THE UNIVERSITY OF CHICAGO

A NETWORK ENCODING MODEL FOR NATURAL FORELIMB MOVEMENTS IN MARMOSET SENSORIMOTOR CORTEX

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES
AND THE PRITZKER SCHOOL OF MEDICINE
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL NEUROSCIENCE

BY
DALTON D. MOORE

CHICAGO, ILLINOIS
DECEMBER 2023

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

Thank you to my wife, Mattie, for consistently supporting and encouraging me; for making our home so cozy; for helping me break up a long work day with a bike ride or a walk with Millie; and for enjoying so many fun trips and tasty restaurants with me.

Thank you to my parents, Darren and Stephanie, for teaching me how to learn and think critically; for providing the resources and environment that allowed me to explore many curiosities and passions; and for endlessly humoring my favorite question as a child – "Why?".

Thank you to all of my family and friends for helping me to enjoy a full life and for always showing interest in my research (or at least in cute pictures of marmosets!).

Thank you to my committee (Nicho Hatsopoulos, Jason MacLean, Dan Margoliash, Matt Kaufman) for all the help, advice, and feedback to help me get to this point. Thank you Nicho for always being available for impromptu meetings, for providing insight and direction throughout my time in the lab, and for allowing me independence and flexibility as I developed my ideas and project.

ABSTRACT

The underlying mechanisms of computation in sensorimotor cortex must be both flexible and robust to support the range of skilled, dynamic forelimb movements observed in natural primate behavior. To understand the full richness of these mechanisms, it is important to study motor behavior in the most natural context possible using analytical tools that account for relationships between movement and both single-unit and population activity. In this work, I demonstrate that accurately capturing naturalistic motor behavior – specifically dynamic forelimb movements during foraging and prey-capture – is both critical and feasible. I leverage natural forelimb movements and pairwise precise spike time structure (represented as a functional network) in marmoset sensorimotor cortex to develop a network encoding model that links single-unit spiking activity with kinematic tuning properties and functional network interactions. I use this model to investigate the computational mechanisms that generate varied and naturalistic motor behavior. I show that a trajectory-based encoding model predicts single-unit activity during naturalistic forelimb movements more accurately than a simpler model. I demonstrate that tuning to kinematics depends on functional interactions between units – particularly on structured strong connections. Finally, I identify a reach-specific functional group that reorganizes to produce dynamic forelimb movements during prey capture. This reach-specific functional group is strongly interconnected and comprises units tightly linked to kinematics with strong, positively correlated preferred trajectories. By examination of the reach-specific functional group and remaining non-specific units, I suggest a potential framework linking single-unit properties to the neural population dynamics that generate movement.

CHAPTER 1

INTRODUCTION

1.1 Computations in Motor Cortex

Any interaction between an organism and the external world requires a motor behavior. As a critical node in the nervous system involved in voluntary motor control, the primary motor cortex (M1) provides the opportunity to investigate neuronal computations and their outputs at the final stage of cortical processing before movement execution. The motor system executes tasks of exquisite precision such as writing, playing musical instruments, or even buttoning a shirt with ease, yet can produce powerful movements like a tennis serve or butterfly stroke using the same neural circuitry. Motor behaviors can be highly consistent and repeatable across this vast behavioral range and the motor system can react quickly and skillfully to external sensory information or perturbations to correct and fine-tune movement. Thus, the underlying mechanisms of computation in M1 must be both flexible and robust to support the full repertoire of human motor behavior. In this work, I combine tuning properties of individual neurons and a functional network approach to investigate the computational mechanisms that generate varied and naturalistic motor behavior.

1.2 Relating M1 neural activity to movement parameters

Studies of primate upper limb control have historically searched for relationships between the activity of individual neurons in M1 and either muscle-related features (such as force) or movement kinematics. In 1968, Evarts used a wrist flexion-extension task with external opposing forces to show that the activity of individual pyramidal tract neurons in M1 was related most directly to exerted force rather than to displacement in a majority of 31 units.[1] In the 1980s, a series of papers by Georgopoulos and colleagues used a center-out

1. Evarts, "Relation of pyramidal tract activity to force exerted during voluntary movement."

reaching task in two and three dimensions to demonstrate a relationship between direction of movement and the mean firing rate of M1 units during the movement.[2,3] These papers modeled direction tuning in individual neurons by fitting a shifted and scaled cosine curve to average firing rates over many trials in each direction. They also introduced the concept of the population vector, in which the direction of movement could be predicted by the sum of preferred directions weighted by firing rate across all direction-tuned units.[4,5] While this concept was instructive for much of the subsequent research in motor control, it treats the output of the motor system as a linear summation operator and ignores any effect of the complex interactions between neurons. A 1992 study partially reconciled the ideas of force versus direction tuning by demonstrating tuning of neurons to the directional change in force rather than to total exerted force.[6] Around this time, the field began to add complexity to the standard center-out movement and force tasks. In this vein, Caminiti et al.[7] and Sergio and Kalaska[8] added nuance to the concept of population vector encoding by demonstrating that directional tuning to movement and isometric force production, respectively, rotates systematically with arm posture. Scott and Kalaska[9] confirmed shifts in directional tuning to movement but, in contrast with Caminiti et al., found no systematic relationship between posture and the population vector. Hocherman and Wise[10] showed that M1 codes for details

2. Georgopoulos *et al.,* "On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex."

3. Schwartz *et al.,* "Primate motor cortex and free arm movements to visual targets in three- dimensional space. I. Relations between single cell discharge and direction of movement".

4. Georgopoulos *et al.,* "Neuronal Population Coding of Movement Direction".

5. Georgopoulos *et al.,* "Primate motor cortex and free arm movements to visual targets in three- dimensional space. II. Coding of the direction of movement by a neuronal population".

6. Georgopoulos *et al.,* "The Motor Cortex and the Coding of Force".

7. Caminiti *et al.,* "Making arm movements within different parts of space: dynamic aspects in the primate motor cortex."

8. Sergio & Kalaska, "Systematic Changes in Motor Cortex Cell Activity With Arm Posture During Directional Isometric Force Generation".

9. Scott & Kalaska, "Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures".

10. Hocherman & Wise, "Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys".

of a movement trajectory as opposed to the net change in position. Moran and Schwartz[11] found that a model incorporating speed along with direction improved on directional tuning alone.

As the field continued to explore the tuning properties of M1, several nuances and dependencies were discovered. For example, multiple studies showed that the consistent directional tuning reported by Georgopoulos and others depended on averaging spiking activity over the duration of movement. In fact, the instantaneous preferred movement direction of neurons in premotor cortex and M1 can change drastically over the course of movement[12,13,14] and between preparatory and movement periods.[15] The study by Sergio et al. also showed that preferred force direction in an isometric force task was more consistent than preferred movement direction in the center-out whole-arm movement task.

Aflalo and Graziano[16] added complexity when they examined tuning of M1 neurons to an array of previously proposed kinematic metrics in unconstrained behavior. They noted that previous studies such as those discussed here related M1 activity to movements that were constrained to a small percentage of possible kinematic variance, and that simple models explaining a large degree of that variance may not extend to less constrained behavior. They showed that direction tuning accounted for only 8% of variance in neural activity while macaques moved their arms freely in the workspace; even the best model – an end-posture model – accounted for only 36% of variance on average. A combination model incorporating hand speed, directional tuning, and final position of the hand and posture of the arm increased the explained variance to 44%, leaving the majority of neural activity

11. Moran & Schwartz, "Motor Cortical Representation of Speed and Direction During Reaching".

12. Churchland & Shenoy, "Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex".

13. Mason *et al.,* "Temporal profile of the directional tuning of the discharge of dorsal premotor cortical cells".

14. Sergio *et al.,* "Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks".

15. Churchland *et al.,* "Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine?"

16. Aflalo & Graziano, "Relationship between unconstrained arm movements and single-neuron firing in the macaque motor cortex."

unexplained. This result, along with the body of work discussed above, suggests that M1 may represent and control a combination of multiple kinematic or muscle-related variables; M1 may also execute unexplored computations that account for some portion of the remaining variance.

A parallel approach to modeling tuning in M1 as a combination of kinematic parameters began with the development of a trajectory encoding model by Hatsopoulos, Xu and Amit in 2007.[17] The authors used a random target pursuit task which produced more kinematic variance than observed in most prior studies, although still within a single plane unlike Aflalo and Graziano.[18] The authors accounted for the shifting of preferred directions of single neurons demonstrated in previous work by relating neural activity to brief movement trajectories that incorporate a range of leading and lagging kinematics relative to neural spiking. They found that a trajectory model incorporating average hand speed, position, and finely sampled movement directions from -100ms before neural activity to +300ms afterward performed best and outperformed a brief trajectory model analogous to direction tuning at a fixed lag of +100 to +150ms. The kinematic trajectory model also outperformed a comparable torque trajectory tuning model for most of the population. The authors note that their recordings were more rostral on the precentral gyrus than work showing a preference for force tuning,[19] which may explain this difference. In a follow-up study,[20] they demonstrated that the preferred trajectories of pairs of neurons could sum linearly when the two neurons fired simultaneously, suggesting a rich variety of complex movement representations could be synthesized in the simultaneous firing of a neural population by adding the trajectory representations of the constituent neurons in the population. Another line of

17. Hatsopoulos *et al.,* "Encoding of Movement Fragments in the Motor Cortex".

18. Aflalo & Graziano, "Relationship between unconstrained arm movements and single-neuron firing in the macaque motor cortex."

19. Sergio *et al.,* "Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks".

20. Hatsopoulos & Amit, "Synthesizing complex movement fragment representations from motor cortical ensembles".

work[21,22] used the same trajectory encoding framework to relate neural activity with grasp and reach-to-grasp behavior, respectively, demonstrating the potential for this approach to explain the neural computations producing increasingly complex and naturalistic behavior. In Chapter 3, I apply this trajectory encoding model to upper-limb kinematics produced by common marmosets during the capture of live moths in an unrestrained, semi-naturalistic environment.

1.3 Computations rely on neuronal interactions and population dynamics

Early work interrogating the computational role of neuronal interactions showed that dynamic synchronization (precisely-timed cross-correlations at zero lag) between neurons can emerge in conjunction with more general aspects of motor control such as the decision to move[23] or cue expectation.[24] Hatsopoulos et al.[25] showed that synchrony varied with movement direction for some neuron pairs in a manner that could not be explained by firing rates alone. Follow-up work[26] demonstrated that pairwise noise correlations carried information beyond that available in firing rates alone and contributed to decoding of movement direction. These studies – constrained by small populations of recorded neurons – suggested an important computational role for neuronal interactions but were unable to explore that role further.

Motivated by these results, by increasing population sizes of recorded neurons, and by the lack of clarity as to the tuning properties of M1 neurons, a contemporary branch of motor con-

21. Saleh *et al.,* "Encoding of coordinated grasp trajectories in primary motor cortex."

22. Saleh *et al.,* "Encoding of coordinated reach and grasp trajectories in primary motor cortex."

23. Vaadia *et al.,* "Dynamics of neuronal interactions in monkey cortex in relation to behavioural events".

24. Riehle *et al.,* "Spike synchronization and rate modulation differentially involved in motor cortical function".

25. Hatsopoulos *et al.,* "Information about movement direction obtained from synchronous activity of motor cortical neurons."

26. Maynard *et al.,* "Neuronal interactions improve cortical population coding of movement direction."

trol work has sought to explain M1 activity as a dynamical system that generates movement through structured dynamics. Seminal work by Churchland et al.[27] showed that oscillatory patterns in single-neuron mean firing rates could be explained by rotational dynamics in a low-dimensional state space representation of population activity. This framework has been used to show that M1 activity remains in an "output-null" subspace of population activity to prevent motion during the preparatory phase, then moves to the "output-potent" subspace and begins to exhibit rotational dynamics as movement begins.[28,29] Moreover, the null and potent spaces seem to be utilized by M1 neural populations more generally to send muscle-like commands downstream while maintaining smooth dynamics.[30] It has also yielded the finding that neural trajectories in M1 are significantly less "tangled" than trajectories in sensory cortices or muscle activity, where low tangling means the future evolution of the neural state is highly dependent on the current state and lends robustness to noise.[31] This study also showed that muscle-related activity could be read out linearly from population dynamics, suggesting a mechanism by which the neural dynamical system might directly control movement in the midst of rotations. The rotational dynamics that dominate M1 population activity seem to be unique – despite similar single-unit modulation in M1 and both upstream areas (supplementary motor area) and downstream readouts (muscle activity), those areas exhibit dynamics with significantly weaker rotations.[32] A recent development termed the Location-Dependent Rotations (LDR) model builds on the work discussed above by allowing for dynamics to unfold in higher dimensional neural space in a task-dependent manner on single trials, rather than in a fixed space across trials and conditions.[33] The au-

27. Churchland *et al.,* "Neural population dynamics during reaching".

28. Elsayed *et al.,* "Reorganization between preparatory and movement population responses in motor cortex".

29. Kaufman *et al.,* "Cortical activity in the null space: Permitting preparation without movement".

30. Stavisky *et al.,* "Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions".

31. Russo *et al.,* "Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response".

32. Lara *et al.,* "Different population dynamics in the supplementary motor area and motor cortex during reaching".

33. Sabatini & Kaufman, "Reach-dependent reorientation of rotational dynamics in motor cortex".

thors demonstrate that neural dynamics occur on a curved manifold, such that frequency of local linear rotational dynamics is conserved while the orientation of the rotational plane depends on the location of the fixed point on that manifold. They also show a near-linear relationship between the fixed point location in neural space and reach kinematics on single trials, suggesting a link between classical kinematic tuning and rotational dynamics. Low-dimensional structure in neural population activity has also provided insight and utility for brain-computer interface research (see Kao et al.[34] and Oby et al.[35] for two examples out of many).

Although dynamical systems approaches have contributed profoundly to our understanding of motor control, it is important to note the contexts in which they may fall short. Rotational dynamics that explain M1 neural activity during movements of the proximal limb cannot explain activity during grasping,[36] due in large part to higher tangling which may reflect higher dependence on extrinsic inputs.[37] This also suggests that dynamical systems may struggle to explain M1 activity during naturalistic tasks that rely on continuous visual and proprioceptive feedback to make online adjustments, such as the reach-to-grasp movements performed during the prey-capture of moths in Chapter 3. Stavisky et al.[38] used low-dimensional representations of output-null and output-potent subspaces to investigate neural trajectories in response to visuomotor feedback but did so for single ballistic perturbations to a well-trained and stereotyped task. Gallego et al.[39] identified consistent low-dimensional latent dynamics as a hallmark of consistent performance of learned behaviors, but this again was for a simple task (cursor movement with a brain-computer interface).

34. Kao *et al.,* "Single-trial dynamics of motor cortex and their applications to brain-machine interfaces".

35. Oby *et al.,* "New neural activity patterns emerge with long-term learning."

36. Suresh *et al.,* "Neural population dynamics in motor cortex are different for reach and grasp".

37. Russo *et al.,* "Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response".

38. Stavisky *et al.,* "Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions".

39. Gallego *et al.,* "Long-term stability of cortical population dynamics underlying consistent behavior".

In contrast, Musall et al.[40] demonstrated that the trial-specific, idiosyncratic movements occurring throughout task performance were necessary to explain a large portion of neuronal activity on single trials; thus, low-dimensional representations that require trial-averaging may not be able to describe the computations occurring during corrective and idiosyncratic movements. Although the LDR model[41] explains much of the single trial variance for curved reaching in two dimensions, to my knowledge neither this model nor any other dynamical systems approach has been applied to movements as varied and complex as those studied in Chapters 2 and 3.

1.4 Moving towards complex and naturalistic motor behaviors

Most of the studies discussed so far relate neural activity to highly constrained and overtrained tasks such as center-out reaching,[42] isometric force application,[43] or even pedaling a wheel.[44] Such behaviors limit the variability in most movement parameters aside from those under investigation and may impose an artificial ceiling on the dimensionality of neural population dynamics.[45] It is clear from the history of motor control work described above that the prevailing model of motor control required more nuance each time the behavioral task was extended to a more complex subset of movements. It would be difficult, or perhaps impossible, to comprehensively account for activity in M1 without studying motor behavior in its full richness. Even in the case of less stereotyped behaviors like random target

40. Musall *et al.,* "Single-trial neural dynamics are dominated by richly varied movements".

41. Sabatini & Kaufman, "Reach-dependent reorientation of rotational dynamics in motor cortex".

42. Georgopoulos *et al.,* "On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex."

43. Evarts, "Relation of pyramidal tract activity to force exerted during voluntary movement."

44. Russo *et al.,* "Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response".

45. Gao & Ganguli, "On simplicity and complexity in the brave new world of large-scale neuroscience".

pursuit,[46] reaching around obstacles,[47] and reach to grasp tasks,[48] primates are restrained in a chair with head-fixation and require weeks of daily training to reliably participate in the experiment. There is evidence to suggest that a significant portion of M1 neurons encode a constrained, well-trained task differently than free behavior[49] and that over-training increases synchrony between M1 neurons.[50] In addition to these results from primate studies, recent work in rodents suggests that motor cortex may not be directly involved in the production of over-trained, skilled movements despite its necessary role in learning the movement.[51] The authors propose that motor cortex trains subcortical circuits over the course of learning, then cedes direct control to those circuits. Seemingly in contrast, Guo et al.[52] showed that rapid, reversible optogenetic inhibition of motor cortex prevented or paused skilled prehension behavior but left untrained movements like grooming unaffected. Emerging work examining the role of subcortical regions in a similar task to Kawai et al. found that lesions to sensorimotor striatum disrupted high-level sequencing of learned behavior but did not disrupt equivalent movements that were visually cued,[53] which means other regions were in control of sensory-guided movements. Taken together, these studies suggest that motor cortex is essential for executing movements that require dexterity (such as reach-to-grasp) or rely on sensory feedback. These results, to the extent they generalize to primate motor control, bring into question the degree to which M1 directly controls over-trained, stereotyped movements at all. Consequently, it is critical to investigate neural control of movement in the most natural, unconstrained behavioral context possible.

46. Hatsopoulos *et al.,* "Encoding of Movement Fragments in the Motor Cortex".

47. Kaufman *et al.,* "Cortical activity in the null space: Permitting preparation without movement".

48. Saleh *et al.,* "Encoding of coordinated reach and grasp trajectories in primary motor cortex."

49. Jackson *et al.,* "Correlations Between the Same Motor Cortex Cells and Arm Muscles During a Trained Task, Free Behavior, and Natural Sleep in the Macaque Monkey".

50. Schieber, "Training and Synchrony in the Motor System".

51. Kawai *et al.,* "Motor Cortex Is Required for Learning but Not for Executing a Motor Skill".

52. Guo *et al.,* "Cortex commands the performance of skilled movement".

53. Mizes *et al.,* "Dissociating the contributions of sensorimotor striatum to automatic and visually-guided motor sequences".

The common marmoset (*Callithrix jacchus*) provides a unique opportunity for studying neural representations of movement during ethological behaviors. Prior work by Walker et al.[54] produced a modular apparatus that allows for semi-automated training and data collection on a variety of naturalistic, goal-directed behavioral tasks. This system leverages marmosets' inclination to engage in a behavioral task sporadically across numerous short bursts of interest. Compared to the classical task paradigms discussed above, prey capture in the apparatus produces a richer set of movements with variable directions, speeds, and amplitudes.

The study of naturalistic behavior presents unique challenges; the first and immediate challenge is accurately and efficiently capturing movement in a less constrained environment. Recent developments in pose estimation technologies leverage deep learning computer vision models to compute frame-by-frame pose estimation. In Chapter 2, I discuss the importance of accurate pose estimation and the performance of DeepLabCut (DLC)[55] compared to marker-based tracking. Crucially, DLC proved capable of reliable pose estimation with relatively little user effort compared to marker-based systems; its demonstrated accuracy precludes the need for implanted or externally attached physical markers that could affect naturalistic behavior. The methods and intuitions developed in Chapter 2 were applied to accurately estimate the shoulder and wrist positions during the prey capture task investigated in Chapter 3.

A second challenge is that naturalistic and untrained behaviors preclude trial-averaging and produce a wealth of movement variability that can be difficult to parse. The models I present in Chapter 3 are designed to leverage the kinematic variability of natural movements in order to produce a motor control model capable of explaining complex, dexterous upper-limb movements on single trials. As marmosets engage in voluntary prey-capture, they initiate internally generated movements that engage motor cortical circuits in their natural

54. Walker *et al.*, "A platform for semiautomated voluntary training of common marmosets for behavioral neuroscience".

55. Mathis *et al.*, "DeepLabCut: markerless pose estimation of user-defined body parts with deep learning".

state, presumably unaltered by over-training of stereotyped movements.

1.5 Population dynamics through the lens of the functional network

Single-neuron approaches to understanding motor control provided a baseline certainty that M1 is instrumental in the generation of upper-limb movements, particularly for skilled and dexterous behaviors. However, the specific tuning properties of single neurons and neural populations remain disputed. The population dynamics approach has provided an alternative view on M1 activity that may transcend the concept of tuning to movement parameters. However, the dimensionality reduction techniques used to characterize these dynamics discard the identity, role, and tuning properties of single neurons, thereby hindering efforts to understand the function of individual neurons within the context of the population. A few studies have attempted to interpret single-unit M1 activity in the context of the population, showing that pairwise spike count correlations provided information about motor behavior beyond what is provided by firing rates alone[56] and improved encoding models that predict single-unit activity.[57] However, spike time correlations which may be more informative of dynamics remain under-studied in M1 neural populations (but see Sundiang et al.,[58] discussed below).

Network science – the study of complex networks – provides an alternative approach to the investigation of neurons and their interactions within the population. In this framework, network activity is summarized in the form of a functional network (FN) that maintains neuron-specific labels while simultaneously capturing all pairwise spike time correlations between neurons. In the work presented here, functional connectivity defined by pairwise

56. Maynard *et al.,* "Neuronal interactions improve cortical population coding of movement direction."

57. Stevenson *et al.,* "Functional Connectivity and Tuning Curves in Populations of Simultaneously Recorded Neurons".

58. Sundiang *et al.,* "Dynamic structure of motor cortical neuron coactivity carries behaviorally relevant information".

statistical relationships in spike timing do not necessarily reflect synaptic connectivity, although there is likely some overlap for local connections.[59] In murine primary visual cortex, FNs can be used to predict single trial activity as an encoding model[60]; to decode single trial visual stimuli[61]; and to generate realistic simulations of single trial activity.[62] Recent work demonstrated that coactivity in groups of source units with functional inputs to a target unit (termed the target unit's "functional group") was the main predictor of single-trial activity in V1 neurons and that the functional group's predictive power depended on the specific topology of the strongest inputs.[63] Another study split V1 neurons into subpopulations that were either tuned or untuned to visual drifting grating stimuli to show that the subpopulations have distinct topological properties, with untuned units occupying central positions in the FN.[64] They also found that decoders utilizing FN topology yield accurate decoding of visual stimuli, and that these decoders depend on the connectivity of both tuned and untuned units. These studies demonstrate the potential for the FN approach to leverage single-neuron properties and population-level interactions together in a cohesive manner.

The FN approach was recently applied to M1 neural activity in non-human primates performing a standard delayed center-out reaching task, showing that FN structure is behaviorally specific and evolves systematically with movement.[65] FNs constructed from closer reach directions were closer in network space and temporal FNs computed in short intervals throughout the trial traversed reach-specific trajectories in a low-dimensional subspace. These trajectories became separable and decodable in the timeframe shortly after the in-

59. Chambers & MacLean, "Multineuronal activity patterns identify selective synaptic connections under realistic experimental constraints".

60. Dechery & MacLean, "Functional triplet motifs underlie accurate predictions of single-trial responses in populations of tuned and untuned V1 neurons".

61. Levy *et al.,* "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding".

62. Sadovsky & MacLean, "Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry".

63. Kotekal & MacLean, "Recurrent interactions can explain the variance in single trial responses".

64. Levy *et al.,* "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding".

65. Sundiang *et al.,* "Dynamic structure of motor cortical neuron coactivity carries behaviorally relevant information".

struction cue and remained so throughout movement initiation and execution. The decoder incorporating both firing rates and temporal FNs provided significant performance gains over the firing rate alone, and the FN-only decoder outperformed an equivalent decoder trained on rate-matched null FNs – both results demonstrate that FN structure contains information beyond that available in instantaneous firing rates. Finally, reciprocity in the FN decreased shortly after the instruction cue, suggesting a transitory reorganization of network topology that may facilitate integration and propagation of external inputs.[66] This change to FN topology may be a complementary representation of the transition from preparatory to movement subspaces identified by the dynamical systems approach.[67,68] FNs provide an opportunity to parse relationships between neural activity and specific features of the network such as reciprocity, in-degree, or membership in functional groups.

1.6 Comparison of M1 to Areas 3a and 3b

To this point I have discussed relationships between upper-limb control and motor cortical activity exclusively. However, in Chapter 3 I investigate kinematic and network encoding models for neurons in M1 and 6dc (pre-motor cortex) and area 3a, as well as some units in area 3b. Area 3b primarily receives and processes cutaneous information from the peripheral nervous system[69,70,71] and has sparser connectivity with M1 and area 3a than with other so-

66. Malonis *et al.,* "M1 dynamics share similar inputs for initiating and correcting movement".

67. Kaufman *et al.,* "Cortical activity in the null space: Permitting preparation without movement".

68. Elsayed *et al.,* "Reorganization between preparatory and movement population responses in motor cortex".

69. Hyvärinen & Poranen, "Receptive field integration and submodality convergence in the hand area of the post-central gyrus of the alert monkey."

70. Iwamura *et al.,* "Rostrocaudal gradients in the neuronal receptive field complexity in the finger region of the alert monkey's postcentral gyrus".

71. Tanji & Wise, "Submodality distribution in sensorimotor cortex of the unanesthetized monkey".

matosensory regions.[72,73,74,75] Area 3a, on the other hand, receives proprioceptive feedback from the periphery[76,77,78] and is highly interconnected with M1.[79,80,81] Area 3a contains pyramidal tract neurons (PTNs) with similar discharge properties and reflex responses to M1 PTNs, suggesting a direct influence on movement[82] in addition to its canonical role receiving and processing ongoing postural and movement-related feedback.[83,84] Similarities between M1 and area 3a extend to representation of proximal upper-limb movement direction and externally applied loads,[85] as well as encoding[86] and decoding[87] of distal limb movements of the wrist and digits. Thus, I expect largely similar results in Chapter 3 for neuron populations in M1 and area 3a, with less similarity between these neurons and those in area 3b. One potential difference is the timescale of neural encoding, such that sensory neurons prefer trajectory samples leading the spike sample window and motor neurons prefer trajectories lagging the spike sample. However, this effect is difficult to parse for natural, continuous

72. Burton & Fabri, "Ipsilateral intracortical connections of physiologically defined cutaneous representations in areas 3b and 1 of macaque monkeys: projections in the vicinity of the central sulcus".

73. Huffman & Krubitzer, "Area 3a: Topographic Organization and Cortical Connections in Marmoset Monkeys".

74. Jones *et al.,* "Intracortical connectivity of architectonic fields in the somatic sensory, motor and parietal cortex of monkeys".

75. Krubitzer & Kaas, "The organization and connections of somatosensory cortex in marmosets".

76. Iwamura *et al.,* "Rostrocaudal gradients in the neuronal receptive field complexity in the finger region of the alert monkey's postcentral gyrus".

77. Phillips *et al.,* "Projection from low-threshold muscle afferents of hand and forearm to area 3a of baboon's cortex".

78. Tanji & Wise, "Submodality distribution in sensorimotor cortex of the unanesthetized monkey".

79. Huerta & Pons, "Primary motor cortex receives input from area 3a in macaques".

80. Huffman & Krubitzer, "Area 3a: Topographic Organization and Cortical Connections in Marmoset Monkeys".

81. Jones *et al.,* "Intracortical connectivity of architectonic fields in the somatic sensory, motor and parietal cortex of monkeys".

82. Fromm & Evarts, "Pyramidal tract neurons in somatosensory cortex: central and peripheral inputs during voluntary movement".

83. Fromm *et al.,* "Sensory response properties of pyramidal tract neurons in the precentral motor cortex and postcentral gyrus of the rhesus monkey".

84. Gardner & Costanzo, "Properties of kinesthetic neurons in somatosensory cortex of awake monkeys".

85. Prud'homme & Kalaska, "Proprioceptive activity in primate primary somatosensory cortex during active arm reaching movements".

86. Goodman *et al.,* "Postural Representations of the Hand in the Primate Sensorimotor Cortex".

87. Okorokova *et al.,* "Decoding hand kinematics from population responses in sensorimotor cortex during grasping".

movements that exhibit long autocorrelations; this is especially true in Chapter 3 because the vast majority of trajectory samples occur entirely in the midst of reaching.

