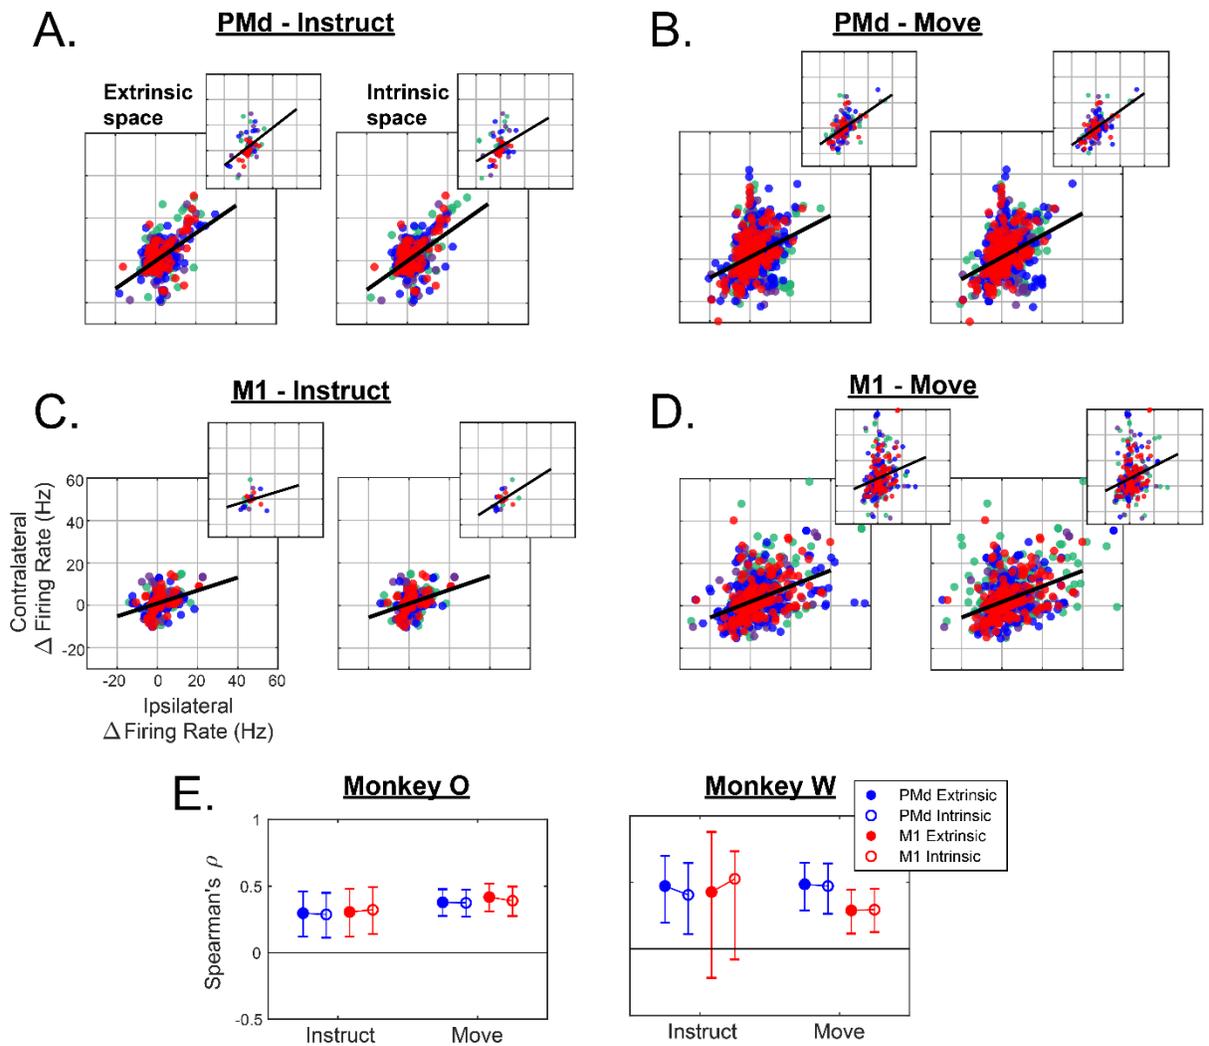


Figure S3.2. Similarity of ipsilateral and contralateral tuning profiles. (A-D) Change in firing rate relative to Rest for ipsilateral and contralateral target reaches. Each datapoint represents a single target for a single unit. The left panel of each pair of plots assigned target labels in an extrinsic reference frame, while the right panels expressed targets in an intrinsic reference frame by mirroring target labels across the midline for one of the two hands. Only units that were tuned for movements of both arms were included. Only the 4 targets that were reachable with both hands when starting from the centered posture are included. Each target is displayed in a different color. Regression lines are plotted in black. Larger plots are monkey O; inset plots are monkey W. (E) Spearman's correlation coefficients for the data displayed in A-D. Datapoints show bootstrapped median \pm 95% confidence intervals.

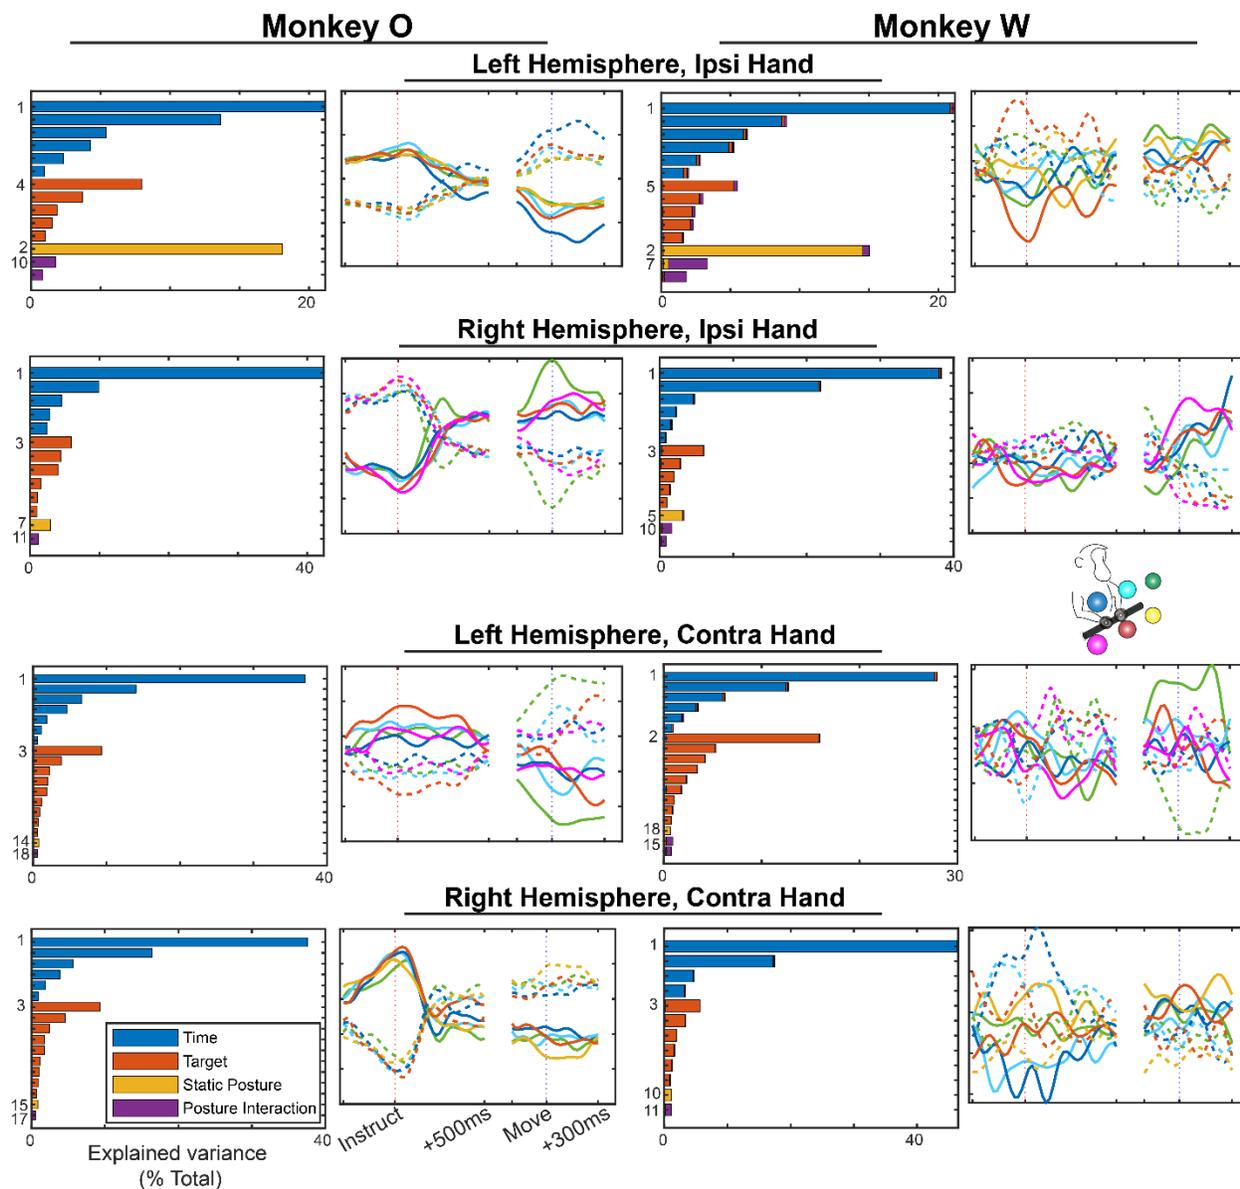


Figure S3.3 dPCA results for PMd. Each pair of plots represents a single dPCA model and includes a bar chart representing the proportion of variance explained by each dPCA component and a projection of the neural activity for all targets and postures onto the top Posture Interaction component. These plots are displayed in the same manner as the examples from Figure 3.6. The left column is monkey O, and the right column is monkey W. Each row corresponds to a specific combination of hemisphere and hand.

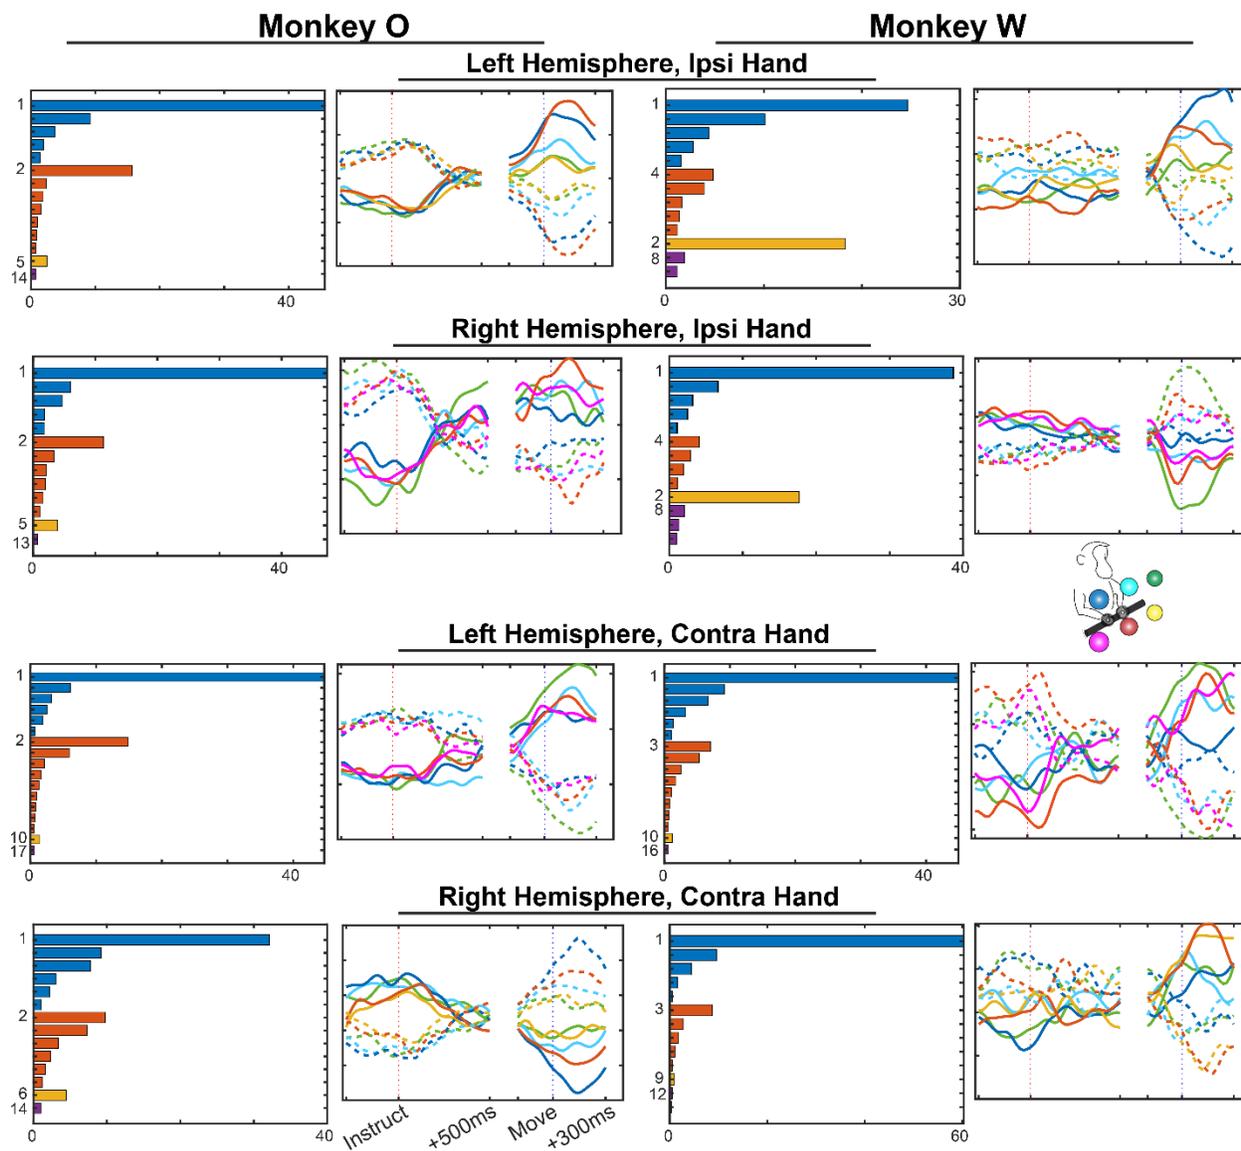


Figure S3.4 dPCA results for M1. This plot is the same as Figure S3.3, except for models trained using data from M1.

Acknowledgements

We thank M. Kitano for help in NHP care and handling, and M. Chen for contributions made in the dPCA analysis. This work was supported by the National Defense Science and Engineering Graduate Fellowship to T.C.D, and a grant from the National Institute of Health (NS097480) to J.M.C.

Chapter 4

Interhemispheric population dynamics during bimanual coordination

4.1 Introduction

Coordinated use of the two arms involves a widely distributed system of neural structures. The two cortical hemispheres represent important nodes in this network, and their communication with each other has shown critical for bimanual coordination. Research with callosotomy patients has revealed that typical patterns of spatial interaction between the two hands are largely abolished in the absence of the corpus callosum, the largest fiber bundle connecting the two cortical hemispheres [79, 88, 132]. The supplementary motor area (SMA) is the cortical brain region that has historically been most tied to bimanual coordination, having the densest transcallosal projections and expressing unique bimanual deficits upon lesion or perturbation [22, 157, 184, 197]. However, the primary motor cortex (M1) and other premotor areas maintain interhemispheric connections as well (including connections with the SMA [184]) and display unique bimanual-related activity patterns [69, 72]. Bilateral activity arising in the M1 and premotor cortices is thus likely to both be influenced by and participate in the computations that underly coordination of the two arms.

Coordination as a general phenomenon requires coupling the dynamics for the two (or more) plants being controlled. This ought to produce overall lower-dimensional dynamics, where the state of each participating plant influences control of the others and constrains behavior to fit the coordinative demands [188]. In the context of bimanual coordination in the motor cortex, it may be expected that control signals directed at one arm (the contralateral arm) would have access to state information

regarding the other arm (the ipsilateral arm). Not only should state information regarding the ipsilateral arm be represented in the motor cortex, but this information would be expected to alter the neural dynamics that give rise to control signals.

Non-invasive neural recording methods have revealed changes in macroscopic coupling between motor cortical hemispheres that are related to bimanual coordination [96, 197]. How this coupling is reflected in the networks of individual neurons is yet unclear. Emerging research has pointed to the importance of cortical population dynamics for directing the evolution of patterned output for motor control [44, 200, 217]. However, there is much yet to be understood about interhemispheric interactions and their relationship with bimanual coordination at this level of description. In M1, intracortical spiking activity from a single hemisphere is predictive of spiking activity in the contralateral hemisphere [8], yet the dynamics of their interaction and any relationship with coordination have yet to be explored.

In the current study, we investigated whether ipsilateral motor activity in the M1 and dorsal premotor cortex (PMd) participates in coupling interhemispheric population dynamics for bimanual coordination. We used a novel behavioral task in monkeys that was designed to provide comparison between unimanual control, independent bimanual control, and coordinated bimanual control. We found that kinematic decoders trained on isolated unimanual movements generalized better for independently controlled bimanual actions than for coordinated bimanual actions, indicating a change in the neural representations of movement across the control modes that impacted both hemispheres. We then fit linear dynamical system (LDS) models to capture the dynamics of motor cortical activity in each task variant. Activity during coordinated control depended less heavily on high frequency oscillatory dynamics than activity during independent bimanual control; however, the dynamics were largely context-invariant on the whole. These results do not support a role for ipsilateral motor cortical activity in altering interhemispheric dynamics for bimanual coordination. Rather, they may be more in line with a hypothesis involving selective inhibition during independent control.

