Imperial College London

Imperial College London

Department of Bioengineering

Constraints on population dynamics across cortical and subcortical sensorimotor circuits

Cecilia Gallego Carracedo

June 2026

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy.

I hereby declare that the work presented in this thesis is my own. Any contributions from other researchers, or previously published materials have been appropriately referenced and acknowledged.

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Behaviour emerges from the coordinated activity across brain regions with distinct architectures. Each brain region produces specific neural signals that make specialised contributions to behaviour. However, the manner in which the distinct regional architectures shape the neural signals which ultimately produce behaviour remains poorly understood. The focus of this thesis was to study the neural signals from brain areas within the sensorimotor system with different architectures, and characterise their similarities and differences with respect to their architecture and functionality.

In the first part of the thesis, we characterised the relationship between local field potentials (LFPs) and neural population dynamics across three distinct regions of the primate sensorimotor cortex during a delayed reaching task. We found that the LFP band power was correlated with the neural population latent dynamics in a frequency-dependent and region-specific manner. The relationship remained stable across movement planning and execution, demonstrating that this relationship reflects architectural constraints rather than the specific neural process.

In the second part, to complement the previous study where the neural constraints were passively studied, we used an operant brain-computer interface (BCI) task in head-fixed mice running on a spherical treadmill to directly probe the constraints of distinct architectures on the volitional modulation of their brain activity. Animals demonstrated volitional modulation of neural activity in primary motor cortex (M1) and dorsolateral striatum (DLS), two highly connected regions within the sensorimotor system. The BCI-driven modulation had a markedly different spatial extent consistent with the architectural properties of each circuit, with localised modulation near BCI channels in DLS, and widespread modulation across M1.

Together, these findings demonstrate that circuit architecture imposes distinct constraints on the neural population signals across the sensorimotor system that shape their contributions to behaviour.

Acknowledgements

First of all, I would like to thank my PhD supervisor, Juan. Thank you for taking a chance on a lost undergrad student and for your supervision all these years. Having you as a mentor has been a great inspiration to become a better scientist. Also, thank you for creating the best lab one can ask for. Good luck in this new chapter of your life, London will miss you.

To our collaborators, Matt, Raeed and Lee. Thank you for your great contributions to my very first published work. It was a pleasure to be able to discuss and learn from all of you.

Nothing of this would have been possible without all the lab members that have joined me in these years, sharing celebrations and also a few tears. To the OGs, Catia, Joanna, Bence, Mostafa and Jimmie, and to the honorary members, Viki, Andrea and Elena, I cannot express how lucky I am to have landed with such an amazing group of people, you made me feel at home from the very first day. And to all the newcomers, Ciara, Genji, Martin, Ioana, Lara, Nicole, Vish, Natalia, Irene, for making what I thought was a perfect lab even better. Thank you for all the evening pubs, weekend picnics and conference trips, you have made these years truly wonderful.

Moving to London to do this PhD has been a great adventure, but it would not have been possible without the support of the people I left back home. Starting with ma beaches, Elena, Ichi, Victor, Irene, Julia, Mario, thank you for your visits, for hosting me, and for giving me so much love during all these years and the plenty before. To my "all doctors" friends Alexia, Julio y Clara, our trips and reunions make me feel like I never left la Carlos, in a good way. To my vecis, Pepa, Almu y Elsa, thank you for always making time for our little gathering, even being far.

To my amazing family. The three best siblings that anyone could ask for Jesu, Javi y Amalia, and the family additions, Sofia y Alex, thank you for your remote (and not that remote) emotional support. And to the best parents, thank you for giving me the space to grow but also staying close when I needed you. I am really grateful for all the little things like sharpening my knives and fixing my pants, but mostly, thank you for taking such good care of me all these years.

To my dogs, Loki and Chuck, although you cannot read this, you deserve this acknowledgement. You are the best dogs anyone could ask for, and your walks and attention have greatly contributed to this work.

And last but probably the most important person to thank, my first husband Josu. I cannot account for all the times you have helped during this PhD, including moving countries. Thank you for all the jokes, love and support. Also, thank you for proofreading this thesis, now any grammatical mistake is OUR mistake.

Contents

Abstract **3**

Acknowledgements **5**

1 Introduction **13**

1.1 From single neurons to neural manifolds 14

1.2 Local field potentials as a window into population dynamics 16

1.3 The sensorimotor system 20

1.3.1 Cortical sensorimotor structures: motor and somatosensory cortices 21

1.3.2 Subcortical sensorimotor structures: basal ganglia 22

1.4 Aims 25

2 Local field potentials reflect cortical population dynamics in a region-specific and frequency-dependent manner **26**

2.1 Abstract 27

2.2 Introduction 27

2.3 Results 30

2.3.1 Behavioural task and neural recordings 30

2.3.2 The LFP-latent dynamics correlation profiles in primary motor cortex are frequency-dependent 33

2.3.3 The LFP-latent dynamics correlation profiles are stable between movement planning and execution 37

2.3.4 The LFP-latent dynamics correlation profiles change between primary motor and premotor cortices 38

2.3.5 The LFP-latent dynamics correlation profiles are different for area 2 of primary somatosensory cortex . . . 40
2.3.6 LFP-latent dynamics correlations are not predicted from single neuron features . . . 42
2.4 Discussion . . . 44
2.4.1 Relation to previous single neuron studies . . . 44
2.4.2 A neural population view is necessary . . . 46
2.4.3 Function as well as circuit biophysics influences the differences between cortical areas . . . 47
2.5 Conclusion . . . 50
2.6 Methods . . . 50
2.6.1 Behavioural task . . . 50
2.6.2 Neural implants and recordings . . . 51
2.6.3 Behavioural epochs . . . 52
2.6.4 LFP processing . . . 52
2.6.5 Neural population latent dynamics . . . 53
2.6.6 Alignment of latent dynamics and LFPs . . . 54
2.6.7 Decoding hand velocity from M1 and area 2 of somatosensory cortex movement activity . . . 55
2.6.8 Predicting reach direction from PMd planning activity . . . 56
2.6.9 Additional analyses including controls . . . 57
2.6.10 Statistics . . . 59

3 A brain-computer interface paradigm reveals different constraints on cortical and striatal activity 75

3.1 Abstract . . . 76
3.2 Introduction . . . 77
3.3 Results . . . 79
3.3.1 Operant Brain Computer Interface task . . . 79
3.3.2 Volitional increase of the firing rate in the BCI channels. . . . 82

3.3.3 BCI-driven firing rate increases reveal distinct constraints on neuronal population activity . . . 85
3.3.4 Animals adopt distinct control strategies for cortical and striatal BCIs 88
3.3.5 Neural manifolds in DLS reorganise over the course of striatal BCI learning, but not M1 manifolds in cortical BCI . . . 90
3.4 Discussion . . . 92
3.4.1 An Operant BCI task as a tool to study neural constraints on volitional modulation . . . 93
3.4.2 Neural population constraints on BCI-induced neuronal changes . . 95
3.4.3 BCI control strategies reveal differences in M1 and DLS . . . 97
3.4.4 Neural manifolds as a static versus dynamic constraint on BCI learning 99
3.5 Conclusion . . . 100
3.6 Methods . . . 101
3.6.1 Subjects . . . 101
3.6.2 Surgeries and regulated procedures . . . 101
3.6.3 Histology . . . 103
3.6.4 Experimental setup . . . 103
3.6.5 Habituation and training protocol . . . 104
3.6.6 Electrophysiological recordings . . . 106
3.6.7 Real-time Brain Computer Interface . . . 107
3.6.8 Data processing . . . 109
3.6.9 Firing rate percentage of change . . . 109
3.6.10 Encoding BCI firing rate from movement . . . 110
3.6.11 Decoding movement from M1 or DLS population activity . . . 111
3.6.12 Neural manifold analysis . . . 111

4 Conclusion **113**

4.1 Summary . . . 113
4.2 Impact . . . 115
4.3 Limitations . . . 116

4.4 Future work 118

Bibliography **120**

List of Figures

1.1 The manifold theory. 15
1.2 Local field potentials. 18
1.3 The sensorimotor system in the brain. 20
1.4 The basal ganglia. 23