1.7 Developing an encoding model incorporating naturalistic kinematics and the functional network

I demonstrate in Chapter 2 that computer vision tools – specifically DLC+Anipose[88,89] – track 3D pose with accuracy comparable to marker-based tracking. DLC+Anipose estimates pose as accurately as expert human labelers and significantly more accurately than novice human labelers. I also characterize the effect of parameter choices on DLC+Anipose accuracy and identify a range of parameter value choices to minimize error and maximize the number of well-tracked frames. These insights were used in Chapter 3 to accurately estimate pose during prey-capture, thereby providing the complex kinematics necessary for the trajectory encoding model.

In Chapter 3 I take steps toward integrating the neuron-centric and network perspectives on computations in motor cortex by developing encoding models that can predict fine-timing spike activity in individual neurons by accounting for temporally precise movement kinematics, neuron-specific properties, and activity in the functional group. The veracity of our results depends on using a network model built on the backbone of a suitably accurate single-unit tuning model. Simple encoding models such as directional tuning at a fixed post-spike delay fail to explain much of the variance in activity during unconstrained natural movements[90] and the majority of tuning models rely on average firing rates through the duration of movement and often across trials. Prior work in the Hatsopoulos lab has demonstrated that a temporally extended, trajectory-based model captures the single-trial tuning

88. Karashchuk *et al.,* "Anipose: A toolkit for robust markerless 3D pose estimation".

89. Mathis *et al.,* "DeepLabCut: markerless pose estimation of user-defined body parts with deep learning".

90. Aflalo & Graziano, "Relationship between unconstrained arm movements and single-neuron firing in the macaque motor cortex."

of M1 neurons more accurately than prior models and can be used to decode instantaneous movement direction with less error.[91,92,93,94] Chartier et al.[95] used the same approach to create an encoding model relating dense ECoG recordings in ventral sensorimotor cortex to articulatory kinematic trajectories during natural, continuous speech production. In Chapter 3, I use this generalized linear model (GLM) framework to elucidate the natural kinematic encoding properties of single motor cortical neurons and their influence on other neurons within an FN shaped by interactions. I identify the best-performing model across a range of lead/lag kinematic sampling windows by evaluating the accuracy of predicted spiking activity for held-out trajectory samples. I then augment the kinematics-only model with coincident and leading network features, computed by the dot product of the FN with coincident or leading activity in functionally connected units. With kinematics, neuron specificity, and network activity combined in a single encoding model, I examine the role and structure of functional interactions and their effect on kinematic tuning.

I find that the trajectory tuning model extends to more naturalistic movements and explains neural activity better than a simpler model. Additionally, the trajectory tuning model without any direct mathematical dependence on the FN links single-unit activity to kinematics most accurately for units that are strongly interconnected within the network. I incorporate network features in the model and show that single-unit spiking activity depends fundamentally on functional interactions – particularly on the topology of strong connections. Finally, I identify a reach-specific functional group that reorganizes to produce dynamic forelimb movements during prey-capture. This reach-specific functional group is strongly interconnected and comprises units tightly linked to kinematics with strong, posi-

91. Hatsopoulos & Amit, "Synthesizing complex movement fragment representations from motor cortical ensembles".

92. Hatsopoulos *et al.,* "Encoding of Movement Fragments in the Motor Cortex".

93. Saleh *et al.,* "Encoding of coordinated grasp trajectories in primary motor cortex."

94. Saleh *et al.,* "Encoding of coordinated reach and grasp trajectories in primary motor cortex."

95. Chartier *et al.,* "Encoding of Articulatory Kinematic Trajectories in Human Speech Sensorimotor Cortex".

tively correlated preferred trajectories. Analysis of the differences between the reach-specific and non-specific groups suggests a potential framework for contextualizing single-unit properties and activity within the generation and evolution of population dynamics.

CHAPTER 2

VALIDATING MARKERLESS POSE ESTIMATION WITH 3D X-RAY RADIOGRAPHY

This work was previously published:

Moore, Dalton D., Jeffrey D. Walker, Jason N. MacLean, and Nicholas G. Hatsopoulos. "Validating markerless pose estimation with 3D X-ray radiography." Journal of Experimental Biology 225, no. 9 (2022): jeb243998.

2.1 Abstract

To reveal the neurophysiological underpinnings of natural movement, neural recordings must be paired with accurate tracking of limbs and postures. Here, we evaluated the accuracy of DeepLabCut (DLC), a deep learning markerless motion capture approach, by comparing it with a 3D X-ray video radiography system that tracks markers placed under the skin (XROMM). We recorded behavioral data simultaneously with XROMM and RGB video as marmosets foraged and reconstructed 3D kinematics in a common coordinate system. We used the toolkit Anipose to filter and triangulate DLC trajectories of 11 markers on the forelimb and torso and found a low median error (0.228 cm) between the two modalities corresponding to 2.0% of the range of motion. For studies allowing this relatively small error, DLC and similar markerless pose estimation tools enable the study of increasingly naturalistic behaviors in many fields including nonhuman primate motor control.

2.2 Introduction

As the study of motor neuroscience progresses toward an emphasis on naturalistic, unconstrained behavior, kinematics must be captured accurately and efficiently. Past research has relied on marker-based systems tracking markers attached to an animal's skin (such as

Vicon, OptiTrack, and PhaseSpace) or surgically implanted radiopaque beads (XROMM).[1] However, these systems are expensive and often impractical with smaller species like mice or marmosets, especially for tracking free or semi-constrained behavior. To solve this problem, multiple groups have developed marker-less pose estimation tools that use deep learning to apply digital markers to recorded video. The most widely used is DeepLabCut (DLC),[2] but alternatives exist.[3,4,5,6] These enable the study of a wider range of behaviors by allowing free movement without the disturbance of physical markers. Furthermore, these tools alleviate the bottleneck of semi-automatic tracking; a well-trained network labels video with accuracy comparable to human labelers[7] and requires minimal hands-on time for subsequent datasets. DLC has been used in many contexts, including tracking eye movements, pupil dilation and hand movements in mice,[8,9,10] estimating 3D pose of freely moving macaques,[11] and even on XROMM video to increase throughput.[12]

DLC accuracy has not been compared to that of marker-based tracking in the context of close-up forelimb tracking that is common in motor control studies. Dunn et al.[13] tested DLC and a geometric deep learning tool (DANNCE) against a rat motion capture dataset, but recording from a small number of cameras in this unconstrained context with significant

1. Brainerd *et al.,* "X-ray reconstruction of moving morphology (XROMM): precision, accuracy and applications in comparative biomechanics research".

2. Mathis *et al.,* "DeepLabCut: markerless pose estimation of user-defined body parts with deep learning".

3. Dunn *et al.,* "Geometric deep learning enables 3D kinematic profiling across species and environments".

4. Graving *et al.,* "DeepPoseKit, a software toolkit for fast and robust animal pose estimation using deep learning".

5. Pereira *et al.,* "Fast animal pose estimation using deep neural networks".

6. Wu *et al., Deep Graph Pose: A semi-supervised deep graphical model for improved animal pose tracking.*

7. Mathis *et al.,* "DeepLabCut: markerless pose estimation of user-defined body parts with deep learning".

8. Sauerbrei *et al.,* "Cortical pattern generation during dexterous movement is input-driven".

9. Siegle *et al., A survey of spiking activity reveals a functional hierarchy of mouse corticothalamic visual areas.*

10. Steinmetz *et al.,* "Distributed coding of choice, action and engagement across the mouse brain".

11. Bala *et al.,* "Automated markerless pose estimation in freely moving macaques with OpenMonkeyStudio".

12. Laurence-Chasen *et al.,* "Integrating XMALab and DeepLabCut for high-throughput XROMM".

13. Dunn *et al.,* "Geometric deep learning enables 3D kinematic profiling across species and environments".

environment- and self-occlusion is beyond the intended use for DLC unless many cameras are used as in Bala et al..[14] Thus, a comparison in the semi-constrained context with a small number of cameras is crucial to confirm whether DLC reliably tracks kinematics with accuracy comparable to human labelers and existing marker-based systems. XROMM provides a useful comparison, as we have shown that the system tracks radiopaque markers with submillimeter precision.[15] To this end, we collect simultaneous recordings with XROMM and RGB video as common marmosets engage in naturalistic foraging, then reconstruct three-dimensional reaching kinematics in a shared coordinate system. We perform filtering, triangulation and optimization steps with Anipose[16] and present the effect of parameter choices on tracking quality. We find that optimized DLC+Anipose tracks position with median absolute error of 0.228 cm (mean absolute error = 0.274 cm), corresponding to 2.0% of the range of marker positions.

2.3 Methods

Subjects

These experiments were conducted with two common marmosets (*Callithrix jacchus*) (an 8-year old, 356g male and a 7-year old, 418g female). All methods were approved by the Institutional Animal Care and Use Committee of the University of Chicago.

Data Collection

The two marmosets were placed together in a 1m x 1m x 1m cage with a modular foraging apparatus attached to the top of the cage, as previously described by.[17] The marmosets were allowed to forage voluntarily throughout recording sessions that lasted 1-2 hours. Recordings

14. Bala *et al.,* "Automated markerless pose estimation in freely moving macaques with OpenMonkeyStudio".

15. Walker *et al.,* "A platform for semiautomated voluntary training of common marmosets for behavioral neuroscience".

16. Karashchuk *et al.,* "Anipose: A toolkit for robust markerless 3D pose estimation".

17. Walker *et al.,* "A platform for semiautomated voluntary training of common marmosets for behavioral neuroscience".

of individual trials were triggered manually with a foot pedal by the experimenters when the marmosets appeared ready to initiate a reach. The manual trigger initiated synchronized video collection by the XROMM system[18] and two visible light cameras, each described in further detail below. We retained all trials that captured right-handed reaches. Marmoset TY produced four useful reaching events containing 5 total reaches and marmoset PT produced 13 reaching events containing 17 reaches.

XROMM

Bi-planar X-ray sources and image intensifiers (90kV, 25mA at 200 fps) were used to track the 3D position of radiopaque tantalum beads (0.5-1 mm, Bal-tec) placed subcutaneously in the arm, hand, and torso. Details of bead implants can be found in,[19] in which the authors also report estimating XROMM marker tracking precision of 0.06 mm based on the standard deviation of inter-marker distances during a recording of a calibration specimen. Marker locations were chosen to approximate the recommendations given by the International Society of Biomechanics for defining coordinate systems of the upper limb and torso in humans.[20] These recommendations were adapted to the marmoset and constrained by surgical considerations. Positions of 13 beads were tracked using a semi-automated process in XMALab[21] following the procedure described there and in the XMALab User Guide (https://bitbucket.org/xromm/xmalab/wiki/Home). Two beads implanted in the anterior torso were ignored for comparison with DLC because corresponding positions on the skin were occluded in nearly every frame captured by visible light cameras.

18. Brainerd *et al.,* "X-ray reconstruction of moving morphology (XROMM): precision, accuracy and applications in comparative biomechanics research".

19. Walker *et al.,* "A platform for semiautomated voluntary training of common marmosets for behavioral neuroscience".

20. Wu *et al.,* "ISB recommendation on definitions of joint coordinate systems of various joints for the reporting of human joint motion–Part II: shoulder, elbow, wrist and hand".

21. Knorlein *et al.,* "Validation of XMALab software for Marker-based XROMM".

DeepLabCut

Two high-speed cameras (FLIR Blackfly S, 200 fps, 1440x1080 resolution) were used to record video for analysis by DLC. The cameras were positioned to optimize visibility of the right upper limb during reaching behavior in the foraging apparatus and to minimize occlusions, while avoiding the path between the X-ray sources and image intensifiers (Fig. 2.1A). The cameras were triggered to record continuous images between the onset and offset of the manual XROMM trigger, with series of images later converted to video for DLC processing. All videos were brightened using the OpenCV algorithm for contrast limited adaptive histogram equalization (CLAHE) prior to labeling. We labeled 11 body parts in DLC – two labels on the torso and three on each of the upper arm, forearm, and hand (Fig. 2.1B). Locations of each label were chosen to be as close as possible to the approximate location of XROMM beads, although concessions had to be made to ensure the location was not occluded consistently in the recordings. We used DLC 2.2 with in-house modifications to produce epipolar lines in image frames that were matched between the two cameras (Fig. 2.1C), which significantly improved human labeling accuracy by correcting gross errors and fine-tuning minor errors. We did not train a network on labels produced without the aid of epipolar lines and therefore cannot evaluate 3D error reduction using epipolar lines. However, we note that labels applied without epipolar lines on the torso were grossly inaccurate – these labels were adjusted by an average of 63 pixels and 57 pixels in camera-1 and camera-2, respectively, after implementation. The other nine labels were adjusted by an average of <1 pixel in camera-1 and 11 pixels in camera-2. This modification has been added as a command line feature in the DLC package (a guide for using epipolar lines can be found at https://deeplabcut.github.io/DeepLabCut/docs/HelperFunctions.html). Aside from this and related changes to the standard DLC process, we followed the steps outlined in Nath et al.[22]

22. Nath *et al.,* "Using DeepLabCut for 3D markerless pose estimation across species and behaviors".

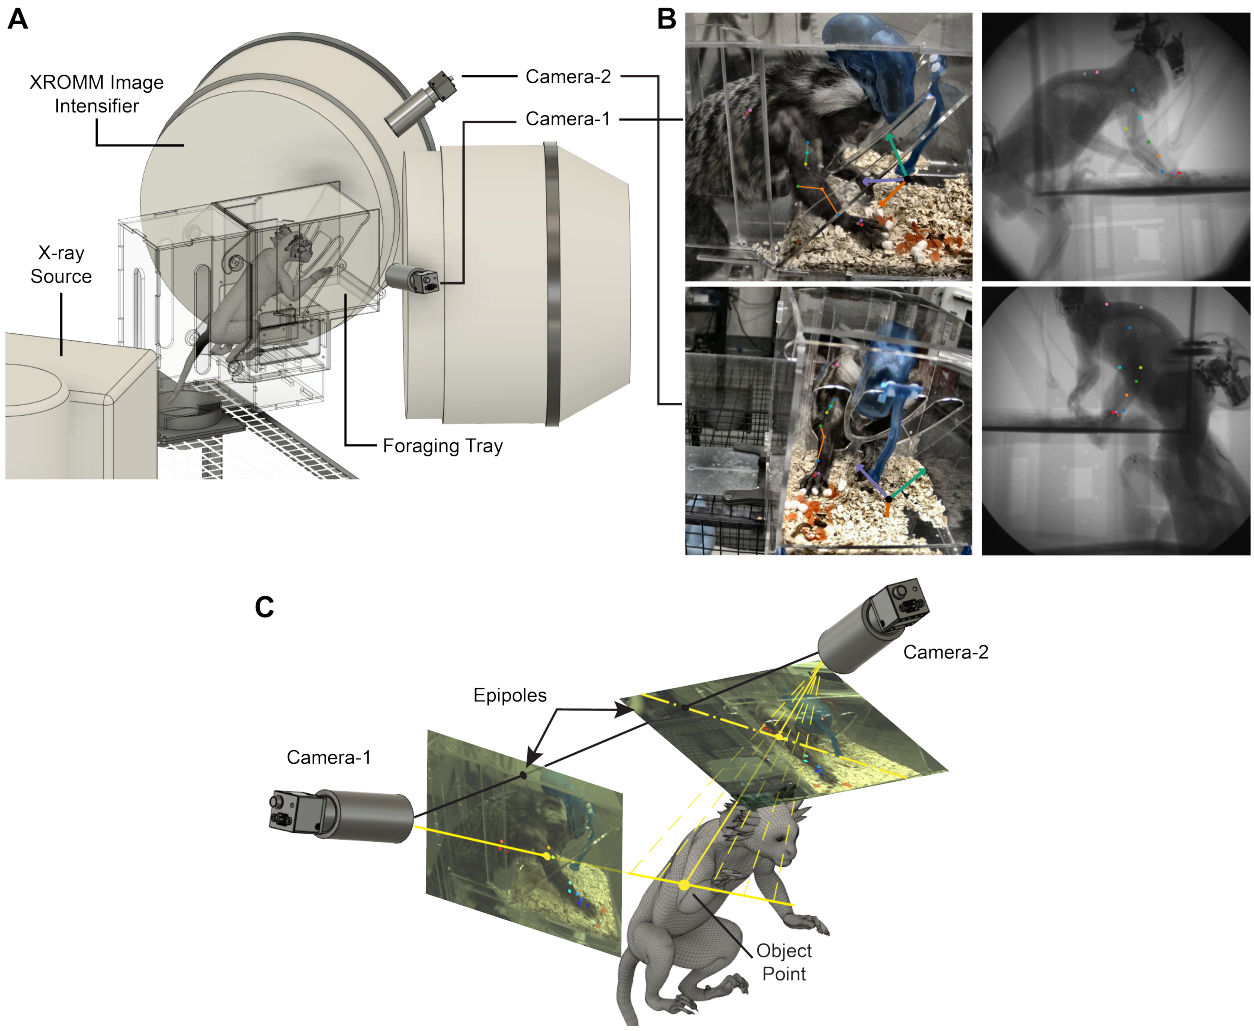

Figure 2.1: Recording apparatus, markers, and labeling with epipolar lines. (A) Marmosets use their upper limb to forage. Blackfly-S cameras and XROMM record foraging behavior simultaneously. **(B)** *Left*: Video frames from camera-1 (top) and camera-2 (bottom). Markers are applied by DLC+Anipose. The coordinate system is shown in green (x), orange (y), and purple (z). *Right*: corresponding XROMM frames. Subcutaneous tantalum beads are overlaid with colors to match corresponding DLC+Anipose labels. **(C)** Epipolar lines improve labeling accuracy. A vector projects from camera-1 through the applied label into 3D space, where it intersects with possible vectors from camera-2 (a subset is shown in dotted yellow). The epipolar line (dot-dash) passes through the epipole in camera-2 and each of the points at which a vector intersects the image plane. A correct label applied along the epipolar line produces accurate triangulation to the object point (correct labels and camera-2 vector shown in solid yellow).

In the first labeling iteration we extracted 100 total frames (50/camera) across the four events for marmoset TY and 254 frames (127/camera) across seven of the 13 events for

marmoset PT, which produced a labeled dataset of 354 frames. These were chosen manually to avoid wasting time labeling frames before and after reaching bouts during which much of the marmoset forelimb was entirely occluded in the second camera angle. An additional 202 frames (101/camera) were extracted using the DLC toolbox with outliers identified by the 'jump' algorithm and frame selection by k-means clustering. We chose the number of frames to extract for each video based on visual inspection of labeling quality and chose the start and stop parameters to extract useful frames that captured reaching bouts. In all cases, frame numbers of extracted frames were matched between cameras to enable the use of epipolar lines. This refinement step resulted in an error reduction of 0.046 cm and percent frames tracked increase of 14.7% after analysis with the chosen Anipose parameters. The final dataset consisted of 278 human-labeled timepoints from 15 of the 17 events and 10,253 timepoints from all 17 events labeled by the network only.

We used the default resnet-50 architecture for our networks with default image augmentation. We trained 3 shuffles of the first labeling iteration with a 0.95 training set fraction and used the first shuffle for the label refinement discussed above. We trained 15 total networks after one round of label refinement – three shuffles each with training fractions of 0.3, 0.5, 0.7, 0.85, and 0.95. Each network was trained for 300,000 iterations starting from the default initial weights. We evaluated each network every 10,000 iterations and selected the snapshot that produced the minimum test error across all labels for further analysis.

We chose the network to use in subsequent analyses by finding the smallest training set size that reached the threshold of human labeling error (discussed next). We then chose the median-performing network of the three shuffles at this training set size for all further analysis.

Human Labeling Error

We selected 134 frames (67/camera) across three events from the same marmoset and session to be relabeled by the original, experienced human labeler and by a second, less

experienced labeler. We used the error between the new and original labels to evaluate whether the networks reached asymptotic performance, defined by the experienced human labeling error.

Calibration

A custom calibration device was built to allow for calibration in both recording domains[23] (instruction manual for small lego cube is located in the XMALab BitBucket). The device was constructed to contain a three-dimensional grid of steel beads within the structure and a two-dimensional grid of white circles on one face of the cube. Calibration of x-ray images was computed in XMALab and calibration of visible light images was computed with custom code using OpenCV. This integrated calibration device, along with the PCA-based alignment procedure described below, ensures that DLC and XROMM tracked trajectories in a common 3D coordinate system. DLC videos were accurately calibrated, with 0.42 pixels and 0.40 pixels of intrinsic calibration error for camera-1 and camera-2, respectively, and 0.63 pixels of stereo reprojection error. XROMM calibration was similarly accurate, with average intrinsic calibration error equal to 0.81 pixels and 1.38 pixels for the two cameras.

Trajectory processing with Anipose

We used Anipose to analyze videos, filter in 2D, triangulate 3D position from 2D trajectories, and apply 3D filters (see Karashchuk et al.[24] for details). For 2D-filtering, we chose to apply a Viterbi filter followed by an autoencoder filter because the authors demonstrate this to be the most accurate combination of 2D filters. For triangulation and 3D filtering, we enabled optimization during triangulation and enabled spatial constraints for each set of three points on the hand, forearm, and upper arm, and for the pair of points on the torso. We identified six Anipose parameters and one post-processing parameter that may affect the final accuracy of DLC+Anipose tracking and ran a parameter sweep to find the

23. Knorlein *et al.,* "Validation of XMALab software for Marker-based XROMM".

24. Karashchuk *et al.,* "Anipose: A toolkit for robust markerless 3D pose estimation".

optimal combination. In 2D filtering, we varied the number of bad points that could be back-filled into the Viterbi filter ("n-back") and the offset threshold beyond which a label was considered to have jumped from the filter. We varied four parameters in 3D processing, including the weight applied to spatial constraints ("`scale_length`") and a smoothing factor ("`scale_smooth`"), the reprojection error threshold used during triangulation optimization, and the score threshold used as a cutoff for 2D points prior to triangulation. We also varied our own post-processing reprojection error threshold that filtered the outputs of DLC+anipose. We tested 3,456 parameter combinations in total, the details of which will be discussed below. We generally chose parameter values centered around those described in Anipose documentation and in Karashchuk et al.[25]

Post-processing of DLC+Anipose trajectories

To process the 3D pose outputs from Anipose, we first used the reprojection error between cameras provided by Anipose to filter out obviously bad frames. We tested two thresholds, 10 and 20 pixels, for 15 of 17 events. We tested much higher thresholds, 25 and 35 pixels, for the final two events of 2019-04-14 because the calibration was poor in these events – we suspect one of the cameras was bumped prior to these events. Next, we deleted brief segments of five or fewer frames and stitched together longer segments separated by fewer than 30 frames. Importantly, we did not have to do any further interpolation to stitch segments together, as Anipose produces a continuous 3D trajectory. Together, these steps remove portions of trajectories captured when the marmoset was chewing or otherwise disengaged from the foraging task and outside of the usable region of interest in camera-2 and combined segments during foraging bouts that were separated only by brief occlusions or minor tracking errors. All steps were performed independently for each label and event.

DLC labels could not be applied to the upper limb and torso in spots corresponding exactly to XROMM bead locations because those locations would often be obstructed from

25. Karashchuk *et al.,* "Anipose: A toolkit for robust markerless 3D pose estimation".

view by the marmoset's own body in one of the camera angles. We therefore applied labels as close as possible to the correct spots and subtracted the average position from each label and bead during post-processing. This removes a constant offset that should not be included in the DLC error calculations.

Despite our best efforts to place DLC and XROMM in the same 3D coordinate system through the calibration process described above, we found the two systems to be slightly misaligned. To fix this, we computed the three principal components across good frames for all DLC+Anipose labels and separately for all XROMM markers, then projected the mean-subtracted DLC+Anipose and XROMM trajectories onto their respective principal components. We found that this brought the coordinate systems into close alignment, such that we could no longer identify any systematic error that could be attributed to misalignment.

Finally, we found that there was a brief delay ranging from 0 to 10 frames between pedal-triggered onset of the XROMM event and the corresponding pedal-triggered TTL pulse initiating the start of the event for the FLIR cameras (and for the pulse ending the event). To adjust for the timing difference, we iterated over a range of possible sample shifts separately for each event to find the shift that minimized the mean absolute error between the DLC+Anipose and XROMM trajectory. We visually inspected each trajectory after the adjustment to ensure the shift was qualitatively accurate.

Evaluation of DLC Performance

We computed the median and mean absolute error between matched trajectories from DLC+Anipose and XROMM for all body parts across all reaching events. We also computed the percent of motion tracked across all labels and all active segments of reaching events. To define active segments, we manually inspected the videos for the first and last frames in each event for which the marmoset was engaged in the task; as mentioned before, the position of camera-2 prevented accurate human labeling when the marmoset was positioned

well behind the partition and the vast majority of these frames are discarded by Anipose and in post-processing.

Statistical Tests

Since the error distributions are right-skewed with long tails of large errors, we use the median error to describe the center of each distribution and the two-tailed Mann-Whitney U-Test to assess statistical significance. The P-values computed with this method are artificially low due to the large sample size (e.g. between 27,630 samples for the three upper arm markers and 11,480 samples for the two torso markers), so we report the correlation effect size defined by the rank-biserial correlation to describe statistical differences between distributions. According to convention, we consider $r < 0.20$ to be a negligible effect.[26]

In order to determine which of the Anipose and post-processing parameters from the parameter sweep significantly affected either the median error or percent of frames tracked, we created two linear regression models using the six parameters and a constant as independent variables and either error or percent tracked as the dependent variable. We tested the effect of individual parameters by calculating the log likelihood ratio Chi-squared test statistic (LR) between the full model and each nested model created by leaving one parameter out at a time (such that each nested model had a constant term and six parameter terms). We computed the p-value of each comparison using a Chi-squared test with two degrees of freedom.

We also created a full interaction model with the seven individual parameter terms and all possible first-order interaction terms. We tested the significance of each term by the same method.

Normalized Error and Fraction of Variance Accounted For

To compute normalized error, we divided the position error by the maximum range of motion for each marker across the dataset. To compute the fraction of variance accounted

26. Cohen, "A power primer."

for, we used the following equation:

$$FVAF = 1 - \frac{\sum [DLC - XROMM]^2}{\sum [XROMM - mean(XROMM)]^2} \tag{2.1}$$

which normalizes the sum of squared DLC error by the XROMM variance and subtracts from one.

2.4 Results & Discussion

Human versus DLC Error

We found that each network reached asymptotic performance within 300,000 iterations (Fig. 2.S1A). DLC error on training images was similar across set sizes, but test error decreased with set size until reaching an asymptote at the 85% training fraction (472 images, 236/camera), with mean error = 7.38 pixels (Fig. 2.S1B). This error matches that of the experienced human labeler (7.53 pixels) and is better than the inexperienced labeler (14.24 pixels). For subsequent analysis we used the median-performing shuffle of the 85% training set networks. As further confirmation of asymptotic performance, we found that the median error of frames that passed the post-processing reprojection error threshold remained constant across training set size while the percent of frames tracked reached an asymptote at the 85% split (Fig. 2.S1C).

Anipose Parameter Selection

By sweeping through a total of 3,456 combinations of six Anipose parameters and one post-processing parameter, we identified four with large, significant effects on the median error linear regression model ($p<0.05$ and log-likelihood ratio LR»50) – pre-triangulation score threshold (LR=4666), post-processing reprojection error threshold (LR=3098), spatial constraint (LR=1881), and smoothing factor (LR=1634; Fig. 2.2A-D). Error decreased linearly with score threshold without reaching an asymptote. It decreased monotonically

with smoothing factor until reaching 6 then continued to decrease at a small, constant rate. Error was lowest for spatial constraint = 2. Anipose reprojection error threshold was technically significant with a comparatively modest effect (LR=23.3). The remaining parameters were not significant (Fig. 2.2E-G).

We identified the same four parameters with large, significant effects on the percent tracked model – score threshold (LR=4689), post-processing reprojection error threshold (LR=4085), spatial constraint (LR=261), and smoothing factor (LR=81.7; Fig. 2.2A-D). Percent tracked decreased linearly with score threshold until dropping steeply at 0.6 and decreased slightly with spatial constraint and smoothing factor. N-back showed a modest effect (LR=8.75) with slightly higher percent tracked for n-back>1, and the remaining parameters were not significant (Fig. 2.2E-G).

The model incorporating first-order interaction terms provided limited information. Significant interaction terms were combinations of significant individual parameters and we found no evidence of nonlinear interactions (Fig. 2.S2). However, the interaction between spatial constraint and smoothing factor helped us choose 2 and 6 for their respective values; these were the smallest weights in the region of negligible error differences highlighted in Fig. 2.2H.