4.2 Results

4.2.1 Task design

The behavioral task in the current study was designed to distinguish true bimanual coordination from both isolated unimanual control and simultaneous control of the two arms with independent action goals. On each trial, monkeys were required to do some combination of two actions: (1) reach towards a target with the left hand, and (2) pull a joystick with the right hand (Figure 4.1). Stimulus presentation on a virtual

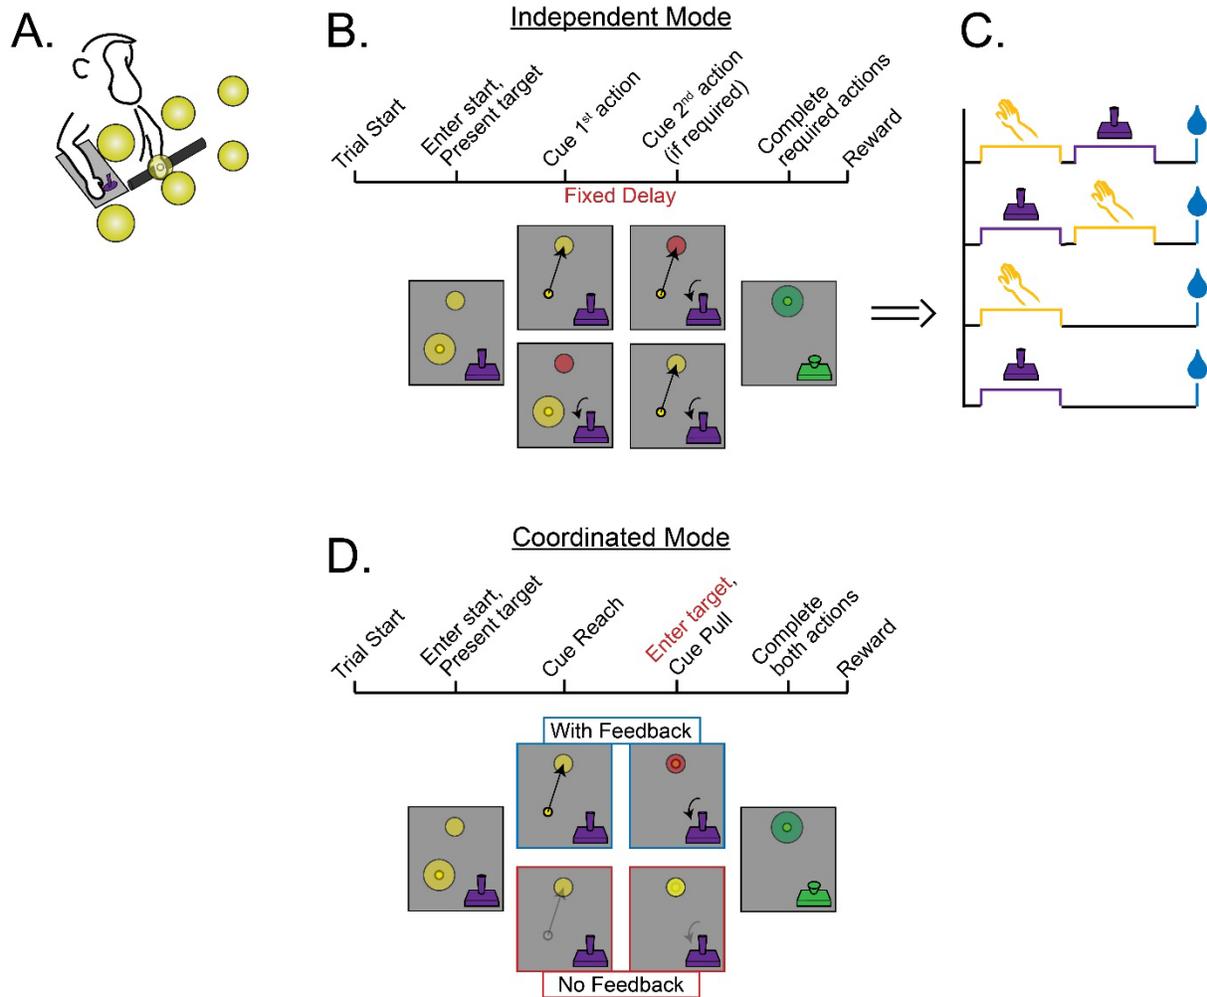


Figure 4.1. Task design. A. The physical setup of the task. Monkeys reached to one of six targets with the left hand and pulled a joystick with the right. Endpoint feedback for the reaching hand (cursor) and all other stimuli were presented on a virtual reality display. Stimulus presentation dictated one of two different control modes. B. Independent Mode control. Monkeys initiated a trial by placing the reaching hand in a central starting target (yellow sphere) and were required to have no deflection of the joystick. A yellow spherical reaching target was then displayed at one of six possible locations. After a randomly sampled delay period, the first action was cued. To indicate a reach cue, the starting target disappeared, and the reaching target enlarged. To indicate a pull cue, the reaching target turned red until the pull was initiated. The second action was then cued at a fixed delay using the same stimulus associations as the first cue. On certain “catch” trials, no second action was cued, and the animal was required to complete the first action without initiating the other. This prevented the monkey from ignoring the second cue and learning to perform the second action at a fixed delay. If all cued actions were initiated in the requisite order and completed, a juice reward was delivered. C. The different stimulus presentations for the Independent Mode resulted in four trial types: Reach-Pull (top row), Pull-Reach (second row), Reach-Catch (third row), Pull-Catch (bottom row). Each row displays a trial timeline with steps indicating the response initiation

window for the two potential actions (yellow: reach, purple: pull). Cartoon water droplets indicate reward delivery. Note that the initiation order was constrained by the task, but movements could be overlapping during their execution. D. Coordinated Mode control. With Feedback trials were initiated in the same manner as the Independent Mode, and cues were presented using the same stimuli; however, the first action on all trials was the reach. In contrast to the fixed timing of the Independent Mode, the timing of the second action (pull) was contingent upon the execution of the reaching movement, with the pull cue presented at the moment the reaching hand entered the reach target. In No Feedback trials, the behavioral constraints of the task were identical to the With Feedback trials; however, as soon as the reach was initiated the reaching cursor became invisible, and the pull cue was never delivered. These No Feedback trials required the monkeys to perform the task using only feed-forward planning and feedback provided by internal/proprioceptive state.

display dictated two different control modes for completion of the bimanual responses: Independent Mode (Figure 4.1B-C) and Coordinated Mode (Figure 4.1D). All trials began with the monkeys placing their reaching hand in a central starting target. Position of the joystick hand was unconstrained so long as it did not move the joystick from its neutral position at the start of the trial, though both monkeys tended to rest the joystick hand in close proximity to the joystick handle between trials. A reaching target was then displayed at one of six locations for a variable delay period during which the monkeys were required to keep both hands still. This marked the last point along the trial timeline where Independent and Coordinated Mode trials were the same.

For Independent Mode trials (Figure 4.1B), an imperative signal was then presented instructing one of the two actions. To cue the reach, the starting target disappeared, and the reach target enlarged. To cue the pull, the reach target turned red. After a fixed delay, the second action was cued; however, on a subset of trials (“catch” trials) the second cue was never provided, and the monkeys were required to withhold the second action. This resulted in four different trial sub-types determined by the order of the two actions and presence of a “catch” (Reach-Pull, Pull-Reach, Reach-Catch, Pull-Catch; Figure 4.1C). The critical distinction of the Independent Mode is that the requirements for each action are fully uncorrelated. Given the range of potential responses and lack of foreknowledge about the trial sub-type, optimal behavior would entail independently responding to each subsequent cue with the appropriate action.

In contrast, Coordinated Mode trials (Figure 4.1D) always instructed the reach action first and there were never catch trials. Rather than presenting the pull cue at a fixed delay, the pull cue was now contingent upon the execution of the reach action as well – as soon as the reaching hand entered its reach target, the pull cue was presented. Not only was there foreknowledge of the action order, but this contingency between the behavioral requirements for the two limbs provided the distinctive “coordinated” nature of these trials. Optimal behavior would therefore entail coordinated use of the two arms, whereby the state of the reaching hand would be integrated into the commands for the pull hand. Notably, this task could be accomplished without any further visual stimuli following the first go-cue (under optimal coordination conditions). We included trials where this was the case (Figure 4.1D “No Feedback”). As soon as the reaching movement was initiated, visual feedback regarding the endpoint position (cursor) was removed, and a pull cue was never delivered. Success on these No Feedback trials required the use of feed-forward planning and/or proprioceptive feedback control, both internally driven processes. Independent and Coordinated Mode trials were performed in fully separate blocks, and a sequence of four auditory tones preceded the Coordinated block to indicate the control mode. This design aimed to produce output behavior that was similar for Independent Reach-Pull trials and Coordinated trials despite the contexts and optimal control policies being quite different.

4.2.2 Behavioral performance and task validation

Two monkeys were trained to perform the behavioral task described above. Attentional demands were high for this task, in particular for the Independent Mode. There were five types of errors for initiated trials: early first action, wrong first action, early second action, late second action, and execution errors (e.g., not holding the reach target for the requisite time). Despite the difficulty of the task, each monkey was able to perform both the Independent and Coordinated Mode trials. We focus here on the Independent Reach-Pull trials and Coordinated trials, as these were the two conditions that were roughly matched in terms of behavior. A full summary of trial outcomes for the two monkeys is provided in Table S4.1 and Table S4.2.

Errors very rarely occurred for the first action. Of all trials that were initiated for the Independent Mode more than 95% advanced past the first action in both monkeys. The most common error type for Reach-First Independent trials was pulling the joystick early (note that for Reach-Catch trials any pull would be considered early). Combined across Reach-Pull and Reach-Catch trials, this accounted for 20% of the trials that had already advanced past the first action in Monkey O, and 5% in Monkey W. The next most common failure mode was pulling the joystick late, accounting for 42% of all Reach-Pull Independent trials that had advanced to the second go-cue in Monkey O, and 3% in Monkey W. However, we note that the response windows differed across the two monkeys. Other execution errors occurred in 0% of Reach-Pull trials for Monkey O and 1% for Monkey W. Overall success rates in the Independent Mode were 72% for Monkey O and 91% for Monkey W.

Errors in the reach component of Coordinated trials were similarly rare, and the most common error type was again completing the pull too early. Pulling the joystick early accounted for a similar proportion of trials as in the Independent Mode for Monkey O (19%) and was slightly greater for Monkey W (11%). The proportion of trials resulting in a late pull was greatly reduced for Monkey O (<1%) but similar for Monkey W (4%). Overall success rates in the Coordinated Mode were 77% for Monkey O and 83% for Monkey W, both having similar rates with or without visual feedback.

Single trial examples of the kinematics in each trial sub-type are provided in Figure S4.1. To analyze the behavioral timing, we timestamped 5 different behavioral events related to the reach and the pull. For the reach, we marked the onset, peak velocity, and target entry. For the pull, we marked the onset of both pull hand movement (“Pull onset”) and joystick deflection (“Joystick onset”), which only truly differed from each other for Monkey W who had an idiosyncratic delay between lifting the pull hand and contacting the joystick. We plotted the distributions of these events aligned to the onset of the first go-cue for Independent Reach-Pull, Coordinated With

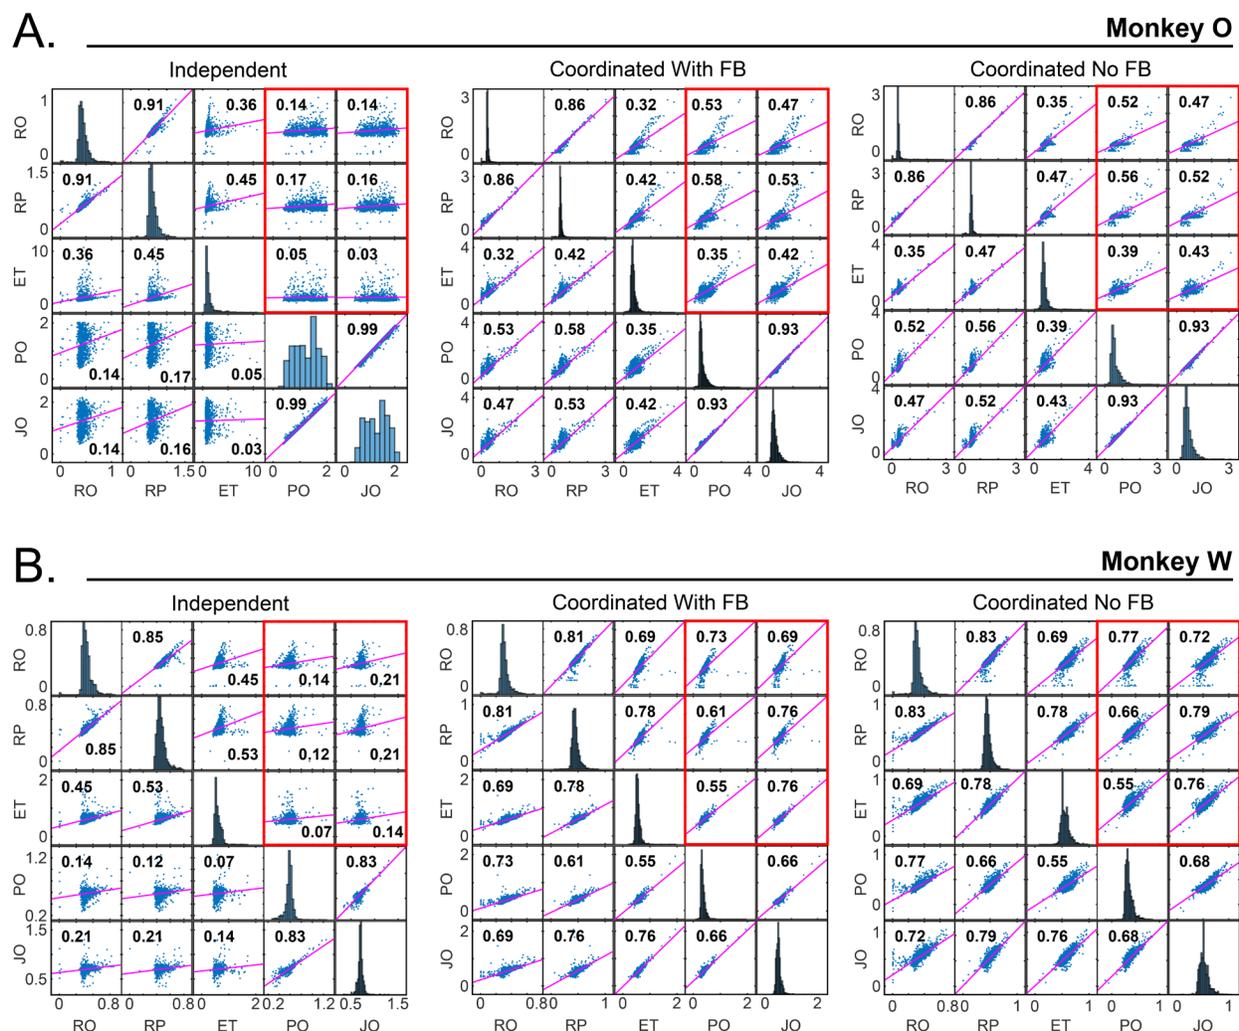


Figure 4.2. Behavioral event correlation plots. The correlations between behavioral event timestamps, aligned to the onset of the reach cue, are displayed for each of three trial types: Independent Reach-Pull trials matched to the coordinated behavior (left column), Coordinated trials with visual feedback (middle column), and Coordinated trials without visual feedback (right column). The cells within each matrix show scatterplots for pairwise comparisons of behavioral events along with a least squares fit line and Spearman's correlation coefficient. Cells along the diagonal show marginal distributions of the timing for each behavioral event. A red box outlines the six comparisons describing relationships between the two movements, e.g. pull onset and reach peak. Abbreviations: RO – reach onset, RP – reach peak, ET – enter target, PO – onset of pull hand movement, JO – onset of joystick deflection. A. Monkey O, B. Monkey W.

Feedback, and Coordinated No Feedback trials in Figure S4.2. One of the goals of the task was to produce distributions of action timestamps that were similarly overlapping for both the Independent and Coordinated Modes. This was largely the case, although both monkeys did perform the Independent trials more sequentially (less overlap in the distributions for reach and pull event times). Reaction times (RT) and movement durations were generally longer during the Independent Mode (Figure S4.2, top row). This was very pronounced for Monkey O. While the peak of Monkey O’s pull RT distribution was around 500ms, he responded with a nearly uniform distribution (Figure S4.2A, top row). Combined with the observation that Monkey O exhibited large proportions of mis-timed pull actions, this indicated that he did not pay close attention to the pull cue on many trials. This was not a problem provided that the reach and pull movement times were not correlated with each other (i.e. coordinated). The primary validation of our behavioral paradigm was therefore to analyze the correlations of reach and pull event times on each trial.

Correlation plots for each of the 5 behavioral events are plotted in Figure 4.2. We were interested specifically in correlations across the two actions (Figure 4.2 red outlined subplots). Correlations between reach and pull timestamps were weak for the Independent Mode in both monkeys. Spearman’s correlation coefficients ranged from 0.03-0.21 across all six pairwise comparisons. As intended, correlations between the two actions were much stronger in the Coordinated Mode. In Monkey O, the strongest correlation was between reach peak and pull onset (With Feedback, $r_{RP,PO}=0.58$; No Feedback, $r_{RP,PO}=0.56$). Similarly, the strongest correlations for Monkey W were between reach peak and joystick onset (With Feedback, $r_{RP,JO}=0.76$; No Feedback, $r_{RP,JO}=0.79$). The task itself constrained these correlations since the response window for the pull movement was aligned to target entry. The joint distributions in Figure 4.2 only include trials where the monkeys completed the joystick pull within that time window, thus it is theoretically possible that these distributions could be groomed by the success criteria to artificially produce strong correlations. However, the high success rates together with the tendency for stronger coupling between pull movements and reach peak (as opposed to target entry) suggest that the coordination profiles were self-generated rather than being an artifact of joint distributions that were groomed by success criteria.

4.2.3 Strength of movement representations in each hemisphere across task conditions

We began by investigating the strength of movement representations in each hemisphere across the different task conditions. Linear kinematic decoders were trained on spiking data from either the hemisphere contralateral to the reaching hand (right

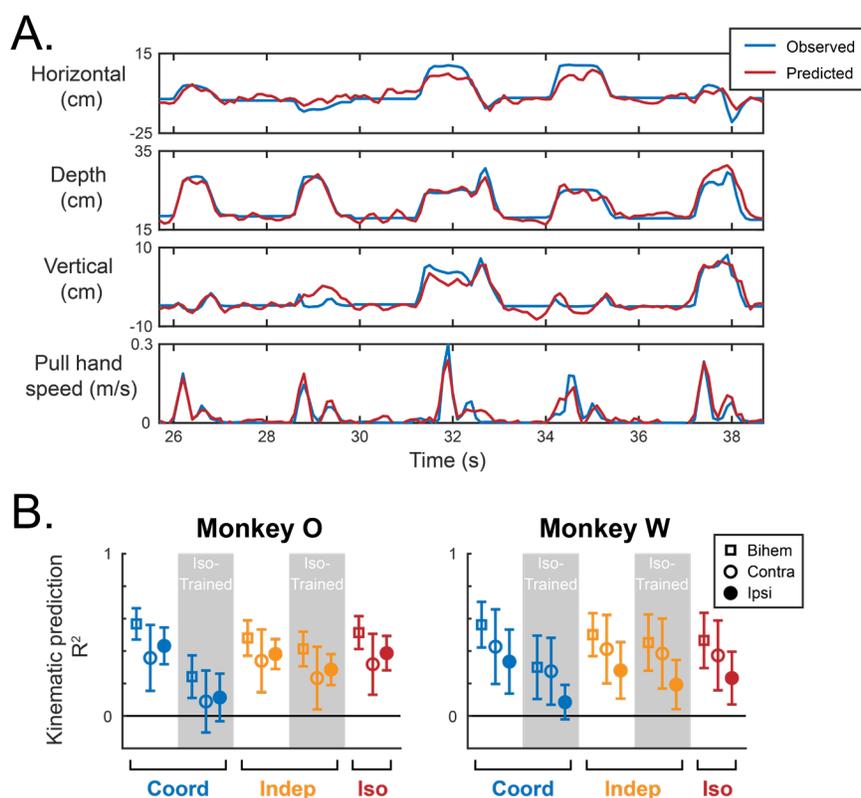


Figure 4.3. Kinematic decoding across control modes. A. Example kinematics and decoded predictions for 5 consecutive trials of the Coordinated Mode task. Linear decoders were trained using neural activity at multiple time lags. This example from Monkey W includes activity from both hemispheres, and predicts the 3D endpoint position of the reaching hand in cartesian coordinates (top 3 subplots) and the speed of the hand pulling the joystick (bottom plot). B. Decoders were assessed for each of three trial sets: coordinated trials (Coord, blue), independent matched trials (Indep, orange), and isolated unimanual movements from the independent control mode (Iso, red). For Monkey O, Iso trials consisted only of the independent catch trials. However, for Monkey W, Iso trials consisted of catch trials as well as the dual-action trials with long delays between the two movements such that they each was performed largely in isolation. Datapoints on white backgrounds were trained and tested using the same trial type and cross-validation. Datapoints on grey backgrounds were trained using Iso trials. Separate models were trained using both hemispheres (Bi-hem, squares), only the hemisphere contralateral to the reaching hand (Contra, open circles), and only the hemisphere ipsilateral to the reaching hand (Ipsi, filled circles). Single R^2 values for each session were taken as the mean over the R^2 for each dimension of the reaching movement. Note that more cells were recorded in the ipsilateral hemisphere, so their raw performance is not readily compared to the contralateral hemisphere. However, changes in performance across model training conditions may be. The data are displayed as mean \pm standard deviation.

hemisphere), ipsilateral to the reaching hand (left hemisphere), or from both hemispheres. Separate decoder models were fit for each of three different task variants: all Coordinated trials, Independent Reach-Pull trials that were matched to the Coordinated behavior, and isolated reach or pull movements from the Independent Mode. Movements were decoded in 3D cartesian space. An example of bihemispheric decoder predictions on five consecutive Coordinated trials is displayed in Figure 4.3A.