2.1 Hypothesis schematic. 31
2.2 Behavioural task and neural recordings. 32
2.3 Similarity between primary motor cortex (M1) latent dynamics and each local field potential (LFP) band during movement execution. 34
2.4 Decoding movement kinematics from primary motor cortex (M1) local field potentials (LFPs). 36
2.5 The primary motor cortex (M1) local field potential (LFP)-latent dynamics correlation profile is preserved between movement planning and execution. 37
2.6 Similarity between dorsal premotor cortex (PMd) latent dynamics and each local field potential (LFP) band during movement preparation. 39
2.7 Similarity between area 2 latent dynamics and each local field potential (LFP) band during somatosensory feedback processing. 41
2.8 Single neuron features do not account for the local field potential (LFP)-latent dynamics correlations. 43
S1 Similarity between primary motor cortex (M1) latent dynamics and each local field potential (LFP) band. 60
S2 Surrogate latent dynamics control. 60

S3 The local field potential (LFP)-latent dynamics correlation profiles are preserved when changing neural manifold dimensionality. 61
S4 Further control analyses to illustrate the frequency-dependent relationship between primary motor cortex (M1) latent dynamics and each local field potential (LFP) band. 63
S5 The relationship between local field potential (LFP) bands does not predict their association with the latent dynamics. 64
S6 The latent local field potential (LFP) signals have a similar association with the neural population latent dynamics as the individual LFP signals do. . . 65
S7 Additional analyses to illustrate movement decoding from primary motor cortex (M1) local field potential (LFP). 66
S8 Decrease in primary motor cortex (M1) firing rates and neural correlations when monkeys are not engaged in the task. 67
S9 Similarity between dorsal premotor cortex (PMd) latent dynamics and each local field potential (LFP) band, and predicting reach direction from PMd LFPs. 68
S10 Similarity of local field potential (LFP) spectral properties across frequency bands and behavioural epochs. 70
S11 Similarity between area 2 latent dynamics and each local field potential (LFP) band, and decoding hand kinematics from area 2 LFPs. 71
S12 Region-specificity of the local field potential (LFP)-latent dynamics correlation profile. 72
S13 Additional data: Single neuron features do not account for the local field potential (LFP)-latent dynamics correlations. 73
S14 Single-neuron population statistics differ across regions, but manifold dimensionality reliably captures population activity. 74

3.1 Operant brain-computer interface (BCI) task schematic. 80
3.2 Real-time brain-computer interface (BCI) firing rate calculation from data. 81
3.3 BCI firing rate increased throughout the session. 82

3.4 Anticipatory licking shows intentional control of the BCI. 84
3.5 Firing rate increase revealed different neural population constraints. 87
3.6 Motion can predict BCI firing rate on M1 but not DLS. 89
3.7 DLS manifolds reorganise over the course of BCI learning, while M1 manifolds remain comparatively stable. 91

Chapter 1

Introduction

The neural basis of behaviour is one of the biggest research questions in neuroscience. Most behaviours are not generated by a single brain region, but emerge from the coordinated activity of distributed circuits. Even an apparently simple motor task like reaching for an object requires the involvement of multiple cortical and subcortical regions contributing to the behaviour, from sensory proprioceptive feedback in the somatosensory cortex, to movement planning and execution in the motor cortex and basal ganglia [1, 2].

Despite their coordinated interaction, different brain regions make specific, specialised contributions to behaviour. Most neuroscience research focuses on how each individual region contributes to behaviour. However, little is known about how the circuitry of each region constrains the produced neural signals, which ultimately generate the behaviour.

Each brain region has its own unique cytoarchitecture, that is, cellular composition, shape, neural density and layer arrangement. Each region also has a distinct connectivity both locally and across brain regions. These structural differences shape the neural activity of each area, constraining it to specific activity patterns that define its behavioural

contributions [3].

Hence, to understand circuit function, we first need to understand the signals they generate. However, neural circuits produce multiple types of signals simultaneously, from the spiking of individual neurons to local field potentials encompassing the coordinated activity of thousands of cells. Each of these signals, spread across a wide variety of spatial and temporal scales, captures different aspects of circuit function and behaviour [4, 5]. Characterising these brain signals and understanding the constraints imposed by different circuit architectures are still fundamental open questions in the field.

1.1 From single neurons to neural manifolds

Among the signals that the neural circuits generate, neuronal spiking is one of the most fundamental and has been the primary focus of efforts to understand how circuits encode behaviour. Early approaches focused on the activity of individual neurons, which were found to show selective tuning to specific movement parameters or sensory stimuli [6]. However, a growing body of evidence suggests that the connectivity of the neural circuit constrains the population into specific covariance patterns, and behaviour is not encoded by individual neurons, but by the coordinated activity of the neural population (Fig. 1.1a), the neural manifold [7, 8, 9].

Under this framework, if the activity of a neural population is represented in a multidimensional space, in which each dimension corresponds to the firing rate of an individual neuron, the collective activity over time would be a trajectory in this neural space. The set of possible trajectories never spans the whole space, but it is restricted to a lower-dimensional hyperplane, which is defined by the covariance patterns imposed by the neural circuit (Fig. 1.1b). This hyperplane is called the neural manifold, and the bases that define it are called the neural modes. In this framework, the individual firing of each neuron is the result of the time-varying activation of each of the neural modes, referred to as latent dynamics (Fig. 1.1c).

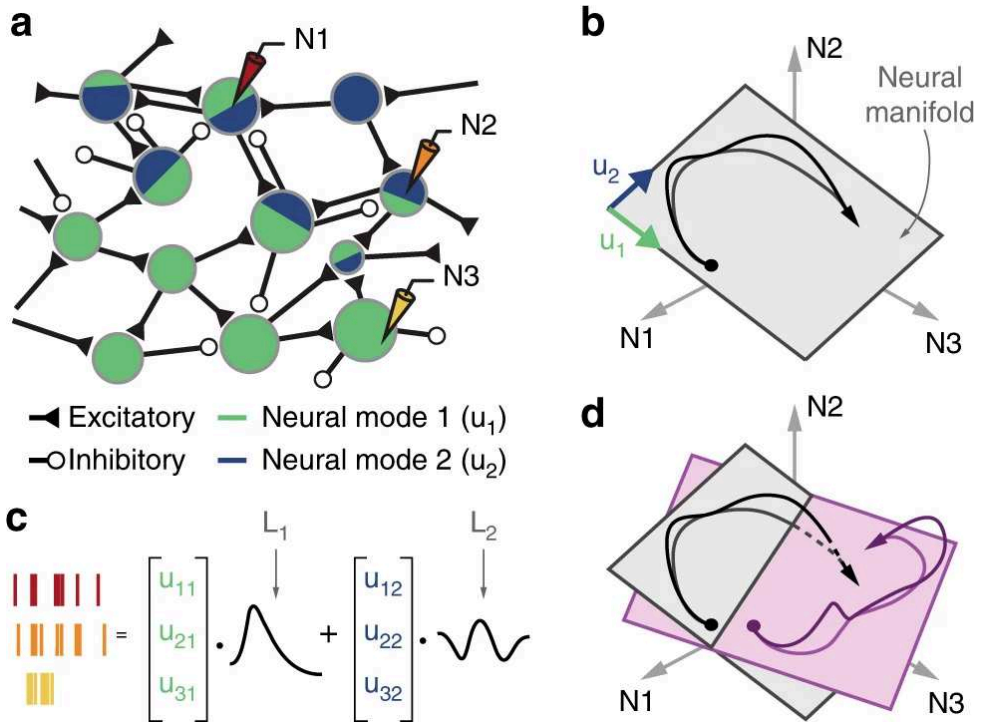

Figure 1.1: The manifold theory. **(a)** The connectivity of the neural circuits results in neural modes whose combined activity explains the specific activity of individual neurons. **(b)** The combined activity of the three neurons recorded in **a** is represented by a trajectory (in black, arrow indicates time direction) mostly confined to a low-dimensional neural manifold (grey plane) spanned by the neural modes (u_1 and u_2 basis vectors). **(c)** The spiking activity of each recorded neuron is the time-varying activation of the two neural modes, the latent dynamics L_1 and L_2. **(d)** Different neural processes of the same brain circuit may span different manifolds. Reproduced from Gallego et al. (2018) [10].