We visualized the effect of score threshold, post-processing reprojection error threshold, and n-back by plotting median error and percent tracked for each parameter set with spatial constraint = 2 and smoothing factor = 6, sorted by score threshold first, post-processing reprojection error threshold second, and n-back third (Fig. 2.2I). We found that `score_threshold` = 0.45 and post-processing reprojection error threshold = 20 provided a balance of low error and high percent tracked, and that n-back>1 improves percent tracked for some combinations. We selected the parameter set that minimized error under these constraints and performed subsequent analyses with `score_threshold`=0.45, `post_reprojection_error_threshold`=20, `scale_smooth`=6, `scale_length`=2,

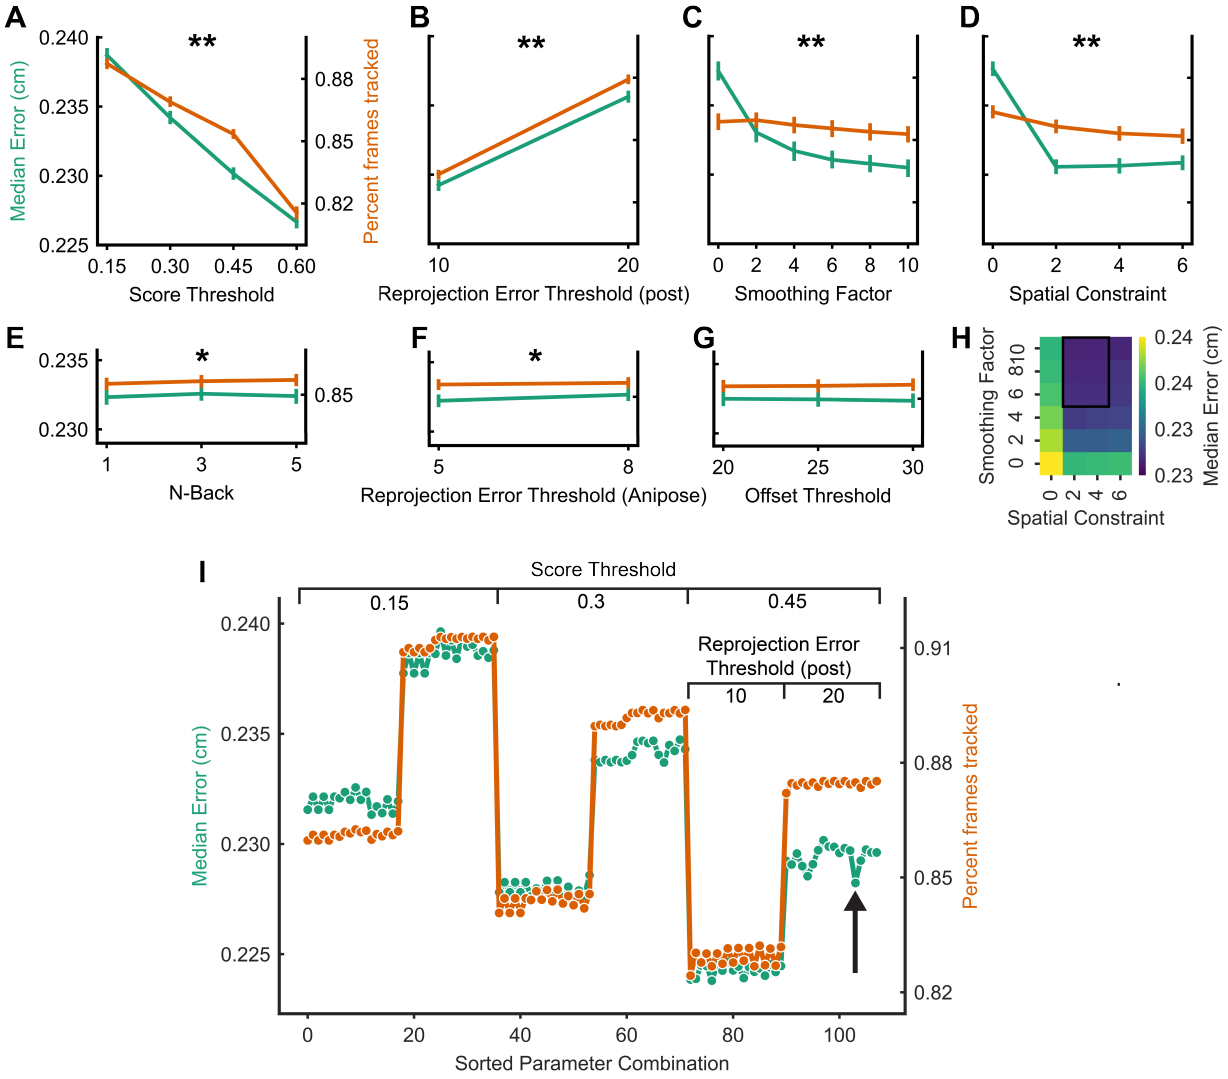

Figure 2.2: Anipose parameter sweep. **(A)** Median error (green) and percent frames tracked (orange) versus score threshold, averaged across the other parameters. Plots show mean±95% confidence interval of median errors. **(B-G)** Same for each of the tested parameters. Plots are labeled with ** to indicate a large effect on both error and percent tracked (log likelihood ratio Chi-squared statistic LR»50, p<0.05) and with * to indicate a modest effect on one or the other (LR<50, p<0.05). **(H)** Median error, represented by cell color, versus smoothing factor and spatial constraint. The combinations with lowest error are boxed. **(I)** Median error and percent tracked for 108 parameter sets with smoothing factor and spatial constraint set to 6 and 2, respectively. Results are sorted along the x-axis by score threshold first, post-processing reprojection error threshold second and n-back third. Arrow indicates the parameter set selected for further analysis.

`n_back`=5, `reproj_error_threshold`=8, and `offset_threshold`=20. We excluded 0.6 from the score threshold due to the large drop in percent tracked and unstable interactions

with other parameters (Fig. 2.S2J).

Tracking Examples

Qualitatively, DLC+Anipose and XROMM forelimb trajectories were nearly identical for most time points and reaches (Fig. 2.3A-B). The 3D view illustrates DLC tracking the overall motion of a foraging bout consisting of a long reach and shorter secondary reach faithfully, with some loss of fine details due to the smoothing factor (Fig. 2.3A). Breaking out x-y-z components demonstrates low position error (median = 0.179 cm; Fig. 2.3B).

Some reaches were not tracked with the same accuracy. Fig. 2.3C shows the medial forearm marker during a reach through the left-hand arm hole, resulting in three brief occlusions, the first two of which are removed by the post-processing reprojection error threshold. In the third, the marker gets "stuck" on the occluding part of the partition around 2 seconds, manifesting as a position oscillation after filtering and smoothing. Fig. 2.3D shows a reach in which a hand marker was occluded by the marmoset's other hand in camera-2 around 1.25 seconds, resulting in brief large errors before an untracked gap. It also shows errors early in the reach resulting from a non-optimal lag adjustment, which could happen if large errors in another label (the torso labels in this case) biased the optimal lag for the event.

Aggregate Error Results

Error distributions were right-skewed with a long tail of large errors, so we used median error as the measure of accuracy. We found that DLC tracked position with median error of 0.228 cm across all markers (Fig. 2.3E). We found that position errors on the torso (median = 0.302 cm) were significantly larger than errors on the forearm (0.222 cm, r=0.26) and upper arm (0.206 cm, r=0.319). We found no significant difference between the median error of 278 timepoints corresponding to the training set and that of the remaining data which were labeled by the network only (test).

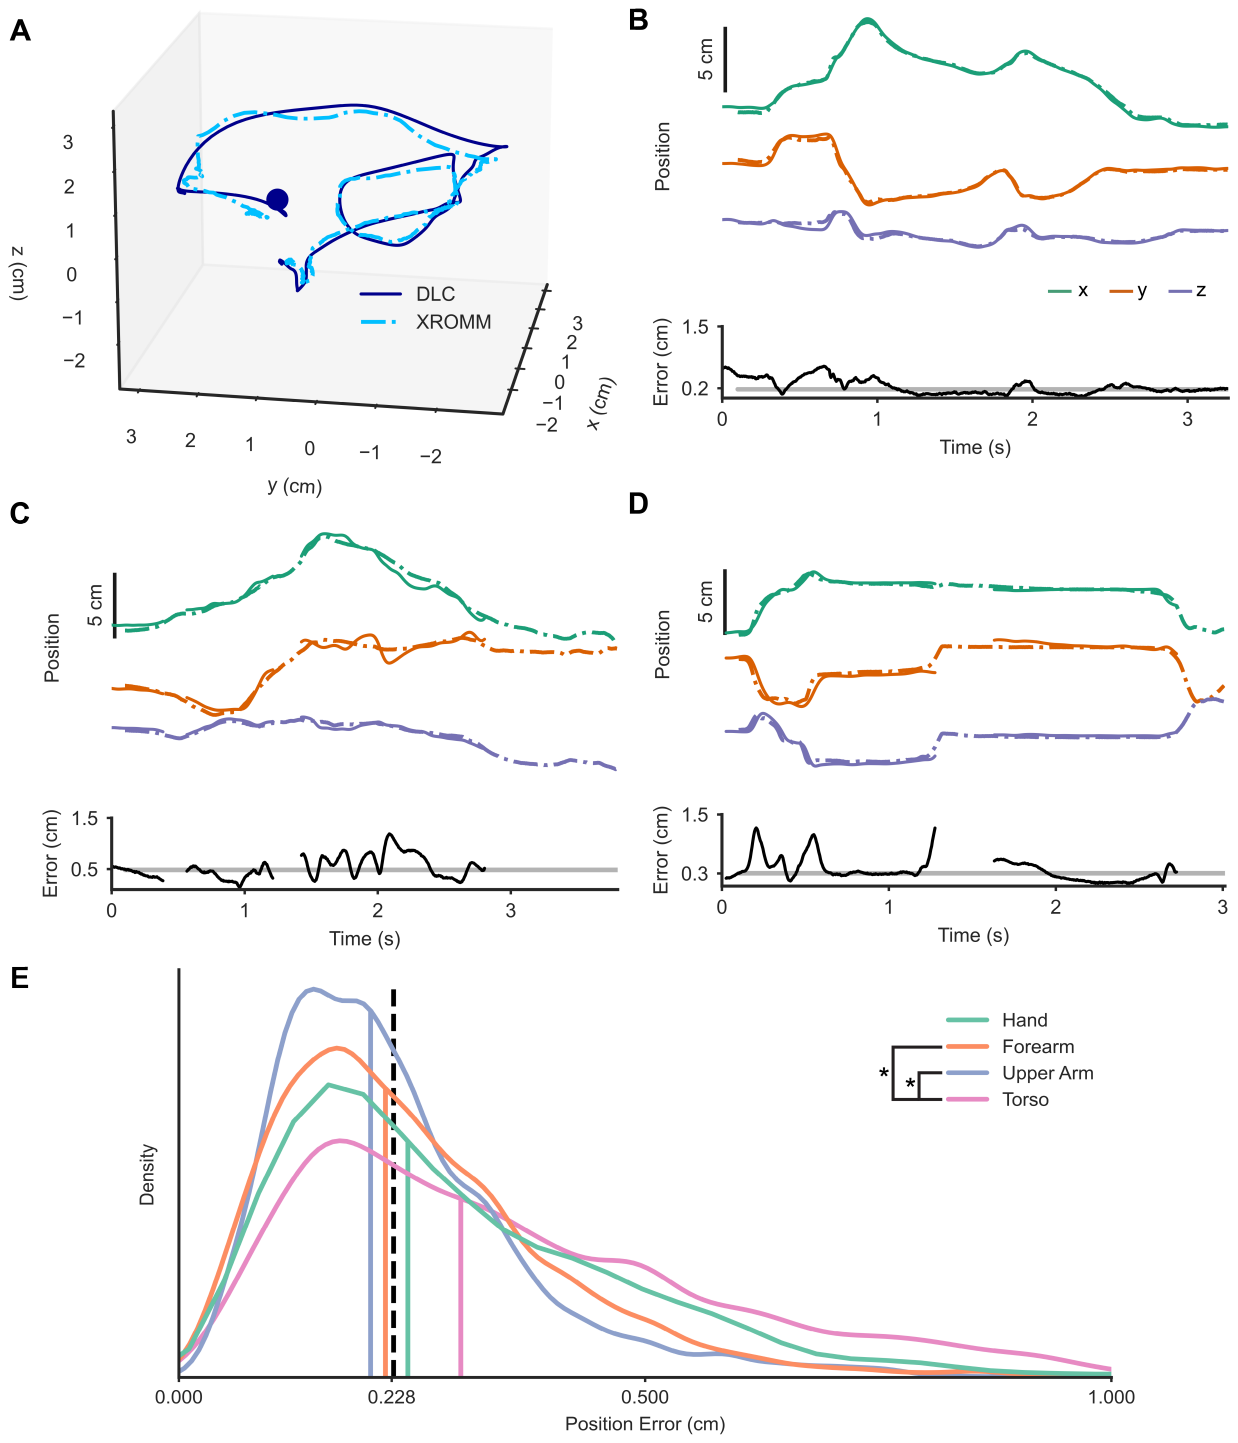
A
B
C
D
E
DLC
XROMM
z (cm)
y (cm)
x (cm)
Position
5 cm
x
y
z
Error (cm)
Time (s)
Density
Position Error (cm)
Hand
Forearm
Upper Arm
Torso
0.000
0.228
0.500
1.000

Figure 2.3: Tracking position with DLC+Anipose and XROMM. (A-B) Example of accurate DLC tracking. **(A)** 3D position of distal hand marker with DLC+Anipose (dark blue, solid) and XROMM (light blue, dot-dash). The movement begins at the blue dot. **(B)** *Top*: Position split into x-y-z coordinates colored to match the axes in Fig. 1B, with DLC+Anipose in solid and XROMM in dot-dash. *Bottom*: Error at each time (black) overlaid on median error for the event (gray). **(C-D)** Examples of DLC+Anipose tracking errors. **(E)** Distribution of errors for the markers of each body segment. Colored vertical lines show median error for each segment and the dashed black line shows median error across all markers. Significant differences between distributions are marked by asterisks in the inset (rank biserial correlation, r>0.2, n_{hand}=24605, $n_{forearm}$=24403, $n_{upperArm}$=27007, n_{torso}=11083).

Median error was equivalent to 2.0% of position range when normalized by the maximum range for each marker. Viewed another way, DLC+Anipose accounted for 97.2% of XROMM position variance. We also computed the pixel error between reprojected trajectories and labeled frames and found the median error to be 3.42 pixels (mean = 4.58 pixels), suggesting that Anipose improved on the 7.38 pixels of error present in the original un-processed network. The level of DLC+Anipose error demonstrated here is sufficient for many purposes, and there is reason to treat this as an upper error bound for well-trained DLC networks.

Assessment of Errors and Limitations

The primary limitation in this study is the viewing angle of DLC camera-2 (Fig. 2.1), which would ideally have been placed in the center of the XROMM image intensifier. The angle resulted in occlusion of torso markers in approximately 34% of relevant timepoints and occlusion of hand markers when the wrist was supinated. We approximated label positions for the supinated hand to maximize percent tracked, but this certainly increased the error. The second limitation is that marker positions were chosen as a compromise between proximity to XROMM bead positions and ease of labeling – for example, the medial forearm marker was placed in a muscular spot where the visual landmark was often ambiguous and resulted in poor re-labeling precision (12.2 pixels error for the experienced labeler), while the upper arm markers were easily located (5.96 pixels). There is also unquantified XROMM error from soft-tissue artifacts as well as skin deformation affecting DLC relative to subcutaneous

XROMM beads. The final limitation is the use of only two cameras – adding more would increase the percentage of frames tracked and improve 3D optimization in Anipose.

Improving Accuracy

DLC is open-source and actively maintained, with frequent implementation of new features. We have contributed one such feature by incorporating epipolar lines to improve labeling precision across cameras for 3D projects (Fig. 2.1C). Epipolar lines simplify the identification of the same landmark in different camera angles, particularly for ambiguous landmarks.

Multiple tools supplement DLC, including Anipose,[27] Pose3d,[28] OpenMonkeyStudio,[29] and a tool developed by Bolaños et al.[30] Anipose and Pose3d implement optimizations for 3D pose estimation python and MATLAB, respectively. OpenMonkeyStudio and the Bolaños et al. method are data augmentation tools – the former uses labels on a subset of 62 camera views to reproject labels on remaining views, while the latter animates a 3D model to render synthetic DLC training data.

Optimal Anipose parameters for our data may not generalize perfectly to different settings. We propose that readers without access to a marker-based system could optimize parameters by minimizing error between reprojected trajectories and held-out human-labeled data. Errors can also be reduced through manual frame-by-frame error correction of DLC-labeled video in XMALab or a similar toolbox.

Comparing DLC and XROMM Workflows

This dataset required 45-50 hours to produce a well-trained network to label 10,531 frames from 17 events (Fig. 2.S3A-B). This compares favorably with the approximately 34-

27. Karashchuk *et al.,* "Anipose: A toolkit for robust markerless 3D pose estimation".

28. Sheshadri *et al.,* "3D reconstruction toolbox for behavior tracked with multiple cameras".

29. Bala *et al.,* "Automated markerless pose estimation in freely moving macaques with OpenMonkeyStudio".

30. Bolaños *et al.,* "A three-dimensional virtual mouse generates synthetic training data for behavioral analysis".

51 hours required for semi-automated XROMM tracking of the same events (Fig. 2.S3C). While XROMM scales poorly to larger datasets, a robust DLC network may label such data with little or no manual intervention. For example, six events that were excluded from the initial set of labeled events required only the 5 hour refinement process, much less than 15 hours of XROMM tracking. Time requirements will depend on context, but the difference in scalability is evident. For work requiring both the precision of XROMM and improved scalability we suggest the procedure described by Laurence-Chasen et al.[31] to apply DLC directly to XROMM data.

Applicability to other pose estimation tools

There are a few alternative pose estimation tools. LEAP[32] and DeepPoseKit[33] use different network architectures to reduce training and inference time. Deep Graph Pose (DPG)[34] uses a similar architecture to DLC with temporal and spatial constraints to eliminate marker jumps. We suspect DPG attenuates high-error frames resulting from marker jumps but would not affect error for well-tracked frames. Based on advances in 3D human pose estimation,[35] DANNCE[36] trains a network on 3D image volumes rather than triangulating and optimizing outputs from a network trained on 2D images (as in DLC+Anipose; also see He et al.[37] and Reddy et al.[38] for subsequent state-of-the-art developments). DANNCE demonstrates higher accuracy than DLC for the same small number of cameras when applied to unconstrained settings in which 3D information is crucial for dealing with occlusions. However, we suspect DANNCE would perform similarly to DLC in semi-constrained settings such as ours

31. Laurence-Chasen *et al.,* "Integrating XMALab and DeepLabCut for high-throughput XROMM".

32. Pereira *et al.,* "Fast animal pose estimation using deep neural networks".

33. Graving *et al.,* "DeepPoseKit, a software toolkit for fast and robust animal pose estimation using deep learning".

34. Wu *et al., Deep Graph Pose: A semi-supervised deep graphical model for improved animal pose tracking.*

35. Iskakov *et al.,* "Learnable Triangulation of Human Pose".

36. Dunn *et al.,* "Geometric deep learning enables 3D kinematic profiling across species and environments".

37. He *et al.,* "Epipolar Transformers".

38. Reddy *et al.,* "TesseTrack: End-to-End Learnable Multi-Person Articulated 3D Pose Tracking".

with just a few cameras. Thus, we expect DLC errors reported here to be representative for existing tools, at least for well-tracked frames in semi-constrained environments.

2.5 Supplementary Figures

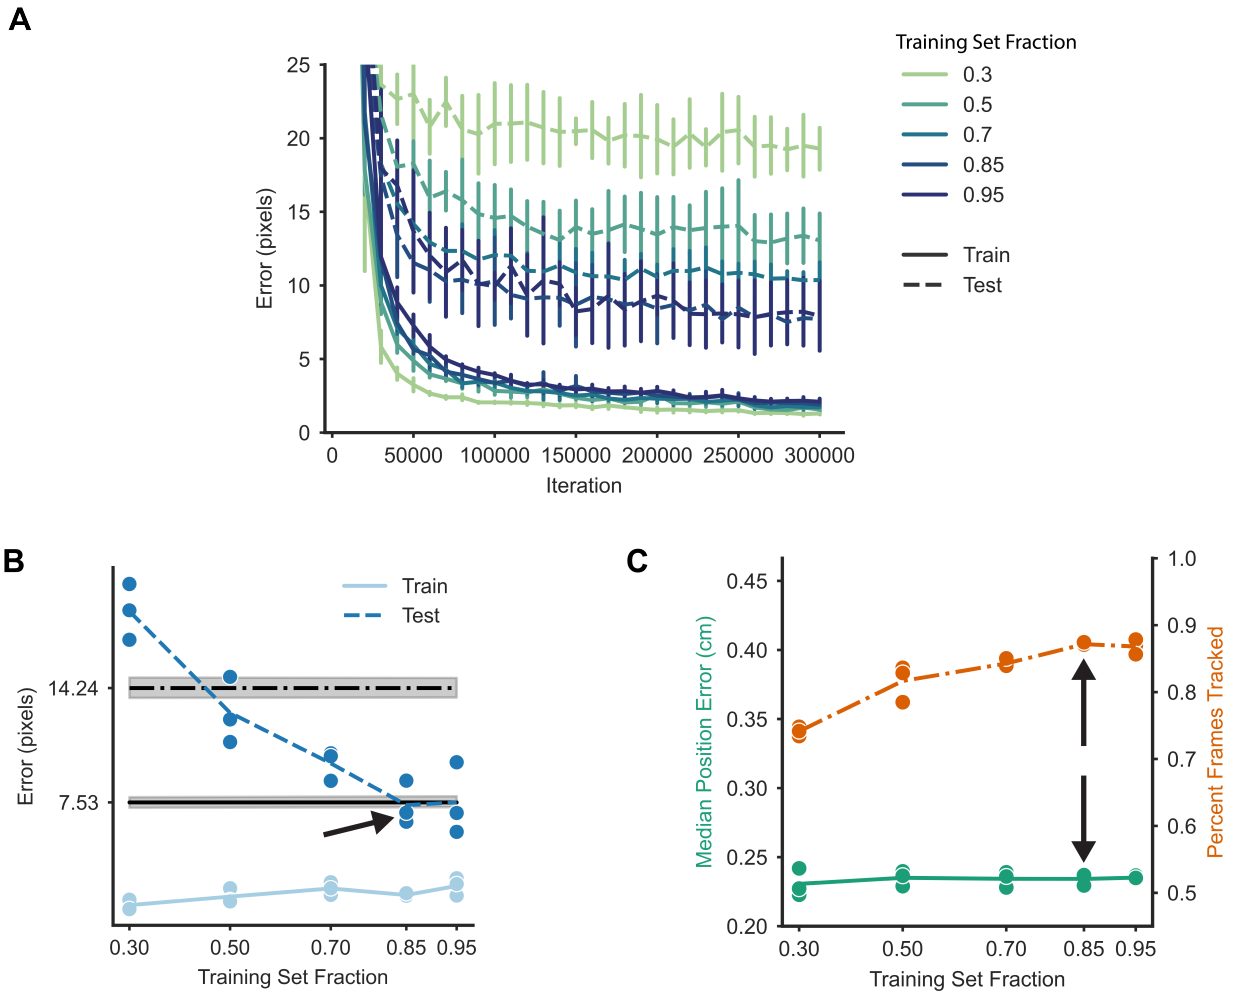

Figure 2.S1: Training and performance of the DLC Network. **(A)** Error as a function of training iteration for the 15 networks analyzed. Solid lines represent training error and dashed represent test error. Coloring varies from light green to dark blue with increasing training set fractions. Each line shows the mean ± standard deviation across three shuffles with that training set fraction. **(B)** Pixel error as a function of training set fraction, evaluated on DLC networks prior to Anipose. Lines represent the mean across shuffles and points indicate individual networks. Train error (solid, light blue) is well below both human labeling errors for all fractions. Test error (dashed, dark blue) drops below inexperienced human labeling error (dot dash, black) at 50% and asymptotes at the experienced human labeling error (solid, black) by the 85% training set. Human labeling errors presented as mean ± 95% confidence interval. Arrow indicates the network used for further analysis. **(C)** Median position error (green, solid) and percent frames tracked (orange, dot dash) as a function of training set fraction. This error was measured at the end of the DLC+Anipose and post-processing pipeline. Median error is not correlated with training set fraction because a reprojection error threshold was applied in post-processing, thus eliminating low-quality frames. Percent of frames tracked, on the other hand, increases with set size because more frames were well-tracked; percent tracked asymptotes at 85%. Arrow indicates the results corresponding to the selected network.

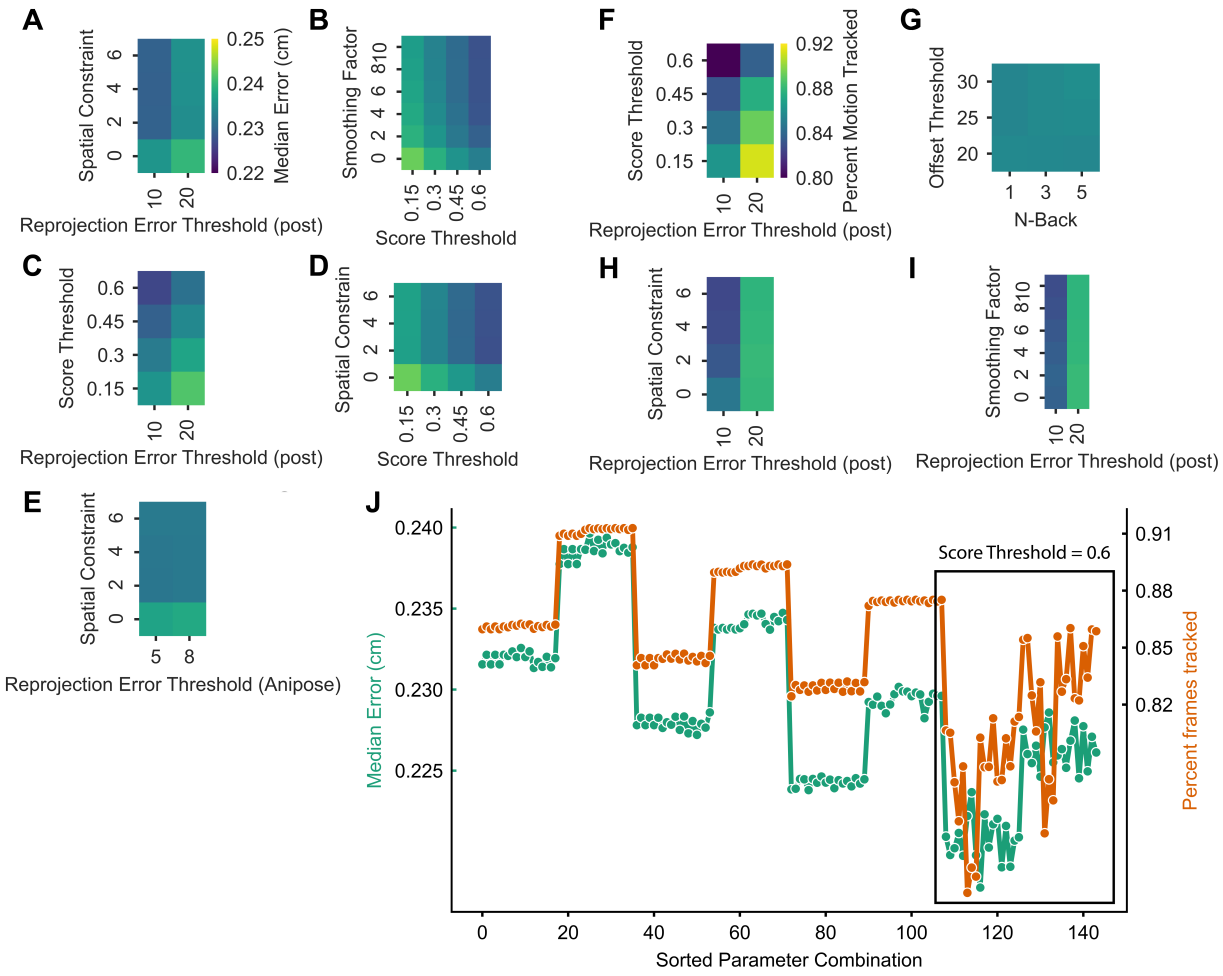

Figure 2.S2: Interaction between anipose parameters. **(A)** Median error, represented by cell color, versus spatial constraint and post-processing reprojection error threshold. Dark blue cells indicate the lowest error. **(B-E)** Median error vs four combinations of parameters. **(F-I)** Similar plots of percent frames tracked vs combinations of parameters, where yellow cells correspond to the highest percent and blue cells to the lowest. Note that the interactions in (A-I), although technically identified as significant, are minor and not evident given the vertical and horizontal striations in the figures. **(J)** Median error and percent tracked for 144 parameter sets including all score threshold values with smoothing factor and spatial constraint set to 6 and 2, respectively. Results are sorted along the x-axis by score threshold first, post-processing reprojection error threshold second and n-back third. The median error and percent tracked vary predictably for score thresholds between 0.15 and 0.45, but combinations with `score_threshold` = 0.6 vary unpredictably with the other parameters.