We focus now on decoding movements of the reaching arm, since the pull movement was very similar on each trial aside from its timing, and it was the state of the reaching arm that governed the response windows in the Coordinated Mode. Reach decoding tended to be strongest in the Coordinated Mode but provided accurate predictions for all trial types using either hemisphere (Figure 4.3B, white backgrounds). This could also be influenced by the greater number of Coordinated trials, which was not accounted for in model training (See Methods 4.4.4, 4.4.5). If the ipsilateral representation of movement provides state information in order to guide coordinated action of the two arms, then we would expect it to be amplified in the Coordinated Mode dataset. To test this specifically, we assessed the performance of decoders trained on isolated unimanual movements (“Iso-trained” decoders) and tested on either the Coordinated trials or the Independent matched trials. The prediction of the bimanual coordination hypothesis is that decoder performance using would be stronger in the Coordinated trials than the Independent matched trials, specifically for the ipsilateral decoders and not those that have access to contralateral activity patterns. This was not the case. In fact, Iso-trained decoders performed worse on the Coordinated trials (Figure 4.3B, grey backgrounds). This was a generic effect impacting each hemisphere to roughly the same degree, i.e., bihemispheric, contralateral, and ipsilateral models were equally affected (two-way permutation-based repeated measures ANOVA, 3 Hemispheric Input x 2 Condition; $ME_{\text{Condition}}: p_{\text{MonkeyO}}=1.0e-5, p_{\text{MonkeyW}}=1.0e-5$; $IE: p_{\text{MonkeyO}}=0.77, p_{\text{MonkeyW}}=0.064$). This means that the features of the bihemispheric neural signals that are leveraged by Iso-trained decoders are altered during Coordinated trials, causing a preferential decrement in performance as compared to the Independent matched trials.

4.2.4 Population dynamics in the ipsilateral vs contralateral hemisphere

If ipsilateral neural activity provides state information regarding one limb in order to coordinate descending drive for the contralateral limb, then one might expect to see alterations in the neural population dynamics when use of the two arms is intended to be interactive. We began investigating this hypothesis by characterizing the dynamics in both hemispheres across a range of behavioral tasks. PCA projections of neural data during both the unimanual reaching task of the previous chapters (Figure 3.2A) and the

bimanual task of the current chapter (Figure 4.4A) reveal prominent rotations through neural state space. Higher dimensions likely include additional dynamical structure. To capture the dynamical behavior of population activity, we fit linear dynamical system (LDS) models to single-trial data, modeled in the form:

$$\mathbf{x}_{t+1} = A\mathbf{x}_t + B \quad (1)$$

The A -matrix therefore relates the current neural state, \mathbf{x}_t , to the future neural state, \mathbf{x}_{t+1} . Eigendecomposition of the A -matrix allows us to isolate separable dynamic modes describing particular aspects of the dynamical behavior, which can be visualized as in Figure 4.4B (more on this in the following section).

We began by assessing the LDS model fits in each of three different tasks: The Coordinated and Independent (Reach-Pull trials only) Modes of the bimanual task, and the isolated unimanual reaching task from the previous chapters. Cross-validated R^2 values were computed for each dataset using the contralateral and ipsilateral hemispheres separately, where contralateral and ipsilateral refers to the relationship to the reaching hand. These values were compared against time-shuffled models that maintained any autoregressive accuracy that was a simple result of data smoothing. These results are plotted in Figure 4.4C. Both contralateral and ipsilateral models had significantly stronger fits than the time-shuffled models in all three tasks (two-sample t-test, $p < 0.05$), suggesting that both hemispheres display significant dynamical behavior. In the unimanual reaching task, both monkeys displayed significantly greater R^2 values for the contralateral models than the ipsilateral models (two-sample t-test; Monkey O: $p = 2.4e-3$, Monkey W: $p = 0.010$). Contralateral models had larger R^2 values on average across both versions of the bimanual task as well; however, the significance of these differences varied (Monkey O: $p_{\text{Coord}} = 0.054$, $p_{\text{Indep}} = 0.42$; Monkey W: $p_{\text{Coord}} = 0.029$, $p_{\text{Indep}} = 0.057$). It is important to note that contralateral reaching signals have greater signal-to-noise overall. Therefore, it is difficult to say whether a greater R^2 for the LDS model predictions would reflect more tightly imposed dynamical behavior, or rather a stronger signal in general.

Performance of model fits was similar across the two control modes of the bimanual task for both the contralateral and ipsilateral hemispheres. Coordinated and Independent R^2 values only had a statistically significant difference in the ipsilateral models of Monkey O (Independent > Coordinated, paired-sample t-test, $p = 1.9e-5$). Although this reached statistical significance, the effect size was quite small ($R^2_{\text{Ipsi,Coord}} = 0.71$, $R^2_{\text{Ipsi,Indep}} = 0.73$). Taken as a whole, this suggests that the strength of dynamical structure in each independent hemisphere is not heavily modulated by the coordinative demands of the task. However, this does not necessarily mean that the dynamical structure itself is the same across both control modes; it is possible that different sets of dynamical rules govern the evolution of activity in each mode while doing so with equal strength. The structure of these dynamics will be further

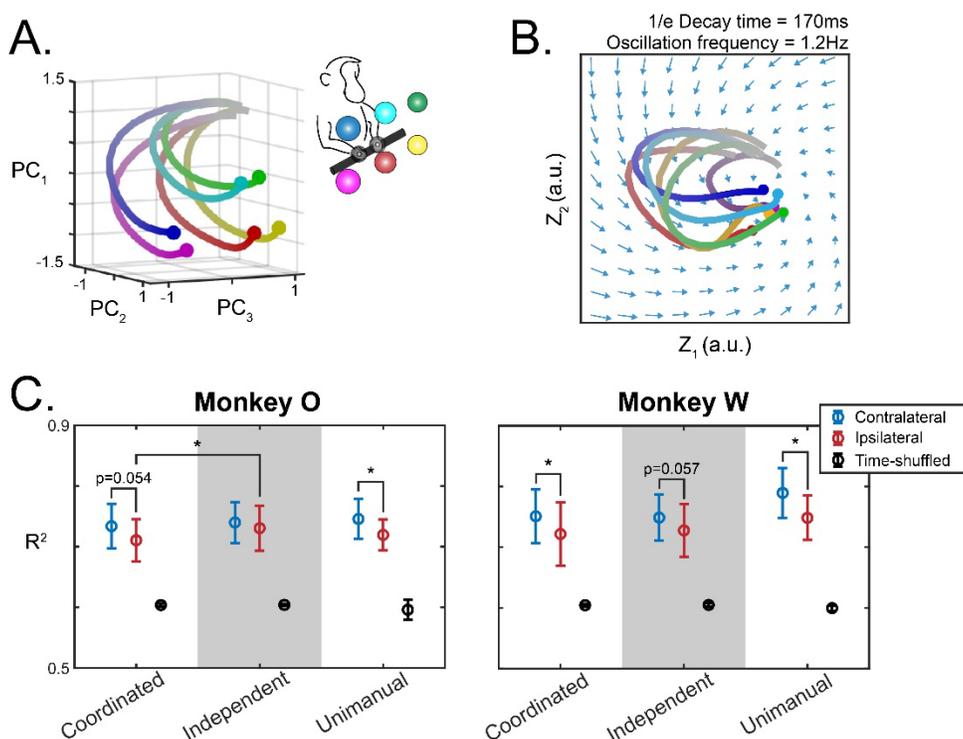


Figure 4.4. Neural population dynamics. A. PCA projections of trial-averaged neural data from the reach phase of Coordinated Mode trials. The data displayed here is from Monkey O and includes PMd and M1 in both hemispheres. Each target is plotted separately according to the color code indicated by the cartoon in the upper right. Line colors fade from grey to their respective color to indicate the advancement of time, extending from 200ms before reach onset to 300ms after. Note the visible rotation through neural state space that completes approximately one full period. B. Linear dynamical system (LDS) models were fit to single trial data (smoothed, but not averaged). The eigenvectors of the resulting A matrices provided separable “dynamic modes.” These eigenvectors may be thought of as slices through neural state space that capture specific linear dynamics. This example from Monkey O projects observed data from a single session onto two eigenvectors that form a complex conjugate pair (the real part of the first and the imaginary part of the second). The arrows in the background illustrate the flow field at different points in neural state space. The effects of this dynamic mode may be summarized by the Decay time (to $1/e$) and Oscillation frequency, which is displayed above the plot. The plotted example is from Monkey O and follows the same coloring rules as the lines in A. C. Cross-validated R^2 values for the fitted LDS models are displayed for both monkeys in each of three different tasks: the Coordinated and Independent (Reach-Pull) Modes of the bimanual task, and the unimanual reaching task in the previous chapters. Model fits are assessed using units from the Contra- and Ipsilateral hemispheres separately and may be compared against a Time-shuffled control that uses units from both hemispheres. The Time-shuffled control shuffled binned spike counts before smoothing the data, and therefore captures artificial dynamics created by the smoothing. Asterisks indicate significance at the $p < 0.05$ level.

investigated in the following sections. In summary, consistent dynamical behavior is observed in populations of neurons within each hemisphere during both coordinated and independently controlled bimanual behavior, as well as during isolated unimanual reaching.

4.2.5 Comparison of dynamical mode properties across coordinated and independent control

In order to examine the dynamical behavior of each dataset, we performed eigendecomposition of the A-matrices and examined dynamics in a latent space defined by the resulting eigenvectors (see Methods). Put simply, this allowed us to analyze independent components of the dynamics in what we refer to as “dynamical modes.” Figure 4.4B displays an example of observed neural data projected onto two such modes, along with the flow field describing their dynamical properties. Each mode may be associated with both a decay time and an oscillation frequency (if rotational).

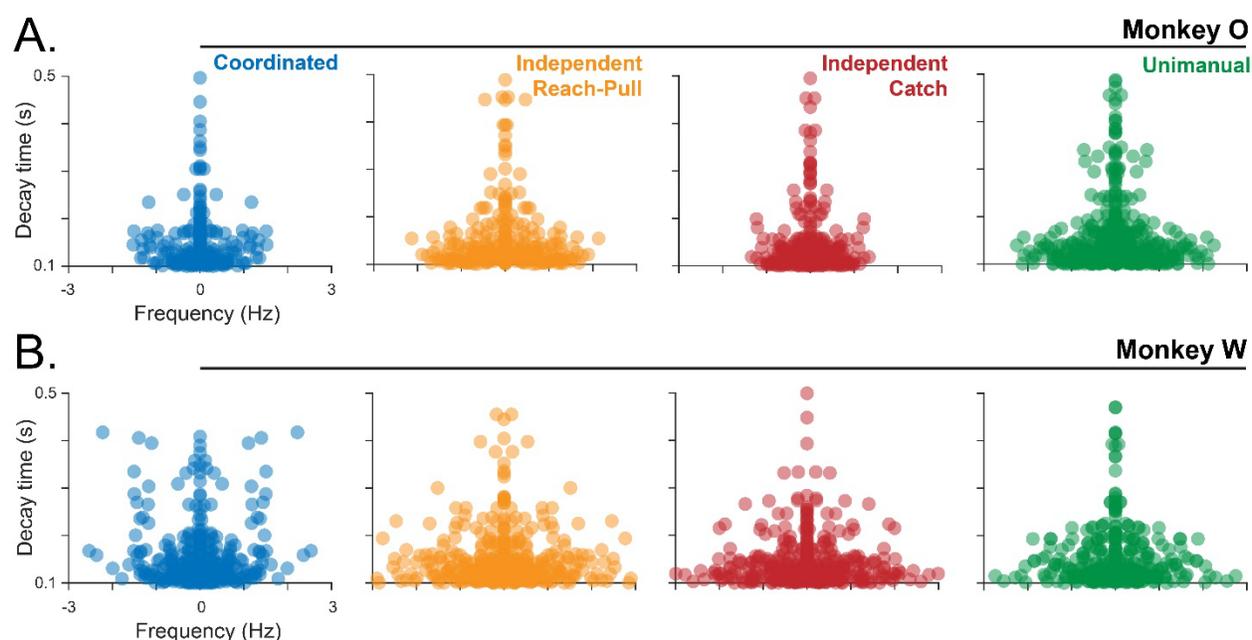


Figure 4.5. Comparison of the oscillatory and decay patterns across behavioral tasks. Each plot displays the decay time and oscillation frequency for all significant dynamic modes identified in each task. Note that oscillatory modes come in conjugate pairs, such that the two will have the same decay time and sign-flipped frequencies. A. Monkey O, B. Monkey W.

We began with an exploratory investigation of the dynamical mode properties in each task variant: Coordinated, Independent Reach-Pull matched trials, Independent Catch trials, and Unimanual reaching from the experiments of the previous chapters. Here, LDS models were fit to bihemispheric activity. We distinguished “significant” modes in each dataset by comparing decay times of observed modes against a null distribution obtained from time-shuffled data and applying a threshold cutoff (see Methods 4.4.6, Figure S4.3). The joint distributions of decay times and oscillation frequencies of significant modes in each task variant are displayed in Figure 4.5. Between the two monkeys, there appeared to be fewer modes at oscillation frequencies above 1.5Hz in the Coordinated Mode. In Monkey W, there was a unique emergence of oscillatory modes in the range of 1-1.5Hz with long delay times. Overall, the distributions did not appear extremely different across the tasks.

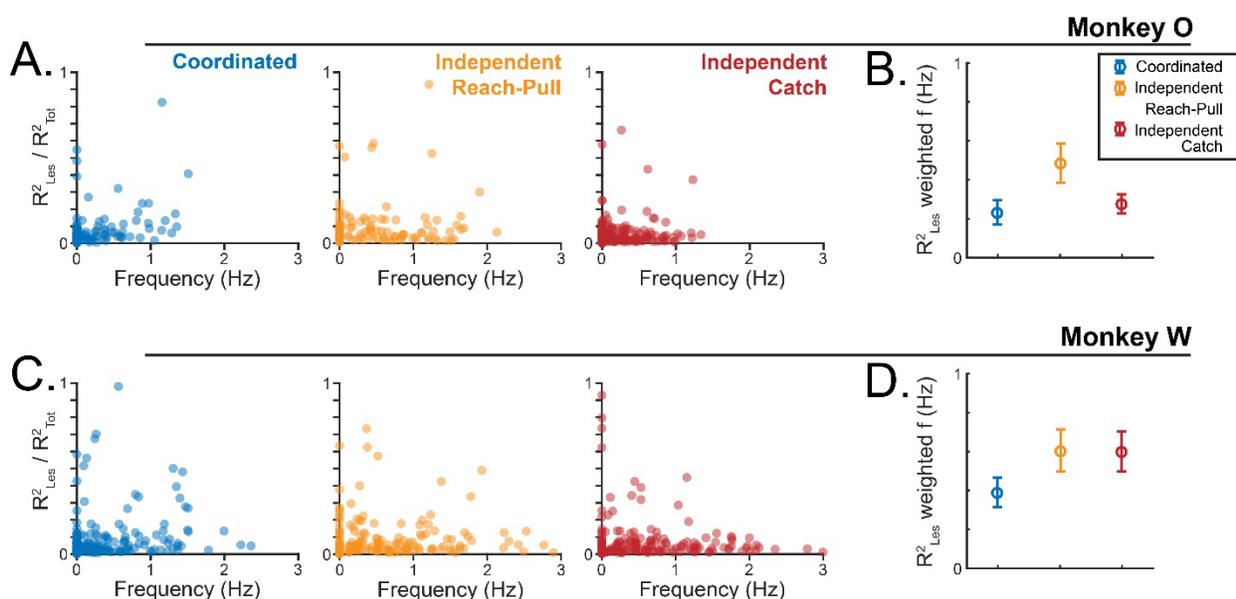


Figure 4.6. Relative importance of oscillatory dynamics for each behavioral context.

A,C. We used a “lesioning” approach to determining the relative importance of dynamical modes with different oscillation frequencies (see Methods 4.X). Each datapoint in these plots represents a single dynamical mode. The ratios plotted on the y-axes represent the reduction in model performance when a single mode is lesioned, as a proportion of the unlesioned model performance. Larger values indicate greater importance of the respective mode for predicting changes in neural activity. Only significant modes are plotted for each of the three trial types. B,D. The data in A,C are summarized by taking an average over all significant modes, weighted by their R^2_{Les} values. This metric provided a coarse summary of the oscillatory behavior of neural population dynamics across task conditions. The plots show mean \pm bootstrapped 95% confidence intervals.

To determine the relative importance of different dynamical modes, we implemented an A-matrix “lesioning” procedure (see Methods 4.4.7). In short, we decomposed the A-matrix, removed a single dynamical mode, then reconstructed the A-matrix from the remaining modes and re-tested performance in predicting future neural activity. The drop in performance, which we call R^2_{Les} , relates the relative importance of the lesioned dynamics for capturing neural variance. We looked specifically at the importance of modes across the spectrum of oscillation frequencies (Figure 4.6). As a summary statistic, we computed the average frequency of all significant modes in each dataset, weighted by their R^2_{Les} . These average frequencies were significantly lower in the Coordinated trials than they were in the Independent matched trials in both monkeys (Figure 4.6B,D – blue vs yellow; permutation test, Monkey O: $p=6.0e-5$, Monkey W: $p=1.8e-3$). This would suggest that non-oscillatory and low-frequency oscillatory dynamics are more dominant during coordinated control. However, the same relationship did not hold for both monkeys when comparing the Coordinated Mode with Independent Catch trials (Figure 4.6B,D – blue vs red). Weighted average dynamical mode frequencies were significantly lower in the Coordinated Mode than the Independent Catch trials for Monkey W (permutation test, $p=1.6e-3$) but not Monkey O ($p=0.26$). It is therefore unclear whether this bias towards lower frequency dynamics is truly a unique feature of the neural activity during coordinated control.

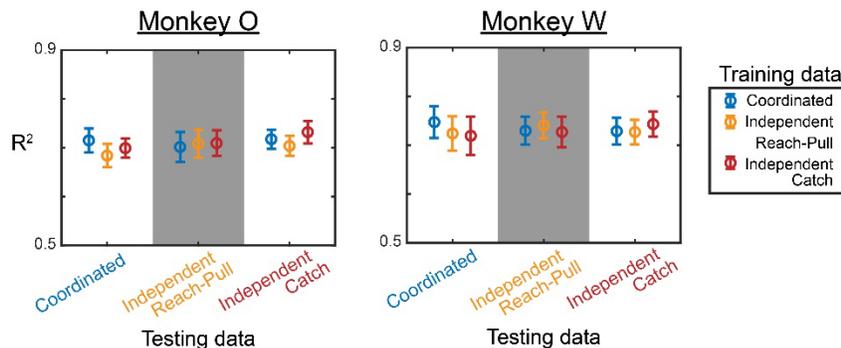


Figure 4.7. Generalization of dynamical behavior across behavioral contexts. LDS models were trained using data from each of the three task conditions, and then tested on each of the three task conditions. Model performance was assessed using the R^2 of predicted changes in neural activity under each training and testing combination, using cross-validation in all cases. Generalization across trial types indicates condition-invariant dynamical behavior, whereas lack of generalization indicates dynamics that are present in one condition and fully unobserved in another. The data are displayed as mean \pm standard deviation.