Two related but distinct concepts underlie this framework. The neural modes are the fixed set of covariance patterns, or basis vectors, that define the manifold; they capture how neurons co-vary with one another, independent of time. The latent dynamics are the time-varying activation of each mode, and capture when and how strongly they are expressed during a given behaviour [7, 9]. This distinction matters because the neural modes are thought to be a comparatively static property of the circuit, set by its anatomy and connectivity, whereas the latent dynamics can change flexibly from trial to trial and behaviour to behaviour, within the constraints the modes impose [8]. Under this view, the dimensionality of the manifold, i.e., the number of neural modes needed to capture most of the population activity, reflects how constrained a circuit's activity is [11, 12]. A low-dimensional manifold indicates that the population is restricted to a small number of co-activation patterns, while a higher-dimensional manifold indicates a richer repertoire of

patterns available to the circuit.

The manifold framework has already transformed the motor control field in multiple studies, from explaining how preparatory and movement execution activity lie in orthogonal manifolds to avoid cross-talk [13, 14], to showing how motor adaptation occurs within the manifold structure [15, 16]. Furthermore, the neural manifolds in motor-related cortices have proved to be stable across long periods of time [17] and even across individuals performing the same behaviour [18]. Overall, the neural manifold has proven relevant for many brain regions and behaviours [7, 8, 19, 20, 21], raising new questions about its role as a fundamental principle of neural computation.

The schematic in Fig. 1.1 shows an example linear manifold, with linear neural modes, although manifolds can also exhibit nonlinear structure, depending on their complexity [22, 23, 24, 25]. The manifold complexity is largely dependent on the circuit cytoarchitecture, which determines the neural connectivity and ultimately constrains the covariance patterns that the neural population can generate [26, 27, 28]. For instance, the highly correlated activity in motor cortex gives rise to a low-dimensional manifold, while the sparser and less correlated activity in subcortical structures like the striatum results in higher-dimensional representations. Hence, during the same behaviour, each brain region shows specific manifold characteristics that are likely affecting its function in the generation of such behaviour.

Beyond the neural population dynamics that define the neural manifold, neural circuits simultaneously generate another set of signals, the local field potentials, that reflect the collective activity of thousands of neurons and often relate to specific brain functions.

1.2 Local field potentials as a window into population dynamics

The local field potentials (LFPs) are among the most widely used signals in neural circuits, as they reflect the extracellular electrical activity of thousands of neurons surrounding

the electrode (Fig. 1.2A). Among the main contributors to the LFPs, synaptic activity is often considered the most prominent. The temporal synchronisation of synaptic currents across many neurons converts their individually negligible contributions into a measurable extracellular signal that is detected by the electrodes [4, 5, 29, 30]. Since LFPs are primarily defined by the synaptic currents in the neural population, the cytoarchitecture of the neural circuit is a key determinant of LFP signals. Consequently, given the biophysical origin of the LFPs, they should be constrained by the same circuit architecture that also shapes the neural manifolds.

The synchronised synaptic currents underlying LFP generation tend to oscillate at specific frequencies that are functionally meaningful. This frequency dependence is one of the defining features of the LFP signals, and has revealed distinct frequency bands associated with specific functions in different brain regions (Fig. 1.2B). For example, in the motor cortex, low (< 5 Hz) and high (> 50 Hz) frequencies carry information about the ongoing movement, while the beta band (12-25 Hz) shows a characteristic suppression during movement onset [31, 32, 33]. The same beta band in the somatosensory cortex reflects afferent inputs from the limb [34]. The frequency dependence reflects oscillations at the circuit level determined by the times of synaptic transmission and the connectivity patterns of the circuit [35, 36].

These frequency-specific oscillations are thought to support neural communication in different ways. Between brain regions, the phase coherence between two LFP signals carries important functional information necessary for inter-area communication (Fig. 1.2C) [37, 38, 39, 40]. According to the communication-through-coherence hypothesis, phase alignment between regions creates recurring windows of high excitability, during which inputs from one region are most effectively integrated by another, thus gating the effective flow of information between circuits [41, 42]. Rhythmic inhibitory activity, particularly in the beta and gamma bands, has similarly been proposed to periodically suppress or release local circuits, structuring which neurons can spike and when, a mechanism known as gating by inhibition [43].

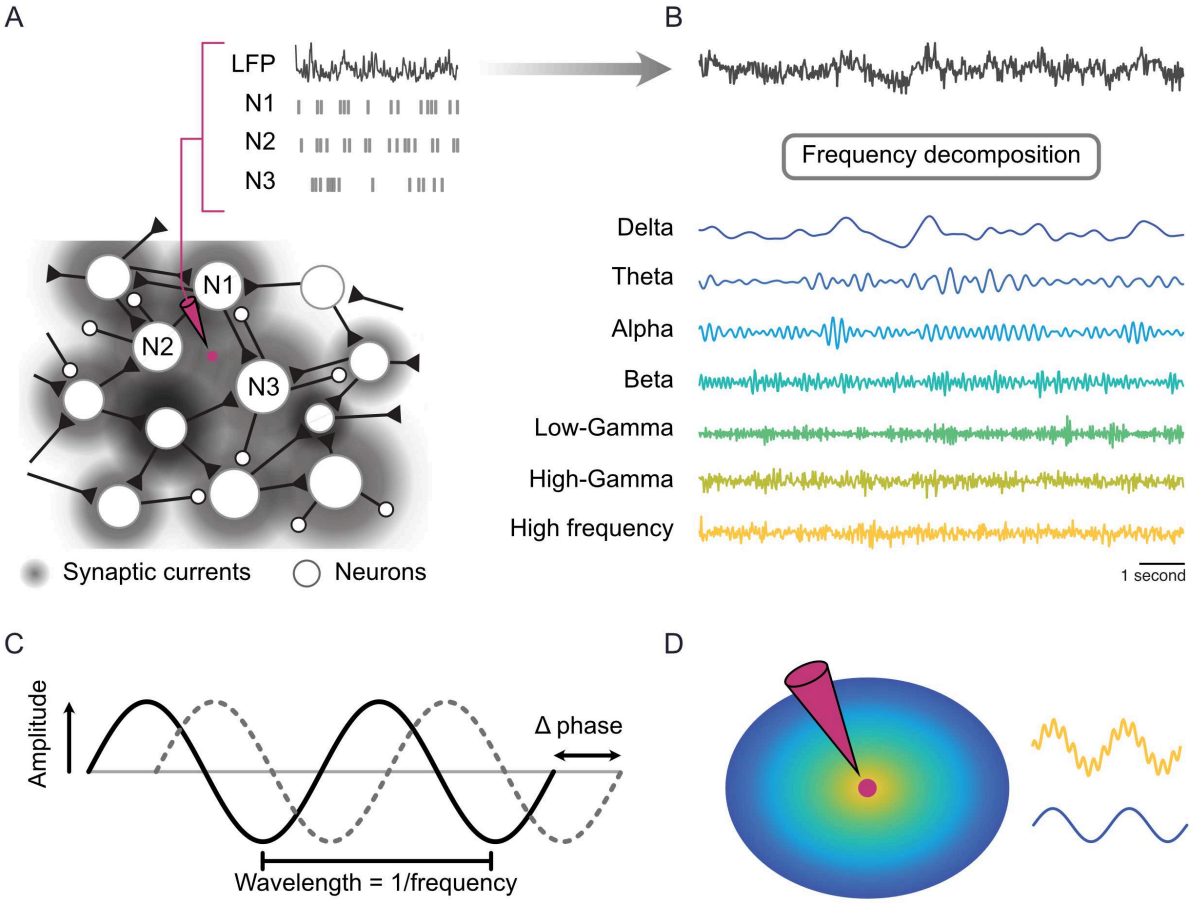

Figure 1.2: Local field potentials description. **(A)** An intracortical electrode records the combined activity of multiple neurons and the local field potential (LFP) in that region. **(B)** A frequency decomposition of the LFP reveals functionally significant frequency bands. Traditional frequency bands are called: delta (0.5-4 Hz), theta (4-8 Hz), alpha (8-12 Hz), beta (12-25 Hz), low-gamma (25-50 Hz), high-gamma (50-100 Hz), high-frequency oscillations (>100 Hz). **(C)** An LFP oscillation can be described by multiple properties: amplitude, frequency and phase difference with respect to other oscillations. These properties reflect the synaptic currents and define the LFPs. **(D)** The spatial reach of the LFP depends on the frequency of the oscillation; both high and low frequencies are found in the vicinity (yellow trace), but only low-frequency oscillations extend beyond that.