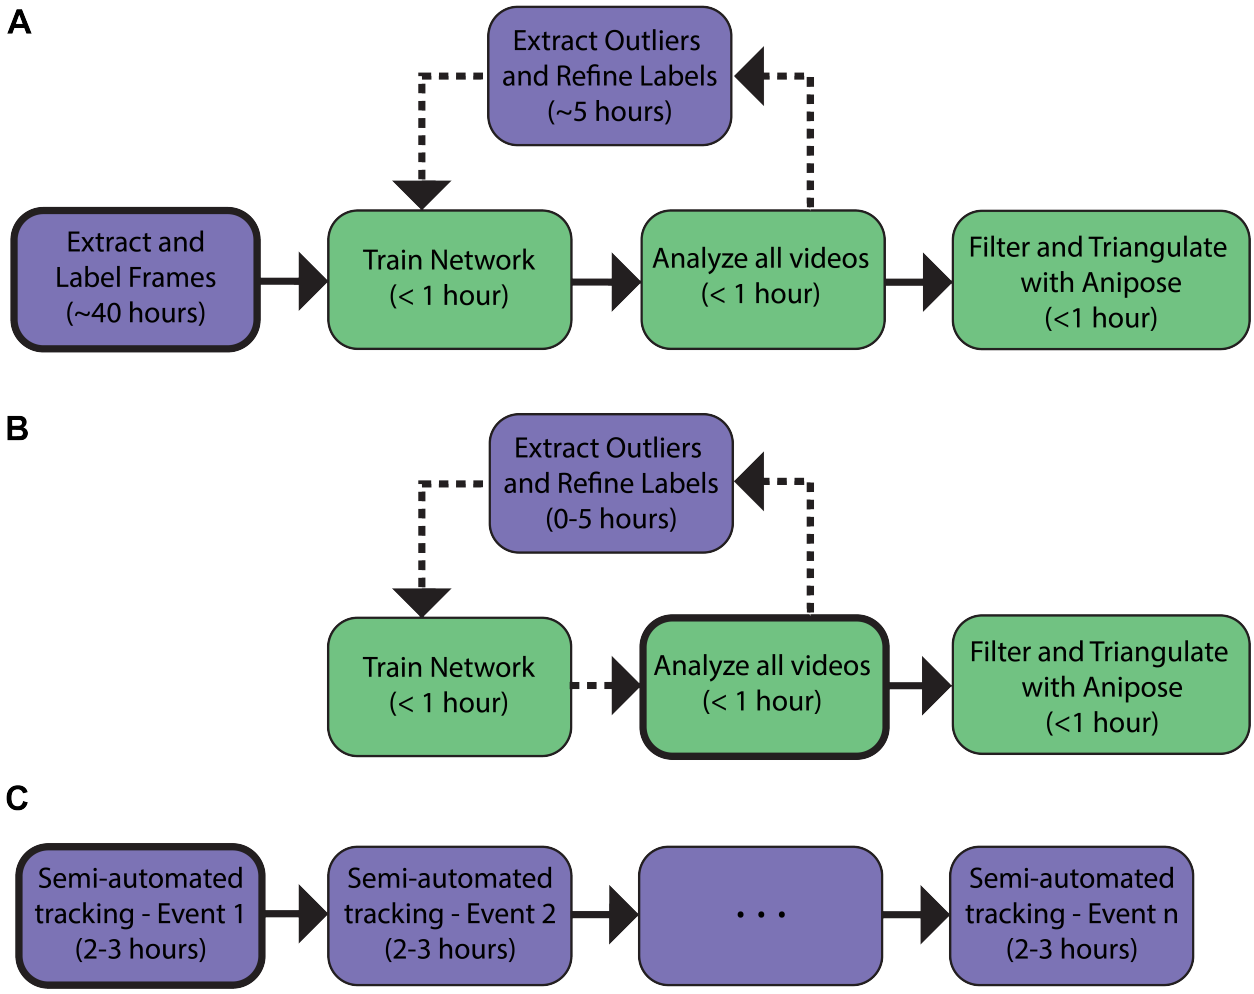

Figure 2.S3: DeepLabCut and XROMM pipelines. (A) The DLC workflow for the first time working with video data in a novel setting. **(B)** DLC workflow for analysis of subsequent data with the well-trained network from (A). **(C)** The standard XMALab workflow. Bold outlines indicate the starting point of each process. Purple steps require significant hands-on work while green steps are primarily computational. All time estimates provided refer to hands-on work either completing manual steps or preparing for computational steps. Dotted lines indicate an optional path in the pipeline. Calibration steps are not very different between DLC and XROMM and are not shown.

XROMM Location	DLC Location	Color
Body of metacarpal 3	Base of metacarpal 3	
Base of metacarpal 4	Proximal to base of metacarpal 3	
Base of metacarpal 2	Base of metacarpal 2	
Flexor carpi ulnaris	Visual landmark on distal forearm	
Brachioradialis	Visual landmark on medial forearm	
Anconeus	Elbow	
Lateral tricep distal	Visual landmark on distal upper arm	
Lateral tricep proximal	Visual landmark on medial upper arm	
Deltoid	Visual landmark on proximal upper arm	
Vertebrae – T4	Visual landmark on torso (anterior)	
Vertebrae – T8	Visual landmark on torso (posterior)	

Table 2.S1: Corresponding XROMM and DLC target locations. XROMM markers were targeted superficial to specific skeletal or muscular locations. DLC labels were chosen to match XROMM locations as close as possible, although we had to adjust marker positions to the location of clear visual landmarks so that consistent human labeling was possible. Colors match the colors in Fig. 2.1.

CHAPTER 3

A REACH-SPECIFIC SUBSET OF NETWORK INTERACTIONS UNDERLIES TUNING TO FORELIMB MOVEMENTS IN SENSORIMOTOR CORTEX

3.1 Introduction

Examining the relationship between population activity in primary motor cortex (M1) and upper limb motor behavior has taken one of two diverging approaches. On the one hand, single neuron tuning properties have been characterized and then combined across the population without considering the interactions between neurons.[1,2] On the other hand, low-dimensional structure in the population dynamics has been linked to movement while ignoring and, in some cases, arguing against the role of single neuron tuning properties.[3,4] Although these divergent methodologies have each proven useful, comparatively few studies have attempted to place single neuron tuning properties in the context of broader M1 population dynamics. Pairwise spike count correlations have been shown to provide information about motor behavior beyond what is provided by firing rates alone[5] and to improve encoding models that predict single-unit activity.[6] Recent work demonstrated that pairwise spike time correlations also carry behaviorally relevant information in M1.[7] A larger of body of work in mostly primary visual cortex has shown that recurrent interactions within local

1. Georgopoulos *et al.,* "Neuronal Population Coding of Movement Direction".

2. Georgopoulos *et al.,* "Primate motor cortex and free arm movements to visual targets in three- dimensional space. II. Coding of the direction of movement by a neuronal population".

3. Churchland *et al.,* "Neural population dynamics during reaching".

4. Shenoy *et al.,* "Cortical Control of Arm Movements: A Dynamical Systems Perspective".

5. Maynard *et al.,* "Neuronal interactions improve cortical population coding of movement direction."

6. Stevenson *et al.,* "Functional Connectivity and Tuning Curves in Populations of Simultaneously Recorded Neurons".

7. Sundiang *et al.,* "Dynamic structure of motor cortical neuron coactivity carries behaviorally relevant information".

functional groups can explain single-trial neural activity,[8] that both tuned and untuned units are essential components of the functional network (FN),[9] and suggested that reliable pairwise correlations, rather than first-order statistical features of spike trains, are the building blocks of coding in visual cortex.[10]

Our understanding of single-unit activity in the context of neuronal interactions has been limited in many cases by the necessity for trial averaging, due historically to wired neural recording systems, difficulty in quantifying complex kinematics, or by the constraints of analysis methods. Constrained and over-trained behavioral tasks such as center-out reaching or pedaling a wheel[11,12] limit the variability in all movement parameters aside from those under investigation and may impose an artificial ceiling on the dimensionality of neural population dynamics.[13] Even in the case of less stereotyped behaviors like random target pursuit,[14] reaching around obstacles,[15] and reach-to-grasp tasks,[16] primates are restrained in a chair with head-fixation and require weeks of daily training to reliably execute an experiment. There is evidence to suggest that a significant portion of M1 neurons encode a constrained, well-trained task differently than free behavior[17] and that over-training increases synchrony between M1 neurons.[18] The extent to which findings discovered in the constrained, trained setting will generalize to natural movements remains unclear.

8. Kotekal & MacLean, "Recurrent interactions can explain the variance in single trial responses".

9. Levy *et al.,* "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding".

10. Levy *et al.,* "A sparse set of spikes corresponding to reliable correlations is highly informative of visual stimulus on single trials".

11. Georgopoulos *et al.,* "On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex."

12. Russo *et al.,* "Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response".

13. Gao & Ganguli, "On simplicity and complexity in the brave new world of large-scale neuroscience".

14. Hatsopoulos *et al.,* "Encoding of Movement Fragments in the Motor Cortex".

15. Kaufman *et al.,* "Cortical activity in the null space: Permitting preparation without movement".

16. Saleh *et al.,* "Encoding of coordinated reach and grasp trajectories in primary motor cortex."

17. Jackson *et al.,* "Correlations Between the Same Motor Cortex Cells and Arm Muscles During a Trained Task, Free Behavior, and Natural Sleep in the Macaque Monkey".

18. Schieber, "Training and Synchrony in the Motor System".

Here we use a wireless neural recording system and computer-vision pose estimation tools to study unrestrained and untrained arm movements executed by the common marmoset (*Callithrix jacchus*) during prey capture of moths.[19,20,21,22] Compared to classical task paradigms, this behavior produces a richer set of movements with variable directions, speeds, and amplitudes and obviates concerns related to overtraining. Instead of assuming a fixed time lag between neural activity and kinematic features,[23,24,25,26,27,28,29,30] we quantify single unit tuning properties using a temporally extended trajectory tuning model.[31] At the same time, we incorporate neuronal interactions using a FN that describes pairwise spike time statistics in the form of a weighted, directed graph. We show that the trajectory tuning model extends to more naturalistic movements and explains neural activity better than a simpler model. We also demonstrated that tuning to kinematics depends on functional interactions between units – particularly on structured strong connections. Finally, we identified a reach-specific functional group that reorganizes to produce the dynamic forelimb movements necessary for prey capture. This reach-specific functional group is strongly interconnected and comprises units tightly linked to kinematics with strong, positively correlated

19. Karashchuk *et al.,* "Anipose: A toolkit for robust markerless 3D pose estimation".

20. Mathis *et al.,* "DeepLabCut: markerless pose estimation of user-defined body parts with deep learning".

21. Moore *et al.,* "Validating markerless pose estimation with 3D X-ray radiography".

22. Walker *et al.,* "Chronic wireless neural population recordings with common marmosets".

23. Aflalo & Graziano, "Partial tuning of motor cortex neurons to final posture in a free-moving paradigm".

24. Aflalo & Graziano, "Relationship between unconstrained arm movements and single-neuron firing in the macaque motor cortex."

25. Caminiti *et al.,* "Making arm movements within different parts of space: dynamic aspects in the primate motor cortex."

26. Evarts, "Relation of pyramidal tract activity to force exerted during voluntary movement."

27. Georgopoulos *et al.,* "On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex."

28. Hocherman & Wise, "Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys".

29. Moran & Schwartz, "Motor Cortical Representation of Speed and Direction During Reaching".

30. Sergio & Kalaska, "Systematic Changes in Motor Cortex Cell Activity With Arm Posture During Directional Isometric Force Generation".

31. Hatsopoulos *et al.,* "Encoding of Movement Fragments in the Motor Cortex".

preferred trajectories.

3.2 Results

Each marmoset was allowed to voluntarily enter and exit the prey-capture apparatus attached to the top of the home enclosure (Fig. 3.1A). Movement was recorded by high-speed cameras that were automatically triggered using a beam-breaker switch (see Methods). Spontaneous behavior in the home enclosure was recorded continuously by a separate camera system, but we do not quantify kinematics or report classes of spontaneous behavior in this work. Spiking activity was recorded throughout the recording session (TY = 175 units, Fig. 3.1B; MG = 73 units, Fig. 3.S1A), during which the monkey was actively reaching for prey (TY = 101 reaches, MG = 55 reaches) or engaged in undirected, spontaneous behavior. Reaches were randomly assigned to *reachSet1* or *reachSet2* and the corresponding spiking activity during reaches and during spontaneous behavior was used to compute functional networks denoted *reachFN1*, *reachFN2*, and *spontaneousFN* (Fig. 3.1C-D; Fig. 3.S1B-D). The reachFNs exhibited structures that were qualitatively very similar to each other and different from the spontaneousFN in both monkeys, which was reflected in the relationship between inter-electrode distance and edge weights for reachFNs versus spontaneousFN (Fig. 3.S1E). We built generalized linear models (GLMs) to predict single-sample spiking activity of individual units given a temporally-extended velocity trajectory and average position of the hand throughout the trajectory, similar to previous work in macaques.[32] We refer to these sets of models and to similar models without the position terms as the *full kinematics* and *trajectory* model, respectively (Fig. 3.1E). We also built models that approached a time-independent preferred direction representation by sampling brief velocity trajectories and average position between +100 and +150 ms with respect to the spike sampling time, which

32. Hatsopoulos *et al.,* "Encoding of Movement Fragments in the Motor Cortex".

is often treated as the optimal lag between a motor cortical neuron's firing and velocity.[33,34] We call these the *short kinematics* (brief trajectory and average position) and *velocity* (brief trajectory only) model (Fig. 3.1F).

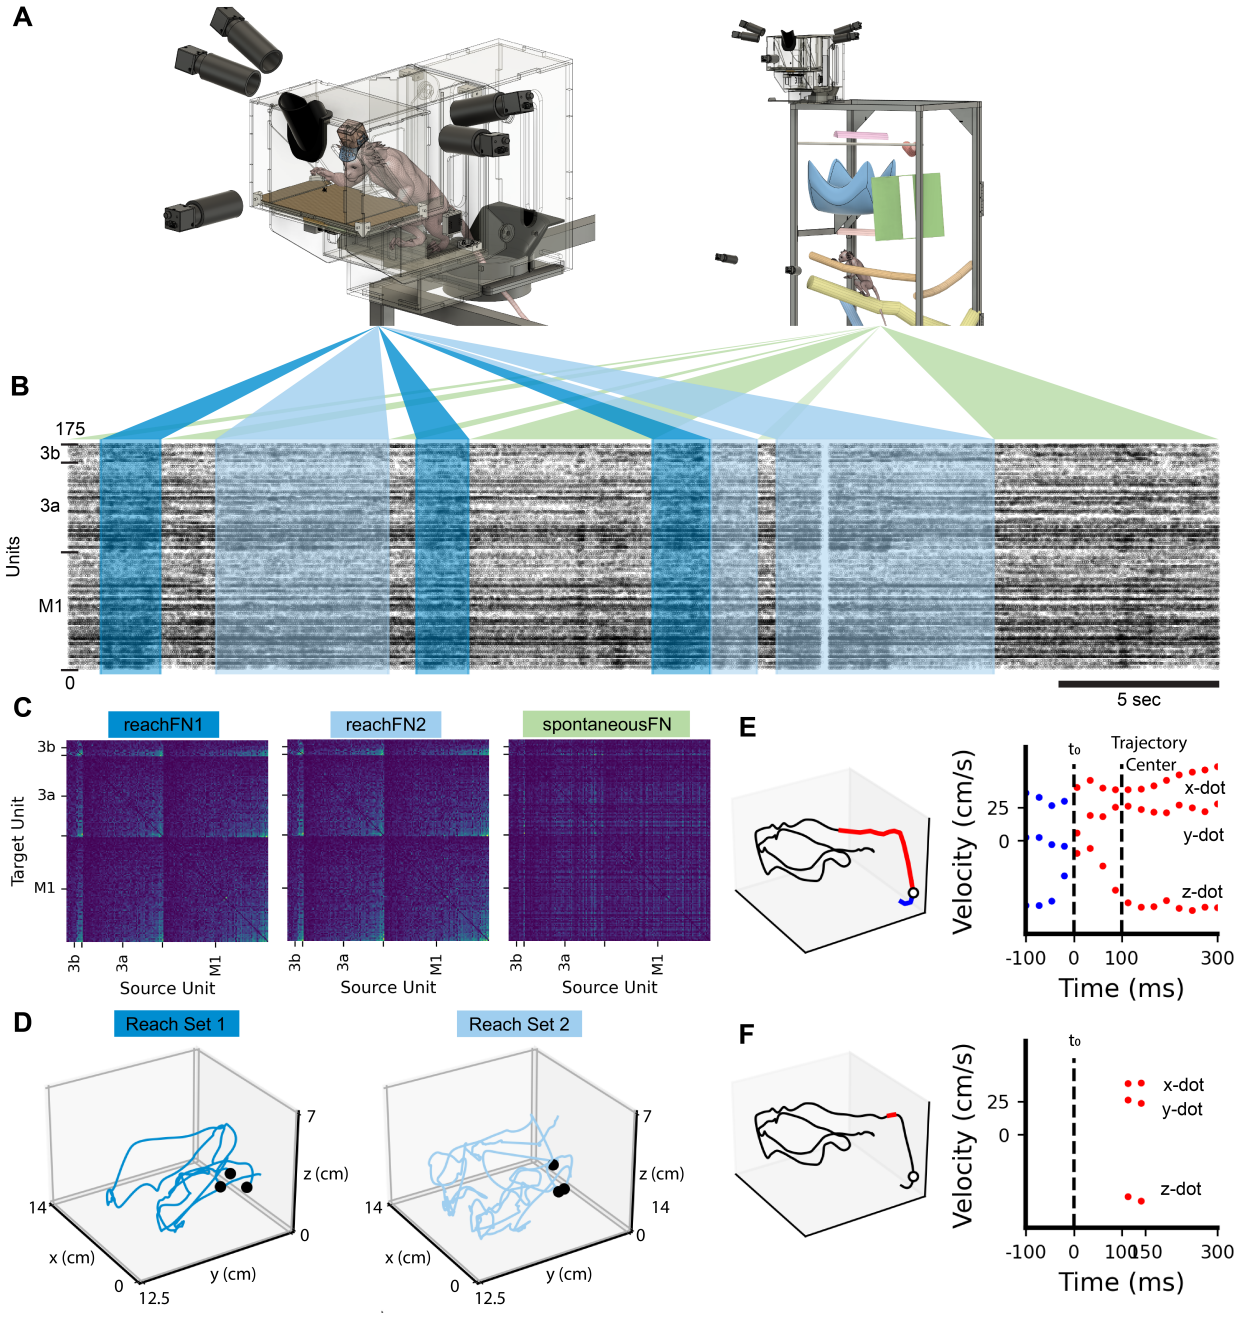

33. Moran & Schwartz, "Motor Cortical Representation of Speed and Direction During Reaching".

34. Paninski *et al.*, "Spatiotemporal Tuning of Motor Cortical Neurons for Hand Position and Velocity".

Figure 3.1: Behavior, data collection and encoding model construction. (A) *Left*: subject capturing prey in the apparatus. Behavioral video data was collected by five (MG) or two (TY) cameras. *Right*: subject behaving freely and spontaneously in the home enclosure. **(B)** A raster plot shows the activity of 175 units recorded from TY across time, ordered vertically by cortical area. Six of 101 reaches are separated into reachSet1 (dark blue) and reachSet2 (light blue). The remaining recording segments correspond to non-reaching, spontaneous behavior (green). **(C)** FNs were computed from activity during either reaches in the corresponding reachSet or during spontaneous behavior. Within each area, units are ordered by the average in-weight to the target unit for reachFN1. The color axis corresponds to the conMI. **(D)** Wrist position corresponding to the reaches in (A), with reach onset at the black dot. **(E)** A representative sample for the trajectory and full kinematics models with $\tau_{lead} = 100\ ms$ and $\tau_{lag} = 300\ ms$. *Left*: wrist position for a single reach (black) overlaid with the spike sample time (white circle) and corresponding trajectory sample including lead (blue) and lag (red) movements subsampled at 40 Hz. *Right*: velocity samples with the spike sample window and trajectory center shown with vertical black dotted lines. **(F)** The corresponding representative sample for the velocity and short kinematics models with $\tau_{lag1} = 100\ ms$ and $\tau_{lag2} = 150\ ms$.

Encoding models incorporating temporally-extended velocity trajectories and average position predict single-unit spiking activity

For each unit in the population, we tested 17 sets of each encoding model with varying sets of lead and lag times in the trajectory samples – the models were 300-500 ms in duration and ranged from entirely lead to entirely lag kinematics. We found that the full kinematics model performed best when it incorporated both lead and lag kinematics, particularly for models containing majority lag kinematics (Fig. 3.2A). We report subsequent results for the [-100, +300] ms model due to its high performance in both monkeys; we also confirmed that results match qualitatively across the multiple lead-lag sets with similar mean performance.

We created two shuffled models for comparison with real data – total shuffle and trajectory shuffle. For the total shuffle model we permuted the kinematic samples to break the relationship between kinematics and spikes. For the trajectory shuffle model we permuted just the trajectory samples, leaving the relationship between average position and spikes unperturbed. We called units for which the real model outperformed the shuffled model for 99 of 100 train/test splits strongly tuned to full kinematics (total shuffle) or to the trajec-

tory specifically (trajectory shuffle) for monkey TY. For monkey MG, we used a threshold of 90/100 to account for the reduced sample count. We found that many units were strongly tuned to kinematics (TY = 112/175 units; MG = 31/73 units), while a subset of these were strongly tuned to the trajectory specifically (TY=62/175; MG=8/73). For results shown in Fig. 3.2B-C, we filtered out units with full kinematics AUC greater than trajectory shuffle AUC for no more than 50% of splits, meaning that they had no sensitivity at all to the

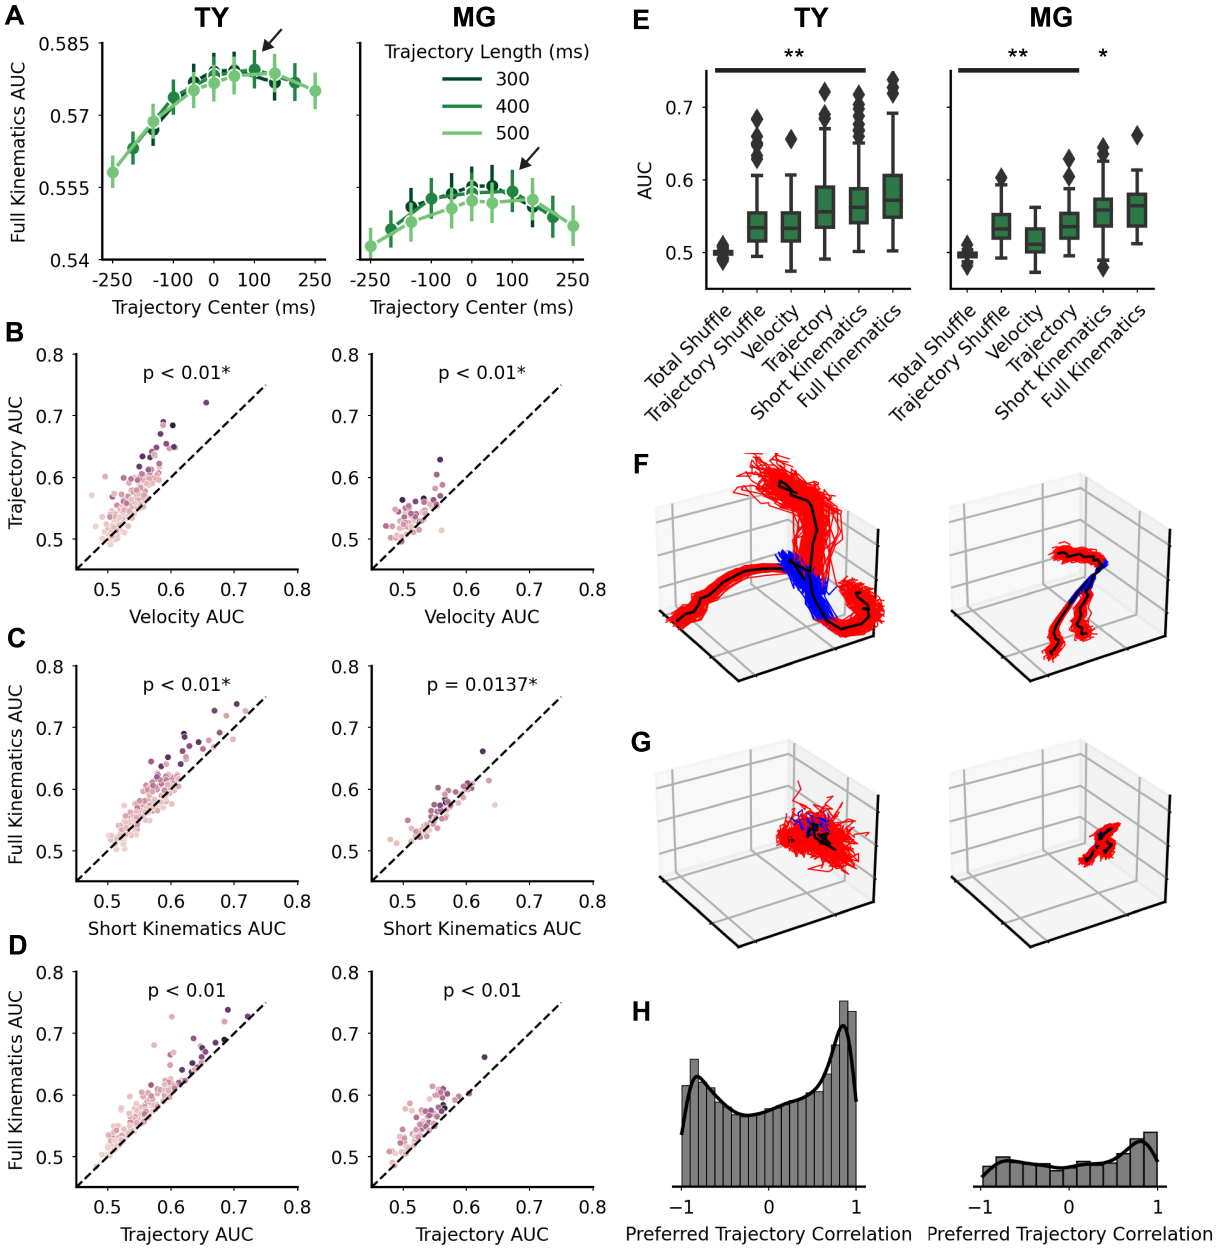

Figure 3.2: Full kinematics encoding model predicts single-unit activity. (A-H) *Left*: monkey TY. *Right*: monkey MG. **(A)** The mean ± sem AUC for the full kinematics model, all lead-lag sets. The trajectory models range from 300 to 500 ms in duration, and from entirely lead-heavy samples at the left end to lag-heavy samples at the right end. The arrow indicates the selection of the [-100, +300] ms trajectory model. **(B)** Scatterplot of the trajectory and velocity model AUC values for each unit, averaged over 100 train/test splits. Each unit's hue corresponds to average in-weight in the FN, which will be explained in further detail in Fig. 3.3. Units above the unity line were predicted better by the trajectory model. P-values are the result of a one-sided sign test. In (B) and (C), we filtered out units with full kinematics AUC greater than trajectory shuffle AUC for no more than 50% of train/test splits, leaving 169/175 units for TY and 61/73 units for MG. **(C)** The full kinematics model predicted activity better than the short kinematics model (TY=139/169 units, $p<0.01$ by sign test; MG=41/61 units, $p=0.0137$). **(D)** The full kinematics model outperforms the trajectory model due to the inclusion of average position terms (TY = 164/175 units, $p<0.01$, MG = 69/73 units, $p<0.01$). **(E)** Summary of model performance. ** indicates $p<0.01$ between the full kinematics model and other models, * indicates $p<0.05$. **(F)** The preferred trajectory pathlets for three of the five units with the highest AUC. Each of 100 train/test splits is shown in blue and red corresponding to lead and lag movements, and the average pathlet is shown in black. **(G)** The pathlets for three of the five units with the lowest AUC. **(H)** Histograms of the Pearson correlation between pathlets for all pairs of units.

details of the trajectory. We were left with 169 units for TY and 61 units for MG.

The trajectory encoding model predicted spike likelihood more accurately than the velocity model (TY = 138/169 units, $p<0.01$, one-sided sign test; MG = 45/61 units, $p<0.01$) and the full kinematics model was more accurate than short kinematics (TY = 139/169 units, $p<0.01$; MG = 41/61 units, $p=0.0137$), confirming previous studies and showing that activity of individual neurons is more closely related to temporally-extended velocity trajectories than to ballistic velocity at a fixed time lag (Fig. 3.2B-C). The full kinematics model was significantly more predictive than the trajectory model (TY = 164/175 units, $p<0.01$, MG = 69/73 units, $p<0.01$; Fig. 3.2D), as well as all the kinematics models we tested (Fig. 3.2E).