To test whether any task-specific differences in the properties of the neural population dynamics actually accounted for a significant amount of the overall variance, we used a generalization procedure with the LDS models. Separate models were trained on the Coordinated trials, Independent Reach-Pull matched trials, and Independent Catch trials. Each model was then tested on all three task variants (using cross-validation), resulting in nine pairwise comparisons. Data from each task tended to be best captured by its natively trained model, yet R^2 values were quite high for models that were generalized across tasks as well (Figure 4.7). The largest generalization cost was only a 4.4% drop in R^2 for Monkey O (Coordinated testing data, Independent Reach-Pull training data) and 3.8% for Monkey W (Coordinated testing data, Independent Catch training data). Taken together, these results suggest that although neural activity in each task may make preferential use of different dynamical properties (Figure 4.6), the dynamics are not categorically different across tasks. Models fit to each dataset still maintain a largely condition-invariant basis set of dynamical modes, such that the majority of neural variance is still captured by generalized models (Figure 4.7).

4.2.6 Interhemispheric interactions

As our final analysis, we investigated whether the neural population dynamics that we have modeled indicate any changes in interhemispheric interactions across control modes. We first asked whether dynamics for the two hemispheres are independent of one another, or whether there exist unique interactions between them that covary with task type. To test this, we created two different LDS models: one which combined activity from both hemispheres and fit a single A-matrix (“Combined” model), and another which fit separate A-matrices for activity from each hemisphere in isolation (“Separate” model). If any dynamics required joint observation of the two hemispheres, then they would be poorly fit by the Separate model. We computed the ratio of R^2 values for the Combined and Separate models, with a value above 1 indicating a benefit from access to interhemispheric dynamics. The Combined to Separate R^2 ratio was not far from 1 for either monkey in the Coordinated Mode (Figure 4.8A,D) and was only significantly greater than it was for the Independent Reach-Pull matched trials in monkey O (paired-sample permutation test, $p_{\text{MonkeyO}}=6.8e-4$, $p_{\text{MonkeyW}}=0.061$).

It is also possible, if not likely, that the Separate models may identify dynamics that are truly interactive across hemispheres purely by observing the component of those dynamics that exists within a single isolated hemisphere. If this were the case, then the R^2 ratios of the two models would not differ substantially, since all dynamical patterns are still accounted for, but the number of significant dynamical modes would be greater in the Separate model since interhemispheric interactions would be “double-logged” as two separate patterns within each hemisphere. We define the number of significant dynamical modes as the “dimensionality” of the dynamics. Although

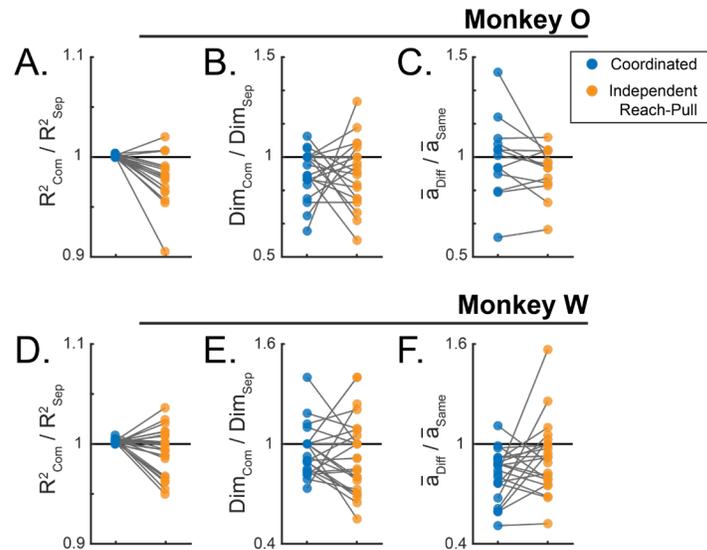


Figure 4.8. Interhemispheric interactions. Each plot displays a different metric of interhemispheric interaction computed for both the Coordinated trials and the Independent Reach-Pull (matched) trials. A. The ratio of R^2 for a combined bihemispheric LDS model over a model that fits two separate A matrices for each hemisphere to predict the same neural activity. The key difference in the two-part model is that each A matrix cannot fit any interactions between units in different hemispheres, and may only identify dynamics which are local to the hemisphere it was trained on. B. A ratio of the number of significant dimensions (or “dynamic modes”) for the same two models in A. If certain dynamical motifs are shared across the two hemispheres, then they may be consolidated and captured with fewer dimensions in the combined model. C. A ratio of the mean coefficient magnitudes in the A matrix that correspond to interactions of units in different hemispheres (Diff) over the same value corresponding to interactions between units of the same hemisphere other than themselves (Same). Connected datapoints in all plots indicate the same population of simultaneously recorded units used in analysis.

dimensionality was reduced by 5-10% in the combined models for both monkeys, the consolidation was not significantly different across Control Modes (Figure 4.8B,E; paired-sample permutation test, $p_{\text{MonkeyO}}=0.99$, $p_{\text{MonkeyW}}=0.36$).

Since the A-matrix describes predictive relationships between the entire population of units, one may specifically analyze the portions that reflect interactions between units from different hemispheres. We computed a metric on the A-matrices that quantified the relative weights assigned to interactions of units in different hemispheres or the same hemisphere (but not the unit itself, which is located along the A-matrix diagonal). Lower values of this ratio indicate dynamics that operate relatively independently for each hemisphere. We again found that these probes of interhemispheric interaction did not vary significantly across control settings (Figure 4.8C,F; paired-sample permutation test, $p_{\text{MonkeyO}}=0.10$, $p_{\text{MonkeyW}}=0.073$). In summary, we did not find any strong evidence that interhemispheric dynamics, as identified through LDS modelling, are adapted to the coordinative demands of the behavioral task.

4.3 Discussion

In the present study we have examined motor cortical population signals in each hemisphere as monkeys engage in a range of novel yet natural bimanual motor tasks, each differing in terms of its coordinative demands. Using population-decoding of reaching kinematics, we found that models trained on isolated unimanual movements generalized more poorly to coordinated bimanual movements than independently controlled bimanual movements. This was true of decoders using contra-, ipsi-, or bilateral populations of units. Furthermore, we fit LDS models to the neural data and performed an exploratory analysis of the dynamical motifs. We identified subtle differences between behavioral contexts yet found that dynamics were largely maintained regardless of the control setting. Finally, interhemispheric interactions in the LDS models were analyzed and showed that there was very little adaptation to the coordinative demands of the task. While ipsilateral activity in the motor cortex may still play some role in coordinating bimanual action, these results show little evidence that it does so flexibly by distinctly modifying dynamics in accord with the behavioral context.

4.3.1 Interhemispheric neural population dynamics

Neurons in the motor cortex are densely interconnected, and their interactions give rise to intrinsic dynamics that may be leveraged for performing a variety of computations [217]. Dynamical motifs such as rotations and decays have been prominently reported in recent works [44, 119] and have been hypothesized to be important features of the population for producing appropriate patterns of activity for driving movement [207]. In this study we have analyzed dynamics across both hemispheres and examined their importance in facilitating bimanual coordination. We

have shown that signals in PMd and M1 of both hemispheres evolve with lawful dynamics, albeit less strictly in the ipsilateral hemisphere, even during isolated unimanual movements (Figure 4.4C). Previous studies have shown that bilateral distribution of motor cortical computations provides robustness to unilateral perturbations, with interactions between the hemispheres supporting recovery of the population state [151]. One critical step that we have taken beyond previous work is to determine whether dynamic coupling may be flexibly altered to suit the coordinative requirements of a task. While we found some subtle differences in the dynamics between Coordinated and Independent control (Figure 4.6), most of the evidence pointed to a largely context-invariant set of dynamics.

However, it should be noted that our investigation of these dynamics was largely exploratory. Our methods could be described as data-driven, rather than approaching the topic with specific, theoretically motivated expectations about how the dynamics could or should be changing. Understanding cortical computations through the framework of neural population dynamics is still in its nascent stages, making it difficult to develop strong expectations, yet examples of detailed theoretical undergirding exist. For example, some researchers have trained recurrent neural networks (RNNs) to perform behavioral tasks and compared features of the resulting artificial dynamics to observed neural responses in order to gain insights about the computations they are involved in [37, 207]. Distinct population geometries and dynamical properties have also been identified in M1 and SMA, theoretically motivated by a role of the SMA in providing contextual information and further supported by artificial network simulations [185]. Our study would benefit from similar theoretical or simulation-based hypothesis development. For example, looking at dynamical modes of a particular frequency that might support bilateral coupling, or analyzing fixed points and stability behavior that are relevant to the constraints placed on bimanual coordination. While our results do not strongly implicate interhemispheric population dynamics or ipsilateral activity patterns in the facilitation of bimanual coordination, it may be that these foundational descriptions have simply not touched the relevant features.

In a similar vein, the method that we use to identify dynamics may also simply not be sensitive enough to identify the important differences. This is particularly relevant when considering how our decoder and dynamics results can be integrated – the decoder analysis did show condition-specificity and could just be the more sensitive probe. Other methods to consider for future dynamics analysis would be latent variable linear models like the Kalman filter [119] or RNN’s [185, 207].

4.3.2 Context-invariant dynamics in PMd and M1

Our results point to a largely context-invariant set of dynamics in the motor cortex. Russo et al., 2020 recently compared the neural dynamics present in M1 and the SMA [185], an area implicated in providing contextual information broadly, and

supporting bimanual coordination in particular. The authors found that activity in the SMA, but not M1, tracked the cycle number in a sequence of identical rhythmic actions by producing helical orbits through neural state space. This helical path could be described as a rotation, similar to those observed in M1, with an orthogonal decaying dimension along its rotational axis. This decaying dimension captured the context (cycle number). Our study may fall in line with the notion of M1 being context invariant; however, it is important to note that the context in our study is quite different. The study by Russo et al. considered the context of time: successive movements in a stereotyped series. Our study considered the context of bilateral state, which contains both spatial and temporal aspects. These different flavors of “context” may be treated differently by the brain.

Indeed, the dichotomy between spatial and temporal aspects of bimanual coordination/interference has proven an important one. In the callosotomy literature, timing constraints on bimanual behavior, such as synchronization of starting and stopping events, appear to be largely unaffected by disconnection of the cortical hemispheres [114, 132, 196, 213]. In contrast, many of the spatial constraints, such as difficulty drawing boxes in orthogonal orientations, are completely abolished [79, 88, 132]. Both M1 single unit responses [72] and interhemispheric interaction patterns in intracortical local field potentials (LFP) [34] have been found to depend on specific non-additive combinations of bimanual behavior, suggesting that computations in M1 do take into account context of the bilateral state in some form. This could be considered spatial context. In contrast, one recent study has suggested that M1 maintains independent representations of the two arms even during bimanual action [55]. While it is possible that the monkeys in our task could have been using purely temporal control policies, the task constraints were essentially spatial. Our current results therefore suggest that the context-invariance of PMd and M1 dynamics extends beyond the temporal domain and into the spatial domain as well.

4.3.3 Inhibition versus coordination

The primary validation of our task was the disruption of intermanual correlation patterns across the two bimanual control modes (Figure 4.2). The intent was to invoke a type of bimanual “motor routine” [96] in the Coordinated Mode that was categorically different from Independent control. However, an alternative cognitive explanation for how the monkeys completed the task may be one of response inhibition. Through this lens, the monkeys may be using similar control in both versions of the task, but with an aspect of inhibition/selection layered on top in the Independent Mode. The Independent trials may therefore be considered as a version of the (selective) stop signal task [76] or a go/no-go task, where it is unknown whether the second action should be released or

withheld until a cue is provided. Foreknowledge that a specific action may need to be withheld has been shown to produce focal corticomotor inhibition [33, 152].

If the Independent trials are being performed with a proactive and selective inhibition, but are otherwise similar to the Coordinated trials, then our result that the dynamics do not drastically change across contexts would not be surprising. Control in these two cases would be fundamentally the same. We did find that decoders trained on isolated unimanual movements generalized more poorly to Coordinated trials than to their Independent counterparts (Figure 4.3B, grey backgrounds). This could be in line with the inhibitory account, since the same inhibition would be present during the isolated unimanual trials (which consisted of Catch trials) as during the Independent Reach-Pull trials. The decoder would likely learn the relationship between neural activity and behavior with this inhibition taken into account, which is only maintained during the other Independent trials. However, that inhibition would be absent in the Coordinated Mode, which would counterintuitively cause a decrement in generalization performance. It is unclear whether this inhibition hypothesis could also explain the subtle differences in oscillatory frequency that we observed between Coordinated and Independent dynamics (Figure 4.6).

There are a few experiments and additional analyses that could help dissect this inhibition hypothesis. First, an additional block of trials performing purely unimanual reaches and another block performing purely unimanual joystick pulls would provide datapoints that isolate the two action components, yet lack the putative inhibitory process since they are outside the context of potential bimanual actions. If decoders still generalized more poorly to Coordinated trials than to their Independent counterparts, then the inhibition hypothesis would be unsupported. However, the hypothesis involving occupancy in different regions of the dynamical landscape would still be valid. Second, an analysis to specifically test for inhibition could be performed. One could compare mean firing rates immediately preceding the time when a cue would be delivered or omitted to assess any suppressive effects. It would likely be expected that this effect would be greatest in the hemisphere contralateral to the limb performing the second action, which in the case of the Reach-Pull Independent trials would be the left hemisphere. However, it should be noted again that the decoding effect we observed did not display a strong lateral bias.

The inhibition hypothesis may have broader applicability to theory on bimanual coordination. This hypothesis would suggest the Coordinated and Independent conditions are fundamentally the same in terms of control, and that Independent control simply consists of suppressing one component of the bimanual action. In a recent study developing shared control-based bimanual robotics, a “vocabulary” of bimanual actions was created by analyzing common patterns of human movement [178]. Each of the patterns identified by the authors consisted of either coordinated movement of the two

arms, or fixation of one arm while the other moved. It is plausible that the brain, and the contribution of the motor cortex, has a rigid default of coordinated dynamics. Attempts at independent movements of the two arms are often clumsy and may involve strategies of alternating rather than simultaneous control. Spatial and temporal interference patterns in bimanual action are well established [87, 124, 129, 130, 131, 135, 173]. The Independent trials in our task were very difficult to perform, for Monkey O in particular (Table S4.1, Figure S4.2A). It may therefore be more appropriate to ask how inhibition plays a role in disrupting coordinated dynamics, rather than how coordinated dynamics emerge from some other baseline.

4.4 Methods

4.4.1 Dataset details

Data from Monkey O were collected across 20 recording sessions. A total of 545 units were recorded from the left hemisphere, and 218 units from the right. Data from Monkey W were collected across 31 recording sessions. A total of 503 units were recorded from the left hemisphere, and 443 units from the right.

4.4.2 Task design

The behavioral task was loosely inspired by [173], which required subjects to open a drawer with one hand and retrieve an item from that drawer using the other. The task required monkeys to perform bimanual actions: reaching a target with the left hand and pulling a joystick with the right. There were two primary modes of the task, the Independent Mode and the Coordinated Mode.

During the Independent Mode, monkeys initiated a trial by placing the reaching hand in a central starting target (yellow sphere, 4cm radius) and were required to have no deflection of the joystick. A reaching target (yellow sphere, 3cm radius) was then displayed at one of six possible locations. These targets were arranged in the same configuration as the reaching task from the previous chapters. After a randomly sampled delay period (uniform, 1-2s), the first action was cued. To indicate a reach cue, the starting target disappeared, and the reaching target enlarged (4cm). To indicate a pull cue, the reaching target turned red until the pull was initiated. The second action was then cued at a fixed delay using the same stimulus associations as the first cue. For Monkey O, this delay period was 1s. For Monkey W, the delay time was either short or long. The short delays were 250ms for 5 of the sessions, and 300ms for the other 26. The long delays were 1.12s. On certain “catch” trials, no second action was cued, and the animal was required to complete the first action without initiating the other. This prevented the monkey from ignoring the second cue and learning to perform the second action at a fixed delay. If all cued actions were initiated in the requisite order and completed, a juice reward was delivered. The different stimulus presentations for the Independent Mode resulted in four trial types: Reach-Pull (top row), Pull-Reach (second row), Reach-Catch (third row), Pull-Catch (bottom row).

Coordinated Mode trials were performed with or without feedback. With Feedback trials were initiated in the same manner as the Independent Mode, and cues were presented using the same stimuli; however, the first action on all trials was the reach. In contrast to the fixed timing of the Independent Mode, the timing of the second action (pull) was contingent upon the execution of the reaching movement, with the pull cue presented at the moment the reaching hand entered the reach target radius. In No

Feedback trials, the behavioral constraints of the task were identical to the With Feedback trials; however, as soon as the reach was initiated the reaching cursor became invisible, and the pull cue was never delivered. These No Feedback trials required the monkeys to perform the task using only feed-forward planning and feedback provided by internal/proprioceptive state. The delay times in the Independent Mode were titrated during training to roughly match the behavior in the Coordinated Mode.

Independent and Coordinated trials were performed in fully separate blocks, typically with the Independent block being completed first in a session. At the beginning of Coordinated blocks, 4 beeps would sound to cue the monkey of the contextual switch. An acclimation period of 75 trials for Monkey O and 35 trials for Monkey W was removed from the start of each block for data analysis to ensure that only trials with steady state behavior were included.

4.4.3 Surgical implantation

All procedures were conducted in compliance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals and were approved by the University of California at Berkeley Institutional Animal Care and Use Committee under protocol ID AUP-2014-09-6720-1. This protocol approval and all surgical methods apply to each chapter of this thesis. Two adult male rhesus monkeys (*Macaca mulatta*) were implanted bilaterally with custom acute recording chambers (Grey Matter Research LLC, Bozeman, MT). Partial craniotomies within the chambers allowed access to the arm regions of dorsal premotor (PMd) and primary motor (M1) cortices in both hemispheres. Localization of target areas was performed using stereotactically aligned structural MRI collected just prior to implantation, alongside a neuroanatomical atlas of the rhesus brain [169].

4.4.4 Linear kinematic decoder

Prior to training kinematic decoders, spiking data were first binned in non-overlapping 100ms bins ranging from 200ms before instruction onset to 500ms after the peak of the return movement. Binned spike counts were then square-root transformed to stabilize variance and smoothed with a 100ms gaussian kernel. These firing rate estimates were then z-scored before model fitting. Linear models were then fit using ridge regression and ridge parameter 1. Kinematics were regressed on neural activity with 5 lagged time bins, i.e., extending 500ms into the past. Decoder performance for models trained and tested in the same trial type was assessed using leave-one-trial-out cross validation. We note that each task variant had a different number of trials, and no

corrections were made for these differing trial counts (see the final paragraph of Methods 4.4.5 for more discussion).

4.4.5 LDS fitting and testing

Prior to fitting LDS models, spiking data were first binned in non-overlapping 20ms bins ranging from 200ms before movement onset to 300ms after movement onset. Binned spike counts were then square-root transformed to stabilize variance and smoothed with a 20ms gaussian kernel. For time-shuffled models, binned spike counts were randomly reordered prior to the smoothing step.

We fit models of the form:

$$\mathbf{x}_{t+1} = A\mathbf{x}_t + B \quad (1)$$

by regressing neural activity on a time-lagged version of itself. Thus, here $\mathbf{x}_{t+1} \in \mathbb{R}^{n \times 24m}$ is the firing rate estimates of n neurons across m trials (24 samples per trial), and $\mathbf{x}_t \in \mathbb{R}^{n \times 24m}$ is the same neural activity one sample behind. Ridge regression with ridge parameter $1e-7$ was used for fitting $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times 1}$. By fitting a B term we allowed the model to find non-zero fixed points, which may be computed as:

$$\mathbf{x}_{FP} = (I - A)^{-1}B \quad (2)$$

Predictions of the one-sample-forward neural signals were obtained using 10-fold cross-validation when models were being evaluated on their own trial type, i.e. iteratively trained on $\mathbf{x}_{t,train}, \mathbf{x}_{t+1,train} \in \mathbb{R}^{n \times 24(m-1)}$, then tested on $\mathbf{x}_{t,test} \in \mathbb{R}^{n \times 24}$. For generalizing models across trial types, a single A_{trial_type} and B_{trial_type} were fit to the entire data that they were meant to represent, $\mathbf{x}_{t+1}, \mathbf{x}_t \in \mathbb{R}^{n \times 24m}$. These parameters were then used to predict the entire data from a different trial type, since those testing data were inherently unused in the model training. Model performance was evaluated using the coefficient of determination, R^2 .