Within a single circuit, different frequency bands likely also reflect distinct biophysical origins. Low-frequency components are thought to arise mainly from slow, spatially extensive synaptic currents reflecting network-wide input, while higher-frequency components, particularly in the gamma range, are more closely tied to the local spiking output of nearby neurons (Fig. 1.2D) [44, 45, 46, 47, 48]. Specific frequency bands may also index distinct behavioural or cognitive states of a circuit, rather than reflecting a single, undifferentiated oscillation. For instance, the beta-band suppression described above has been proposed to

reflect the maintenance of the current sensorimotor state, becoming suppressed only once a new movement is about to be generated [32, 33]. Finally, recent modelling work has shown that shared, low-dimensional modes can capture behaviourally relevant dynamics common to both low and high frequency motor cortical LFPs and the spiking population activity, offering a further route through which frequency-specific oscillations may relate to the coordinated activity of a neural population [49].

Beyond basic science, LFPs have demonstrated a crucial clinical role. Compared to single neuron recordings, LFPs are more robust to signal degradation over time and easier to acquire, making them a preferred signal for clinical applications that target the brain [50, 51]. In epilepsy, abnormal high-frequency LFP activity, such as fast ripples in the hippocampus and entorhinal cortex, has been proposed as a biomarker for seizure-generating tissue, enabling targeted stimulation to reduce seizure frequency [52, 53]. Another well-studied clinical application is in Parkinson's disease, where the beta band of the LFPs in the subthalamic nucleus (STN) is pathologically elevated [54, 55]. Deep-brain stimulation (DBS) triggered by the pathological beta LFP has shown greater efficiency in tremor reduction compared to conventional open-loop DBS [56, 57]. In motor control, a wide range of LFP signals, mainly the power in low and high-frequency bands, have been used to control or inform BCI performance [58, 59, 60].

Although the LFPs mostly reflect the synaptic currents that also drive the neural population dynamics, and both are constrained by the cytoarchitectural properties of the neural circuit, the relationship between LFPs and neural population dynamics remains largely unexplored. Most studies focus on relating the properties of LFPs to the activity of single neurons and specific cell types [61, 44, 62, 63]; however, a population-level comparison may reveal a more fundamental relationship, given the common constraints between LFP and neural population dynamics [45, 49, 64].

The sensorimotor system provides an ideal system to test these questions because it comprises multiple regions with very distinct and well characterised cytoarchitectural properties, spanning from cortical areas like primary motor and premotor cortex to

subcortical structures like the striatum. Furthermore, the neural population dynamics and the LFPs in these regions have been extensively studied during movement, an ideal scenario to explore how the constraints imposed by the circuit cytoarchitecture affect the relationship between brain signals within and across areas, and their functionality.

1.3 The sensorimotor system

The sensorimotor system is one of the best characterised in systems neuroscience. It spans multiple regions with distinct cytoarchitectures, from laminar structures in cortical areas to unstratified organisation in the basal ganglia [1, 2]. The anatomy and connectivity of these regions have been extensively studied and provide a rich architectural diversity to investigate how the circuit properties shape the relationship between LFPs and neural population dynamics, and to what extent neural activity is constrained in cortical and subcortical structures, limiting their contribution to behaviour.

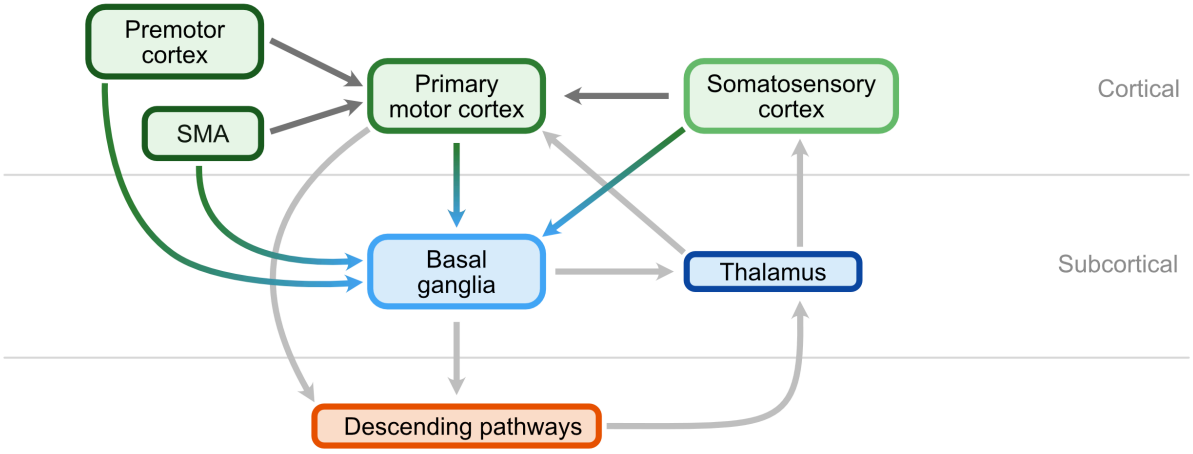

Figure 1.3: The sensorimotor system in the brain. Box diagram of the main sensorimotor system regions in the brain and their main connections, separated by cortical and subcortical structures. The dark grey lines correspond to cortico-cortical connections, and the green/blue lines to cortico-striatal connections. Other relevant anatomical connections such as the descending projection from primary motor cortex to brainstem and spinal cord (descending pathways), the cortico-basal-thalamic loop, and the afferent feedback to somatosensory cortex through thalamus are represented in light grey and were out of the scope of the project. SMA: Supplementary motor area.

Fig. 1.3 indicates the main contributors to motor control in both cortical and subcortical regions, together with the main connections between them. The basal ganglia have been

simplified as a single node; the collection of nuclei and its connectivity can be seen in Fig. 1.4.

1.3.1 Cortical sensorimotor structures: motor and somatosensory cortices

The sensorimotor cortex comprises a collection of areas with distinct but complementary roles in the planning and execution of movement. It is organised by layers with distinct cell types and input/output connections that serve different computational roles [3, 65]. The thickness of each of the layers varies significantly between cortical regions. For example, layer IV, which receives thalamic inputs, is very prominent in somatosensory cortex and virtually absent both in primary motor cortex (M1) and premotor cortex [66, 67]. In contrast, layer V, with downstream projections to subcortical structures, is particularly developed in M1 [68, 69] and premotor cortex has a particularly developed layer III, consisting mainly of long-range cortical connections, reflecting its role in integrating inputs from multiple cortical areas [70].

The sensorimotor cortex can be separated into somatosensory cortex, which processes proprioceptive sensory information during an ongoing movement; and motor cortex, which integrates sensorimotor information to plan and execute voluntary movement [1, 2]. The primary motor cortex (M1) is the primary output of the motor cortex, with direct projection to the spinal cord and to the basal ganglia via layer V [68, 69, 71].

M1 population activity has been widely studied. It has been well described by a low-dimensional manifold in a variety of behaviours [7, 13, 14, 15], and has been shown to remain stable across long periods of time and individuals performing the same behaviour [17, 18]. M1 LFPs are also well characterised; in fact, the movement-related changes in specific frequency bands are among the best studied LFP phenomena [31, 32, 33, 59, 72]. This combination makes it an ideal starting point to characterise the relationship between these two signals.

Anterior to M1, the supplementary motor area (SMA) is located medially and the premotor cortex laterally. Specifically, the premotor cortex is a high-order region of the motor cortex mainly involved in movement planning [14, 73]. It receives inputs from many cortical and subcortical structures and projects to M1 via its prominent layer III [70]. The difference in layer organisation and connections makes M1-PMd comparison ideal for testing the region specificity of the LFP-neural population relationship.

Within the subdivisions of somatosensory cortex, area 2 is of particular interest, as it integrates proprioceptive information about the state of the limb [67, 74]. Its activity is primarily driven by sensory feedback, not motor commands; that is, neurons fire in relation to the state of the limb during and after movement. Its granular cytoarchitecture, with a prominent layer IV, contrasts sharply with the agranular motor cortical regions [75]. Together, M1, premotor cortex and area 2 of primary somatosensory cortex form a complementary set of sensorimotor regions with distinct cytoarchitectural and functional properties that provide the ideal framework to evaluate whether their population dynamics and LFPs are stable or circuit-dependent.