We treated the coefficients of the velocity trajectory terms in the full kinematics model as instantaneous velocities and integrated to obtain the preferred position trajectory, or pathlet. Units with high AUC values were tuned to high amplitude preferred pathlets that

were consistent across train/test splits, while units with low AUC values had no discernible preferred pathlet (Fig. 3.2F-G). We computed the Pearson correlation between pathlets for all pairs of units and found that preferred pathlets tended to be strongly correlated and skewed toward positive correlations (Fig. 3.2H). The full kinematics model incorporating a temporally-extended velocity trajectory and average position over the movement sample was the best predictor of single-unit spiking activity, and individual units exhibited distinct preferred trajectories that tended to be strongly correlated with a preference for positive correlations.

Activity in the functional group improves the encoding model and relates to kinematic tuning

We incorporated the FN in each target unit's encoding model by taking the dot product of all source unit spikes with the edge weights from the source units to the target unit (Eq. 3.6) for the leading bin and the coincident bin – this resulted in two network features for each target unit model. We took this approach rather than the alternative – fitting all pairwise coupling coefficients in the GLM – for three reasons. First, using conMI values as the coupling coefficients to produce network features constrained the encoding models to strict dependence on finely timed spike-train statistics measured directly from the neural data rather than on maximum likelihood computations. Prior work has shown that couplings taken from an FN approached optimality in a similar model.[35] Second, our approach allowed us to probe the influence of specific FN features and subsets and observe a direct result in model performance. Third, adding just two terms to each target unit model rather than 73 or 175 terms reduces the likelihood of overfitting the model. Our final kinematics+reachFN model contained the trajectory features with 48 $x-y-z$ velocity terms, the average position features with three terms, and the coincident and leading network feature terms (Eq. 3.5). Note that we use the term functional group to denote subsets of the FN for which functional

35. Dechery & MacLean, "Functional triplet motifs underlie accurate predictions of single-trial responses in populations of tuned and untuned V1 neurons".

interactions were studied in the GLM (in the unmodified kinematics+reachFN model, the functional group is the entire set of source units to the target unit).

The inclusion of network activity improves the model significantly across the population (TY = 169/175 units, p<0.01; MG = 73/73 units, p<0.01, one-sided sign test; Fig. 3.3A). Improvements were similar across cortical areas (Fig. 3.S2A-B). We found that performance of the full kinematics models increased with increasing average in-weight in the functional group (TY: Pearson correlation r = 0.67; MG: r = 0.51; Fig. 3.3B). Additionally, model

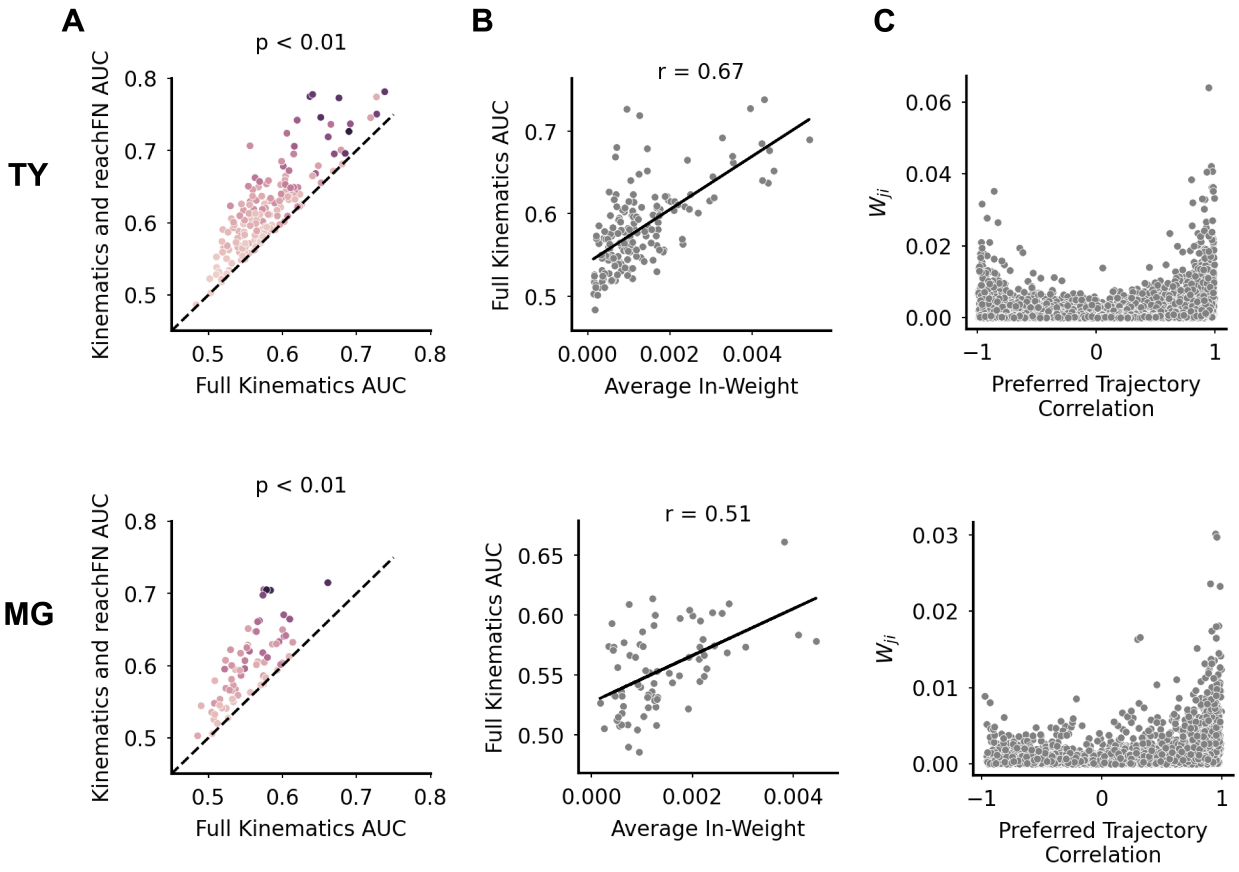

Figure 3.3: Network features improve the encoding model and first order FN structure is linked to kinematic tuning. (A-C) *Top*: monkey TY. *Bottom*: monkey MG. **(A)** Adding the reachFN network features to the full kinematics model improves prediction of single-unit activity (TY = 169/175 units, p<0.01, one-sided sign-test; MG = 73/73 units, p<0.01). Each unit's hue corresponds to average in-weight in the FN. **(B)** Performance of the full kinematics model (which contains no network feature terms) increases with average in-weight to the unit in the FN (TY: r = 0.67, Pearson correlation; MG: r = 0.51). **(C)** A scatterplot of all edge weights versus the pairwise preferred trajectory correlation.

improvement from adding network features was positively correlated with average in-weight but not with performance in the full kinematics model (TY: r = 0.66, MG: r = 0.49; Fig. 3.S2C-D). While we expect that a model incorporating functional group activity would improve with stronger inputs, it is striking that the units most strongly tuned to the full kinematics model (without network features) also received the stronger functional group inputs. Units with strongly correlated preferred trajectories also tended to have strong functional connections (Fig. 3.3C). The relationship between average full kinematics AUC in a pair of units and their preferred trajectory correlation is weak (Fig. 3.S2E), suggesting that the link between strong weights and strong trajectory correlations did not solely result from the relationship between full kinematics AUC and average in-weight. None of these findings are the result of variations in unit firing rate or waveform signal-to-noise ratio (Fig. 3.S3A-F). Taken together, these results demonstrate that functional interactions influence single-unit activity, that a unit's tuning to kinematics increases with the degree to which it is interconnected with the surrounding network, and that strongly connected units exhibit similar (or to a lesser extent, opposite) tuning properties.

The topology of strong functional interactions underlies accurate prediction of single-sample activity

Previous work in murine visual cortex has shown the strongest connections in FN are most informative of single-trial activity.[36,37] To test the importance of specific topological features in strongly connected functional groups in sensorimotor cortex, we identified the top N% of strongest edge weights in the FN. Using the method described by Kotekal and MacLean, we first permuted the weights amongst unchanged target-source unit pairs, which preserves the original functional groups (Fig. 3.4A). Network features computed from the permuted weights FN were then substituted into the kinematics+reachFN model. We did

36. Dechery & MacLean, "Functional triplet motifs underlie accurate predictions of single-trial responses in populations of tuned and untuned V1 neurons".

37. Kotekal & MacLean, "Recurrent interactions can explain the variance in single trial responses".

not re-train the GLM with permuted network features, but rather computed the change in AUC values on training data compared to the kinematics+reachFN model. Re-training the GLM would optimize model coefficients to minimize model error, which would obscure the effect of permuted weights. We also selected functional groups of matched size at random to determine whether functional groups of strong connections were of particular importance. Across 100 permutations for monkey TY (1 for each training set as in previous sections), we found that permuting weights in the connection strength functional group resulted in significantly greater information loss than permuting weights randomly for functional group sizes of greater than 15% and less than 100% of the network ($p < 0.01$, two-sided sign test; Fig. 3.4B). The effect of removing FN terms entirely from the GLM was significantly greater than the effect of weight permutations for set sizes up to 40%, beyond which shuffling functional groups including the next 50% of weaker connections had a similar effect as removing the network features entirely. However, increasing the set size did increase the amount of information lost, and shuffling 95% or 100% of the FN even degraded the model beyond simple removal of the terms.

We observed a less pronounced effect of permuted weights in monkey MG. The strongly connected functional group carried more information than the random functional group for multiple, but not all group sizes ($p < 0.01$; Fig. 3.4B). Unlike TY, the effect of removing the FN terms from the model exceeded the effect of permuted weights for all functional group sizes – even the 100% permuted FN. Although these results in MG demonstrate that the strength functional groups – and the FN more broadly – contributed to prediction of single-unit activity, they also suggest that a portion of the information available in the network terms came from overall population activity rather than the specific structure of functional interactions.

Next we permuted the target unit of the strongest (or random) edges, thereby generating new functional groups and allowing us to determine the sensitivity of predictive power to

the underlying strong edges. Unlike permuting the weights, which left the functional group of inputs to each target unit unaltered, this procedure permuted edges across all functional groups (Fig. 3.4A). As before, prediction loss was assessed on training data for 100 permutations. Permuting edges produced qualitatively similar results to permuting weights, although with heightened importance of the strongly connected functional groups for even the smallest group sizes (Fig. 3.4C). Strength-based permutation of edges had a consistently larger effect than random permutation for group sizes up to 75% in TY and 40% in MG, with the effect of edge permutation matching the effect of entirely removing FN terms for

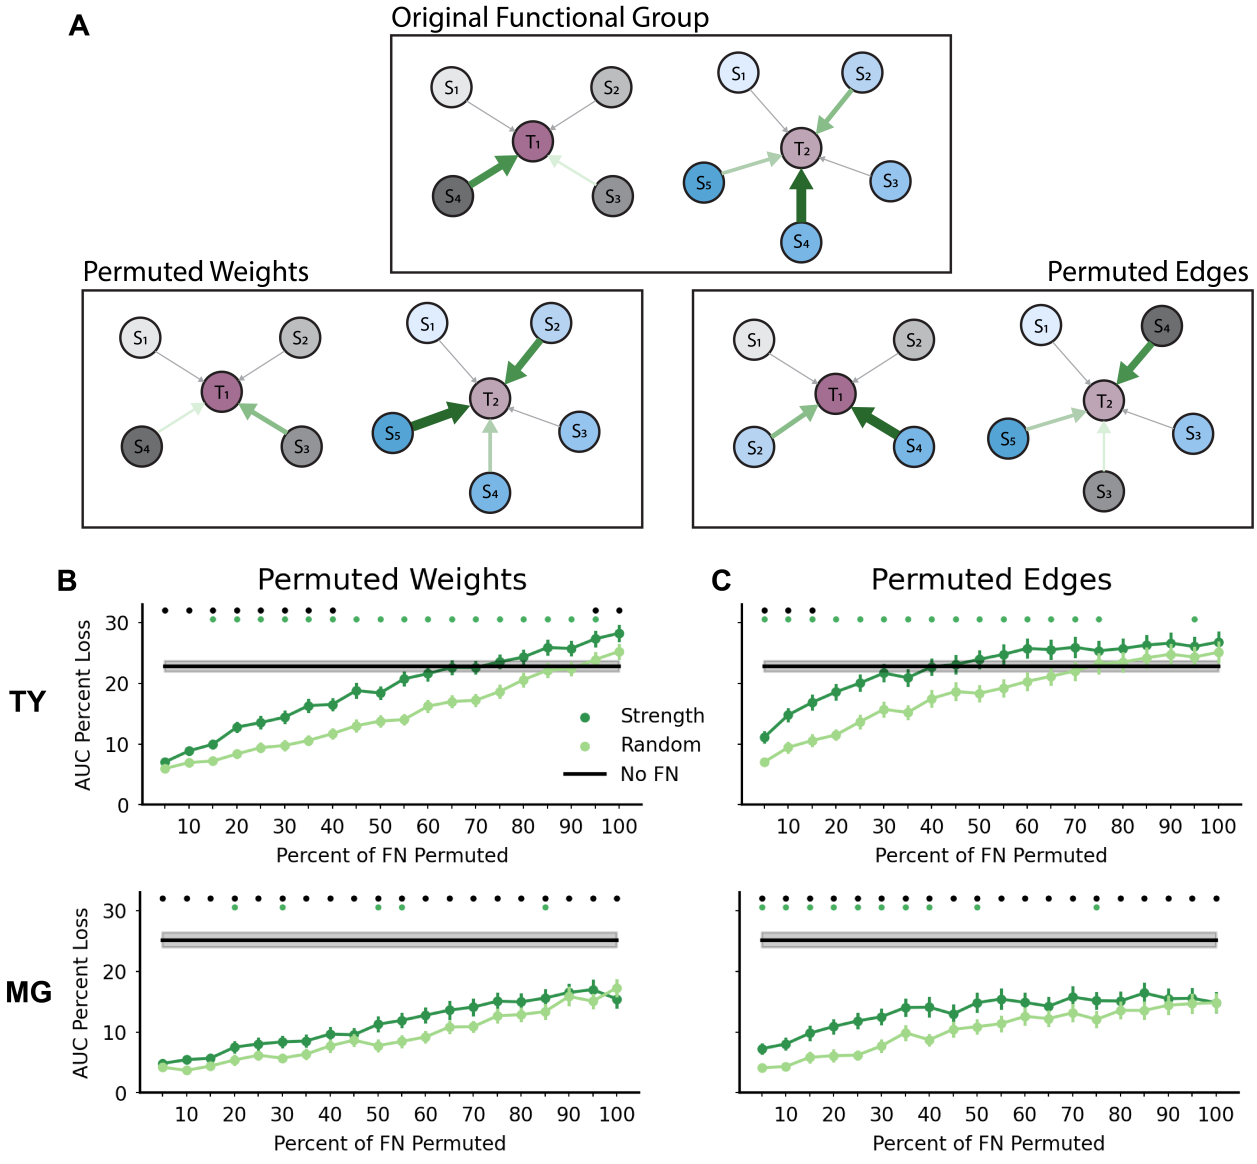

Figure 3.4: The topology of strong functional interactions underlies accurate prediction of single-sample activity. (A) Diagram of permutation methods. *Top*: two original functional groups of source units (S_n) and their input edge weights to target units (T_n). Green arrows indicate membership of this source-edge pair in the strongest N% of edge weights, with larger and darker arrows indicating a larger weight. Gray arrows constitute the remaining edges in the FN. *Lower-left*: for permuted weights, the strongest edges were permuted freely, even across functional groups, while the source-target unit pairs were unchanged. *Lower-right*: for permuted edges, the source-edge pair was held constant, but the target unit receiving input from the pair was permuted freely. This resulted in entirely new functional groups. Note that for simplicity this diagram does not depict the existence or permutation of reciprocal connections or the interconnectivity of functional groups. **(B)** *Top*: The percent AUC loss on the training set (mean ± sem for 100 shuffles) resulting from permuting the strongest N% of weights (dark green) or a random N% (light green) versus the size of the permuted functional group for monkey TY. Green circles above the plot indicate functional group sizes for which permuting the strongest weights resulted in significantly greater loss than permuting random weights ($p < 0.01$, two-sided median test). Black circles indicate sizes for which the effect of removing the network feature terms entirely was significantly different than permuting the strong weights ($p < 0.01$). *Bottom*: same for monkey MG. **(C)** Same as (B) for the effect of permuting edges in the original functional groups.

strongly connected functional groups larger than 15% in TY. As before, permutations of edges never matched the effect of removing FN terms for MG.

The study in visual cortex[38] permuted weights or edges for the top 25% of edge weights – here we observe significant differences between the strong connection and random functional groups across multiple group sizes. However, the difference observed for TY between perturbing or removing the FN terms suggests that information carried in functional interactions is focused in the 15-40% of strongest connections (Fig. 3.4B,C). The results from MG are less conclusive for this comparison to the findings in visual cortex.

The finding that strong functional interactions underlie accurate prediction by the kinematics+reachFN model lends support to our observations that strong connections correlate with tuning to kinematics and that units with highly correlated preferred pathlets tend to be more strongly connected. Taken together, these provide evidence that the specific structure of interactions captured by the FN are informative of the underlying dynamics involved in

38. Kotekal & MacLean, "Recurrent interactions can explain the variance in single trial responses".

kinematic production in the sensorimotor neural population, particularly for strongly connected functional groups.

A reach-specific functional group reorganizes during prey-capture

Although these results and prior work demonstrate the importance of strongly connected functional groups in the FN, connection strength may not be the only critical factor in functional group organization. We next took an approach based on model performance that was agnostic to any specific feature of the FN. We trained a kinematics and spontaneous network model by computing network feature terms from the dot product of the spontaneousFN with the spiking activity associated with trajectory samples. We then tested this model on network features computed with reachFNs to compare generalization of the kinematics+spontaneousFN model against the performance of the original kinematics+reachFN model. Essentially, we asked whether functional interactions computed during spontaneous behavior were informative during prey-capture. We found that the kinematics+reachFN model significantly outperformed kinematics+spontaneousFN generalization across the population for TY ($p < 0.01$, one-sided sign test) but not for MG ($p = 0.0503$; Fig. 3.5A,F). However, we observed that the kinematics+spontaneousFN model generalized well for most units – in fact, there was only a small subset of the population for which it clearly could not generalize (TY=19/175, MG=6/73). We call this subset the reach-specific functional group and the remaining units non-specific.

We isolated the reach-specific and non-specific functional groups for reachFN1 and spontaneousFN (Fig. 3.5B-C,G-H). We found that the difference between the reachFN1 and spontaneousFN was much larger for the reach-specific group compared to both the non-specific and full networks ($p < 0.01$, two-sided median test; Fig. 3.5D,I). Reach-specific reachFN1 also had a significantly stronger distribution of edge weights than the full network, while the non-specific reachFN1 had a significantly weaker distribution ($p < 0.01$; Fig. 3.5E,J). The pattern was reversed for the spontaneousFN in TY, for which the reach-specific

group had significantly weaker edges ($p < 0.01$). This did not hold for MG ($p > 0.05$). The non-specific group had a statistically weaker distribution of edge weights compared to the full FN in both monkeys ($p < 0.01$). These and subsequent results are consistent for reachFN2 – we don't display or report them because they are essentially identical. The reach-specific functional group which reorganized during prey-capture was also more tightly linked to forelimb movement. We found that the full kinematics model more accurately predicted spiking activity for units in the reach-specific functional group than for non-specific units

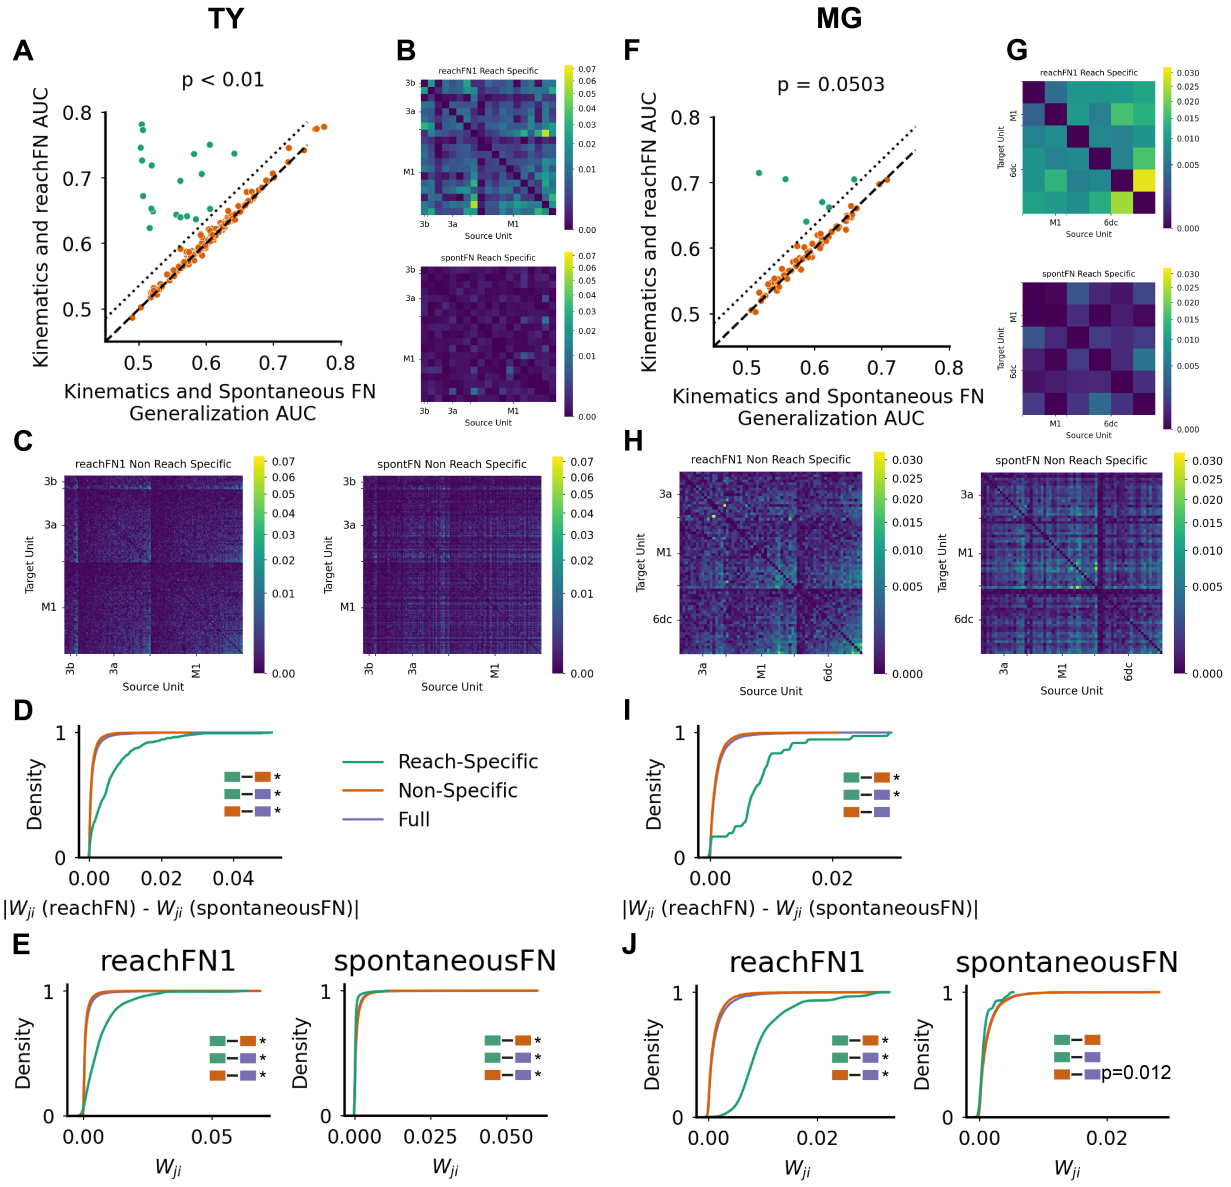

Figure 3.5: A reach-specific functional group reorganizes during prey-capture. **(A)** The kinematics and reachFN model versus the network encoding model trained with spontaneousFN network features and tested on reachFN network features, monkey TY. Units above the dotted line (threshold = 0.03 AUC away from the unity line) are defined as the reach-specific functional group (green, 19/175 units), and the remaining are non-specific units (orange, 156/173 units). **(B)** The reach-specific functional group, with reachFN1 on top and spontaneousFN on bottom. **(C)** The non-specific reachFN1 and spontaneousFN. **(D)** The cumulative distribution of the absolute difference in edge weights between reachFN1 and spontaneousFN for the reach-specific, non-specific, and full (purple) groups. Distribution comparisons are inset as colored pairs of boxes. For D and E, the * indicates significantly different distributions with p<0.01 (two-sided median test), remaining p-values less than 0.05 are written, and a blank next to the comparison inset indicates no significant difference. **(E)** The cumulative distribution of edge weights for reachFN1 (*left*) and spontaneousFN (*right*). **(F-J)** Corresponding results for MG, with 6/73 units in the reach-specific functional group.

(TY: $p < 0.01$, MG: $p = 0.03$, median test), and that the distribution of full kinematics AUC values was higher than in the full set of units as well (TY: $p < 0.01$, MG: $p = 0.03$; Fig. 3.6A). We also found that the preferred trajectories were more strongly and positively correlated within the reach-specific functional group than within the non-specific or full FNs ($p < 0.01$; Fig. 3.6B). To rule out the possibility that more strongly positive preferred trajectory correlations could be explained simply by the higher distribution of AUC values, we also tested unit pairs within the non-specific and full FNs for which both units had full kinematics AUCs exceeding the lowest full kinematics AUC in the reach-specific functional group. Pairs in the reach-specific group still exhibited more strongly positive preferred trajectory correlations than these AUC-matched pairs in both the non-specific and full FNs ($p < 0.01$; Fig. 3.S4A). Similarly, AUC-matched comparisons of the edge weight distributions shown in Fig. 3.5E,J produced the same results (Fig 3.S4B-C), with the exception that for monkey MG the distributions of edge weights were not significantly different between the reach-specific, non-specific, or full FNs. We note that reach-specific units did not comprise all of the units performing well in the full kinematics or full kinematics model, but were intermixed with well-explained non-specific units (Fig. 3.S4D).

Finally, we examined the location of reach-specific and non-specific units on the cortical

map as determined by intra-cortical microstimulation (ICMS) and passive tactile and somatic

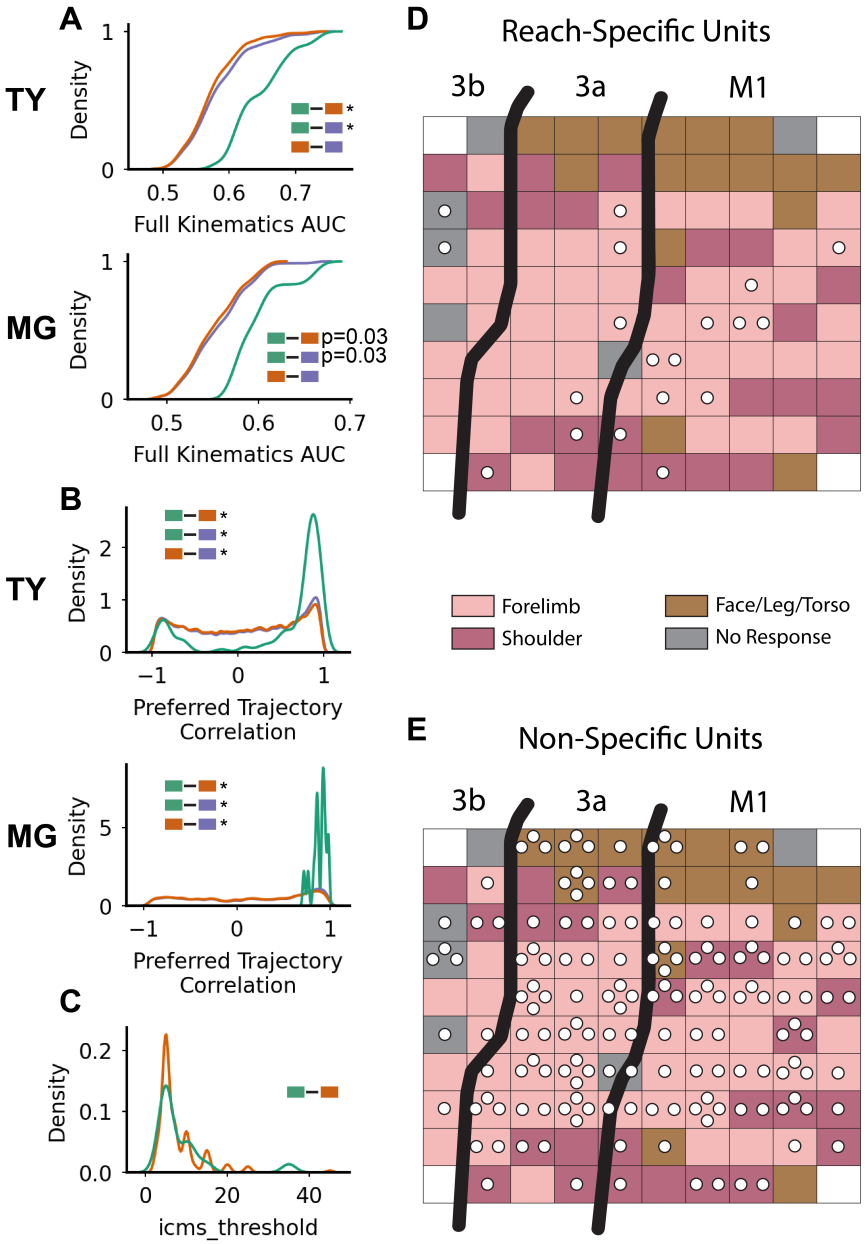

Figure 3.6: Tuning properties and cortical location of the reach-specific functional group. (A) Cumulative distribution of AUC values in the full kinematics model for TY (*top*) and MG (*bottom*), separated by reach-specific vs non-specific/full groups. Legend and statistics are denoted as in Fig 3.5. **(B)** Distributions of preferred trajectory correlation. **(C)** Cortical location of reach-specific units (white circles) on the array. Each box in the grid represents an array channel and is color-coded by the body-part associated with the channel through ICMS and/or passive tactile or somatic stimulation. **(D)** Cortical location of non-specific units.

stimulation while the marmoset was in a restful state (see Methods). We found that ICMS thresholds were not different between groups, ruling this out as an explanation for our findings (Fig. 3.6C). We discovered that for monkey TY the units in the reach-specific functional group were located entirely on channels linked to the shoulder or forelimb, with most units residing within M1 or within 3a near the boundary with M1 (Fig. 3.6D). The location of these units is striking in contrast with the non-specific units which were spread across the cortical map (Fig. 3.6E). Importantly, many channels on which reach-specific units were found also recorded non-specific units, and many non-specific units were located on other channels linked to forelimb movement or sensory feedback. Our preliminary examination of the location of reach-specific units for monkey MG was less conclusive (Fig. 3.S5A-B). Based on passive somatic and tactile stimulation, we have drawn estimated boundaries for the forelimb regions of area 3a, M1, and premotor cortex. We have not yet tested the effects of ICMS. Based on this mapping and expected channel response properties to ICMS, five of the six reach-specific units were located in the forelimb region and were spread across the cortical areas (Fig. 3.S5A). Non-specific units were spread across the array (Fig. 3.S5B).