We note that there were different numbers of trials in each task variant and no adjustments were made to account for this during model training. These models should be improved upon by adopting a cross-validation protocol that sets a fixed training sample size to be used for all task variants so that sample size does not have any effect on the results. However, our datasets are likely all large enough to where any effects due to differing trial counts would be minimal, yet this is not known for certain.

4.4.6 Characterization of dynamical modes

Eigendecomposition was performed on the A matrices fit to each dataset using the *eig* function in MATLAB. This decomposes the A matrix into a set of eigenvectors

contained in the matrix T and their associated eigenvalues located along the diagonal of matrix L , such that:

$$AT = TL \quad (3)$$

and

$$A = TLT^{-1} \quad (4)$$

The LDS expressed in equation 1 may therefore be re-expressed as:

$$x_{t+1} = TLT^{-1}x_t + B \quad (5)$$

By pre-multiplying both sides by T^{-1} and substituting $T^{-1}x$ with a new variable z , we may now observe the dynamics of a new matrix L acting on latent variable z :

$$z_{t+1} = Lz_t + T^{-1}B \quad (6)$$

Recall that L is a diagonal matrix of eigenvalues, making each dimension independent of one another.

The dynamical behavior of these separable “dynamical modes” can be characterized by their eigenvalues, which may be complex numbers in the form $l = a + bi$, where a is the real component and b is the imaginary component. We used the eigenvalues for each dynamical mode to compute the $\frac{1}{e}$ decay time as:

$$decay\ time = \frac{-(\text{bin width})}{\ln(\sqrt{a^2 + b^2})} \quad (7)$$

and the oscillation frequency as:

$$f = \frac{\tan^{-1}\left(\frac{b}{a}\right)}{2\pi(\text{bin width})} \quad (8)$$

where the bin width just provides the change in time for a single sample.

Decay times were used to determine “significance” of dynamic modes. Since the time-shuffled modes resulted solely from noise-fitting and data smoothing, their decay times should be fast. These values therefore represent a distribution of dynamic modes reflecting features that we are not interested in, and can be used to apply a cutoff for filtering dynamic modes that capture features that we are interested in. A cutoff below 100ms was used to filter out unwanted modes, which eliminated nearly the entire time-shuffled distribution (Figure S4.3). This left only modes that were persistent, i.e., not quickly decaying, in the dataset.

Please see [134] for an excellent python notebook tutorial from which most of these methods were adapted.

4.4.7 Procedure for “lesioning” dynamical modes

In order to determine the relative importance of individual dynamic modes, we performed a “lesioning” procedure which selectively eliminated specific eigenvectors from the A-matrix. In short, we decomposed the A-matrix, removed a single dynamical mode, then reconstructed the A-matrix from the remaining modes and re-tested performance in predicting future neural activity. The drop in performance, which we call R^2_{Les} , relates the relative importance of the lesioned dynamics for capturing neural variance. More precisely, given eigendecomposition in equation (4), we followed the protocol:

1. Remove column i from T , yielding $T_i \in \mathbb{R}^{n \times (n-1)}$
2. Remove row and column i from L , yielding $L_i \in \mathbb{R}^{(n-1) \times (n-1)}$
3. Compute new $T_i^{-1} \in \mathbb{R}^{(n-1) \times n}$ by taking the pseudo-inverse of T_i
4. Reconstruct new $A_i \in \mathbb{R}^{n \times n} = T_i L_i T_i^{-1}$
5. Reconstruct data matrix $\hat{x}_{t+1,i} = A_i x_t + B$

4.5 Supplementary Figures

Monkey O	Indep Total	Reach -First	Reach -Pull	Reach -Catch	Pull -First	Pull -Reach	Pull -Catch		Coord Total	With FB	No FB
Initiated	7123	3685	1523	1415	3128	1421	1517		6387	4566	1771
Early 1 st action	209 3%								187 3%		
Wrong 1 st action	101 1%								22 <1%		
Early 2 nd action	937 13%	747 20%			190 6%				1238 19%	856 19%	382 24%
Late 2 nd action	687 10%	633 17%	633 42%		54 2%	54 4%			8 <1%	6 <1%	2 <1%
Execution error	42 1%	0 0%	0 0%	0 0%	42 1%	42 3%	0 0%		0 0%	0 0%	0 0%
Success	5147 72%	2305 63%	890 58%	1415 100%	2842 91%	1421 94%	1421 100%		4932 77%	3704 81%	1228 76%

Table S4.1. Monkey O trial outcome proportions. Trial outcomes were coded according to five different error types, or success. Percentages reflect the proportion of initiated trials. The “Initiated” trials here refer to those that advanced past key milestones along the trial timeline. In the “Indep Total” and “Coord Total” columns, this refers to all trials where a first action (whether correct or incorrect) was initiated. In the “Reach-First,” “Pull-First,” “With FB,” and “No FB” columns, initiated trials are those where the first action was correctly initiated. In the “Reach-Pull” and “Pull-Reach” columns, initiated trials are those where the trial advanced past the go-cue for the second action. The only error type represented in the catch trial columns was execution error – the table does not discriminate between early 2nd actions on dual-action trials and un-cued 2nd actions on dual-action trials, which are both represented in the “Reach-First” or “Pull-First” columns as “Early 2nd action”.

Monkey W	Indep total	Reach -First	Reach -Pull	Reach -Catch	Pull -First	Pull -Reach	Pull -Catch		Coord total	With FB	No FB
Initiated	14214	7104	5433	1290	6809	5339	1306		13491	5596	7687
Early 1 st action	136 1%								89 1%		
Wrong 1 st action	165 1%								119 1%		
Early 2 nd action	545 4%	381 5%			164 2%				1506 11%	670 12%	836 11%
Late 2 nd action	236 2%	179 3%	179 3%		57 1%	57 1%			571 4%	314 6%	257 3%
Execution error	157 1%	72 1%	72 1%	0 0%	85 1%	85 2%	0 0%		44 <1%	44 1%	0 0%
Success	12975 91%	6472 91%	5182 95%	1290 100%	6503 96%	5197 97%	1306 100%		11162 83%	4568 82%	6594 86%

Table S4.2. Monkey W trial outcome proportions. Same as Table S4.1, except for Monkey W. Note that for Monkey W there were both long and short delay dual-action Independent trials, which is why there are more dual-action trials than catch trials (See Methods 4.4.2).

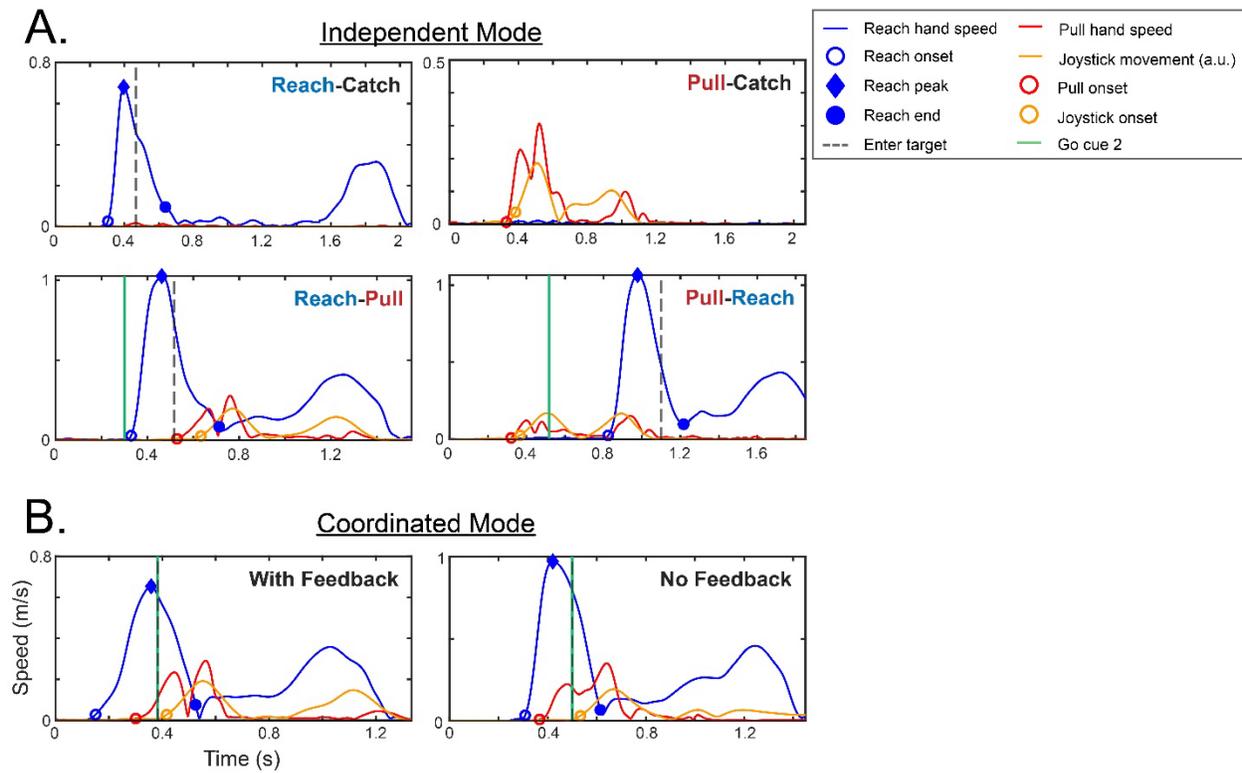


Figure S4.1 Single trial examples of behavioral kinematics. For each trial sub-type, an example single trial showing the speed of both hands along with the deflection of the joystick is displayed. Vertical lines indicate task events, and markers placed along the kinematic traces indicate behavioral events, e.g. the time of peak speed for the reaching hand (blue diamond). All examples are from Monkey W. A. Independent Mode trial sub-types, B. Coordinated Mode trial sub-types.

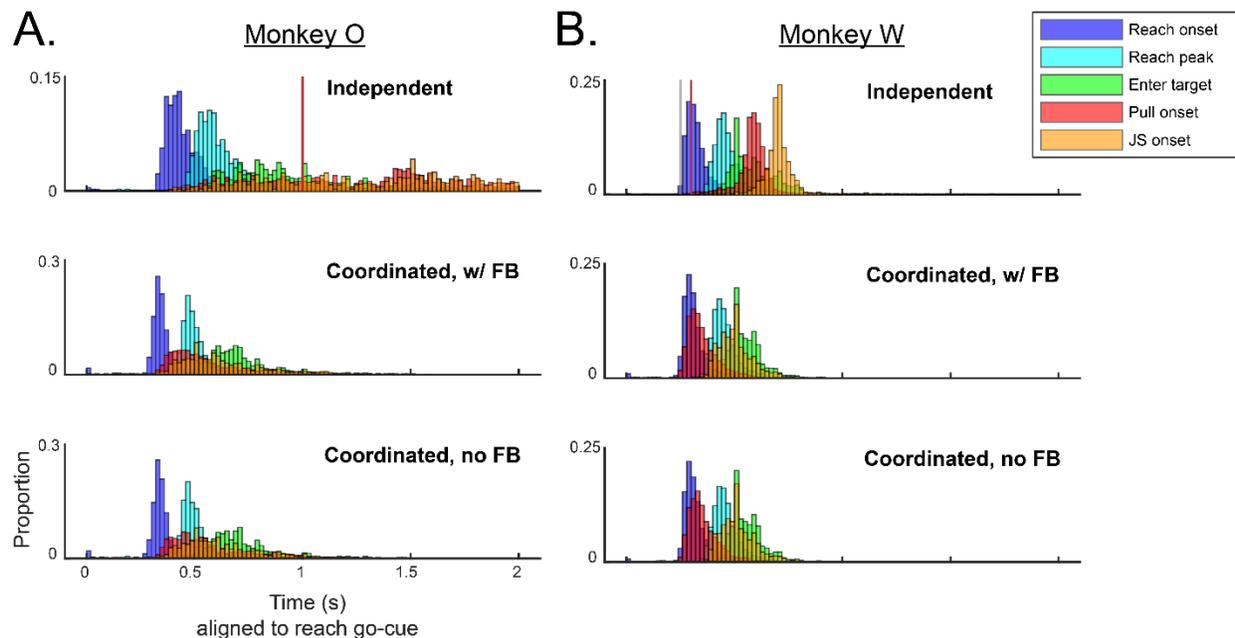


Figure S4.2 Distributions of behavioral events along the trial timeline. Each subplot displays the timing distributions for all behavioral events aligned to the reach cue. Plots are made for only the Independent Reach-Pull trials that were matched to the Coordinated behavior (top row), the Coordinated trials with visual feedback (middle row), and Coordinated trials without visual feedback (bottom row). Red lines in the top row show the time of the pull cue, which was at a fixed delay from the reach cue in Independent Reach-Pull trials (1s for Monkey O, 0.3s for Monkey W). In 5/31 sessions for Monkey W, an earlier pull cue was used, as indicated by the grey line (0.25s). A. Monkey O, B. Monkey W.

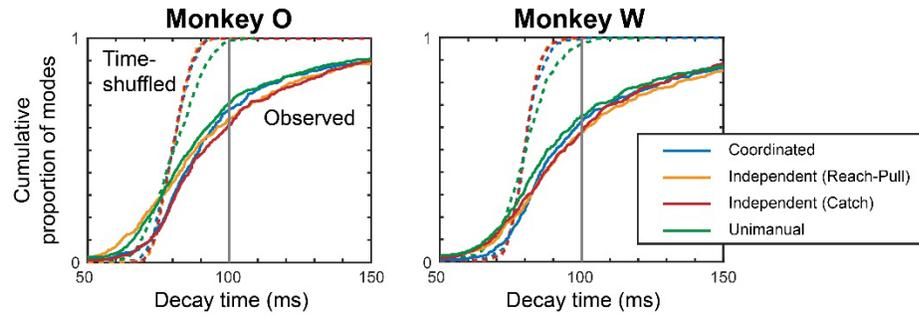


Figure S4.3. Determination of significant dynamic modes. These plots display the cumulative distributions of decay times associated with the dynamic modes fit in each dataset. Each line represents modes collected across all recording sessions. Dotted lines represent time-shuffled data. The time-shuffled data serve as a null distribution over decay time values since they are the sole result of noise-fitting and data smoothing, both of which should result in fast decay times. Vertical grey lines at 100ms indicate the cutoff for “significant” dynamic modes. Nearly the entire time-shuffled distribution lies to the left of this line, conservatively filtering out only modes from the observed data that were persistent (i.e., did not quickly decay) in the datasets. Note that these models produce as many dynamic modes as there are recorded units, so it should be expected that a small proportion of those would be considered significant.

Acknowledgements

We thank M. Kitano for help in NHP care and handling, and P. Khanna for helpful discussion of dynamical systems methodology. This work was supported by the National Defense Science and Engineering Graduate Fellowship to T.C.D, and a grant from the National Institute of Health (NS097480) to J.M.C.

Chapter 5

Conclusion

It is reasonable to assume that all movements we make with our arms are in fact bimanual. Some are more obvious than others: typing on a keyboard, tying your shoes, swinging a bat. Yet, even movements that seem to involve just one arm communicate forces across the body that may impact the contralateral limb. While a significant body of work coming from clinical neurology and experimental neuroscience have made clear that the motor cortex is principally involved in driving muscle activity on the opposite side of the body, the presence of activity related to bilateral movements may not be all too surprising when considering how we use our arms during daily activity. In this thesis, we have considered multiple hypotheses regarding laterality in the motor cortex and investigated them using a range of bimanual motor tasks. Here, we summarize the contributions of our work and discuss their implications for future avenues of research.

5.1 Summary of contributions

5.1.1 Independence and interaction of unimanual reaching signals

In Chapter 2 we investigate the organizational properties of population signals in the motor cortex as they evolve across the process of motor preparation and execution, with specific interest in laterality. This study can be thought of as linking and expanding upon two recent observations made in the field.

First, it has been shown that pre-movement activity develops within a neural subspace that is roughly orthogonal to the subspace spanned by peri-movement activity [80, 122]. A dominant theory surrounding motor cortical function posits that this allows preparatory neural activity to develop towards an optimal initial state without causing movement itself [5, 40, 43]. Second, two recent studies have found that motor cortical activity related to each arm resides in orthogonal neural subspaces as well [8, 108]. In

the same way that preparatory activity may refrain from causing movement, activity related to one arm may refrain from causing movement in the other arm. Together, these studies suggest a constant reorganization of the population statistics to serve different computational functions and target different effectors. How then can we integrate these ideas to gain a holistic understanding of how unimanual motor commands develop within the motor cortex? At what point in the process do signals for the two arms diverge? Precisely what organizational structure results in this divergence, and is there any heterogeneity in these properties?

We find that divergence of arm-specific neural subspaces occurs gradually throughout preparation and movement. The two phases did not appear categorically different from one another; rather, the same organizational principles that produced subspace separation were present in both phases but to different extents. Divergence of these subspaces was the unique result of units that were strongly dedicated to one arm or the other. This was commonly the contralateral arm. It has been previously observed that many units in the motor cortex exhibit modulation during movements of either arm [8, 45, 108], which has led to a notion that signals for the two arms are mixed within units [8]. This is a puzzling notion given our understanding of contralateral bias in the motor cortex. We found that units which were more strongly modulated also tended to have stronger arm preference, which resolves these two ideas by positing that bilaterally active units exist in high quantity yet still account for a relatively small portion of the population variance.

We additionally found that signals which were mixed within units did not distinguish between the two arms even in terms of their population covariance patterns. This component of the population response did contain behaviorally specific information such that reach targets could be classified from its activity. Interestingly, we did not find strong differences between PMd and M1 regarding these laterality properties. That is not to say that they were not different in general, as PMd was clearly more active during the preparatory period, but conditioned upon the units being active at all they were not readily distinguished by the metrics we assessed (e.g. during movement). Together, our results from Chapter 2 suggest a heterogeneity in the population responses within the motor cortex, consisting primarily of an arm-dedicated component yet also containing a component that is fundamentally bilateral in nature.

5.1.2 Integration of bilateral state for unilateral motor commands

The distributed component of Chapter 2 was highly relevant to the question asked in Chapter 3. Is bilateral state integrated into unilateral motor commands? We designed an experiment that tested the sensitivity of unimanual motor signals to posture of the stationary hand. We considered two main scenarios in which this effect would

exist. First, we considered whether covert motor commands for the stationary hand might exist in the motor cortex ipsilateral to the active hand. This activity would in fact reflect what the stationary hand would do if it were selected for movement. Second, we considered whether the state of the stationary hand might influence the neural activity associated with movements of the active arm for the purpose of facilitating bimanual interactions.

Our primary finding was that decoders trained to predict reaching movements were counterintuitively sensitive to the posture of the stationary hand, which was not directly involved in the action. The errors that these decoders made tended to miss in the same direction as the shift in posture of the stationary hand, which is a prediction that follows from the covert motor command hypothesis. However, we obtained mixed support for this hypothesis as response properties for the two hands were not strongly correlated in any particular reference frame, which would be another prediction.