Beyond the cortex, motor control also involves subcortical structures that receive direct inputs from several cortical structures, modulating and transmitting the motor output to descending pathways.

1.3.2 Subcortical sensorimotor structures: basal ganglia

The two main subcortical structures involved in motor control are the basal ganglia and the thalamus. Particularly, the basal ganglia are a collection of interconnected subcortical nuclei that integrate cortical signals and modulate motor output (Fig. 1.4) [76, 77]. Although anatomically well characterised, the functionality of this region and its nuclei is still a very debated topic [78, 79].

The main input to the basal ganglia is the striatum, an integrative region that incorporates relevant inputs from multiple regions of the cortex [76, 81]. These strong cortical projections are remarkably characterised and topographically organised, with the sensorimotor cortex

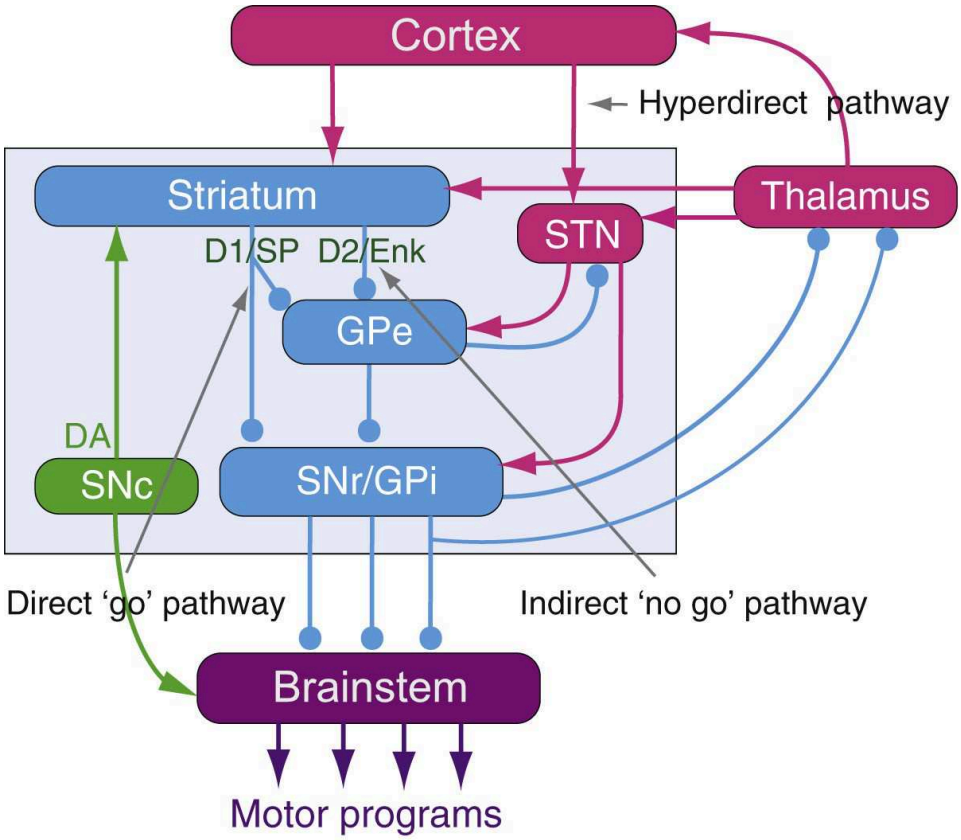

Figure 1.4: The basal ganglia. The blue square contains the nuclei of the basal ganglia. The lines show anatomical connections, with arrow endings representing excitatory connections and circle endings, inhibitory. GPe/GPi: globus pallidus externus/internus; SNc/SNr: substantia nigra pars compacta/reticulata; STN: subthalamic nucleus. Slightly modified from [80].

primarily projecting to the dorsolateral striatum (DLS), while other high-level cognitive regions such as the association cortices and the limbic system mainly project to dorsomedial and ventral striatum, respectively [71, 82, 83, 84].

Unlike cortical laminar cytoarchitecture, the cellular anatomy of DLS is not geometrically structured. Instead, it predominantly consists of medium spiny neurons (MSN), presenting a hyperpolarised resting membrane potential that requires strong convergent input to generate spiking activity [81, 85, 86]. Consequently, neurons in DLS usually have sparser and less correlated activity compared to other cortical regions [22, 87, 88]. This structural contrast has direct consequences for the neural manifold each region can support. The layered organisation of cortex, with its stereotyped local microcircuitry, promotes strong, structured correlations among nearby neurons, constraining cortical population activity

to a comparatively low-dimensional manifold, as introduced above. DLS, lacking this layered structure and instead relying on convergent input from many, often functionally diverse, cortical afferents to drive MSN firing, shows reduced local correlation among its neurons, resulting in sparser activity and a higher-dimensional manifold than that of cortex. The direct strong connection between M1 and DLS, together with their significant architectural differences, therefore make these regions an ideal setting to test the neural activity patterns they can generate given the distinct circuit constraints.

The striatum projects to different ganglia through the alternating direct and indirect pathways (Fig. 1.4) [76, 77, 81]. When the neurons in the direct pathway are activated, they inhibit the globus pallidus internus (GPi) and substantia nigra pars reticulata (SNr), the main output of the basal ganglia. The inhibition of basal ganglia output releases the previously inhibited thalamus and brainstem pathways to produce a given movement. On the other hand, when the indirect pathway is activated, it inhibits globus pallidus externus (GPe), which usually inhibits the basal ganglia output. When GPe is inhibited, the output nuclei inhibition is released, making the nuclei active and suppressing any downstream activation, preventing any movement from happening. The combination of these two pathways, mediated by dopamine generated in substantia nigra pars compacta (SNc), regulates movement execution [81]. Additionally, the subthalamic nucleus (STN) serves as an extra controller of movement gating through a hyperdirect pathway from cortex, activating GPe or the basal ganglia output according to the movement necessities [89, 90].

Despite the extensive characterisation, the function of the basal ganglia beyond simple movement gating, especially the striatal contributions to movement and the relationship between DLS and the cortical areas, remains debated. Some studies suggest that DLS primarily reflects cortical inputs, mirroring cortical population dynamics [91]. Others propose that DLS makes independent functional contributions that go beyond simply reflecting cortical activity [79, 92, 93].

Testing the individual functionality of two brain regions is challenging as it cannot be resolved by passively recording the brain activity. To directly probe whether M1 and

DLS operate under different constraints, a causal perturbation is needed. Brain-computer interfaces (BCIs) provide a unique solution, in which the perturbation in the neural activity comes from internally generated processes provoked by the volition of the subjects to control the BCI [94, 95]. Importantly, the generated perturbations with this paradigm follow the natural constraints of the neural circuits [15, 16, 96]. Using an operant BCI task to directly compare the capability for volitional modulation in M1 and DLS and the activity changes induced by the change could provide unique insights into whether these connected but architecturally distinct areas are subjected to specific circuit constraints.

1.4 Aims

Given the presented gaps, the purpose of this thesis is to understand how the biophysical and architectural properties of different brain regions shape their neural activity. We address this gap by studying multiple neural signals from several connected brain regions. We chose the sensorimotor system for this purpose since it is an extensively characterised set of brain regions with a very rich architectural diversity. We present the results in two complementary studies described in the following chapters. In Chapter 2, we show that specific frequencies of LFPs in cortical areas of the sensorimotor system reflect neural population dynamics that are dependent on the cytoarchitecture of their region and invariant for different behaviours. In Chapter 3, we actively perturb the neural population dynamics of a cortical and subcortical neural population with an operant BCI and study how the induced brain activity is directly related to the constraints imposed by the neural population organisation. Together, these two studies address complementary aspects of the same fundamental question: how circuit architecture shapes neural signals and constrains their modulation.

Chapter 2

Local field potentials reflect cortical population dynamics in a region-specific and frequency-dependent manner

The local field potentials and the neural population dynamics are both shaped by the coordinated synaptic currents within a neural circuit, yet how circuit architecture constrains both signals and their relationship remains largely uncharacterised. In this chapter, we characterised simultaneously recorded LFP and neural population activity in three cytoarchitecturally distinct regions of the primate sensorimotor cortex during a delayed reaching task.