3.3 Discussion

We have shown that precisely timed single-unit activity during forelimb movements can be predicted accurately by a network plus kinematics encoding model incorporating functional interactions, a temporally-extensive velocity trajectory and the average wrist position taken over the trajectory. We also demonstrated that tuning to kinematics depends fundamentally on functional interactions between units – particularly on structured strong connections. This builds on past work in macaque motor cortex demonstrating that movement related information is present in pairwise spike count correlations,[39,40] and provides complementary

39. Maynard *et al.,* "Neuronal interactions improve cortical population coding of movement direction."

40. Stevenson *et al.,* "Functional Connectivity and Tuning Curves in Populations of Simultaneously Recorded Neurons".

insights to recent work showing that the structure of fine-timing spike correlations in a FN contains movement-related information and evolves systematically over the course of behavior.[41] Finally, we identified a reach-specific functional group that reorganizes its functional interactions to produce the dynamic forelimb movements necessary for prey capture.

The trajectory encoding model predicts activity across sensorimotor cortex during naturalistic behavior

Given evidence that neural activity recorded in association with highly constrained and over-trained tasks may not generalize completely to naturalistic, unrestrained behavior,[42,43,44,45] it was not guaranteed that the results of our full kinematics model would match those of Hatsopoulos et al.[46] However, as in that study, we found that the trajectory model predicted single-unit activity more accurately than a velocity model and that the full kinematics model was most accurate for trajectories including a range of lead and lag kinematics. In fact, the same [-100, +300] ms model that performed best for M1 units in macaques performing a random-target pursuit task was also amongst the best-performing models here. Interestingly, the model accurately predicted units in area 3a and part of area 3b, as well as in motor cortex; this aligns with recent studies demonstrating similar encoding[47] and decoding[48] of distal limb movements of the wrist and digits for units in M1 and area 3a. We also found that significantly tuned units exhibited distinct and consistent preferred trajectories in addition to average position tuning. Since the trajectory tuning model

41. Sundiang *et al.,* "Dynamic structure of motor cortical neuron coactivity carries behaviorally relevant information".

42. Aflalo & Graziano, "Relationship between unconstrained arm movements and single-neuron firing in the macaque motor cortex."

43. Gao & Ganguli, "On simplicity and complexity in the brave new world of large-scale neuroscience".

44. Jackson *et al.,* "Correlations Between the Same Motor Cortex Cells and Arm Muscles During a Trained Task, Free Behavior, and Natural Sleep in the Macaque Monkey".

45. Schieber, "Training and Synchrony in the Motor System".

46. Hatsopoulos *et al.,* "Encoding of Movement Fragments in the Motor Cortex".

47. Goodman *et al.,* "Postural Representations of the Hand in the Primate Sensorimotor Cortex".

48. Okorokova *et al.,* "Decoding hand kinematics from population responses in sensorimotor cortex during grasping".

extended to accurate prediction of spiking activity during naturalistic behavior, it served as a useful tool for understanding additional contributions gained from the functional network.

Strong functional interactions underlie kinematic tuning properties

We found that the inclusion of network features in the kinematics+reachFN model significantly increased predictive power over the full kinematics model, and that performance of the full kinematics model increased with stronger average functional inputs – despite no direct link built into the kinematics-only model. Furthermore, we demonstrated that the specific structure of the strongest 15-40% of edge weights contained most of the information available in the FN. This agrees with a similar study in murine visual cortex which demonstrated that the topology of the functional group containing the largest 25% of edge weights was critical to the performance gained by incorporating network features.[49]

Another study in visual cortex found that both tuned and untuned units play essential roles in the functional network, and that untuned units were central to the structure of the network.[50] This is in contrast with two results presented here: that strongly interconnected units tended to be more tuned to kinematics, and that members of the reach-specific functional group were more strongly connected to each other and more tightly linked to kinematics. This suggests that untuned units may play a different role in sensorimotor cortex than in visual cortex. We posit that this difference is related to the generation of temporally smooth population dynamics that are necessary for production of motor behavior.[51,52,53] We discuss this possibility in further detail below.

49. Kotekal & MacLean, "Recurrent interactions can explain the variance in single trial responses".

50. Levy *et al.,* "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding".

51. Churchland *et al.,* "Neural population dynamics during reaching".

52. Russo *et al.,* "Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response".

53. Shenoy *et al.,* "Cortical Control of Arm Movements: A Dynamical Systems Perspective".

A reach-specific functional group for forelimb movement

We identified a small subset of the population, the reach-specific functional group, for which the kinematics+spontaneousFN model could not generalize to match the kinematics+reachFN model. When we compared this group to the non-specific and full FNs, we discovered that the reach-specific functional group was more strongly interconnected, contained pairs of units with strong, positively correlated preferred trajectories, and exhibited markedly different structures in the reachFNs and the spontaneousFN. In contrast, the structure of interactions between non-specific units was fairly consistent across spontaneous and reaching behavior. Additionally, the reach-specific functional group was more strongly tuned to forelimb kinematics than the non-specific units.

The novel finding that a reach-specific functional group, located primarily in forelimb M1 and 3a, reorganizes to produce dynamic forelimb movements during prey capture suggests intriguing possibilities for contextualizing single-unit properties within the framework of population dynamics. The functional interactions making up the non-specific group were consistent across dynamic reaching and the range of spontaneous motor behaviors in the home enclosure. We speculate that these units exhibiting static interaction patterns may constitute the population of neurons that generate low-dimensional and rotational dynamics described across the motor control literature.[54,55,56,57,58] Consistent pairwise interactions would contribute to dynamics that evolve in a predictable fashion and could produce neural trajectories with low tangling.[59] The reach-specific group, on the other hand, might reorganize its pattern of interactions in order to read out motor commands for movement production. Alternatively, the reach-specific group could be involved in setting the fixed

54. Churchland *et al.,* "Neural population dynamics during reaching".

55. Kaufman *et al.,* "Cortical activity in the null space: Permitting preparation without movement".

56. Lara *et al.,* "Different population dynamics in the supplementary motor area and motor cortex during reaching".

57. Russo *et al.,* "Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response".

58. Shenoy *et al.,* "Cortical Control of Arm Movements: A Dynamical Systems Perspective".

59. Russo *et al.,* "Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response".

point about which rotational dynamics unfold in neural space, a mechanism suggested as a link between classical tuning and dynamics in recent work.[60] Furthermore, it may be that the reach-specific functional group contributes preferentially to activity in the output-potent subspace,[61,62] while the non-specific group contributes preferentially to the output-null subspace. This is difficult to demonstrate with certainty, as it would require recordings of muscle activity in marmosets (which is technically challenging from a surgical perspective) or recordings of both goal-directed and spontaneous motor behavior in macaques with EMG recordings. As an alternative explanation, work by Dann et al.[63] suggests that reach-specific units identified here could be tightly coupled by preferred trajectory correlations and by strong interconnectivity in order to facilitate communication between areas.

The differences between reach-specific and non-specific groups aligns with recent work demonstrating that reliable pairwise correlations, rather than first-order statistical features of spike trains, are the building blocks of coding in visual cortex.[64] The fact that the non-specific group did not reorganize between spontaneous and prey-capture behavior implies a weaker dependence on the structure of interactions and a stronger dependence on source unit activity (first order features), while the reorganization of the reach-specific functional group demonstrates a dependence on precisely structured interactions for units strongly tuned to kinematics.

Future work could leverage the temporal network approach used by Sundiang et al.[65] to create a separate FN for each trajectory sample. Each FN would be truly specific to the reach segment and might provide insight into the reorganization of reach-specific functional

60. Sabatini & Kaufman, "Reach-dependent reorientation of rotational dynamics in motor cortex".

61. Elsayed *et al.,* "Reorganization between preparatory and movement population responses in motor cortex".

62. Kaufman *et al.,* "Cortical activity in the null space: Permitting preparation without movement".

63. Dann *et al.,* "Uniting functional network topology and oscillations in the fronto-parietal single unit network of behaving primates".

64. Levy *et al.,* "A sparse set of spikes corresponding to reliable correlations is highly informative of visual stimulus on single trials".

65. Sundiang *et al.,* "Dynamic structure of motor cortical neuron coactivity carries behaviorally relevant information".

groups on a fine timescale. Additionally, classification of spontaneous motor behavior would provide context to the spontaneousFN and would allow for multiple FNs computed during behavioral classes that are more or less related to forelimb reach-to-grasp.

3.4 Methods

Subjects

These experiments were conducted with two common marmosets, *Callithrix jacchus* (a 10 year old, 370 g male designated TY and an 8 year old, 350 g female designated MG). All methods were approved by the Institutional Animal Care and Use Committee of the University of Chicago.

Data Collection – Behavior and Cameras

A custom-built, modular apparatus designed for prey-capture and other goal-directed tasks[66] was attached at the top of the home enclosure. Subjects were allowed to enter the apparatus and engage in prey-capture voluntarily. When the subject was prepared and alert, the experimenter dispensed a single moth into the apparatus to initiate a prey-capture episode. The next prey was dispensed when the previous prey was either captured and eaten or had escaped the apparatus (a rare occurrence). The subject was engaged in spontaneous behavior in the home enclosure or in the apparatus when not directly alert and responsive to the prey-capture task. Spontaneous behavior includes, but is not limited to: leaping, hanging, grooming of self or pair housed monkey, ambulating or resting in the home enclosure and visual exploration or chewing/fine manipulation of prey in between capture episodes in the apparatus. A recording session lasted approximately 90 minutes for TY and 2.5 hours for MG.

High-speed cameras (Blackfly S, 200 frames/sec, 1440×1080 resolution; Teledyne FLIR)

66. Walker *et al.*, "A platform for semiautomated voluntary training of common marmosets for behavioral neuroscience".

were used to record video for pose estimation by DLC. For marmoset TY, two cameras were positioned to optimize visibility of the left upper limb and recorded at 150 fps. For marmoset MG, five cameras recorded at 200 fps – two cameras for each side view and one front-facing camera to improve coverage for reaches occluded in the side views. Image acquisition was triggered intermittently by the marmoset activating a beam-break sensor when approaching the partition within the apparatus. Additional cameras were used to record behavior in the home enclosure from a wider angle, but we do not attempt to classify or quantify behavior in the home enclosure in this work.

Data Collection – Array and Neural Recording

Each subject was implanted with a 96-channel Utah Array (Blackrock Microsystems, Salt Lake City, UT) using stereotaxic coordinates[67] to target the forelimb area of the primary motor cortex in the right (TY) or left (MG) hemisphere. The surgical procedure is described in detail by.[68] Neural data was collected using a Blackrock Cereplex Exilis, which is houses a digital amplifier, wireless transmitter and Li-ion rechargeable battery capable of powering 90 minutes of continuous recording in a compact headstage. A quick-connect solution designed in-house facilitated the removal/attachment and charging cycle with minimal experimenter intervention and disruption to the marmosets' natural behaviors (see Walker et al., which describes many of the design concepts which were adapted for the Exilis headstage). Additional battery life was provided in the MG recordings by a detachable external battery circuit designed in-house. Data is transmitted to 8 receiving antennas, then processed by additional products from Blackrock Microsystems. All data was recorded as a raw 30 kHz signal without filtering.

Data Collection – ICMS and receptive field mapping

Mapping of the arrays by ICMS and passive sensory stimulation was conducted at night

67. Burish *et al.,* "Microstimulation and architectonics of frontoparietal cortex in common marmosets (Callithrix jacchus)".

68. Walker *et al.,* "Chronic wireless neural population recordings with common marmosets".

during quiet restfulness. For ICMS, a medium-strength current was used to identify muscle groups for which stimulation evoked a twitch or movement. Then the current was reduced to identify the specific muscle target of the channel and the lowest current that could evoke a response. For passive stimulation, the skin was lightly brushed with a cotton-tipped applicator (tactile) or the muscle body was palpated (proprioceptive). This was repeated 20 or more times for each muscle/body region using a 5sec on/5 sec off protocol. Peri-stimulus time histograms were computed from the repetitions and mapped onto the array for manual inspection by the experimenters. Cortical area boundaries were estimated from the composite maps of ICMS and receptive field mapping, in accordance with cortical maps from.[69]

Data Processing – Spike Sorting

Spike sorting was performed on raw neural data at 30 kHz using Spike Interface.[70] Ironclust (https://github.com/flatironinstitute/ironclust) was the primary sorter, with Spyking-Circus[71] and waveclus[72] used to cross-reference for consistent units that were identified across all sorters. Inter-spike interval, autocorrelation, and signal-to-noise ratio thresholds were applied to automatically classify units as single units, multi-unit activity, noise, or unclassified. Automatically sorted spikes were manually curated using phy (https://github.com/cortex-lab/phy).

Data Processing – Kinematics

A comprehensive DeepLabCut network (DLC)[73] with Resnet-50 base architecture was trained on 2,343 labeled images from TY and MG recording sessions. Three labels were

69. Burish *et al.,* "Microstimulation and architectonics of frontoparietal cortex in common marmosets (Callithrix jacchus)".

70. Buccino *et al.,* "Spikeinterface, a unified framework for spike sorting".

71. Yger *et al.,* "A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo".

72. Chaure *et al.,* "A novel and fully automatic spike-sorting implementation with variable number of features".

73. Mathis *et al.,* "DeepLabCut: markerless pose estimation of user-defined body parts with deep learning".

applied to each side of the wireless headstage housing (for head-tracking in future work), three to the upper limb on each side (shoulder, elbow and wrist), and three to corners of the apparatus to establish a coordinate system – a total of 15 labeled points. The model was iteratively trained and refined until pose estimation was consistent throughout all significant upper-limb movements in the apparatus prey capture space. Anipose was used to apply the DLC network to videos and to perform subsequent 2D filtering and 3D calibration and triangulation, as described in Moore et al.[74] and Karashchuk et al.[75]

We added multiple steps of post-processing to reduce the effect of brief tracking lapses in outputs from the well-trained DLC network (such as spurious jumps and brief occlusions). Each step was applied independently to all markers and video events. We first filtered out timepoints with reprojection error greater than 20 pixels (35 for MG) and fewer than two cameras tracking the label, leaving only well-tracked segments of kinematic data. We then removed very brief segments shorter than 50ms, then filled tracking data back in for inter-segment gaps shorter than 200ms. Most lapses were fixed or removed correctly after these simple steps. Next, we removed any remaining marker jumps using the reprojection error and position data together. Then we replaced any brief inter-segment gaps that remained with either original tracking data if that data matched a linear interpolation closely, or with the interpolation. Finally, we trimmed long segments with a high percentage of interpolated data at the beginning and end of video events – these segments corresponded to the marmoset maintaining an occluded posture in the back of the apparatus before or after prey-capture activity. The processed data was smoothed with a 3rd order, 70ms Savitsky-Golay filter. We identified reaching segments by finding y-position peaks that indicated extension of the upper-limb into the prey-capture space and assessing when hand speed crossed a threshold before and after each peak.

74. Moore *et al.,* "Validating markerless pose estimation with 3D X-ray radiography".

75. Karashchuk *et al.,* "Anipose: A toolkit for robust markerless 3D pose estimation".

Trajectory encoding model

To build the encoding model for each neuron, we used a generalized linear model that estimated the set of coefficients to maximize the likelihood of correctly predicting the spike count within a 10 ms spike sampling window given the corresponding sample of kinematics. Each kinematic sample was the monkey's hand velocity trajectory and average position over a kinematic sampling window of length $\tau_{lead} + \tau_{lag}$, with the kinematic and paired spike samples centered at t_0. Samples of spikes and kinematics were extracted every 30 ms throughout reaching segments, resulting in around 8,000 samples for monkey TY (8,149 for [-100, +300] ms model) and 4,000 samples for monkey MG (4,250 for [-100, +300] ms model). Instantaneous velocity trajectories were down-sampled from 150 or 200 Hz down to 40 Hz based on the observation that the power spectrum of velocities mostly fell below 25 Hz. Each velocity sample can be formalized as:

$$\hat{v}_{t_0} = (\vec{v}(t_0 - \tau_{lead}),\ \ldots\ ,\ \vec{v}(t_0 + n\Delta t),\ \ldots\ ,\ v(t_0 + \tau_{lag}) \tag{3.1}$$

where $\vec{v}(t_0 + n\Delta t)$ is the instantaneous 3-dimensional velocity vector at time $t_0 + n\Delta t$ and $\Delta t = 25ms$. The full kinematics model relates these terms and the average position vector $\vec{p}$ to the conditional spike intensity of target unit i:

$$P\left(spike_i\left(t_0\right) \middle|\, \hat{v}_{t_0},\ \vec{p}\right) = \exp\left(\gamma + \vec{k} \cdot \hat{v}_{t_0} + \vec{c} \cdot \vec{p}\right) \tag{3.2}$$

where $\vec{k}$ is referred to as the preferred velocity trajectory of the neuron and $\vec{c}$ is the vector of coefficients for the average position terms. The preferred path or "pathlet" of the neuron is obtained by integrating $\vec{k}$ in time as if it were a vector of 3-dimensional velocities.

GLMs were implemented using the Python statsmodels package and trained using ridge regression (L1=0) with a light penalty of $\alpha = 0.00001$. Results of the full kinematics model were consistent for α values between 0.01 and 0.000001. All results of the kine-

matics+reachFN model were consistent for α values 0.00001 and 0.000001, and most results were consistent for α as large as 0.01.

Area under the receiver operating characteristic curve

To assess the predictive power of each encoding model, we computed receiver operating characteristic (ROC) curves that quantified the relationship between the probability of correctly predicting a spike (hit probability) versus the probability of incorrectly predicting a spike when it was absent (false positive probability). We used 90% of spike-kinematics sample pairs to train the model for each unit and used 10% as held-out test data. We predicted the conditional spike intensity from the encoding model given the 10% of held-out kinematics samples and compared predictions to the held-out spike samples. We then computed the area under the ROC curve (AUC), which measures the predictive value of the model and can vary from 0.5 (no predictive power) to 1.0 (perfect prediction). The train/test split was sampled randomly 100 times, resulting in 100 encoding models for each unit.

Functional Network computation

We created weighted, directed FNs by computing pairwise spike time statistics between recorded units. We binned the recorded spike trains into 10-ms bins, assigning a value of 1 if at least one spike occurred in that bin, and 0 otherwise. We then computed the confluent mutual information (conMI) between the binned spike trains.[76] ConMI tells us how much information we gain from the firing state of a source unit i at time t about the firing state of target unit j in the same time bin t and the consecutive bin, $t+1$:

$$conMI = \sum_{j(t)\in\{0,1\}} \sum_{i(t)\in\{0,1\}} p\left(j\left(t\right), i\left(\hat{t}\right)\right) \cdot \log_2 \left[\frac{p\left(j\left(t\right), i\left(\hat{t}\right)\right)}{p\left(i\left(t\right)\right) \cdot p\left(j\left(\hat{t}\right)\right)}\right], \tag{3.3}$$

76. Chambers *et al.,* "Ensemble stacking mitigates biases in inference of synaptic connectivity".

$$\text{where } i\left(\hat{t}\right) = \begin{cases} 1, & \text{if } i(t) = 1 \text{ OR } i(t+1) = 1 \\ 0, & \text{otherwise} \end{cases} \tag{3.4}$$

We computed two FNs during reaching segments from a separate 50% of reaching segments chosen at random. We paired the kinematics from reachSet1 with the FN computed during reachSet2, and vice versa, to eliminate the concern that kinematics and the FN might co-vary within a single sample. We also computed the FN during the remainder of the session in which the marmoset was behaving in an undirected, spontaneous manner either in the apparatus or the home enclosure (spontaneousFN). Each FN was represented as a square matrix of directed edge weights (the conMI) between nodes (units), with target units represented along the rows and source units along the columns.

Incorporating the functional network into the encoding model

Activity in the functional group was incorporated in the encoding model by computing the sum of input weights times spiking activity:

$$P\left(spike_i\left(t_0\right)\middle|\hat{v}_{t_0}, \vec{p}, F^0, F^1\right) = \exp\left(\gamma + \vec{k}\cdot\hat{v}_{t_0} + \vec{c}\cdot\vec{p} + \beta_0 F^0 + \beta_1 F^1\right) \tag{3.5}$$

$$F^0 = \sum_j w_{ji} s_j^0 \; and \; F^1 = \sum_j w_{ji} s_j^1 \tag{3.6}$$

where w_{ji} is the edge weight (the conMI value) from source unit j to target unit i and s_j^0 and s_j^1 are spike activity in the source unit at times t and $t-1$. Thus, F^0 and F^1 are the coincident and leading network features and β_0 and β_1 are the corresponding scaling terms fit in the GLM.

Statistical Tests

We used three statistical tests to evaluate significance of results. We used the sign test for

paired tests, including comparisons of unit performances in two encoding models and comparing FN permutations. We used the median test to compare distributions for reach-specific, non-specific and full FNs. We chose the median test because it is conservative and valid for distributions with different sample sizes and variance. We used the Pearson correlation to quantify preferred trajectory correlations and to evaluate the correlation between pairs of features (for example, full kinematics AUC vs average in-weight). For correlations, we considered $|r| < 0.2$ to be uncorrelated, $|r|$ in $[0.2, 0.5]$ to be weakly or moderately correlated, and $|r| > 0.5$ to be strongly correlated.

3.5 Supplementary Figures

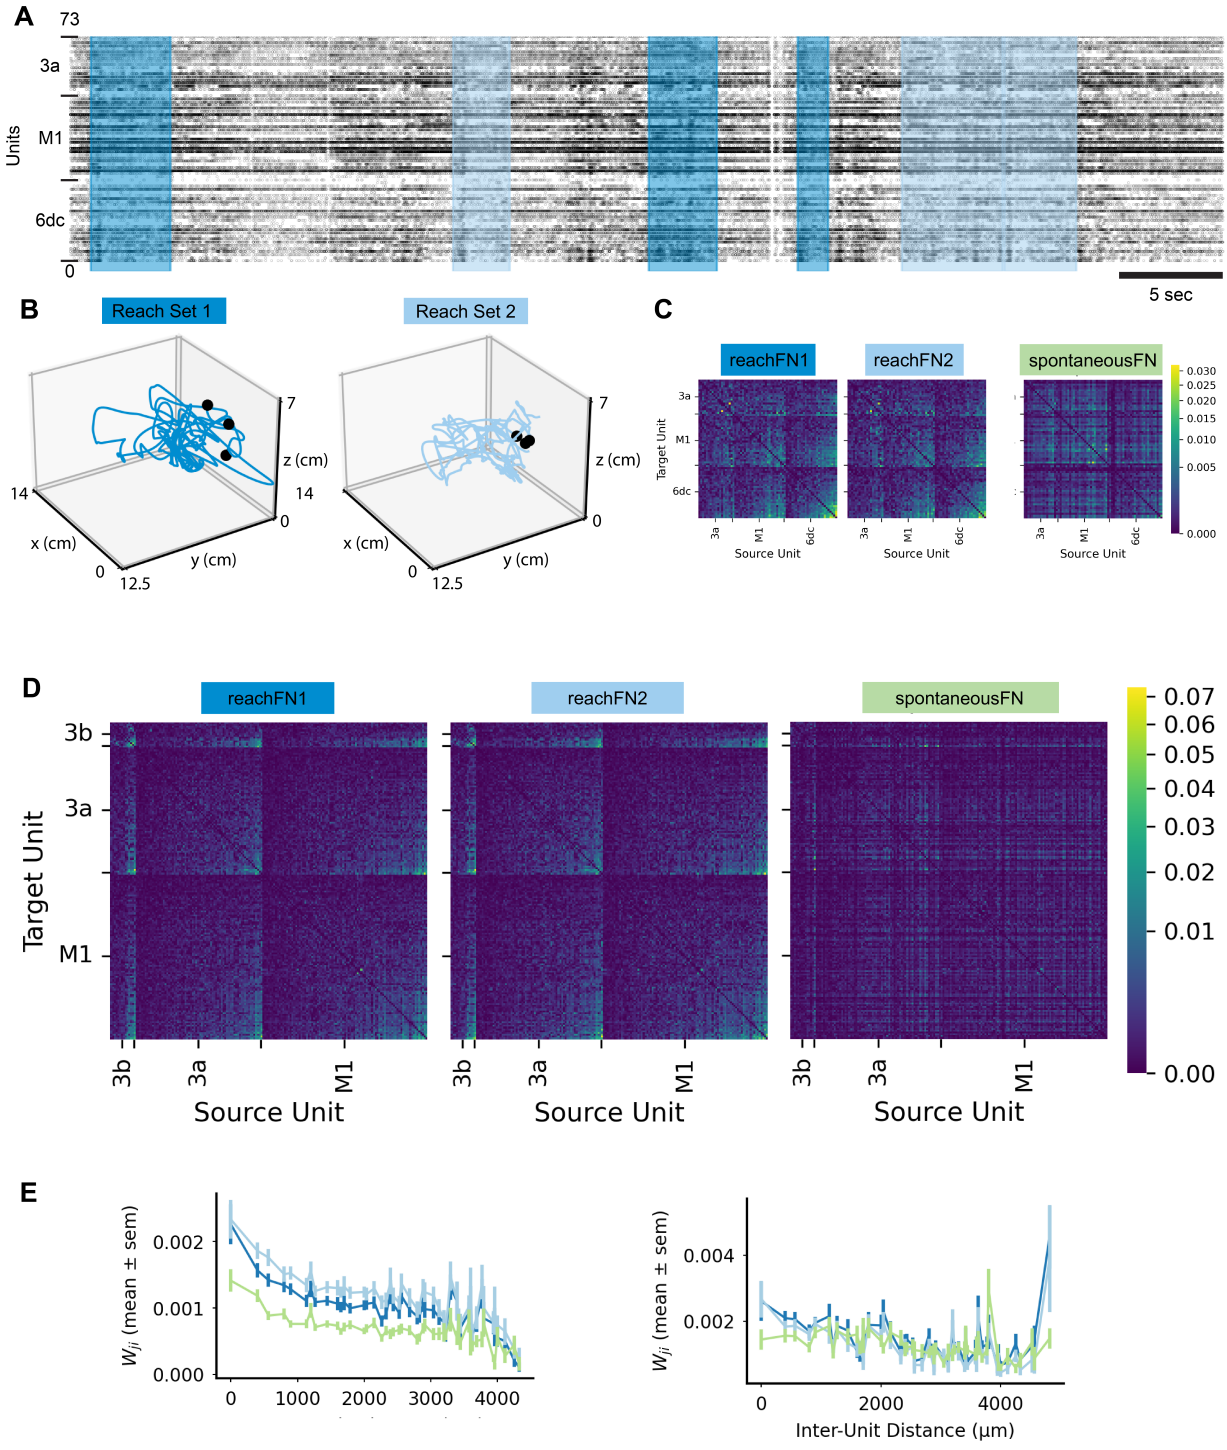

Figure 3.S1: MG data and full size FNs. (A-E) Data presentation and legends identical to Fig. 3.1. **(A)** Raster plot from MG recording, with preliminary estimates of cortical area boundaries used to group units. **(B)** Reaches associated with raster. **(C)** FNs computed for MG. **(D)** Full-size versions of FNs computed for TY. **(E)** Edge weights (mean ± sem) versus inter-electrode distance in the computed FNs for TY (*left*) and MG (*right*).