An alternative interpretation centers on bimanual coordination or interaction. M1 has been shown to integrate information across multiple joints of a single limb in order to produce contextually appropriate feedback responses [177]. It may be the case that M1 and its connected secondary motor areas may perform a similar computation that integrates information across the two arms. One version of this hypothesis may focus on manipulatory actions, where commands provided to one arm need to know the state of the other when acting on a common object (e.g., opening a drawer to remove an item). Another version may draw stronger analogy to multi-joint integration. While the two arms can move with some level of independence, they are still mechanically coupled through the torso similar to how the joints of a single arm are connected through direct linkages. Movements that on the surface appear to isolate a single arm may in fact require coordinated whole-body control to counter the rotational forces communicated through the torso. Our results also lend strong evidence to these bimanual hypotheses, yet are limited by the lack of EMG recordings.

5.1.3 Interhemispheric population dynamics involved in bimanual coordination

Chapter 4 presents a novel task for distinguishing unimanual action from coordinated or independent (but simultaneous) bimanual action. Using this design, we asked if, and how, the population dynamics in each hemisphere adapt to different bimanual coordination requirements.

We found that kinematic decoders trained on isolated unimanual movements generalized poorly to coordinated bimanual trials, suggesting that the relationship between neural activity and behavior differed across control modes. The effect was not specific to ipsilateral activity, however. Despite this difference, we found little evidence that the neural population dynamics were adapting to the coordination requirements of

each task variant. These two primary results suggest that there may be a context-invariant dynamical landscape while the neural state in different contexts simply occupies different regions of that landscape. For example, the initial state may be seeded differently but obey the same flow-field (this specific idea has not been tested in the current work).

Our results may also point to a qualitatively different account of how independent and coordinated patterns of behavior are produced. Rather than flexibly switching between two control modes where the limbs are either treated as a unified plant or two independent plants, control may be always subject to certain bimanual constraints (i.e. coordinated). The presence of decorrelated behavior for the two arms may be the product of a separate process of response selection/gating, where the bimanual motor plan is adapted *post-hoc* by selectively withholding or modifying certain components. These processes would likely be distributed across the broader motor system and not simply reflected in motor cortical activity.

5.2 Future directions

A valuable complement to the work of this thesis would be comparison with activity in the SMA, a region commonly implicated in bilateral control. In particular, it would be useful to know how this area interacts with PMd/M1 during the behavioral tasks of Chapter 4. Our investigation of the interhemispheric dynamics across different bimanual behaviors (which may be expanded upon with further analysis) showed limited adaptation to the coordination requirements. Observing the interactions with both the ipsi- and contralateral SMA may offer new insights into the distributed computations underlying flexible adaptation of bimanual behavior.

The investigation of neural population dynamics in Chapter 4 was quite exploratory. Future work should be done to develop strong theoretical hypotheses about precisely how dynamics may change to create different coordinated profiles of behavior. A powerful tool for this would be recurrent neural networks (RNNs). Training an RNN to perform similar task-switching and observing the behavior of the model may provide key hypotheses that would guide improved testing in the biological analogs.

Our results from Chapter 2 suggest that there is a limited amount of independent ipsilateral activity in PMd and M1. These areas have historically been common targets for control signals in brain-machine interfaces (BMIs). An area of translational relevance for the work in this thesis is applying BMIs to restore mobility in patients with hemispheric stroke. With signals in one hemisphere no longer available, resulting in hemiparesis, one would hope to use a BMI to allow the intact hemisphere to do two jobs: isolate signals for reinnervating the paretic limb while continuing natural control of

the other. The weak independent ipsilateral signals that we observe present a critical engineering consideration. In order to achieve high quality control signals, many channels of recording may be required to “find” an appropriate amount of activity in a single hemisphere that is independent of the signals involved in contralateral control. BMI performance would also likely benefit from decoder training protocols that explicitly learn two independent readouts, e.g. training a decoder to not only create intended movements but to create no movement (learn a null space) during movements of the other limb. This would prevent the shared (or “distributed”) component of the population activity from causing inadvertent movement of the BMI during movement of the intact limb. Additionally, one could design a decoder that leverages the shared signal. To distinguish this activity from contralateral control, it would require a form of gating mechanism – for example, a two-stage decoder that first classifies which hand control signals belong to and then outputs a continuous motor signal. The independent ipsilateral activity could serve this gating function in addition to direct control.

An additional BMI consideration in light of our Chapter 4 results is whether a bimanual BMI using only control signals in PMd/M1 would be capable of producing the type of coordinated arm usage that occurs in natural control. Our results showed limited contextual information about the coordination requirements specific to each bimanual task variant. This study would be nicely complemented by a BMI equivalent testing whether flexible switching between coordinated patterns of behavior may be produced from control signals located purely in the same areas where we recorded.

5.3 Publications and presentations resulting from thesis

Main thesis work

1. **Dixon TC**, Merrick CM, Wallis JD, Ivry RB, Carmena JM. 2020. Hybrid dedicated and distributed coding in PMd/M1 provides separation and interaction of bilateral arm signals. *bioRxiv*. doi: 10.1101/2020.09.23.310664. *under review*.
2. **Dixon TC**, Merrick CM, Knight RT, Ivry RB, Carmena JM. 2019. Emergence of lateralized population activity in motor cortex across instructed-delay and execution phases of reaching. Poster presented at: The Society for Neuroscience 49th Annual Meeting; 2019 October 19-23; Chicago, IL.
3. **Dixon TC**, Merrick CM, Knight RT, Ivry RB, Carmena JM. 2019. Bilateral representation of reaching movements in motor cortex. Poster presented at: The

9th International IEEE EMBS Conference on Neural Engineering; 2019 March 20-23; San Francisco, CA.

4. **Dixon TC**, Merrick CM, Knight RT, Ivry RB, Carmena JM. 2018. Stability and independence of the ipsilateral representation of reaching movements in motor cortex. Poster presented at: The Society for Neuroscience 48th Annual Meeting; 2018 November 3-7; San Diego, CA.

Other

1. Merrick CM, **Dixon TC**, Lin JJ, Chang EF, Carmena JM, Knight RT, Ivry RB. 2019. Kinematic encoding model of ECoG activity in the contralateral and ipsilateral hemisphere during a reaching task. Poster, not presenting author: The Society for Neuroscience 49th Annual Meeting; 2019 October 19-23; Chicago, IL.
2. Kim HE, Morehead JR, Parvin DE, Moazzezi R, **Dixon TC**, Ivry RB. 2016. Behavioral changes induced by sensory prediction errors in the absence of task performance errors: Implications for models of sensorimotor adaptation. Poster, not presenting author: The Society for Neuroscience 46th Annual Meeting; 2016 November 12-16; San Diego, CA.

References

5. Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV. 2011. Single-trial neural correlates of arm movement preparation. *Neuron*, 71(3), 555–564. doi: 10.1016/j.neuron.2011.05.047
6. Alagona G, Delvaux V, Gérard P, De Pasqua V, Pennisi G, Delwaide PJ, Nicoletti F, Maertens De Noordhout A. 2001. Ipsilateral motor responses to focal transcranial magnetic stimulation in healthy subjects and acute-stroke patients. *Stroke*, 32(6), 1304–1309. doi: 10.1161/01.STR.32.6.1304
7. Albe-Fessard D, Liebeskind J. 1966. Origine des messages somato-sensitifs activant les cellules du cortex moteur chez le singe. *Experimental Brain Research*, 1(2), 127–146. doi: 10.1007/BF00236866
8. Ames KC, Churchland MM. 2019. Motor cortex signals for each arm are mixed across hemispheres and neurons yet partitioned within the population response. *eLife*, 8, 1–36. doi: 10.7554/eLife.46159
9. Anderson MJ, Ter Braak CJF. 2003. Permutation tests for multi-factorial analysis of variance. *Journal of Statistical Computation and Simulation*, 73(2), 85–113. doi: 10.1080/00949650215733
10. Athalye VR, Costa RM, Ganguly K, Costa RM, Carmena JM. 2017. Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control. *Neuron*, 93(4), 955–970.e5. doi: 10.1016/j.neuron.2017.01.016
11. Baker SN. 2011. The primate reticulospinal tract, hand function and functional recovery. *Journal of Physiology*, 589(23), 5603–5612. doi: 10.1113/jphysiol.2011.215160
12. Baker SN, Zaaimi B, Fisher KM, Edgley SA, Soteropoulos DS. 2015. Pathways mediating functional recovery. *Progress in Brain Research* (1st ed., Vol. 218). Elsevier B.V. doi: 10.1016/bs.pbr.2014.12.010
13. Bastian A, Schöner G, Riehle A. 2003. Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience*, 18(7), 2047–2058. doi: 10.1046/j.1460-9568.2003.02906.x
14. Bennett KMB, Lemon RN. 1996. Corticomotoneuronal contribution to the fractionation of muscle activity during precision grip in the monkey. *Journal of Neurophysiology*, 75(5), 1826–1842. doi: 10.1152/jn.1996.75.5.1826
15. Belhaj-Saïf A, Cheney PD. 2000. Plasticity in the distribution of the red nucleus output to forearm muscles after unilateral lesions of the pyramidal tract. *Journal of Neurophysiology*, 83(5), 3147–3153. doi: 10.1152/jn.2000.83.5.3147
16. Berlot E, Prichard G, O’Reilly J, Ejaz N, Diedrichsen J. 2019. Ipsilateral finger representations in the sensorimotor cortex are driven by active movement

- processes, not passive sensory input. *Journal of Neurophysiology*, 121(2), 418–426. doi: 10.1152/jn.00439.2018
17. Bernstein NA. 1967. *The Co-ordination and Regulation of Movements*. Pergamon Press, Oxford.
 18. Biasiucci A, Leeb R, Iturrate I, Perdikis S, Al-Khodairy A, Corbet T, Schnider A, Schmidlin T, Zhang H, Bassolino M, Viceic D, Vuadens P, Guggisberg AG, Millán, JDR. 2018. Brain-actuated functional electrical stimulation elicits lasting arm motor recovery after stroke. *Nature Communications*, 9(1), 1–13. doi: 10.1038/s41467-018-04673-z
 19. Bostan AC, Strick PL. 2018. The basal ganglia and the cerebellum: Nodes in an integrated network. *Nature Reviews Neuroscience*, 19(6), 338–350. doi: 10.1038/s41583-018-0002-7
 20. Brinkman J, Kuypers HGJM. 1973. Cerebral control of contralateral and ipsilateral arm, hand and finger movements in the split-brain rhesus monkey. *Brain*, 93, 653–674. doi: 10.1093/brain/96.4.653
 21. Brinkman C, Porter R. 1979. Supplementary motor area in the monkey: Activity of neurons during performance of a learned motor task. *Journal of Neurophysiology*, 42(3), 681–709. doi: 10.1152/jn.1979.42.3.681
 22. Brinkman C. 1981. Lesions in supplementary motor area interfere with a monkey's performance of a bimanual coordination task. *Neuroscience Letters*, 27(3), 267–270. doi: 10.1016/0304-3940(81)90441-9
 23. Brinkman C. 1984. Supplementary motor area of the monkey's cerebral cortex: Short- and long-term deficits after unilateral ablation and the effects of subsequent callosal section. *Journal of Neuroscience*, 4(4), 918–929. doi: 10.1523/jneurosci.04-04-00918.1984
 24. Brodal P. 1980a. The cortical projection to the nucleus reticularis tegmenti pontis in the rhesus monkey. *Experimental Brain Research*, 38(1), 19–27. doi: 10.1007/BF00237926
 25. Brodal P. 1980b. The projection from the nucleus reticularis tegmenti pontis to the cerebellum in the rhesus monkey. *Experimental Brain Research*, 38(1), 29–36. doi: 10.1007/BF00237927
 26. Buch E, Weber C, Cohen LG, Braun C, Dimyan MA, Ard T, Mellinger J, Caria A, Soekadar S, Fourkas A, Birbaumer N. 2008. Think to move: A neuromagnetic brain-computer interface (BCI) system for chronic stroke. *Stroke*, 39(3), 910–917. doi: 10.1161/STROKEAHA.107.505313
 27. Buford JA, Davidson AG. 2004. Movement-related and preparatory activity in the reticulospinal system of the monkey. *Experimental Brain Research*, 159(3), 284–300. doi: 10.1007/s00221-004-1956-4
 28. Bundy DT, Wronkiewicz M, Sharma M, Moran DW, Corbetta M, Leuthardt EC. 2012. Using ipsilateral motor signals in the unaffected cerebral hemisphere as a

- signal platform for brain-computer interfaces in hemiplegic stroke survivors. *Journal of Neural Engineering*, 9(3). doi: 10.1088/1741-2560/9/3/036011
29. Bundy DT, Souders L, Baranyai K, Leonard L, Schalk G, Coker R, Moran DW, Huskey T, Leuthardt EC. 2017. Contralesional Brain-Computer Interface Control of a Powered Exoskeleton for Motor Recovery in Chronic Stroke Survivors. *Stroke*, 48(7), 1908–1915. doi: 10.1161/STROKEAHA.116.016304
 30. Bundy DT, Szrama N, Pahwa M, Leuthardt EC. 2018. Unilateral, 3D arm movement kinematics are encoded in ipsilateral human cortex. *Journal of Neuroscience*, 38(47), 10042–10056. doi: 10.1523/JNEUROSCI.0015-18.2018
 31. Burle B, Bonnet M, Vidal F, Possamaï CA, Hasbroucq T. 2002. A transcranial magnetic stimulation study of information processing in the motor cortex: relationship between the silent period and the reaction time delay. *Psychophysiology*, 39(2), 207–217. doi: 10.1017/S0048577202010077
 32. Byblow WD, Summers JJ, Lewis GN, Thomas J. 2002. Bimanual coordination in Parkinson's disease: Deficits in movement frequency, amplitude, and pattern switching. *Movement Disorders*, 17(1), 20–29. doi: 10.1002/mds.1281
 33. Cai W, Oldenkamp CL, Aron AR. 2011. A proactive mechanism for selective suppression of response tendencies. *Journal of Neuroscience*, 31(16), 5965–5969. doi: 10.1523/JNEUROSCI.6292-10.2011
 34. Cardoso de Oliveira S, Gribova A, Donchin O, Bergman H, Vaadia E. 2001. Neural interactions between motor cortical hemispheres during bimanual and unimanual arm movements. *European Journal of Neuroscience*, 14(11), 1881–1896. doi: 10.1046/j.0953-816X.2001.01801.x
 35. Carey JR, Kimberley TJ, Lewis SM, Auerbach EJ, Dorsey L, Rundquist P, Ugurbil K. 2002. Analysis of fMRI and finger tracking training in subjects with chronic stroke. *Brain*, 125(4), 773–788. doi: 10.1093/brain/awf091
 36. Catsman-Berrevoets CE, Kuypers HGJM. 1976. Cells of origin of cortical projections to dorsal column nuclei, spinal cord, and bulbar medial reticular formation in the rhesus monkey. *Neuroscience Letters*, 3, 245–252. doi: 10.1016/0304-3940(76)90050-1
 37. Chaisangmongkon W, Swaminathan SK, Freedman DJ, Wang XJ. 2017. Computing by Robust Transience: How the Fronto-Parietal Network Performs Sequential, Category-Based Decisions. *Neuron*, 93(6), 1504-1517.e4. doi: 10.1016/j.neuron.2017.03.002
 38. Chen R, Yung D, Li JY. 2003. Organization of ipsilateral excitatory and inhibitory pathways in the human motor cortex. *Journal of Neurophysiology*, 89(3), 1256–1264. doi: 10.1152/jn.00950.2002
 39. Cheney PD, Fetz EE. 1980. Functional classes of primate corticomotoneuronal cells and their relation to active force. *Journal of Neurophysiology*, 44(4), 773–791. doi: 10.1152/jn.1980.44.4.773

40. Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV. 2006a. Neural Variability in Premotor Cortex Provides a Signature of Motor Preparation. *Journal of Neuroscience*, 26(14), 3697–3712. doi: 10.1523/JNEUROSCI.3762-05.2006
41. Churchland MM, Santhanam G, Shenoy KV. 2006b. Preparatory Activity in Premotor and Motor Cortex Reflects the Speed of the Upcoming Reach. *Journal of Neurophysiology*, 96(6), 3130–3146. doi: 10.1152/jn.00307.2006
42. Churchland MM, Shenoy KV. 2007. Delay of Movement Caused by Disruption of Cortical Preparatory Activity. *Journal of Neurophysiology*, 97(1), 348–359. doi: 10.1152/jn.00808.2006
43. Churchland MM, Cunningham JP, Kaufman M., Ryu SI, Shenoy KV. 2010. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron*, 68(3), 387–400. doi: 10.1016/j.neuron.2010.09.015
44. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. 2012. Neural population dynamics during reaching. *Nature*, 487(7405), 51–56. doi: 10.1038/nature11129
45. Cisek P, Crammond DJ, Kalaska JF. 2003. Neural activity in primary motor and dorsal premotor cortex in reaching tasks with the contralateral versus ipsilateral arm. *J Neurophysiol*, 89(2), 922–942. doi: 10.1152/jn.00607.2002
46. Cisek P, Kalaska JF. 2004. Neural correlates of mental rehearsal in dorsal premotor cortex. *Nature*, 431(7011), 993–996. doi: 10.1038/nature03005
47. Cisek P, Kalaska JF. 2005. Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. *Neuron*, 45(5), 801–814. doi: 10.1016/j.neuron.2005.01.027
48. Cisek P. 2006. Integrated neural processes for defining potential actions and deciding between them: A computational model. *Journal of Neuroscience*, 26(38), 9761–9770. doi: 10.1523/JNEUROSCI.5605-05.2006
49. Cisek P, Puskas GA, El-Murr S. 2009. Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience*, 29(37), 11560–11571. doi: 10.1523/JNEUROSCI.1844-09.2009
50. Cisek P, Kalaska JF. 2010. Neural Mechanisms for Interacting with a World Full of Action Choices. *Annual Review of Neuroscience*, 33(1), 269–298. doi: 10.1146/annurev.neuro.051508.135409
51. Cramer SC. 2008a. Repairing the human brain after stroke: I. Mechanisms of spontaneous recovery. *Annals of Neurology*, 63(3), 272–287. doi: 10.1002/ana.21393
52. Cramer SC. 2008b. Repairing the human brain after stroke. II. Restorative therapies. *Annals of Neurology*, 63(5), 549–560. doi: 10.1002/ana.21412
53. Crammond DJ, Kalaska JF. 1996. Differential relation of discharge in primary motor cortex and premotor cortex to movements versus actively maintained