The work presented in this chapter constitutes an article published in the eLife journal [64]. My contributions to this work included conceptualisation, analysis and validation of the data, methodology and custom software development, and writing the original manuscript.

2.1 Abstract

The spiking activity of populations of cortical neurons is well described by the dynamics of a small number of population-wide covariance patterns, whose activation we refer to as 'latent dynamics'. These latent dynamics are largely driven by the same correlated synaptic currents across the circuit that determine the generation of local field potentials (LFPs). Yet, the relationship between latent dynamics and LFPs remains largely unexplored. Here, we characterised this relationship for three different regions of primate sensorimotor cortex during reaching. The correlation between latent dynamics and LFPs was frequency-dependent and varied across regions. However, for any given region, this relationship remained stable throughout the behaviour: in each of primary motor and premotor cortices, the LFP-latent dynamics correlation profile was remarkably similar between movement planning and execution. These robust associations between LFPs and neural population latent dynamics help bridge the wealth of studies reporting neural correlates of behaviour using either type of recordings.

2.2 Introduction

Researchers and clinicians often monitor brain activity using implanted electrodes to understand how the brain drives behaviour and to improve clinical outcomes [50, 51, 97, 98]. Modern multielectrode arrays allow the spiking activity of hundreds of neurons to be observed simultaneously [99]. These same electrodes can also capture the lower frequency local field potentials (LFPs) that result from synaptic currents summed across many thousands of neurons [4, 5, 29, 30, 46]. Uncovering a relationship between the collective dynamics of populations of single neurons and those of larger scale LFPs will help bridge studies looking at each of these signals in isolation and advance our understanding of the human brain.

The spiking activity of neural populations can be well characterised by dynamics of relatively few covariance patterns [7, 8], which we refer to as the 'latent dynamics' [17,

100]. Studying these latent dynamics has provided insight into aspects of animal behaviour that was not apparent from the activity of single neurons alone, including decision making [20, 21], movement planning [13, 14, 73], learning [15, 16, 101, 102, 103], the control timing [73, 104, 105], and production of consistent behaviour [17].

The latent dynamics are constrained by the neural covariance structure, which is shaped in part by circuit connectivity [15, 16, 26, 27, 106]. Currents through these same connections as well as other biophysical properties of the circuit are the main contributors to the generation of the LFPs [4, 5, 29, 30, 46] (Fig. 2.1A). In particular, correlations in the synaptic input currents across the population yield changes in LFP power at specific frequency bands, which often relate to specific brain functions [31, 32, 59, 60, 61, 72, 107, 108, 109, 110, 111, 112, 113, 114, 115, 116, 117]. For example, in the motor cortices, movement-related information lies primarily at low and high frequencies [31, 59, 72, 112] (i.e., <5 Hz and >50 Hz), yet, the clearest LFP signature of imminent movement initiation is a marked decrease in activity in the 15–30 Hz band [32, 33, 34, 113]. In somatosensory cortical areas 2 and 3a, this same 15–30 Hz band contains signatures potentially related to afferent input from the limb [115].

This relationship between LFPs and behaviour takes two related forms. Specific LFP bands show power changes time-locked to particular behavioural events: the 15–30 Hz suppression described above occurs specifically around movement initiation, rather than being constantly present throughout the task. More directly, LFP power in certain bands can be used to predict, or decode, ongoing behaviour: both low- and high-frequency LFPs allow accurate prediction of muscle activity [31] and movement kinematics [59, 72] on a moment-to-moment basis, approaching the accuracy achieved using the spiking activity of the underlying population. This LFP-behaviour relationship also has a pathological counterpart: in Parkinson's disease, exaggerated beta-band oscillations in the basal ganglia correlate directly with motor symptoms such as tremor, and detecting this pathological beta activity in real time now forms the basis of closed-loop deep brain stimulation [56, 57]. Given this strong and varied correlation to behaviour, it seems likely that these

region-specific LFP phenomena relate to the neural population dynamics that more directly drive it.

Previous studies investigated the relationship between single neuron activity and LFPs [61, 44, 62, 113, 116, 118, 119]. However, we lack a systematic description of how the different LFP bands relate to the latent dynamics reflecting the coordinated activity of the neural populations driving behaviour. Here, we report on the region-specific relationship between the LFPs and the latent dynamics by addressing four hypotheses. First, since both LFP [4, 5, 29, 30, 46] and latent dynamics [15, 16, 26] are likely shaped by the circuit biophysics (Fig. 2.1B), we expect to find a robust relationship between these two signals. Second, we anticipate finding fundamental differences across LFP bands since only a specific subset of them is correlated with behaviour for any given brain region [31, 32, 59, 60, 61, 72, 107, 108, 109, 110, 111, 112, 113, 114, 115, 116, 117, 120]. That is, the relationship between LFPs and latent dynamics should be frequency-dependent. Third, we expect that as long as circuit biophysics remain unchanged, the LFP-latent dynamics associations will remain similarly stable. Thus, the various associations between motor cortical LFPs and latent dynamics should not change in the short time scale spanning the preparation and execution of a movement. Lastly, since different sensorimotor cortical areas have important differences in inputs and function [1, 2, 121, 122], as well as cytoarchitecture [65, 70], we hypothesise that the relationships between LFPs and latent dynamics should be region-specific.

We tested these four hypotheses using intracortical recordings from three different regions of the primate sensorimotor cortex during the same reaching behaviour: (1) primary motor cortex (M1), the main cortical output that controls movement execution [123, 124, 125, 126]; (2) dorsal premotor cortex (PMd), a region that integrates inputs from many structures and is largely involved in movement planning [73, 127, 128, 129]; and (3) area 2 of somatosensory cortex, where multimodal proprioceptive and cutaneous information about the state of the limb is processed [67, 74, 75]. Our results show that the relationship between the LFPs and the latent dynamics from the same cortical region

is indeed frequency-dependent and varies across cortical regions. Yet, the region-specific LFP-latent dynamics 'correlation profiles' remain remarkably stable across the evolving processes mediating movement planning and execution [73, 130, 131]. Finally, we show that these relationships do not trivially arise from correlations between single unit firing rates and LFP recordings from the same intracortical electrode; instead, they reflect the coordinated activity of neural populations. This holds both during movement execution and as animals perform more 'abstract computations' related to movement planning.

These results, which are critical for bridging studies of the sensorimotor system that use LFPs and neural population activity, paint a picture in which LFP bands relate to the latent dynamics in a stable, region-specific, and frequency-dependent manner. This picture has one unexpected feature: some LFP bands are strongly correlated with latent dynamics yet carry little information about the animal's movement kinematics in M1 and area 2, or the direction of the upcoming reach in PMd. Given that low-dimensional latent dynamics have been found across many brain regions [7, 19], our findings are likely to translate beyond the sensorimotor system. Such translation may be especially insightful for regions commonly studied using LFP rhythms, such as the hippocampus [132].

2.3 Results

2.3.1 Behavioural task and neural recordings

We trained four monkeys (C, M, H, and L) to perform an instructed delay centre-out reaching task using a manipulandum (Fig. 2.2A) (Methods). The monkeys started each trial by holding a cursor in the central target before one of the eight outer targets was presented. After a variable delay period, an auditory 'go cue' instructed them to move the cursor towards the presented target. A liquid reward was given after the monkeys had held the cursor in the target for a given period of time. For the monkeys with implants in area 2 (Monkeys H and L), both the delay and holding periods were omitted.

Each monkey was implanted with one or two 96-channel microelectrode arrays (Fig. 2.2B).