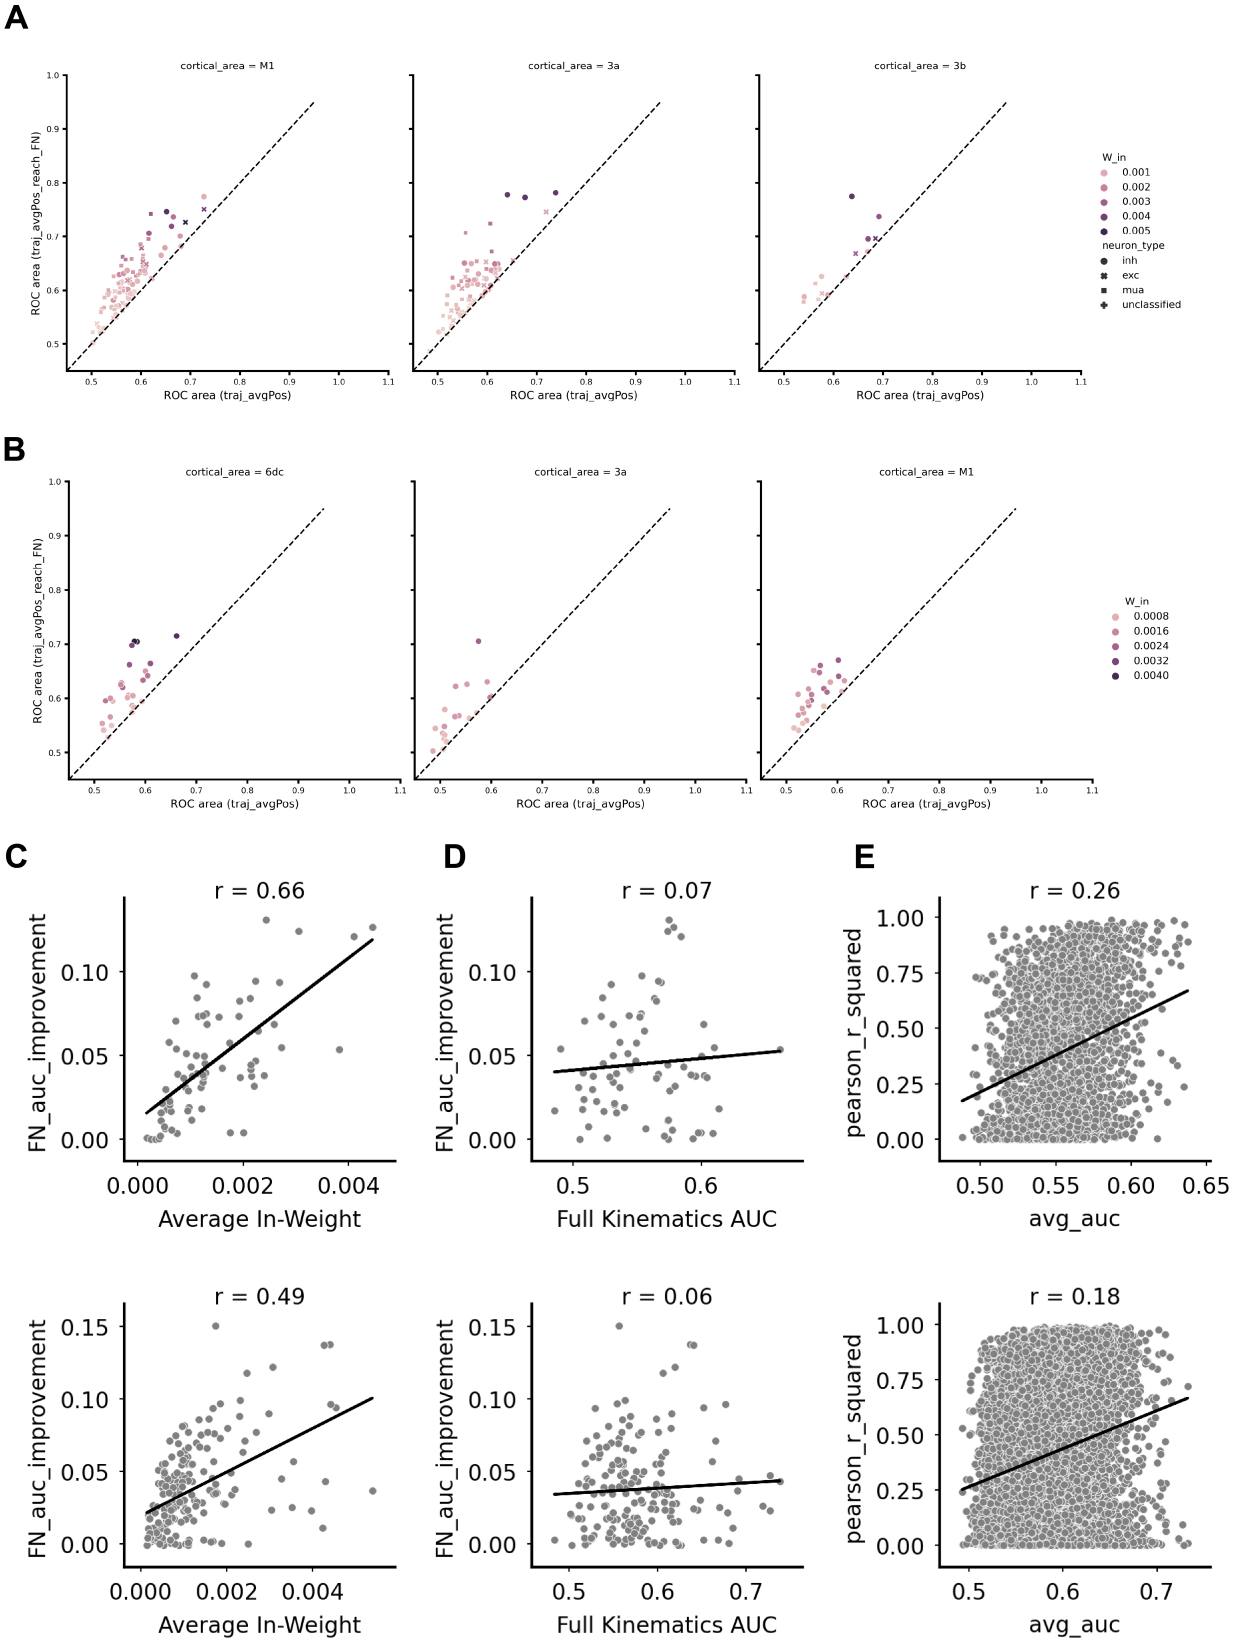
A
cortical_area = M1
cortical_area = 3a
cortical_area = 3b
ROC area (traj_avgPos_reach_FN)
ROC area (traj_avgPos)
W_in
0.001
0.002
0.003
0.004
0.005
neuron_type
inh
exc
mua
unclassified
B
cortical_area = 6dc
cortical_area = 3a
cortical_area = M1
ROC area (traj_avgPos_reach_FN)
ROC area (traj_avgPos)
W_in
0.0008
0.0016
0.0024
0.0032
0.0040
C
D
E
r = 0.66
r = 0.07
r = 0.26
FN_auc_improvement
Average In-Weight
Full Kinematics AUC
pearson_r_squared
avg_auc
r = 0.49
r = 0.06
r = 0.18

Figure 3.S2: Further results of FNs and the kinematics+reachFN model. (A-B) The kinematics+reachFN vs full kinematics model, split by cortical area for TY (A) and MG (B). **(C)** Scatterplots showing a strong positive correlation between improvement in AUC from the full kinematics model to the kinematics+reachFN model and average in-weight to the unit for TY (*top*) and MG (*bottom*). **(D)** Same for FN-driven AUC improvement and the original performance of the full kinematics model, showing no correlation. **(E)** Scatterplots of pairwise preferred trajectory correlation versus average AUC of the unit pair, showing a very weak positive correlation.

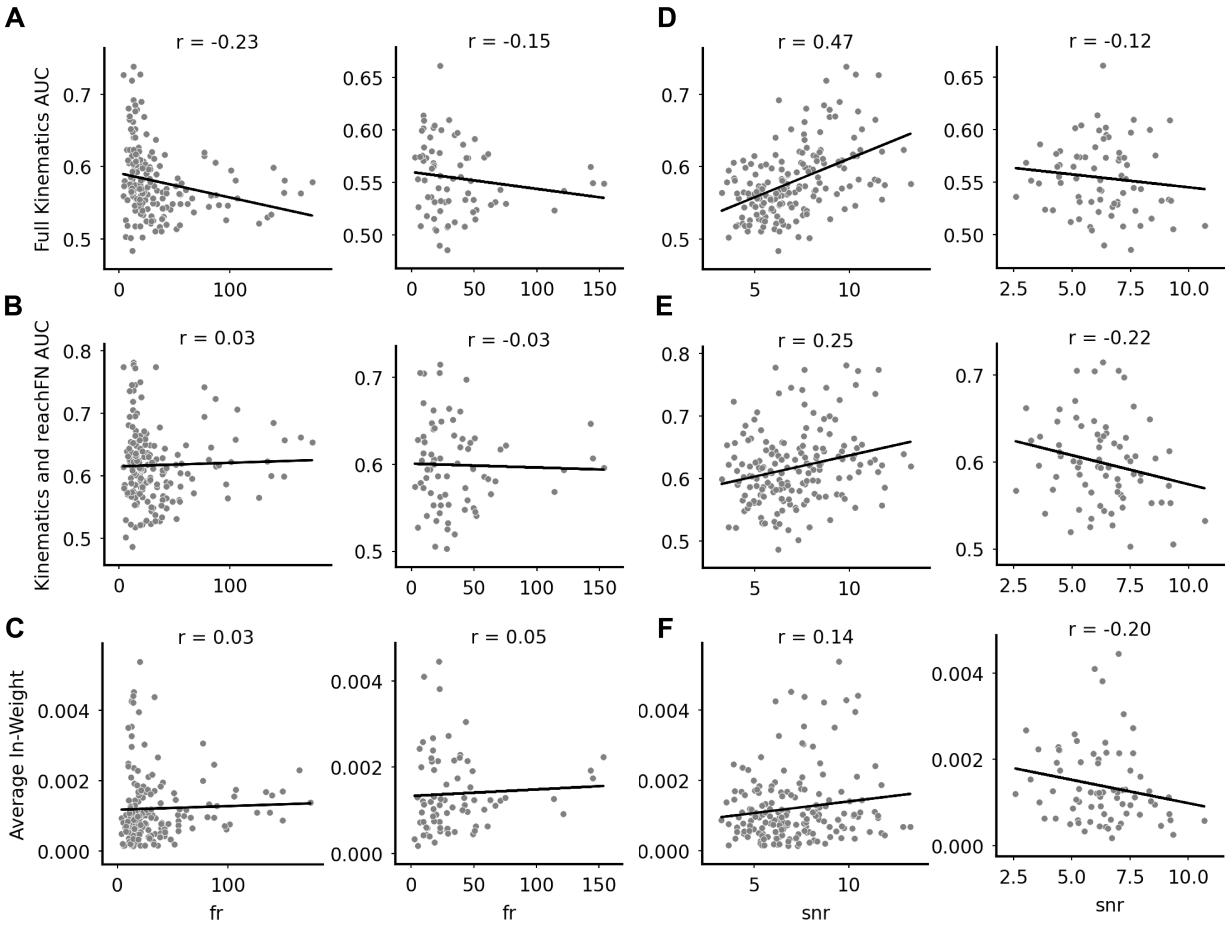

Figure 3.S3: Checks of potential confounds. (A) Scatterplots showing weak negative correlation between AUC of the full kinematics model and firing rate for TY (*left*) and MG (*right*). **(B)** No correlation for kinematics+reachFN model versus firing rate. **(C)** No correlation for average in-weight versus firing rate. **(D)** Positive correlation between full kinematics AUC and signal to noise ratio (SNR) for TY, no correlation for MG. **(E)** Weak positive correlation between kinematics+reachFN model and SNR for TY, weak negative correlation for MG. **(F)** Weak/no correlations between average in-weight and SNR.

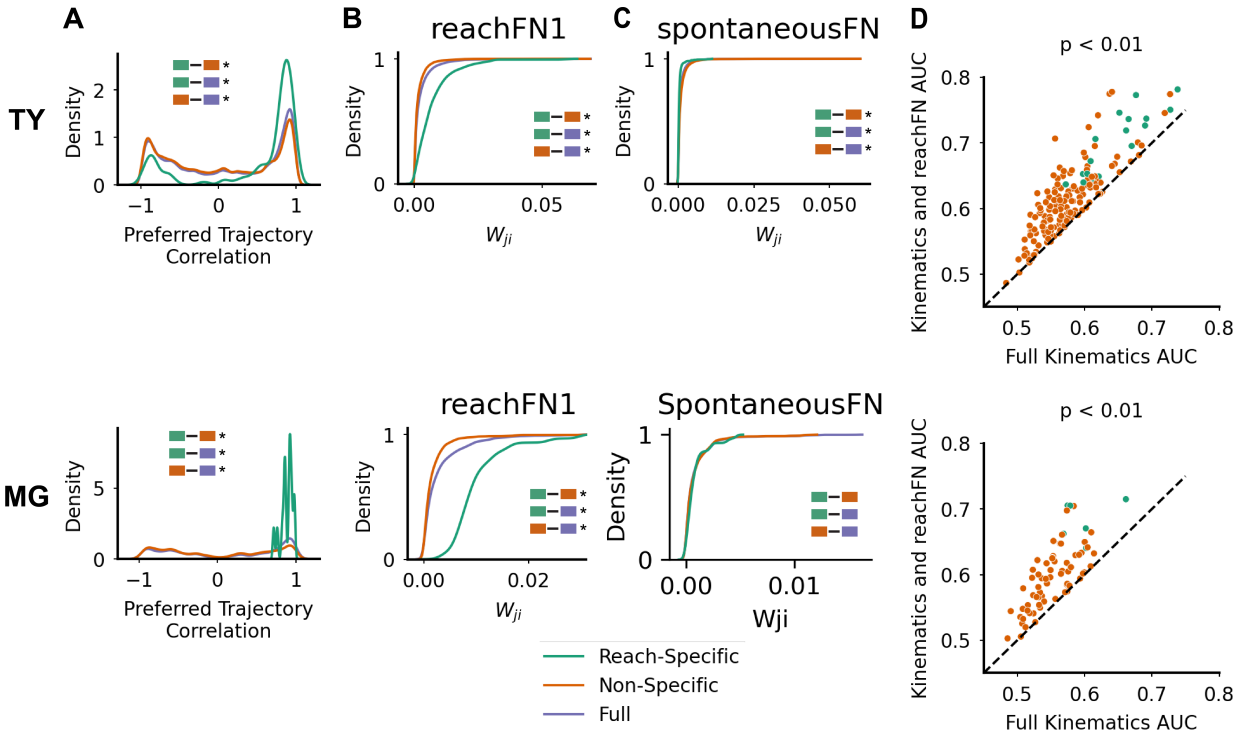

Figure 3.S4: Additional reach- versus non-specific functional group results. (A) Distributions of preferred trajectory correlations for pairs of units in the full and non-specific groups that each had kinematics AUC values greater than the lowest AUC value in the reach-specific group, compared to the reach-specific group. **(B-C)** Similar AUC-matched cumulative distributions of edge weight in reachFN1 and spontaneousFN. **(D)** Scatterplots from Fig. 3.3A, displayed with hue designating functional group membership.

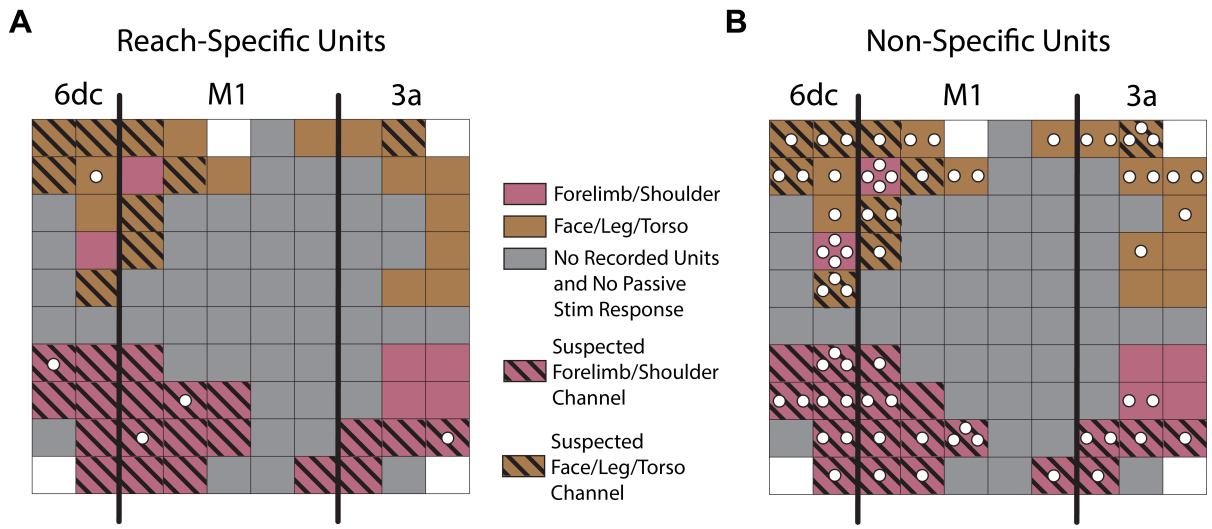

Figure 3.S5: Cortical location of reach- and non-specific units for MG. (A-B) Same as Fig. 3.6C-D for monkey MG. Channels with black stripes are hypothesized ICMS results based on estimated boundaries and passive stimulation results. The estimated channels are based on the location of somatic stimulation results and the estimated cortical region boundaries derived from those results. These maps are preliminary.

CHAPTER 4

GENERAL DISCUSSION

The underlying mechanisms of computation in sensorimotor cortex must be both flexible and robust to support the full repertoire of skilled motor behavior available to humans and non-human primates. To understand the richness of these mechanisms, it is important to study motor behavior in the most natural context possible. Furthermore, single-unit activity should be examined in relation to movement and in the context of population activity in an interacting network. In this work, I demonstrate that accurately capturing naturalistic motor behavior - specifically dynamic forelimb movements during foraging and prey-capture - is both critical and feasible. I leverage natural forelimb movements and pairwise precise spike time structure (measured using confluent mutual information and represented as a functional network) in marmoset sensorimotor cortex to develop a network encoding model that links single-unit spiking activity with kinematic tuning properties and functional network interactions. I use this model to investigate the computational mechanisms that generate varied and naturalistic motor behavior.

I demonstrate in Chapter 2 that computer vision tools (specifically DLC+Anipose)[1,2] track 3D pose with accuracy comparable to marker-based tracking. This is crucial for studying motor behavior of the common marmoset because they are less amenable to long, repetitive training and recording sessions in a constrained setting but can be engaged by other methods; in fact, prior work in our lab has demonstrated that marmosets will complete hundreds of reaches if allowed to enter the behavioral apparatus voluntarily to participate in a foraging or prey-capture task.[3] I characterize the effect of parameter choices on DLC+Anipose accuracy and identify a range of parameter value choices to minimize error and maximize the

1. Karashchuk *et al.,* "Anipose: A toolkit for robust markerless 3D pose estimation".

2. Mathis *et al.,* "DeepLabCut: markerless pose estimation of user-defined body parts with deep learning".

3. Walker *et al.,* "A platform for semiautomated voluntary training of common marmosets for behavioral neuroscience".

number of well-tracked frames. I use the optimal parameter set in Chapter 3 to accurately estimate pose during prey-capture, thereby providing the complex kinematics necessary for our trajectory encoding model.

In Chapter 3, I develop an encoding model to predict single-unit spiking activity in 10ms windows from a temporally-extensive velocity trajectory of the wrist, average position during the trajectory, and network features derived from the spiking activity of source units passed through the functional network. I show that the trajectory tuning model extends to more naturalistic movements and explains neural activity better than a simpler model. Additionally, the trajectory-based encoding model without any direct mathematical dependence on the FN links single-unit activity to kinematics most accurately for units that are strongly interconnected within the network. I show that the structure of the functional group of source units is informative of spike activity in the target unit, and that the information is focused in the 15-40% of strongest connections. Finally, I identify a reach-specific functional group for which the functional connectivity is reorganized during prey-capture. This reach-specific functional group is strongly interconnected and comprises units tightly linked to kinematics with strong, positively correlated preferred trajectories. This functional group exhibits distinct interaction patterns during prey capture and spontaneous behavior, while the non-specific group interacts similarly across the spectrum of spontaneous and goal-directed behavior. The reorganization or static nature of these sub-populations in primate sensorimotor cortex may have implications for understanding the generation of smooth neural trajectories during movement production.

4.1 Analysis of Neuron Properties in the Network

This work demonstrates the utility of a comprehensive encoding model incorporating kinematic and network features – I leverage dynamic, variable forelimb reaching movements and pairwise functional interactions on a fine timescale to identity important structural features

of the neural circuit and the role these features play in moment-to-moment spiking activity. However, I only scratch the surface of the information available in the functional network. Every unit in the FN can be characterized according to many properties - these include cortical area, degree of kinematic tuning, laminar depth (if recorded with a different array technology, discussed below), and neuron type (narrow-spiking, wide-spiking or multi-unit activity).

I address the cortical area of units tangentially in Chapter 3. Units strongly tuned to kinematics are distributed across sensorimotor areas, and for one monkey the reach-specific functional group is preferentially located in M1 and in 3a near the boundary with M1. Further analysis may provide insight into the structure of functional interactions within and across cortical areas. I plan to train and compare prediction performance of encoding models using network features computed from exclusively intra- or inter-area connections, then probe the importance of network topology as in Fig. 3.4. Additionally, I will examine relationships between cortical area, functional edge weights, and pairwise preferred trajectory correlations to potentially add context to the results presented in Chapter 3. Finally, work by Dann et al.[4] shows that modularity in the functional network links large groups of interconnected units in a single cortical area to smaller groups in other areas, suggesting a mechanism for information flow between areas. It may be that members of the reach-specific functional group described in this work, which are tightly coupled by preferred trajectory correlations and strong edge weights, participate in reach-specific modules to facilitate inter-area communication.

I show in Chapter 3 that the degree to which a unit is tuned to kinematics can be partially explained by the strength of functional inputs from the network. Furthermore, units in the reach-specific functional group are more strongly tuned to kinematics than non-specific units. This suggests strongly tuned units are free to reorganize for generation of specific movements, while weakly tuned and untuned units, which tend to interact similarly across the range of

4. Dann *et al.*, "Uniting functional network topology and oscillations in the fronto-parietal single unit network of behaving primates".

behavior studied here, may contribute to maintenance of the neural state or to smooth transitions between states. The finding that strongly tuned units tend to be more strongly connected within the circuit disagrees with results in visual cortex showing that untuned units occupy central positions in the topology and act as crucial hubs for propagating information through the network.[5] It is plausible that untuned neurons play distinct roles in sensorimotor and visual cortex. These areas perform different computations, and areas or behaviors with different computational constraints exhibit distinct population dynamics.[6,7] On the other hand, a systematic examination of tuned and untuned sub-populations within the FN and their impact on prediction of target unit activity may reveal similarities between visual and sensorimotor cortex that are not yet apparent.

Similarly, narrow-spiking (NS) and wide-spiking (WS) units may play different roles in the generation of complex kinematics and contribute differently within the functional network. Using neuron classifications estimated by the waveform analysis described in Snyder et al.,[8] I have found in preliminary analyses of TY data (Chapter 3) that NS (putative inhibitory) neurons seem to be over-represented in the set of units that were best predicted by the full kinematics model and the kinematics+reachFN model (Fig. 4.1A), and in the reach-specific functional group (Fig. 4.1B). However, the distribution of AUC values was not significantly different between the NS and WS (putative excitatory) populations in either the full kinematics or kinematics+reachFN model (Fig. 4.1C). I have yet to investigate whether NS neurons are significantly over-represented in the reach-specific functional group or whether they exhibit differences in FN structure. I have also not ruled out the possibility that these results are confounded by significantly higher waveform signal-to-noise ratios for NS neurons (Fig. 4.1D). To test for structure in the tuning properties of each class in a more

5. Levy *et al.,* "Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding".

6. Lara *et al.,* "Different population dynamics in the supplementary motor area and motor cortex during reaching".

7. Suresh *et al.,* "Neural population dynamics in motor cortex are different for reach and grasp".

8. Snyder *et al.,* "Dynamics of excitatory and inhibitory networks are differentially altered by selective attention".

nuanced fashion, we could examine the preferred trajectory correlations between NS-NS, NS-WS, and WS-WS pairs. I hypothesize that NS-WS pairs will contain a greater degree of strong anti-correlations than expected by chance, particularly for pairs on nearby channels. This would be consistent with putative inhibitory-excitatory interactions, but may not bear out because we have no way of measuring direct synaptic connectivity and because our recordings contain an extremely sparse sample of inhibitory-excitatory interactions. Then we could investigate the role of each connection type in the functional network by permuting

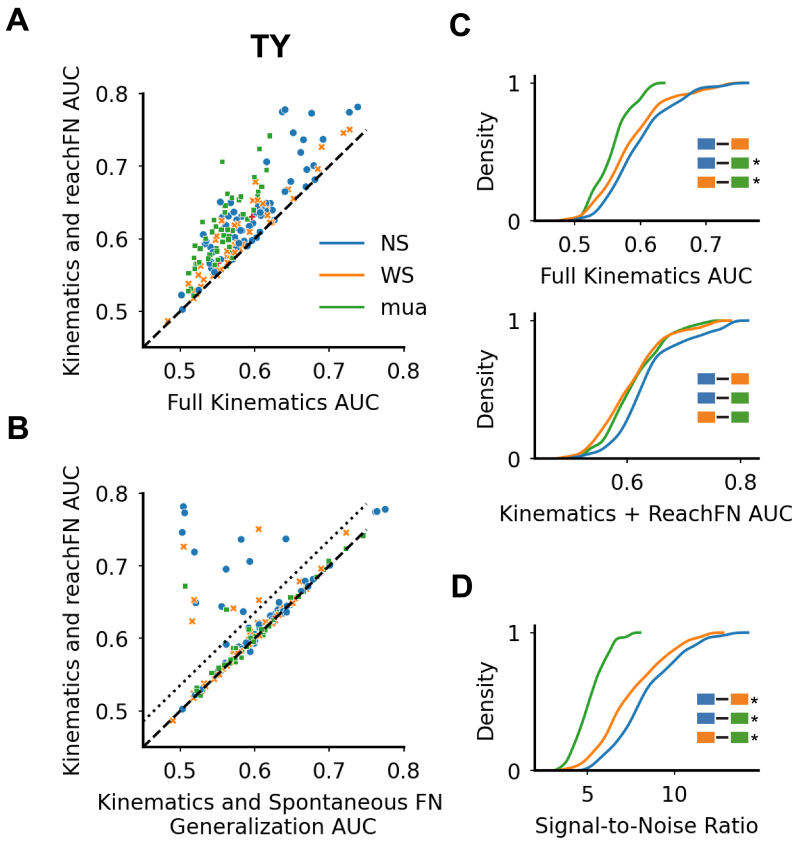

Figure 4.1: Looking at narrow-spiking and wide-spiking neurons. **(A)** Recreation of Fig. 3.3A with unit hue representing neuron classification. **(B)** Recreation of Fig 3.5A. **(C)** Cumulative distributions of full kinematics (*top*) and kinematics and reachFN (*bottom*) AUC values, split by neuron classification. Significant differences indicated by * ($p < 0.05$). **(D)** Cumulative distribution of waveform signal-to-noise ratios.

FN topology (as in Fig. 3.4) according to NS-NS, NS-WS, and WS-WS connectivity rather than by the strongest edges. In this analysis, it is important to permute within functional

groups of matched edge weight distributions to control for the demonstrated relationship between edge weight and predictive power of both the full kinematics model and the kinematics+reachFN model.