- postures during a reaching task. *Exp Brain Res*, 108(1):45-61. doi: 10.1007/bf00242903
54. Crammond DJ, Kalaska JF. 2000. Prior information in motor and premotor cortex: Activity during the delay period and effect on pre-movement activity. *Journal of Neurophysiology*, 84(2), 986–1005. doi: 10.1152/jn.2000.84.2.986
 55. Cross KP, Heming EA, Cook DJ, Scott SH. 2020. Maintained Representations of the Ipsilateral and Contralateral Limbs during Bimanual Control in Primary Motor Cortex. *Journal of Neuroscience*, 40(35), 6732–6747. doi: 10.1523/JNEUROSCI.0730-20.2020
 56. Cunningham JP, Yu BM. 2014. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11), 1500–1509. doi: 10.1038/nn.3776
 57. Dąbrowska PA, Voges N, von Papen M, Ito J, Dahmen D, Riehle A, Brochier T, Grün S. 2020. On the complexity of resting state spiking activity in monkey motor cortex. *bioRxiv*. doi: 10.1101/2020.05.28.121095
 58. Dancause N. 2006. Vicarious function of remote cortex following stroke: Recent evidence from human and animal studies. *The Neuroscientist*, 12(6), 489–499. doi: 10.1177/1073858406292782
 59. Davidson AG, Buford JA. 2006. Bilateral actions of the reticulospinal tract on arm and shoulder muscles in the monkey: Stimulus triggered averaging. *Experimental Brain Research*, 173(1), 25–39. doi: 10.1007/s00221-006-0374-1
 60. Davidson AG, Schieber MH, Buford JA. 2007. Bilateral spike-triggered average effects in arm and shoulder muscles from the monkey pontomedullary reticular formation. *Journal of Neuroscience*, 27(30), 8053–8058. doi: 10.1523/JNEUROSCI.0040-07.2007
 61. Dayan P, Abbott LF. 2001. Theoretical neuroscience: computational and mathematical modeling of neural systems. Cambridge, MA: MIT Press
 62. de Lacoste MC, Kirkpatrick JB, Ross ED, 1985. Topography of the Human Corpus Callosum, *Journal of Neuropathology & Experimental Neurology*, 44(6):578–591. doi: 10.1097/00005072-198511000-00004
 63. Diedrichsen J, Hazeltine E, Kennerley S, Ivry RB. 2001. Moving to directly cued locations abolishes spatial interference during bimanual actions. *Psychological Science*, 12(6), 493–498. doi: 10.1111/1467-9280.00391
 64. Diedrichsen J, Nambisan R, Kennerley SW, Ivry RB. 2004. Independent on-line control of the two hands during bimanual reaching. *European Journal of Neuroscience*, 19(6), 1643–1652. doi: 10.1111/j.1460-9568.2004.03242.x
 65. Diedrichsen J, Verstynen T, Lehman SL, Ivry RB. 2005. Cerebellar involvement in anticipating the consequences of self-produced actions during bimanual movements. *Journal of Neurophysiology*, 93(2), 801–812. doi: 10.1152/jn.00662.2004

66. Diedrichsen J. 2007. Optimal Task-Dependent Changes of Bimanual Feedback Control and Adaptation. *Current Biology*, 17(19), 1675–1679. doi: 10.1016/j.cub.2007.08.051
67. Diedrichsen J, Gush S. 2009. Reversal of bimanual feedback responses with changes in task goal. *Journal of Neurophysiology*, 101(1), 283–288. doi: 10.1152/jn.90887.2008
68. Diedrichsen J, Shadmehr R, Ivry RB. 2010. The coordination of movement: optimal feedback control and beyond. *Trends in Cognitive Sciences*, 14(1), 31–39. doi: 10.1016/j.tics.2009.11.004
69. Diedrichsen J, Wiestler T, Krakauer JW. 2013. Two distinct ipsilateral cortical representations for individuated finger movements. *Cerebral Cortex*, 23(6), 1362–1377. doi: 10.1093/cercor/bhs120
70. di Pellegrino G, Fadiga L, Fogassi L, Gallese V, Rizzolatti G. 1992. Understanding motor events: a neurophysiological study. *Experimental Brain Research*. doi: 10.1007/BF00230027
71. Dixon TC, Merrick CM, Wallis JD, Ivry RB, Carmena JM. 2020. Hybrid dedicated and distributed coding in PMd/M1 provides separation and interaction of bilateral arm signals. *bioRxiv*. doi: 10.1101/2020.09.23.310664
72. Donchin O, Gribova A, Steinberg O, Bergman H, Vaadia E. 1998. Primary motor cortex is involved in bimanual coordination. *Nature*, 395(6699), 274–278. doi: 10.1038/26220
73. Dum RP, Strick PL. 1991. The origin of corticospinal projections from the premotor areas in the frontal lobe. *Journal of Neuroscience*, 11(3), 667–689. doi: 10.1523/jneurosci.11-03-00667.1991
74. Duque J, Ivry RB. 2009. Role of corticospinal suppression during motor preparation. *Cerebral Cortex*, 19(9), 2013–2024. doi: 10.1093/cercor/bhn230
75. Duque J, Lew D, Mazzocchio R, Olivier E, Ivry RB. 2010. Evidence for Two Concurrent Inhibitory Mechanisms during Response Preparation. *Journal of Neuroscience*, 30(10), 3793–3802. doi: 10.1523/JNEUROSCI.5722-09.2010
76. Duque J, Greenhouse I, Labruna L, Ivry RB. 2017. Physiological Markers of Motor Inhibition during Human Behavior. *Trends in Neurosciences*, 40(4), 219–236. doi: 10.1016/j.tins.2017.02.006
77. Dushanova J, Donoghue J. 2010. Neurons in primary motor cortex engaged during action observation. *European Journal of Neuroscience*, 31(2), 386–398. doi: 10.1111/j.1460-9568.2009.07067.x
78. Eccher M. 2014. Corpus Callosum. In *Encyclopedia of the Neurological Sciences* (Second Edition), 1, 867–868. doi: 10.1016/B978-0-12-385157-4.01137-4
79. Eliassen JC, Baynes K, Gazzaniga MS. 1999. Direction information coordinated via the posterior third of the corpus callosum during bimanual movements. *Experimental Brain Research*, 128(4), 573–577. doi: 10.1007/s002210050884

80. Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP. 2016. Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications*, 13239. doi: 10.1038/ncomms13239
81. Erlhagen W, Schöner G. 2002. Dynamic field theory of movement preparation. *Psychological Review*, 109(3), 545–572. doi: 10.1037/0033-295X.109.3.545
82. Evarts EV. 1968. Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of Neurophysiology*, 31(1), 14–27. doi: 10.1152/jn.1968.31.1.14
83. Felten DL, O'Banion MK, Maida MS. 2016. Motor Systems. *Netter's Atlas of Neuroscience*, 391–420. doi: 10.1016/b978-0-323-26511-9.00015-1
84. Fetz EE, Finocchio DV, Baker MA, Soso MJ. 1980. Sensory and motor responses of precentral cortex cells during comparable passive and active joint movements. *Journal of Neurophysiology*, 43(4), 1070–1089. doi: 10.1152/jn.1980.43.4.1070
85. Fetz EE, Shupe LE, Murthy VN. 1990. Neural networks controlling wrist movements. *Proceedings of 1990 International Joint Conference on Neural Networks*, (II), 675–679. doi: 10.1109/ijcnn.1990.137778
86. Fetz EE. 1992. Are movement parameters recognizably coded in the activity of single neurons? *Movement Control*. doi: 10.1017/cbo9780511529788.008
87. Franz EA, Zelaznik HN, McCabe G. 1991. Spatial topological constraints in a bimanual task. *Acta Psychologica*, 77(2), 137–151. doi: 10.1016/0001-6918(91)90028-X
88. Franz EA, Eliassen JC, Ivry RB, Gazzaniga MS. 1996a. Dissociation of spatial and temporal coupling in the bimanual movements of callosotomy patients. *Psychological Science*, 7(5), 306–310. doi: 10.1111/j.1467-9280.1996.tb00379.x
89. Franz EA, Ivry RB, Helmuth LL, Franz L. 1996b. Reduced timing variability in patients with unilateral cerebellar lesions during bimanual movements. *Journal of Cognitive Neuroscience*, 8(2), 107–118. doi: 10.1162/jocn.1996.8.2.107
90. Franz EA, Waldie KE, Smith MJ. 2000. The effect of callosotomy on novel versus familiar bimanual actions: A neural dissociation between controlled and automatic processes? *Psychological Science*, 11(1), 82–85. doi: 10.1111/1467-9280.00220
91. Frolov AA, Mokienko O, Lyukmanov R, Biryukova E, Kotov S, Turbina L, Nadareyshvily G, Bushkova Y. 2017. Post-stroke rehabilitation training with a motor-imagery-based brain-computer interface (BCI)-controlled hand exoskeleton: A randomized controlled multicenter trial. *Frontiers in Neuroscience*, 11(JUL). doi: 10.3389/fnins.2017.00400
92. Ganguly K, Secundo L, Ranade G, Orsborn A, Chang EF, Dimitrov DF, Wallis JD, Barbaro NM, Knight RT, Carmena JM. 2009. Cortical representation of ipsilateral arm movements in monkey and man. *Journal of Neuroscience*, 29(41), 12948–12956. doi: 10.1523/JNEUROSCI.2471-09.2009

93. Gazzaniga MS. 1989. Organization of the human brain. *Science*, 245(4921), 947 LP – 952. doi: 10.1126/science.2672334
94. Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. 1982. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of Neuroscience*, 2(11), 1527–1537. doi: 10.1523/jneurosci.02-11-01527.1982
95. Georgopoulos AP, Schwartz AB, Kettner RE. 1986. Neuronal population coding of movement direction. *Science*, 233(4771), 1416–1419. doi: 10.1126/science.3749885
96. Gerloff C, Andres FG. 2002. Bimanual coordination and interhemispheric interaction. *Acta Psychologica*. doi: 10.1016/s0001-6918(02)00032-x
97. Gibson BYAR, Houk JC, Kohlerman NJ. 1985. Magnocellular red nucleus activity during different types of limb movement in the macaque monkey. *J. Physiol.*, 358, 527–549. doi: 10.1113/jphysiol.1985.sp015565
98. Gielen CC, Ramaekers L, van Zuylen EJ. 1988. Long-latency stretch reflexes as co-ordinated functional responses in man. *The Journal of Physiology*, 407(1), 275–292. doi: 10.1113/jphysiol.1988.sp017415
99. Graziano MSA, Patel KT, Taylor CSR. 2004. Mapping from motor cortex to biceps and triceps altered by elbow angle. *Journal of Neurophysiology*, 92(1), 395–407. doi: 10.1152/jn.01241.2003
100. Grefkes C, Fink GR. 2011. Reorganization of cerebral networks after stroke: New insights from neuroimaging with connectivity approaches. *Brain*, 134(5), 1264–1276. doi: 10.1093/brain/awr033
101. Grossberg S. 1973. Contour Enhancement, Short Term Memory, and Constancies in Reverberating Neural Networks. *Studies in Applied Mathematics*, 52(3), 213–257. doi: 10.1002/sapm1973523213
102. Haggard P, Hutchinson K, Stein J. 1995. Patterns of coordinated multi-joint movement. *Experimental Brain Research*, 107(2), 254–266. doi: 10.1007/BF00230046
103. Haith AM, Pakpoor J, Krakauer JW. 2016. Independence of Movement Preparation and Movement Initiation. *Journal of Neuroscience*, 36(10), 3007–3015. doi: 10.1523/JNEUROSCI.3245-15.2016
104. Haken H, Kelso JAS, Bunz H. 1985. A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, 51(5), 347–356. doi: 10.1007/BF00336922
105. Hatem SM, Saussez G, della Faille M, Prist V, Zhang X, Dispa D, Bleyenheuft Y. 2016. Rehabilitation of motor function after stroke: A multiple systematic review focused on techniques to stimulate upper extremity recovery. *Frontiers in Human Neuroscience*, 10(SEP2016), 1–22. doi: 10.3389/fnhum.2016.00442

106. Hatsopoulos NG, Suminski AJ. 2011. Sensing with the motor cortex. *Neuron*, 72(3), 477–487. doi: 10.1016/j.neuron.2011.10.020
107. Helmuth LL, Ivry RB. 1996. When Two Hands Are Better Than One: Reduced Timing Variability during Bimanual Movements. *Journal of Experimental Psychology: Human Perception and Performance*, 22(2), 278–293. doi: 10.1037/0096-1523.22.2.278
108. Heming EA, Cross KP, Takei T, Cook DJ, Scott SH. 2019. Independent representations of ipsilateral and contralateral limbs in primary motor cortex. *eLife*, 8(2002), 1–26. doi: 10.7554/eLife.48190
109. Hongo T, Jankowska E, Lundberg A. 1969. The rubrospinal tract. I. Effects on alpha-motoneurons innervating hindlimb muscles in cats. *Experimental Brain Research*, 7(4), 344–364. doi: 10.1007/BF00237320
110. Hoshi E, Tanji J. 2002. Contrasting neuronal activity in the dorsal and ventral premotor areas during preparation to reach. *Journal of Neurophysiology*, 87(2), 1123–1128. doi: 10.1152/jn.00496.2001
111. Hoshi E, Tanji J. 2007. Distinctions between dorsal and ventral premotor areas: anatomical connectivity and functional properties. *Current Opinion in Neurobiology*, 17(2), 234–242. doi: 10.1016/j.conb.2007.02.003
112. Hummel FC, Cohen LG. 2006. Non-invasive brain stimulation: a new strategy to improve neurorehabilitation after stroke? *The Lancet Neurology*, 5(8), 708–712. doi: 10.1016/s1474-4422(06)70525-7
113. Ifft PJ, Shokur S, Li Z, Lebedev MA, Nicolelis MAL. 2013. A Brain-Machine Interface Enables Bimanual Arm Movements in Monkeys. *Science Translational Medicine*, 5(210), 1-13. doi: 10.1126/scitranslmed.3006159
114. Ivry RB, Hazeltine E. 1999. Subcortical locus of temporal coupling in the bimanual movements of a callosotomy patient. *Human Movement Science*, 18, 345-375. doi: 10.1016/S0167-9457(99)00014-7
115. Ivry RB, Richardson TC. 2002. Temporal control and coordination: the multiple timer model. *Brain and Cognition*, 48(1), 117–132. doi: 10.1006/breg.2001.1308
116. Ivry RB, Diedrichsen J, Spence RM, Hazeltine E, Semjen A. 2004. A cognitive neuroscience perspective on bimanual coordination and interference. In: *Interlimb Coordination*, edited by Swinnen S, Duysens J. Boston, MA: Kluwer Academic, p. 259 –295.
117. Johnson KA, Cunnington R, Bradshaw JL, Phillips JG, Lansak R, Rogers MA. 1998. Bimanual co-ordination in Parkinson's disease. *Brain*, 121(4), 743–753. doi: 10.1093/brain/121.4.743
118. Kakei S, Hoffman DS, Strick PL. 1999. Muscle and movement representations in the primary motor cortex. *Science*, 285(5436), 2136–2139. doi: 10.1126/science.285.5436.2136

119. Kao JC, Nuyujukian P, Ryu SI, Churchland MM, Cunningham JP, Shenoy KV. 2015. Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nature Communications*, 6(May). doi: 10.1038/ncomms8759
120. Kaufman MT, Churchland MM, Santhanam G, Yu BM, Afshar A, Ryu SI, Shenoy KV. 2010. Roles of monkey premotor neuron classes in movement preparation and execution. *Journal of Neurophysiology*, 104(2), 799–810. doi: 10.1152/jn.00231.2009
121. Kaufman MT, Churchland MM, Shenoy KV. 2013. The roles of monkey M1 neuron classes in movement preparation and execution. *Journal of Neurophysiology*, 110(4), 817–825. doi: 10.1152/jn.00892.2011
122. Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. 2014. Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci*, 17(3), 110–111. doi: 10.1038/nn.3643
123. Kaufman MT, Seely JS, Sussillo D, Ryu SI, Shenoy KV, Churchland MM. 2016. The largest response component in the motor cortex reflects movement timing but not movement type. *eNeuro*, 3(4). doi: 10.1523/ENEURO.0085-16.2016
124. Kazennikov O, Wicki U, Corboz M, Hyland B, Palmeri A, Rouiller EM, Wiesendanger M. 1994. Temporal structure of a bimanual goal-directed movement sequence in monkeys. *The European Journal of Neuroscience*, 6(2), 203–210. doi: 10.1111/j.1460-9568.1994.tb00262.x
125. Kazennikov O, Hyland B, Corboz M, Babalian A, Rouiller EM, Wiesendanger M. 1999. Neural activity of supplementary and primary motor areas in monkeys and its relation to bimanual and unimanual movement sequences. *Neuroscience*, 89(3), 661–674. doi: 10.1016/S0306-4522(98)00348-0
126. Keizer K, Kuypers HGJM. 1989. Distribution of corticospinal neurons with collaterals to the lower brain stem reticular formation in monkey (*Macaca fascicularis*). *Experimental Brain Research*, 311–318. doi: 10.1007/bf00248864
127. Kelly RM, Strick, PL. 2003. Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate. *Journal of Neuroscience*, 23(23), 8432–8444. doi: 10.1523/jneurosci.23-23-08432.2003
128. Kelly RM, Strick PL. 2004. Macro-architecture of basal ganglia loops with the cerebral cortex: use of rabies virus to reveal multisynaptic circuits. *Prog Brain Res*, 143:449-59. doi: 10.1016/s0079-6123(03)43042-2
129. Kelso JAS, Southard DL, Goodman D. 1979. On the coordination of two-handed movements. *Journal of Experimental Psychology: Human Perception and Performance*, 5(2), 229–238. doi: 10.1037/0096-1523.5.2.229
130. Kelso JAS, Putnam CA, Goodman D. 1983. On the space-time structure of human interlimb co-ordination. *The Quarterly Journal of Experimental Psychology Section A*, 35(2), 347–375. doi: 10.1080/14640748308402139

131. Kelso JAS. 1984. Phase transitions and critical behavior in human bimanual coordination. *American Journal of Physiology - Regulatory Integrative and Comparative Physiology*, 15(6). doi: 10.1152/ajpregu.1984.246.6.r1000
132. Kennerley SW, Diedrichsen J, Hazeltine E, Semjen A, Ivry RB. 2002. Callosotomy patients exhibit temporal uncoupling during continuous bimanual movements. *Nature Neuroscience*, 5(4), 376–381. doi: 10.1038/nn822
133. Kermadi I, Liu Y, Tempini A, Calciati E, Rouiller EM. 1998. Neuronal activity in the primate supplementary motor area and the primary motor cortex in relation to spatio-temporal bimanual coordination. *Somatosensory and Motor Research*, 15(4), 287–308. doi: 10.1080/08990229870709
134. Khanna P. 2020. Dynamics Notebook [Source code].
https://github.com/pkhanna104/dynamics_notebook.
135. Klapp ST. 1979. Doing two things at once: The role of temporal compatibility. *Memory & Cognition*, 7(5), 375–381. doi: 10.3758/BF03196942
136. Kobak D, Brendel W, Constantinidis C, Feierstein CE, Kepecs A, Mainen ZF, Xue-Lian Q, Romo R, Uchida N, Machens CK. 2016. Demixed principal component analysis of neural population data. *eLife*, 5(APRIL2016), 1–36. doi: 10.7554/eLife.10989
137. Künzle H. 1975. Bilateral projections from precentral motor cortex to the putamen and other parts of the basal ganglia. An autoradiographic study in *Macaca fascicularis*. *Brain Research*, 88(2), 195–209. doi: 10.1016/0006-8993(75)90384-4
138. Kuypers HGJM, Lawrence DG. 1967. Cortical projections to the red nucleus and the brain stem in the rhesus monkey. *Brain Research*, 4, 151–188. doi: 10.1016/0006-8993(67)90004-2
139. Kuypers HGJM. 1981. Anatomy of the Descending Pathways. *Comprehensive Physiology*, 597–666. doi: 10.1002/cphy.cp010213
140. Lacroix S, Havton LA, McKay H, Yang H, Brant A, Roberts J, Tuszynski MH. 2004. Bilateral Corticospinal Projections Arise from Each Motor Cortex in the Macaque Monkey: A Quantitative Study. *Journal of Comparative Neurology*, 473(2), 147–161. doi: 10.1002/cne.20051
141. Lashley KS. 1933. Integrative functions of the cerebral cortex. *Physiol. Rev.*, 13: 1–42. doi: 10.1152/physrev.1933.13.1.1
142. Lawrence DG, Kuypers HGJM. 1968a. The functional organization of the motor system in the monkey: I. The effects of bilateral pyramidal lesions. *Brain*, 91(1), 1–14. doi: 10.1093/brain/91.1.1
143. Lawrence DG, Kuypers HGJM. 1968b. The functional organization of the motor system in the monkey: II. The effects of lesions of the descending brain-stem pathways. *Brain*, 91(1), 15–36. doi: 10.1093/brain/91.1.15