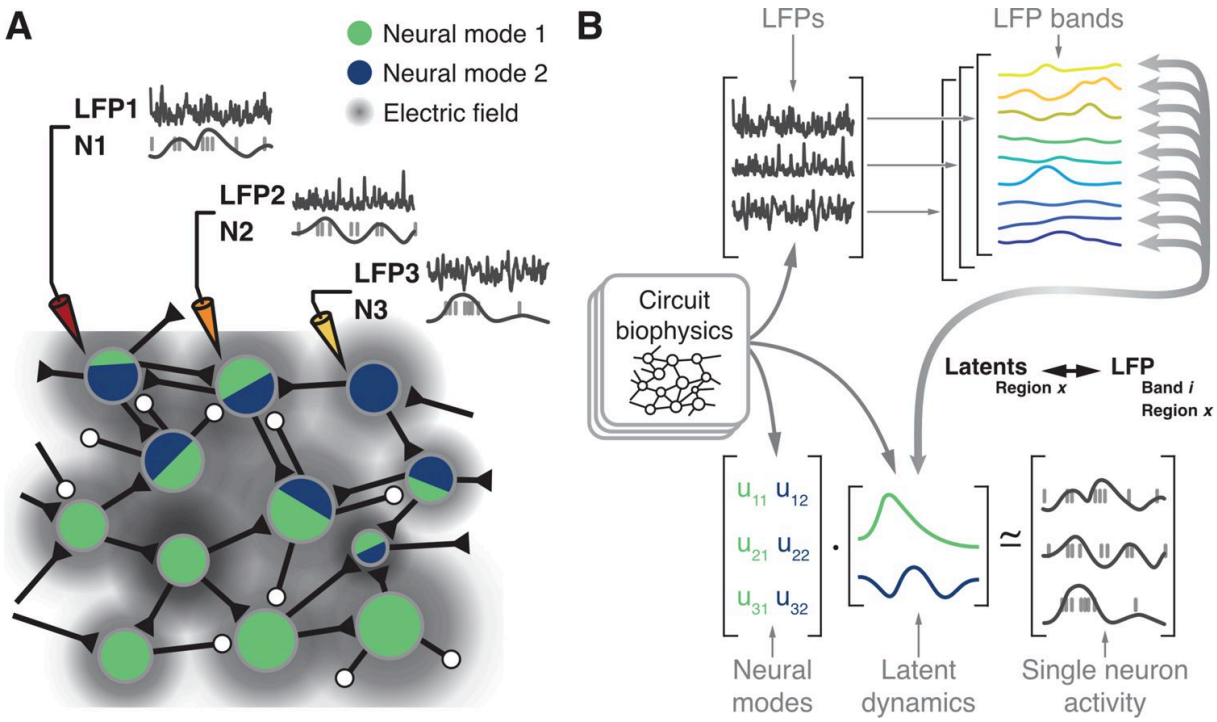

Figure 2.1: Hypothesis. (A) Microelectrode arrays record both the spiking activity of single neurons (denoted as N1, N2, N3) and the local field potentials at these same sites (denoted as LFP1, LFP2, LFP3) of a cortical region X. **(B)** Synchronisation of synaptic currents within this circuit generate LFPs in different bands (top). Circuit connectivity and other biophysical properties also constrain the coordinated latent dynamics of the N1, N2, N3 population (bottom). We thus hypothesise that, for each brain region, there will be a frequency-dependent association between LFPs and latent dynamics that should remain stable while circuit biophysics remain stable.

Monkeys C and M had dual implants in the arm regions of M1 and PMd, whereas Monkeys H and L had a single implant in the arm region of area 2. We also analysed a second dataset from the other (right) hemisphere of Monkey C, which we recorded in a different set of experiments; we denote this dataset as Monkey C_R, to distinguish it from the previous dual-region dataset from the same animal, which we denote as Monkey C_L.

We simultaneously recorded neural spiking and LFP on each electrode. The spiking signals were manually sorted to identify putative single neurons (Fig. 2.2D). To compute the neural population latent dynamics, we constructed a high-dimensional neural state space in which the smoothed firing rate of each neuron was represented on a different axis; this way, the state of a population of N neurons at a given time t corresponds to one point in an N-dimensional state space. We used principal component analysis (PCA) to find the lower-dimensional neural manifold spanning the dominant population-wide