The methods and analyses presented in Chapter 3 are perfectly suited for examining the way in which functional interactions are structured by cortical depth. Given that excitatory pathways from L2/3 to L5 contain the strongest and most stable projections between lamina in M1,[9] I hypothesize that L5 neurons would be more reliant on inputs from L2/3 than other network subsets to explain their activity. We also know that sensory-related cortical and thalamic inputs project primarily to L2/3 of M1, while L5 of motor cortex contains the vast majority of pyramidal tract neurons projecting to the spinal cord.[10,11] I hypothesize that L2/3 units would be predicted best by models using trajectories that extend further back in time, thereby encoding kinematics leading the neural response. In contrast, I expect preferred trajectories of L5 neurons (from which all data in Chapter 3 was recorded) to extend further into the future to encode upcoming motor output (as suggested by Fig. 3.2A). However, this could be difficult to determine conclusively; I originally hypothesized that I would observe similar differences in the preferred trajectories of motor versus sensory units in Chapter 3 and was unable to identify any conclusive differences in the lead/lag set of preferred trajectories across cortical areas. Investigations into laminar structure in M1 could be conducted in the future by implanting an N-Form array (Modular Bionics) instead of the Utah array used in Chapter 3. The N-Form array allows for custom placement of multiple electrode sites along the depth of each shank, and we have an array ready for implantation that would record from approximately three depths in L2/3 and two depths in L5 on each of 16 shanks.

9. Hooks *et al.,* "Laminar Analysis of Excitatory Local Circuits in Vibrissal Motor and Sensory Cortical Areas".

10. Ibid.

11. Shinoda & Kakei, "Distribution of terminals of thalamocortical fibers originating from the ventrolateral nucleus of the cat thalamus".

4.2 Investigate temporal FNs

In Chapter 3, we computed two functional networks from two sets of reaches across the entire session (reachFNs) and clearly demonstrated that they captured interactions and structure that were not present in the FN computed from spontaneous behavior (spontaneousFN). However, Sundiang et al.[12] showed that temporal FNs capture cortical dynamics in finer detail and can be used to generate insights into the timescale at which motor cortex processes the requirements of individual trials and reorganizes its interactions to produce different movements. In that study, the authors computed the temporal FN by computing confluent mutual information from spiking activity in 200-ms windows centered on each consecutive 10-ms bin. I will modify this method to compute an FN associated with each trajectory sample, such that for the [-100, +300] ms trajectories studied in depth in Chapter 3 the FN window would be 400-ms in duration and sampled at 30-ms intervals. It would be interesting to investigate systematic changes in the structure of the temporal FN and how such changes relate to details of the kinematics, particularly for the reach-specific functional group.

To do so, we could start by training a kinematics+temporalFN model that computes network features using the sample-specific FN in place of reachFNs. I expect this model would predict unit activity more accurately because of the fine-grained information available in the temporal FN. Next, we could perform a similar generalization experiment similar to analysis that revealed the reach-specific functional group (see Fig. 3.5), training on network features from the spontaneousFN or reachFNs and testing on the kinematics+temporalFN model. I expect to find that the reach-specific functional group is conserved in this analysis. I also expect to find an additional functional group that is particularly dependent on the information captured only by the temporalFN. We could then focus on this sample-specific functional group and the reach-specific group to investigate reorganization of the FN across time. It would be particularly intriguing to find that the sample-specific and reach-specific

12. Sundiang *et al.,* "Dynamic structure of motor cortical neuron coactivity carries behaviorally relevant information".

functional groups reorganize in a way that is linked to details of the movement – perhaps to the trajectory sample's tortuosity or mean speed. Structural details may be revealed by first order analysis of something like in-weight, or may require more nuanced methods such as the graph alignment score or UMAP-based methods described by Sundiang et al.[13]

4.3 Leveraging the details of spontaneous behavior

As demonstrated in Figs. 3.5 and 3.6, functional interactions measured during spontaneous behavior are sufficient to predict activity during reaching for many, but not all, units in the population. In Chapter 3, we considered all moments that the marmoset was not directly engaged in the prey-capture task to belong to the broad class of spontaneous behavior. However, such behavior is highly variable and includes but is not limited to ambulating, leaping, climbing, hanging, grooming of self, grooming of partner in the enclosure, and awake restfulness. Although we did not attempt to classify spontaneous behavior into these categories or to quantify movement at all in Chapter 3, there is a wealth of information in these behaviors that could add nuance to our findings. Behaviors that require coordinated forelimb movements such as climbing and partner grooming may exhibit functional network structure that is more similar to the reachFNs. On the other hand, activities such as hanging, self-grooming and resting may exhibit even larger differences in the FN than the spontaneousFN presented here. An investigation of functional groups underlying each of these varied behaviors would be valuable. Additionally, understanding the structure of interactions during each behavior would provide context to the reach-specific functional group and non-specific units. It may be that a different functional group is engaged preferentially for each class of behavior; alternatively, as we argue in Chapter 3, structured interactions within the non-specific group might be conserved across behaviors while the reach-specific group reorganizes based on the demands of each behavior.

13. Sundiang *et al.,* "Dynamic structure of motor cortical neuron coactivity carries behaviorally relevant information".

4.4 Single-unit activity and functional interactions within the framework of population dynamics

It is difficult to place single-unit tuning properties in the context of low-dimensional population dynamics. However, our finding that a reach-specific functional group located primarily in forelimb M1 and 3a reorganizes to produce dynamic forelimb movements during prey capture suggests intriguing possibilities for contextualizing single-unit properties within the framework of population dynamics. In Chapter 3, we speculate that the non-specific units that interacted consistently from spontaneous behavior to reaching could contribute to the generation of low-dimensional and rotational dynamics that evolve in a predictable fashion with low tangling and govern activity in the output-null subspace.[14,15,16,17] The reach-specific functional group, on the other hand, may govern dimensions of the neural trajectory that can be directly read out by downstream targets and thereby contribute preferentially to the output-potent subspace. Testing this directly would require the experiment conducted in Chapter 3 to be repeated with the addition of electromyographic recordings from marmoset forelimb muscles. Alternatively, our lab and others have frequently recorded muscle activity in behaving macaque monkeys. To my knowledge, however, no studies have reported unrestrained and naturalistic reaching behavior, richly varied spontaneous behavior, and simultaneous neural recordings together in macaques. Fortunately, it is becoming feasible to record rich, unconstrained behavior in macaques,[18] which in turn provides the opportunity to apply our methods to macaque experiments with the addition of EMG recordings.

Alternatively, emerging work demonstrates that neural dynamics occur on a curved manifold such that frequency of local linear rotational dynamics is conserved while the orientation

14. Churchland *et al.,* "Neural population dynamics during reaching".

15. Elsayed *et al.,* "Reorganization between preparatory and movement population responses in motor cortex".

16. Kaufman *et al.,* "Cortical activity in the null space: Permitting preparation without movement".

17. Russo *et al.,* "Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response".

18. Bala *et al.,* "Automated markerless pose estimation in freely moving macaques with OpenMonkeyStudio".

of the rotational plane depends on the location of the fixed point on that manifold.[19] In this study, the location of the fixed point was linked to reach kinematics on single trials. It is possible that the reach-specific functional group determines the location of the fixed point in neural space, around which the non-specific units produce rotational neural dynamics to generate movement. Since the reach-specific group was significantly more tuned to velocity and average position in Chapter 3, an experiment showing this to be true would be a compelling reconciliation between classical tuning and dynamics.

There remain substantial barriers to understanding single-unit activity in the context of broader population dynamics. There are a limited number of labs in the world that can record single unit electrophysiology in non-human primates along with both spontaneous behavior and unrestrained, goal-directed forelimb movements – either with marmosets or macaques. Even as this becomes a possibility in our lab and others thanks to wireless neural recording devices and computer vision pose estimation tools, wrangling the immense datasets resulting from such experiments is daunting. Furthermore, many of the computational tools used to study population dynamics depend on trial-averaging, which can be difficult or impossible in the study of natural behavior. However, the work I present here - along with work by Sundiang et al.[20] – demonstrates that pairwise functional interactions computed on a fine timescale provide a window into the neural dynamics in primate sensorimotor cortex underlying natural forelimb movements. By combining these methods with dynamical systems methods that can be applied to single trials such as the LDR-model described by Sabatini and Kaufman,[21] it may be possible to contextualize single-unit activity and tuning properties within a dynamical systems framework.

19. Sabatini & Kaufman, "Reach-dependent reorientation of rotational dynamics in motor cortex".

20. Sundiang *et al.,* "Dynamic structure of motor cortical neuron coactivity carries behaviorally relevant information".

21. Sabatini & Kaufman, "Reach-dependent reorientation of rotational dynamics in motor cortex".

BIBLIOGRAPHY

1. Aflalo, T. N. & Graziano, M. S. A. Partial tuning of motor cortex neurons to final posture in a free-moving paradigm. *Proceedings of the National Academy of Sciences* **103,** 2909–2914. ISSN: 0027-8424 (Feb. 2006).

2. Aflalo, T. N. & Graziano, M. S. A. Relationship between unconstrained arm movements and single-neuron firing in the macaque motor cortex. *The Journal of Neuroscience* **27,** 2760–80. ISSN: 1529-2401 (Mar. 2007).

3. Bala, P. C. *et al.* Automated markerless pose estimation in freely moving macaques with OpenMonkeyStudio. *Nature Communications 2020 11:1* **11,** 1–12. ISSN: 2041-1723 (Sept. 2020).

4. Bolaños, L. A. *et al.* A three-dimensional virtual mouse generates synthetic training data for behavioral analysis. *Nature Methods* **18,** 378–381. ISSN: 1548-7091 (Apr. 2021).

5. Brainerd, E. L. *et al.* X-ray reconstruction of moving morphology (XROMM): precision, accuracy and applications in comparative biomechanics research. *Journal of Experimental Zoology Part A: Ecological Genetics and Physiology* **313,** 262–279. ISSN: 19325223 (June 2010).

6. Buccino, A. P. *et al.* Spikeinterface, a unified framework for spike sorting. *eLife* **9,** 1–24. ISSN: 2050084X (Oct. 2020).

7. Burish, M. J., Stepniewska, I. & Kaas, J. H. Microstimulation and architectonics of frontoparietal cortex in common marmosets (Callithrix jacchus). *The Journal of Comparative Neurology* **507,** 1151–1168. ISSN: 00219967 (Mar. 2008).

8. Burton, H. & Fabri, M. Ipsilateral intracortical connections of physiologically defined cutaneous representations in areas 3b and 1 of macaque monkeys: projections in the vicinity of the central sulcus. *The Journal of Comparative Neurology* **355,** 508–538. ISSN: 0021-9967 (1995).

9. Caminiti, R, Johnson, P. B. & Urbano, A. Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. *The Journal of Neuroscience* **10,** 2039–58. ISSN: 0270-6474 (July 1990).

10. Chambers, B., Levy, M., Dechery, J. B. & MacLean, J. N. Ensemble stacking mitigates biases in inference of synaptic connectivity. *Network Neuroscience* **2,** 60–85. ISSN: 2472-1751 (Mar. 2018).

11. Chambers, B. & MacLean, J. N. Multineuronal activity patterns identify selective synaptic connections under realistic experimental constraints. *Journal of Neurophysiology* **114,** 1837–1849. ISSN: 0022-3077 (Sept. 2015).

12. Chartier, J., Anumanchipalli, G. K., Johnson, K. & Chang, E. F. Encoding of Articulatory Kinematic Trajectories in Human Speech Sensorimotor Cortex. *Neuron* **98,** 1042–1054. ISSN: 0896-6273 (June 2018).

13. Chaure, F. J., Rey, H. G. & Quian Quiroga, R. A novel and fully automatic spike-sorting implementation with variable number of features. *Journal of Neurophysiology* **120,** 1859–1871. ISSN: 15221598 (Oct. 2018).

14. Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron* **68,** 387–400. ISSN: 0896-6273 (Nov. 2010).

15. Churchland, M. M. & Shenoy, K. V. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *Journal of Neurophysiology* **97,** 4235–4257. ISSN: 00223077 (June 2007).

16. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487** (2012).

17. Cohen, J. A power primer. *Psychological Bulletin* **112,** 155–159. ISSN: 1939-1455 (1992).

18. Dann, B., Michaels, J. A., Schaffelhofer, S. & Scherberger, H. Uniting functional network topology and oscillations in the fronto-parietal single unit network of behaving primates. *eLife* **5.** ISSN: 2050084X (Aug. 2016).

19. Dechery, J. B. & MacLean, J. N. Functional triplet motifs underlie accurate predictions of single-trial responses in populations of tuned and untuned V1 neurons. *PLOS Computational Biology* **14** (ed Beck, J.) e1006153. ISSN: 1553-7358 (May 2018).

20. Dunn, T. W. *et al.* Geometric deep learning enables 3D kinematic profiling across species and environments. *Nature Methods* **18,** 564–573. ISSN: 1548-7105 (Apr. 2021).

21. Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications* **7,** 1–15. ISSN: 2041-1723 (Oct. 2016).

22. Evarts, E. V. Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of Neurophysiology* **31,** 14–27. ISSN: 00223077 (1968).

23. Fromm, C., Wise, S. P. & Evarts, E. V. Sensory response properties of pyramidal tract neurons in the precentral motor cortex and postcentral gyrus of the rhesus monkey. *Experimental Brain Research* **54,** 177–185. ISSN: 00144819 (Feb. 1984).

24. Fromm, C. & Evarts, E. V. Pyramidal tract neurons in somatosensory cortex: central and peripheral inputs during voluntary movement. *Brain Research* **238,** 186–191. ISSN: 0006-8993 (Apr. 1982).

25. Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A. & Miller, L. E. Long-term stability of cortical population dynamics underlying consistent behavior. *Nature Neuroscience* **23,** 260–270. ISSN: 1546-1726 (Jan. 2020).

26. Gao, P. & Ganguli, S. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology* **32,** 148–155. ISSN: 0959-4388 (June 2015).

27. Gardner, E. P. & Costanzo, R. M. Properties of kinesthetic neurons in somatosensory cortex of awake monkeys. *Brain Research* **214,** 301–319. ISSN: 0006-8993 (June 1981).

28. Georgopoulos, A. P., Ashe, J., Smyrnis, N. & Taira, M. The Motor Cortex and the Coding of Force. *Science* **256,** 1692–1695. ISSN: 0036-8075 (Sept. 1992).

29. Georgopoulos, A. P., Kalaska, J. F., Caminiti, R & Massey, J. T. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of Neuroscience* **2,** 1527–37. ISSN: 0270-6474 (Nov. 1982).

30. Georgopoulos, A. P., Kettner, R. E. & Schwartz, A. B. Primate motor cortex and free arm movements to visual targets in three- dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience* **8,** 2928–2937. ISSN: 0270-6474 (Aug. 1988).

31. Georgopoulos, A. P., Schwartz, A. B. & Kettner, R. E. Neuronal Population Coding of Movement Direction. *Science* **233,** 1416–1419. ISSN: 00368075 (1986).

32. Goodman, J. M. *et al.* Postural Representations of the Hand in the Primate Sensorimotor Cortex. *Neuron* **104,** 1000–1009. ISSN: 0896-6273 (Dec. 2019).

33. Graving, J. M. *et al.* DeepPoseKit, a software toolkit for fast and robust animal pose estimation using deep learning. *eLife* **8,** e47994. ISSN: 2050084X (2019).

34. Guo, J. Z. *et al.* Cortex commands the performance of skilled movement. *eLife* **4.** ISSN: 2050084X (Dec. 2015).

35. Hatsopoulos, N. G., Ojakangas, C. L., Paninski, L & Donoghue, J. P. Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proceedings of the National Academy of Sciences of the United States of America* **95,** 15706–11. ISSN: 0027-8424 (Dec. 1998).

36. Hatsopoulos, N. G. & Amit, Y. Synthesizing complex movement fragment representations from motor cortical ensembles. *Journal of Physiology-Paris* **106,** 112–119. ISSN: 0928-4257 (May 2012).

37. Hatsopoulos, N. G., Xu, Q. & Amit, Y. Encoding of Movement Fragments in the Motor Cortex. *Journal of Neuroscience* **27,** 5105–5114. ISSN: 0270-6474 (2007).

38. He, Y., Yan, R., Fragkiadaki, K. & Yu, S. I. *Epipolar Transformers* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (May 2020), 7779–7788.

39. Hocherman, S. & Wise, S. Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys. *Experimental Brain Research* **83,** 285–302. ISSN: 0014-4819 (Jan. 1991).

40. Hooks, B. M. *et al.* Laminar Analysis of Excitatory Local Circuits in Vibrissal Motor and Sensory Cortical Areas. *PLoS Biology* **9** (ed Petersen, C. C.) e1000572. ISSN: 1545-7885 (Jan. 2011).

41. Huerta, M. F. & Pons, T. P. Primary motor cortex receives input from area 3a in macaques. *Brain Research* **537,** 367–371. ISSN: 0006-8993 (Dec. 1990).

42. Huffman, K. J. & Krubitzer, L. Area 3a: Topographic Organization and Cortical Connections in Marmoset Monkeys. *Cerebral Cortex* **11,** 849–867. ISSN: 1047-3211 (Sept. 2001).

43. Hyvärinen, J. & Poranen, A. Receptive field integration and submodality convergence in the hand area of the post-central gyrus of the alert monkey. *The Journal of Physiology* **283,** 539. ISSN: 14697793 (Oct. 1978).

44. Iskakov, K., Burkov, E., Lempitsky, V. & Malkov, Y. *Learnable Triangulation of Human Pose* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (May 2019), 7718–7727. ISBN: 9781728148038.

45. Iwamura, Y., Tanaka, M., Sakamoto, M. & Hikosaka, O. Rostrocaudal gradients in the neuronal receptive field complexity in the finger region of the alert monkey's postcentral gyrus. *Experimental Brain Research* **92,** 360–368. ISSN: 0014-4819 (Jan. 1993).

46. Jackson, A., Mavoori, J. & Fetz, E. E. Correlations Between the Same Motor Cortex Cells and Arm Muscles During a Trained Task, Free Behavior, and Natural Sleep in the Macaque Monkey. *Journal of Neurophysiology* **97,** 360–374. ISSN: 0022-3077 (2007).

47. Jones, E. G., Coulter, J. D. & Hendry, S. H. Intracortical connectivity of architectonic fields in the somatic sensory, motor and parietal cortex of monkeys. *Journal of Comparative Neurology* **181,** 291–347. ISSN: 1096-9861 (Sept. 1978).

48. Kao, J. C. *et al.* Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nature Communications* **6,** 7759. ISSN: 2041-1723 (Nov. 2015).

49. Karashchuk, P. *et al.* Anipose: A toolkit for robust markerless 3D pose estimation. *Cell Reports* **36,** 109730. ISSN: 2211-1247 (Sept. 2021).

50. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: Permitting preparation without movement. *Nature Neuroscience* **17,** 440–448. ISSN: 10976256 (2014).

51. Kawai, R. *et al.* Motor Cortex Is Required for Learning but Not for Executing a Motor Skill. *Neuron* **86,** 800–812. ISSN: 0896-6273 (May 2015).

52. Knorlein, B. J., Baier, D. B., Gatesy, S. M., Laurence-Chasen, J. D. & Brainerd, E. L. Validation of XMALab software for Marker-based XROMM. *Journal of Experimental Biology* **219,** 3701–3711. ISSN: 00220949 (Dec. 2016).

53. Kotekal, S. & MacLean, J. N. Recurrent interactions can explain the variance in single trial responses. *PLOS Computational Biology* **16,** e1007591. ISSN: 1553-7358 (2020).

54. Krubitzer, L. A. & Kaas, J. H. The organization and connections of somatosensory cortex in marmosets. *The Journal of Neuroscience* **10,** 952. ISSN: 02706474 (Mar. 1990).

55. Lara, A. H., Cunningham, J. P. & Churchland, M. M. Different population dynamics in the supplementary motor area and motor cortex during reaching. *Nature Communications* **9,** 2754. ISSN: 2041-1723 (Dec. 2018).

56. Laurence-Chasen, J. D., Manafzadeh, A. R., Hatsopoulos, N. G., Ross, C. F. & Arce-Mcshane, F. I. Integrating XMALab and DeepLabCut for high-throughput XROMM. *Journal of Experimental Biology* **223.** ISSN: 14779145 (Sept. 2020).

57. Levy, M., Guo, J. K. & MacLean, J. N. A sparse set of spikes corresponding to reliable correlations is highly informative of visual stimulus on single trials. *bioRxiv,* 2022.01.24.477564 (Jan. 2022).

58. Levy, M., Sporns, O. & MacLean, J. N. Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding. *Cell Reports* **31,** 107483. ISSN: 2211-1247 (Apr. 2020).

59. Malonis, P. J. *et al.* M1 dynamics share similar inputs for initiating and correcting movement. *bioRxiv,* 2021.10.18.464704 (Oct. 2021).

60. Mason, C. R., Johnson, M. T., Fu, Q. G., Gomez, J. E. & Ebner, T. J. Temporal profile of the directional tuning of the discharge of dorsal premotor cortical cells. *NeuroReport* **9,** 989–995. ISSN: 0959-4965 (Apr. 1998).

61. Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience* **21,** 1281–1289. ISSN: 1097-6256 (Aug. 2018).

62. Maynard, E. M. *et al.* Neuronal interactions improve cortical population coding of movement direction. *The Journal of Neuroscience* **19,** 8083–93. ISSN: 1529-2401 (Sept. 1999).

63. Mizes, K. G. C., Lindsey, J., Escola, G. S. & Ölveczky, B. P. Dissociating the contributions of sensorimotor striatum to automatic and visually-guided motor sequences. *bioRxiv,* 2022.06.13.495989 (Nov. 2022).

64. Moore, D. D., Walker, J. D., MacLean, J. N. & Hatsopoulos, N. G. Validating markerless pose estimation with 3D X-ray radiography. *Journal of Experimental Biology* **225.** ISSN: 14779145 (May 2022).

65. Moran, D. W. & Schwartz, A. B. Motor Cortical Representation of Speed and Direction During Reaching. *Journal of Neurophysiology* **82,** 2676–2692. ISSN: 0022-3077 (Nov. 1999).

66. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience* **22,** 1677–1686. ISSN: 1546-1726 (Sept. 2019).

67. Nath, T. *et al.* Using DeepLabCut for 3D markerless pose estimation across species and behaviors. *Nature Protocols* **14,** 2152–2176. ISSN: 17502799 (July 2019).

68. Oby, E. R. *et al.* New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences of the United States of America* **116,** 15210–15215. ISSN: 1091-6490 (July 2019).

69. Okorokova, E. V., Goodman, J. M., Hatsopoulos, N. G. & Bensmaia, S. J. Decoding hand kinematics from population responses in sensorimotor cortex during grasping. *Journal of Neural Engineering* **17,** 046035. ISSN: 1741-2552 (Aug. 2020).

70. Paninski, L., Fellows, M. R., Hatsopoulos, N. G. & Donoghue, J. P. Spatiotemporal Tuning of Motor Cortical Neurons for Hand Position and Velocity. *Journal of Neurophysiology* **91,** 515–532. ISSN: 00223077 (Jan. 2004).

71. Pereira, T. D. *et al.* Fast animal pose estimation using deep neural networks. *Nature Methods* **16,** 117–125. ISSN: 15487105 (Jan. 2019).

72. Phillips, C. G., Powell, T. P. & Wiesendanger, M. Projection from low-threshold muscle afferents of hand and forearm to area 3a of baboon's cortex. *The Journal of Physiology* **217,** 419–446. ISSN: 1469-7793 (Sept. 1971).

73. Prud'homme, M. J. & Kalaska, J. F. Proprioceptive activity in primate primary somatosensory cortex during active arm reaching movements. *Journal of Neurophysiology* **72,** 2280–2301. ISSN: 00223077 (1994).

74. Reddy, N. D., Guigues, L., Pischulini, L., Eledath, J. & Narasimhan, S. *TesseTrack: End-to-End Learnable Multi-Person Articulated 3D Pose Tracking* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2021), 15190–15200.

75. Riehle, A., Grün, S., Diesmann, M. & Aertsen, A. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* **278,** 1950–1953. ISSN: 00368075 (Dec. 1997).

76. Russo, A. A. *et al.* Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron* **97,** 953–966. ISSN: 0896-6273 (Feb. 2018).

77. Sabatini, D. A. & Kaufman, M. T. Reach-dependent reorientation of rotational dynamics in motor cortex. *bioRxiv,* 2021.09.09.459647 (July 2023).

78. Sadovsky, A. J. & MacLean, J. N. Scaling of Topologically Similar Functional Modules Defines Mouse Primary Auditory and Somatosensory Microcircuitry. *Journal of Neuroscience* **33,** 14048–14060. ISSN: 0270-6474 (Aug. 2013).

79. Saleh, M., Takahashi, K., Amit, Y. & Hatsopoulos, N. G. Encoding of coordinated grasp trajectories in primary motor cortex. *The Journal of Neuroscience* **30,** 17079–90. ISSN: 1529-2401 (Dec. 2010).

80. Saleh, M., Takahashi, K. & Hatsopoulos, N. G. Encoding of coordinated reach and grasp trajectories in primary motor cortex. *The Journal of Neuroscience* **32,** 1220–32. ISSN: 1529-2401 (Jan. 2012).

81. Sauerbrei, B. A. *et al.* Cortical pattern generation during dexterous movement is input-driven. *Nature* **577,** 386–391. ISSN: 14764687 (Jan. 2020).

82. Schieber, M. H. Training and Synchrony in the Motor System. *The Journal of Neuroscience* **22,** 5277–81 (2002).

83. Schwartz, A. B., Kettner, R. E. & Georgopoulos, A. P. Primate motor cortex and free arm movements to visual targets in three- dimensional space. I. Relations between single cell discharge and direction of movement. *Journal of Neuroscience* **8,** 2913–2927. ISSN: 0270-6474 (Aug. 1988).

84. Scott, S. H. & Kalaska, J. F. Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. *Journal of Neurophysiology* **73,** 2563–2567. ISSN: 00223077 (1995).

85. Sergio, L. E., Hamel-Pâquet, C. & Kalaska, J. F. Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. *Journal of Neurophysiology* **94,** 2353–2378. ISSN: 00223077 (2005).

86. Sergio, L. E. & Kalaska, J. F. Systematic Changes in Motor Cortex Cell Activity With Arm Posture During Directional Isometric Force Generation. *Journal of Neurophysiology* **89,** 212–228. ISSN: 0022-3077 (Jan. 2003).

87. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience* **36,** 337–359. ISSN: 0147-006X (2013).

88. Sheshadri, S., Dann, B., Hueser, T. & Scherberger, H. 3D reconstruction toolbox for behavior tracked with multiple cameras. *Journal of Open Source Software* **5,** 1849. ISSN: 2475-9066 (Jan. 2020).

89. Shinoda, Y. & Kakei, S. Distribution of terminals of thalamocortical fibers originating from the ventrolateral nucleus of the cat thalamus. *Neuroscience Letters* **96,** 163–167. ISSN: 0304-3940 (Jan. 1989).

90. Siegle, J. H. *et al. A survey of spiking activity reveals a functional hierarchy of mouse corticothalamic visual areas* Oct. 2019.

91. Snyder, A. C., Morais, M. J. & Smith, M. A. Dynamics of excitatory and inhibitory networks are differentially altered by selective attention. *Journal of Neurophysiology* **116,** 1807–1820. ISSN: 15221598 (Oct. 2016).

92. Stavisky, S. D., Kao, J. C., Ryu, S. I. & Shenoy, K. V. Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions. *Neuron* **95,** 195–208. ISSN: 0896-6273 (July 2017).

93. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576,** 266–273. ISSN: 14764687 (Dec. 2019).

94. Stevenson, I. H. *et al.* Functional Connectivity and Tuning Curves in Populations of Simultaneously Recorded Neurons. *PLoS Computational Biology* **8** (ed Behrens, T.) e1002775. ISSN: 1553-7358 (Nov. 2012).

95. Sundiang, M., Hatsopoulos, N. G. & MacLean, J. N. Dynamic structure of motor cortical neuron coactivity carries behaviorally relevant information. *Network Neuroscience,* 1–18 (Mar. 2023).

96. Suresh, A. K. *et al.* Neural population dynamics in motor cortex are different for reach and grasp. *eLife* **9,** 1–16. ISSN: 2050084X (Oct. 2020).

97. Tanji, J. & Wise, S. P. Submodality distribution in sensorimotor cortex of the unanesthetized monkey. *Journal of Neurophysiology* **45,** 467–481. ISSN: 00223077 (1981).

98. Vaadia, E *et al.* Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* **373,** 515–518. ISSN: 1476-4687 (Feb. 1995).

99. Walker, J. D., Pirschel, F., Gidmark, N., MacLean, J. N. & Hatsopoulos, N. G. A platform for semiautomated voluntary training of common marmosets for behavioral neuroscience. *Journal of Neurophysiology* **123,** 1420–1426. ISSN: 0022-3077 (Apr. 2020).

100. Walker, J. D. *et al.* Chronic wireless neural population recordings with common marmosets. *Cell Reports* **36,** 109379. ISSN: 2211-1247 (July 2021).

101. Wu, A. *et al. Deep Graph Pose: A semi-supervised deep graphical model for improved animal pose tracking* Aug. 2020.

102. Wu, G. *et al.* ISB recommendation on definitions of joint coordinate systems of various joints for the reporting of human joint motion–Part II: shoulder, elbow, wrist and hand. *Journal of Biomechanics* **38,** 981–992. ISSN: 0021-9290 (May 2005).

103. Yger, P. *et al.* A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo. *eLife* **7.** ISSN: 2050084X (Mar. 2018).