144. Lee BY, Zhu XH, Li X, Chen W. 2019. High-resolution imaging of distinct human corpus callosum microstructure and topography of structural connectivity to cortices at high field. *Brain Structure and Function*, 224(2), 949–960. doi: 10.1007/s00429-018-1804-0
145. Lemon RN, Hanby JA, Porter R. 1976. Relationship between the activity of precentral neurones during active and passive movements in conscious monkeys. *Proceedings of the Royal Society of London - Biological Sciences*, 194(1116), 341–373. doi: 10.1098/rspb.1976.0083
146. Lemon RN. 1993. Cortical control of the primate hand. *Experimental Physiology*, 78(3), 263–263. doi: 10.1113/expphysiol.1993.sp003686
147. Lemon RN. 2008. Descending pathways in motor control. *Annual Review of Neuroscience*, 31(Cm), 195–218. doi: 10.1146/annurev.neuro.31.060407.125547
148. Lemon RN. 2016. Cortical projections to the red nucleus and the brain stem in the rhesus monkey. *Brain Research*, 1645, 28–30. doi: 10.1016/j.brainres.2016.01.006
149. Leocani L, Cohen LG, Wassermann EM, Ikoma K, Hallett M. 2000. Human corticospinal excitability evaluated with transcranial magnetic stimulation during different reaction time paradigms. *Brain*, 123(6), 1161–1173. doi: 10.1093/brain/123.6.1161
150. Li N, Chen T, Guo ZV, Gerfen CR, Svoboda K. 2015. A motor cortex circuit for motor planning and movement. *Nature*, 519(7541), 51–56. doi: 10.1038/nature14178
151. Li N, Daie K, Svoboda K, Druckmann S. 2016. Robust neuronal dynamics in premotor cortex during motor planning. *Nature*, 537(7618), 122. doi: 10.1038/nature18623
152. Majid DSA, Cai W, George JS, Verbruggen F, Aron AR. 2012. Transcranial magnetic stimulation reveals dissociable mechanisms for global versus selective corticomotor suppression underlying the stopping of action. *Cerebral Cortex*, 22(2), 363–371. doi: 10.1093/cercor/bhr112
153. Mark RF, Sperry RW. 1968. Bimanual coordination in monkeys. *Experimental Neurology*, 21(1), 92–104. doi: 10.1016/0014-4886(68)90036-8
154. Matsunami K, Hamada I. 1981. Characteristics of the ipsilateral movement-related neuron in the motor cortex of the monkey. *Brain Research*, 204(1), 29–42. doi: 10.1016/0006-8993(81)90649-1
155. Maxwell DJ, Soteropoulos DS. 2020. The mammalian spinal commissural system: Properties and functions. *Journal of Neurophysiology*, 123(1), 4–21. doi: 10.1152/jn.00347.2019
156. McCall AA, Miller DM, Yates BJ. 2017. Descending influences on vestibulospinal and vestibulosympathetic reflexes. *Frontiers in Neurology*, 8(MAR), 1–15. doi: 10.3389/fneur.2017.00112

157. McGuire PK, Bates JF, Goldman-Rakic PS. 1991a. Interhemispheric Integration: I. Symmetry and Convergence of the Corticocortical Connections of the Left and the Right Principal Sulcus (PS) and the Left and the Right Supplementary Motor Area (SMA) in the Rhesus Monkey. *Cerebral Cortex*, 1(5), 390–497. doi: 10.1093/cercor/1.5.390
158. McGuire PK, Bates JF, Goldman-Rakic PS. 1991b. Interhemispheric integration: II. Symmetry and convergence of the corticostriatal projections of the left and the right principal sulcus (PS) and the left and the right supplementary motor area (SMA) of the rhesus monkey. *Cerebral Cortex*, 1(5), 408–417. doi: 10.1093/cercor/1.5.408
159. McMillan S, Ivry RB, Byblow WD. 2006. Corticomotor excitability during a choice-hand reaction time task. *Experimental Brain Research*, 172(2), 230–245. doi: 10.1007/s00221-005-0331-4
160. Mechsner F, Kerzel D, Knoblich G, Prinz W. 2001. Perceptual basis of bimanual coordination. *Nature*, 414(6859), 69–73. doi: 10.1038/35102060
161. Mewes K, Cheney PD. 1991. Facilitation and suppression of wrist and digit muscles from single rubromotoneuronal cells in the awake monkey. *Journal of Neurophysiology*, 66:6, 1965-1977. doi: 10.1152/jn.1991.66.6.1965
162. Montgomery LR, Herbert WJ, Buford JA. 2013. Recruitment of ipsilateral and contralateral upper limb muscles following stimulation of the cortical motor areas in the monkey. *Experimental Brain Research*, 230(2), 153–164. doi: 10.1007/s00221-013-3639-5
163. Moran DW, Schwartz AB. 1999. Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology*, 82(5), 2676–2692. doi: 10.1152/jn.1999.82.5.2676
164. Murase N, Duque J, Mazzocchio R, Cohen LG. 2004. Influence of interhemispheric interactions on motor function in chronic stroke. *Annals of Neurology*, 55(3), 400–409. doi: 10.1002/ana.10848
165. Ochiai T, Mushiake H, Tanji J. 2002. Effects of image motion in the dorsal premotor cortex during planning of an arm movement. *Journal of Neurophysiology*, 88(4), 2167–2171. doi: 10.1152/jn.2002.88.4.2167
166. Ono T, Shindo K, Kawashima K, Ota N, Ito M, Ota T, Mukaino M, Fujiwara T, Kimura A, Liu M, Ushiba J. 2014. Brain-computer interface with somatosensory feedback improves functional recovery from severe hemiplegia due to chronic stroke. *Frontiers in Neuroengineering*, 7(July), 1–8. doi: 10.3389/fneng.2014.00019
167. Pandya DN, Karol EA, Heilbronn D. 1971. The topographical distribution of interhemispheric projections in the corpus callosum of the rhesus monkey. *Brain Research*, 32(1), 31–43. doi: 10.1016/0006-8993(71)90153-3

168. Papaioannou S, Dimitriou M. 2021. Goal-dependent tuning of muscle spindle receptors during movement preparation. *Science Advances*, 7(9), 1–14. doi: 10.1126/sciadv.abe0401
169. Paxinos G, Huang XF, Toga AW. 2000. The Rhesus Monkey Brain in Stereotaxic Coordinates. Academic Press.
170. Penfield W, Boldrey E. 1937. Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*, 389–443. doi: 10.1093/brain/60.4.389
171. Penfield W, Welch K. 1951. The supplementary motor area of the cerebral cortex: A clinical and experimental study. *AMA Archives of Neurology and Psychiatry*, 66(3), 289–317. doi: 10.1001/archneurpsyc.1951.02320090038004
172. Penfield W, Jasper H. 1954. Epilepsy and the Functional Anatomy of the Human Brain, Little Brown, Boston.
173. Perrig S, Kazennikov O, Wiesendanger M. 1999. Time structure of a goal-directed bimanual skill and its dependence on task constraints. *Behavioural Brain Research*, 103(1), 95–104. doi: 10.1016/S0166-4328(99)00026-1
174. Pompeiano O. 1972. Spinovestibular Relations: Anatomical and Physiological Aspects. *Progress in Brain Research*, 37(C), 263–296. doi: 10.1016/S0079-6123(08)63907-2
175. Preilowski BFB. 1972. Possible contribution of the anterior forebrain commissures to bilateral motor coordination. *Neuropsychologia*, 10(3), 267–277. doi: 10.1016/0028-3932(72)90018-8
176. Preilowski B. 1975. Bilateral Motor Interaction: Perceptual-Motor Performance of Partial and Complete “Split-Brain” Patients. *Cerebral Localization*, 115–132. doi: 10.1007/978-3-642-66204-1_9
177. Pruszynski JA, Kurtze I, Nashed JY, Omrani M, Brouwer B, Scott SH. 2011. Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature*, 478(7369), 387–390. doi: 10.1038/nature10436
178. Rakita D, Mutlu B, Gleicher M, Hiatt LM. 2019. Shared control-based bimanual robot manipulation. *Science Robotics*, 4(30). doi: 10.1126/scirobotics.aaw0955
179. Rathelot JA, Strick PL. 2009. Subdivisions of primary motor cortex based on cortico-motoneuronal cells. *Proceedings of the National Academy of Sciences*, 106(3), 918–923. doi: 10.1073/PNAS.0808362106
180. Rich E, Wallis JD. 2016. Decoding subjective decisions from orbitofrontal cortex supplementary. *Nature Neuroscience*, 19(7), 973–980. doi: 10.1038/nn.4320
181. Riehle A, Requin J. 1993. The predictive value for performance speed of preparatory changes in neuronal activity of the monkey motor and premotor cortex. *Behavioural Brain Research*, 53(1–2), 35–49. doi: 10.1016/S0166-4328(05)80264-5

182. Rosenbaum DA. 1980. Human movement initiation: Specification of arm, direction, and extent. *Journal of Experimental Psychology: General*, 109(4), 444–474. doi: 10.1037/0096-3445.109.4.444
183. Rosenzweig ES, Brock JH, Culbertson MD, Lu P, Moseanko R, Edgerton VR, Havton LA, Tuszynski MH. 2009. Extensive spinal decussation and bilateral termination of cervical corticospinal projections in rhesus monkeys. *J Comp Neurol*, 513(2), 151–163. doi: 10.1002/cne.21940
184. Rouiller EM, Babalian A, Kazennikov O, Moret V, Yu XH, Wiesendanger M. 1994. Transcallosal connections of the distal forelimb representations of the primary and supplementary motor cortical areas in macaque monkeys. *Experimental Brain Research*, 102(2), 227–243. doi: 10.1007/BF00227511
185. Russo AA, Khajeh R, Bittner SR, Perkins SM, Cunningham JP, Abbott LF, Churchland MM. 2020. Neural Trajectories in the Supplementary Motor Area and Motor Cortex Exhibit Distinct Geometries, Compatible with Different Classes of Computation. *Neuron*, 107(4), 745–758.e6. doi: 10.1016/j.neuron.2020.05.020
186. Schaechter JD, Kraft E, Hilliard TS, Dijkhuizen RM, Benner T, Finklestein SP, Rosen BR, Cramer SC. 2002. Motor Recovery and Cortical Reorganization after Constraint-Induced Movement Therapy in Stroke Patients: A Preliminary Study. *Neurorehabilitation and Neural Repair*, 16(4), 326–338. doi: 10.1177/154596830201600403
187. Schmammann JD, Pandya DN. 1997. Anatomic organization of the basilar pontine projections from prefrontal cortices in rhesus monkey. *Journal of Neuroscience*, 17(1), 438–458. doi: 10.1523/jneurosci.17-01-00438.1997
188. Scholz JP, Schöner G. 1999. The uncontrolled manifold concept: Identifying control variables for a functional task. *Experimental Brain Research*, 126(3), 289–306. doi: 10.1007/s002210050738
189. Scott SH, Kalaska JF. 1997a. Reaching movements with similar hand paths but different arm orientations. I. Activity of individual cells in motor cortex. *Journal of Neurophysiology*, 77(2), 826–852. doi: 10.1152/jn.1997.77.2.826
190. Scott SH, Sergio LE, Kalaska JF. 1997b. Reaching movements with similar hand paths but different arm orientations. II. Activity of individual cells in dorsal premotor cortex and parietal area 5. *Journal of Neurophysiology*, 78(5), 2413–2426. doi: 10.1152/jn.1997.78.5.2413
191. Scott SH. 2003. The role of primary motor cortex in goal-directed movements: Insights from neurophysiological studies on non-human primates. *Current Opinion in Neurobiology*, 13(6), 671–677. doi: 10.1016/j.conb.2003.10.012
192. Scott SH. 2004. Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5(7), 532–544. doi: 10.1038/nrn1427

193. Sergio LE, Hamel-Pâquet C, Kalaska JF. 2005. Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. *Journal of Neurophysiology*, 94(4), 2353–2378. doi: 10.1152/jn.00989.2004
194. Serrien DJ, Wiesendanger M. 2000. Temporal control of a bimanual task in patients with cerebellar dysfunction. *Neuropsychologia*, 38(5), 558–565. doi: 10.1016/S0028-3932(99)00116-5
195. Serrien DJ, Wiesendanger M. 2001a. A higher-order mechanism overrules the automatic grip-load force constraint during bimanual asymmetrical movements. *Behavioural Brain Research*, 118(2), 153–160. doi: 10.1016/S0166-4328(00)00317-X
196. Serrien DJ, Nirkko AC, Wiesendanger M. 2001b. Role of the corpus callosum in bimanual coordination: a comparison of patients with congenital and acquired callosal damage. 1897-905. *Eur J Neurosci*, 14(11). doi: 10.1046/j.0953-816x.2001.01798.x
197. Serrien DJ, Strens LHA, Oliviero A, Brown P. 2002. Repetitive transcranial magnetic stimulation of the supplementary motor area (SMA) degrades bimanual movement control in humans. *Neuroscience Letters*, 328(2), 89–92. doi: 10.1016/S0304-3940(02)00499-8
198. Shaikhouni A, Donoghue JP, Hochberg LR. 2013. Somatosensory responses in a human motor cortex. *Journal of Neurophysiology*, 109(8), 2192–2204. doi: 10.1152/jn.00368.2012
199. Shen L, Alexander GE. 1997. Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. *Journal of Neurophysiology*, 77(3), 1195–1212. doi: 10.1152/jn.1997.77.3.1195
200. Shenoy KV, Sahani M, Churchland, MM. 2013. Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience*, 36(1), 337–359. doi: 10.1146/annurev-neuro-062111-150509
201. Soma S, Saiki A, Yoshida J, Ríos A, Kawabata M, Sakai Y, Isomura Y. 2017. Distinct laterality in forelimb-movement representations of rat primary and secondary motor cortical neurons with intratelencephalic and pyramidal tract projections. *Journal of Neuroscience*, 37(45), 10904–10916. doi: 10.1523/JNEUROSCI.1188-17.2017
202. Soteropoulos DS, Edgley SA, Baker, SN. 2011. Lack of evidence for direct corticospinal contributions to control of the ipsilateral forelimb in monkey. *Journal of Neuroscience*, 31(31), 11208–11219. doi: 10.1523/JNEUROSCI.0257-11.2011
203. Stein BM, Carpenter MB. 1967. Central projections of portions of the vestibular ganglia innervating specific parts of the labyrinth in the rhesus monkey. *American Journal of Anatomy*, 120(2), 281–317. doi: 10.1002/aja.1001200205

204. Steinberg O, Donchin O, Gribova A, Cardoso de Oliveira S, Bergman H, Vaadia E. 2002. Neuronal populations in primary motor cortex encode bimanual arm movements. *European Journal of Neuroscience*, 15(8), 1371–1380. doi: 10.1046/j.1460-9568.2002.01968.x
205. Suminski AJ, Tkach DC, Hatsopoulos NG. 2009. Exploiting multiple sensory modalities in brain-machine interfaces. *Neural Networks*, 22(9), 1224–1234. doi: 10.1016/j.neunet.2009.05.006
206. Suminski AJ, Tkach DC, Fagg AH, Hatsopoulos NG. 2010. Incorporating Feedback from Multiple Sensory Modalities Enhances Brain-Machine Interface Control. *Journal of Neuroscience*, 30(50), 16777–16787. doi: 10.1523/JNEUROSCI.3967-10.2010
207. Sussillo D, Churchland MM, Kaufman MT, Shenoy KV. 2015. A neural network that finds a naturalistic solution for the production of muscle activity. *Nat Neurosci*, 116(8), 1477–1490. doi: 10.1038/nn.4042
208. Swinnen SP, Wenderoth N. 2004. Two hands, one brain: Cognitive neuroscience of bimanual skill. *Trends in Cognitive Sciences*, 8(1), 18–25. doi: 10.1016/j.tics.2003.10.017
209. Tanji J, Evarts EV. 1976. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *Journal of Neurophysiology*, 39(5), 1062–1068. doi: 10.1152/jn.1976.39.5.1062
210. Tanji J, Okano K, Sato KC. 1988. Neuronal activity in cortical motor areas related to ipsilateral, contralateral, and bilateral digit movements of the monkey. *Journal of Neurophysiology*, 60(1), 325–343. doi: 10.1152/jn.1988.60.1.325
211. Tkach D, Reimer J, Hatsopoulos NG. 2007. Congruent activity during action and action observation in motor cortex. *Journal of Neuroscience*, 27(48), 13241–13250. doi: 10.1523/JNEUROSCI.2895-07.2007
212. Todorov E, Jordan MI. 2002. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11), 1226–1235. doi: 10.1038/nn963
213. Tuller B, Kelso JAS. 1989. Environmentally-specified patterns of movement coordination in normal and split-brain subjects. *Experimental Brain Research*, 75(2), 306–316. doi: 10.1007/BF00247936
214. van Meer MPA, van Der Marel K, Wang K, Otte WM, El Bouazati S, Roeling TAP, Viergever MA, van der Sprenkel JWB, Dijkhuizen RM. 2010. Recovery of sensorimotor function after experimental stroke correlates with restoration of resting-state interhemispheric functional connectivity. *Journal of Neuroscience*, 30(11), 3964–3972. doi: 10.1523/JNEUROSCI.5709-09.2010
215. Verstynen T, Diedrichsen J, Albert N, Aparicio P, Ivry B. 2005. Ipsilateral Motor Cortex Activity During Unimanual Hand Movements Relates to Task Complexity. *Journal of Neurophysiology*, 1209–1222. doi: 10.1152/jn.00720.2004

216. Verstynen T, Ivry RB. 2011. Network dynamics mediating ipsilateral motor cortex activity during unimanual actions. *Journal of Cognitive Neuroscience*, 23(9), 2468–2480. doi: 10.1162/jocn.2011.21612
217. Vyas S, Golub MD, Sussillo D, Shenoy KV. 2020. Computation through Neural Population Dynamics. *Annual Review of Neuroscience*, 43, 249–275. doi: 10.1146/annurev-neuro-092619-094115
218. Wang W, Chan S, Heldman D, Moran D. 2007. Motor cortical representation of position and velocity during reaching. *Journal of Neurophysiology*, 4258–4270. doi: 10.1152/jn.01180.2006
219. Weinrich M, Wise SP, Mauritz KH. 1984. A neurophysiological study of the premotor cortex in the rhesus monkey. *Brain*, 107(2), 385–414. doi: 10.1093/brain/107.2.385
220. Welford AT. 1968. *Fundamentals of Skill*. Methuen, London.
221. Wiesendanger M, Serrien DJ. 2004. The quest to understand bimanual coordination. *Progress in Brain Research*, 143(03), 491–505. doi: 10.1016/S0079-6123(03)43046-X
222. Wilkins KB, Yao J, Owen M, Karbasforoushan H, Carmona C, Dewald JPA. 2020. Limited capacity for ipsilateral secondary motor areas to support hand function post-stroke. *Journal of Physiology*, 598(11), 2153–2167. doi: 10.1113/JP279377
223. Willett FR, Deo DR, Avansino DT, Rezaii P, Hochberg LR, Henderson JM, Shenoy KV. 2020. Hand Knob Area of Premotor Cortex Represents the Whole Body in a Compositional Way. *Cell*, 181(2), 396-409.e26. doi: 10.1016/j.cell.2020.02.043
224. Yeh F, Panesar S, Fernandes D, Meola A, Yoshino M, Fernandez-Miranda J, Vettel JM, Verstynen T. 2018. Population-Averaged Atlas of the Macroscale Human Structural Connectome and Its Network Topology. *Neuroimage*, (178), 57–68. doi: 10.1016/j.neuroimage.2018.05.027
225. Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. 2009. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, 102(1), 614–635. doi: 10.1152/jn.90941.2008
226. Ziemann U, Ishii K, Borgheresi A, Yaseen Z, Battaglia F, Hallett M, Cincotta M, Wassermann EM. 1999. Dissociation of the pathways mediating ipsilateral and contralateral motor-evoked potentials in human hand and arm muscles. *Journal of Physiology*, 518(3), 895–906. doi: 10.1111/j.1469-7793.1999.0895p.x