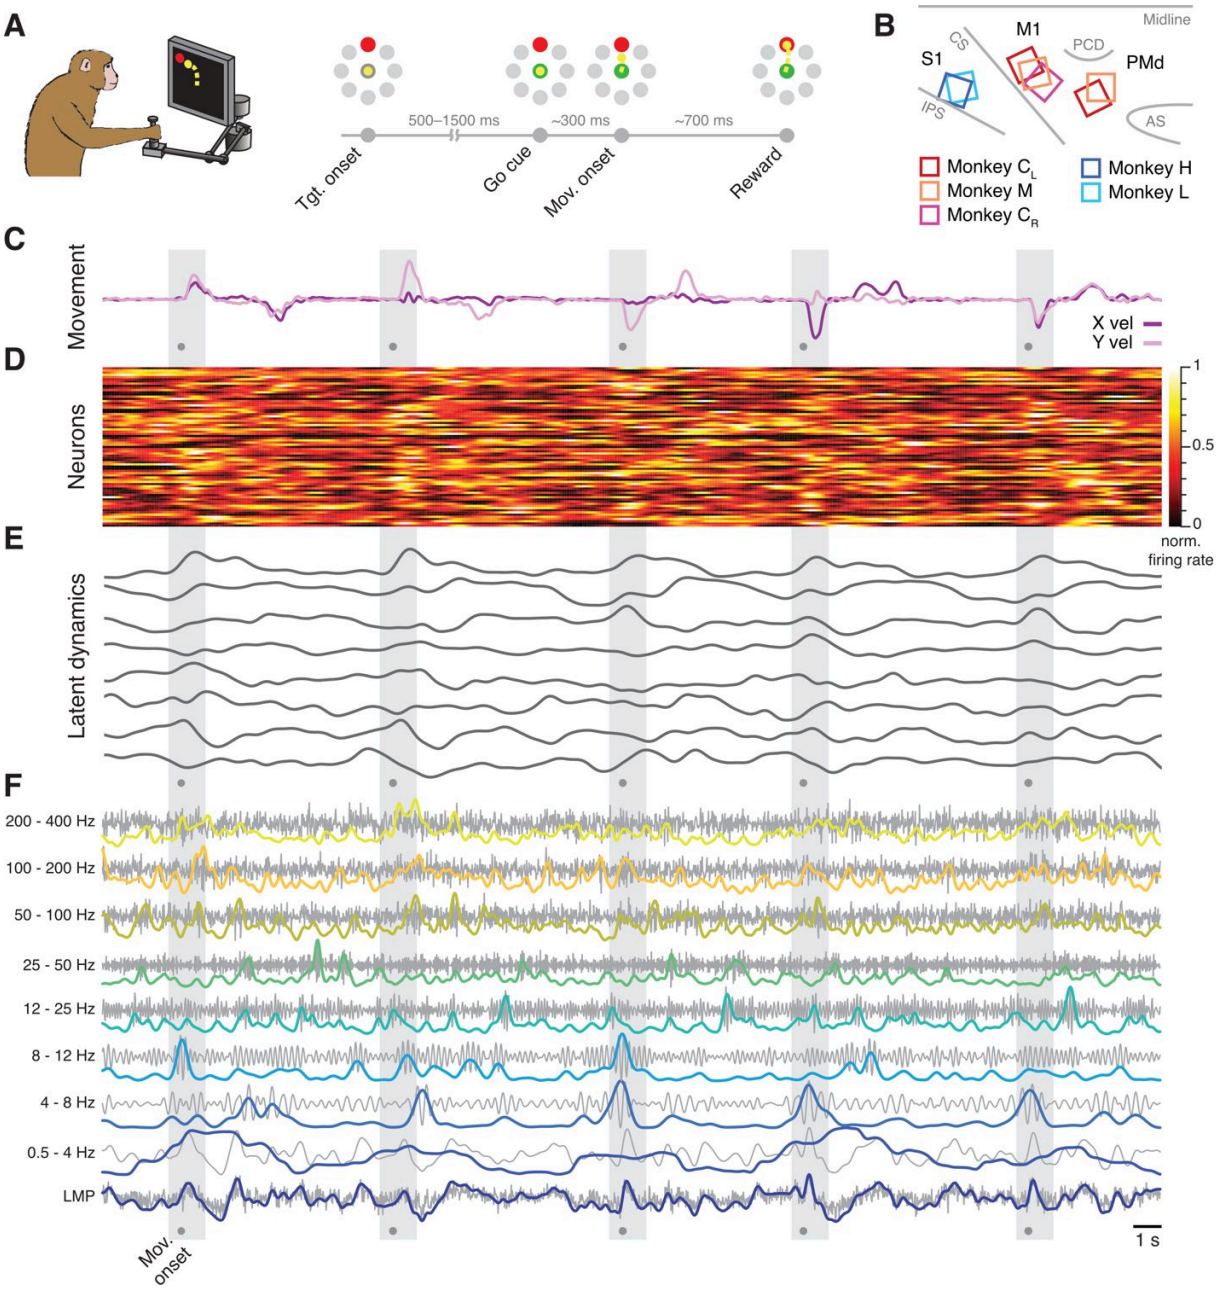

Figure 2.2: Behavioural task and neural recordings. (A) Monkeys performed a reaching task using a planar manipulandum. **(B)** Approximate locations of all seven arrays; each colour is one set of implants (legend). IPS, intraparietal sulcus; CS, central sulcus; AS, arcuate sulcus; PCD, precentral dimple. **(C)** Hand velocity along the horizontal (X) and vertical (Y) axes during five trials from a representative session from Monkey C_L. Shaded grey areas, reach epoch for each trial; dots, movement onset. **(D)** Example neural recordings showing the activity of 84 simultaneously recorded putative single neurons. **(E)** Latent dynamics corresponding to the neural firing rates in D. **(F)** Example local field potential (LFP) recordings. Each row shows the LFP activity filtered in one of the nine frequency bands we studied (grey) along with its power (colour). A, B have been adapted from Figure 2A and Figure 3A from [17].

activity patterns [7, 9] choosing the following manifold dimensionalities based on previous studies [14, 17, 101]: 10 dimensions for M1, 15 for PMd, and 8 dimensions for area 2. We computed the latent dynamics by projecting the single neuron firing rates into each of the axes (principal components) that defined a given neural manifold (Fig. 2.2E).

We calculated the LFP power in eight standard frequency bands [31, 59, 72, 112]: 0.5–4 Hz, 4–8 Hz, 8–12 Hz, 12–25 Hz, 25–50 Hz, 50–100 Hz, 100–200 Hz, and 200–400 Hz for each electrode from which we could record at least one neuron. We also computed the local motor potential (LMP), which captures fluctuations in total LFP power [31, 59] (Fig. 2.2F). For full details, see Methods section.

2.3.2 The LFP-latent dynamics correlation profiles in primary motor cortex are frequency-dependent

We first investigated the similarity between the latent dynamics in M1 and each of the LFP frequency bands defined above. We used canonical correlation analysis (CCA), a method that quantifies the similarity between two sets of signals by finding the linear transformations that maximise their correlation [17, 133, 134] (Methods); we refer to this process simply as 'alignment'. For each session, we aligned the latent dynamics with the LFP power in each frequency band (Fig. 2.3A). We performed this process separately for each recorded LFP signal, which yielded nine distributions of canonical correlation coefficients (CCs), one per frequency band, with as many samples as electrodes containing identified neurons (shown as individual data points in Fig. 2.3B).

Fig. 2.3B shows one representative session for each M1 monkey. We found that the relationship between LFPs and latent dynamics was both consistent across subjects and frequency-dependent: there were clear correlations between LFP and latent dynamics in the low and high bands (median correlation: 0.3–0.6), whereas the correlations approached zero in the mid-range bands (8–50 Hz) (coloured distributions in Fig. 2.3B; Fig. S1 summarises all datasets). Importantly, these frequency-dependent LFP-latent dynamics correlations were also stable across sessions (Fig. S1).

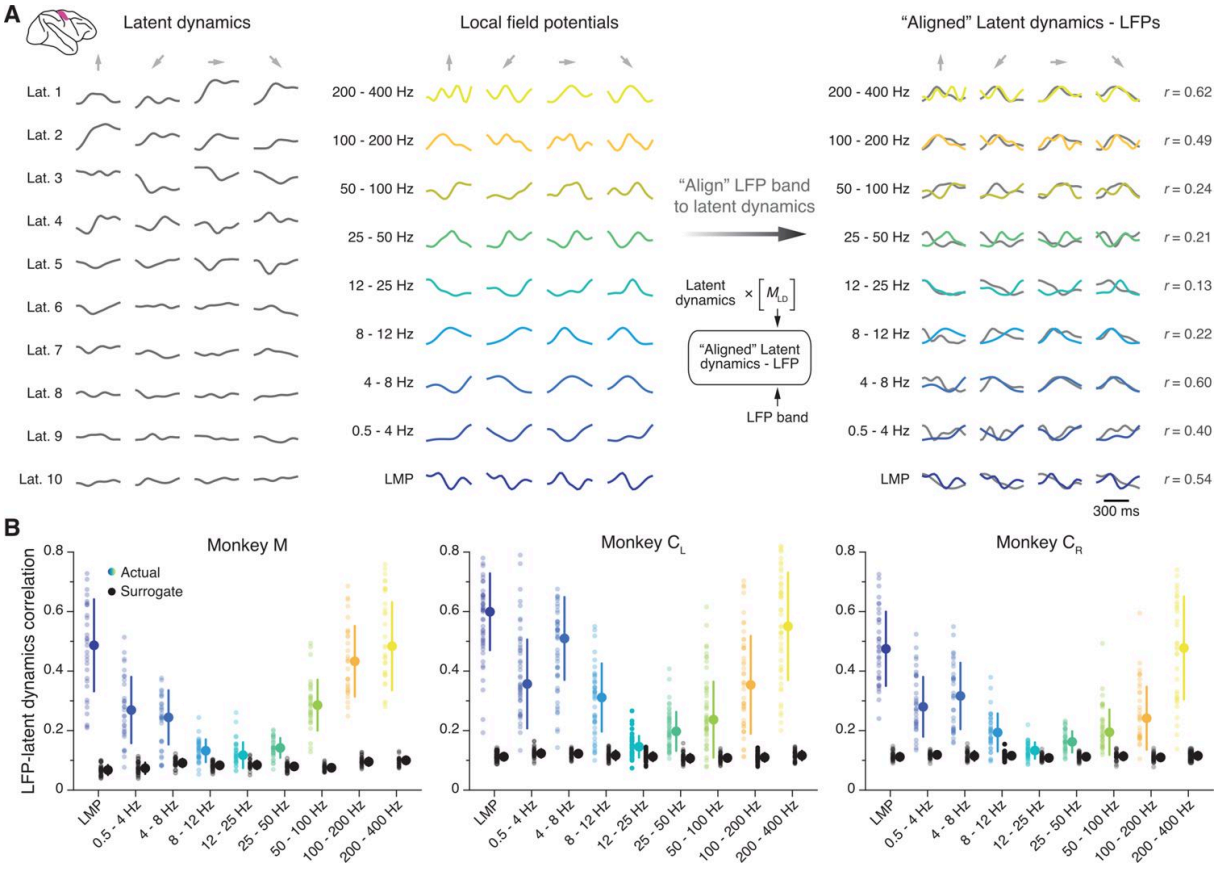

Figure 2.3: Similarity between primary motor cortex (M1) latent dynamics and each local field potential (LFP) band during movement execution. **(A)** Left: Example M1 latent dynamics during four reaches to different targets (direction indicated by the arrows above each column); figure shows top 10 dimensions of the latent dynamics. Middle: Example LFP power in each of the nine bands we study during the same three trials. Right: Canonical correlation analysis (CCA) 'alignment' finds strong similarities between certain LFP bands and the latent dynamics. r, correlation coefficient. From one session from Monkey C_L. **(B)** Correlation between each LFP band and the latent dynamics (coloured markers) during one representative session from each M1 monkey. Black markers show the control correlation values obtained after generating surrogate neural activity using tensor maximum entropy (TME); note that at low and high LFP frequencies, the actual correlations are much larger than the surrogate correlations. Error bars, median ± s.d. (n=31, 50, 34 for Monkeys M, C_L, and C_R, respectively).

We devised a control analysis to verify that the larger correlations of the low and high frequency bands captured a significant relationship between LFPs and latent dynamics. Using tensor maximum entropy (TME) [135], we generated surrogate neural firing rates that both lay on the same neural manifold – that is, preserved the covariance across neurons – and had temporal statistics similar to the actual data (examples in Fig. S2) (Methods). We reasoned that despite their spectral similarity, the correlations between the LFPs and the surrogate latent dynamics should be much lower than those between the LFPs and

the actual latent dynamics. As predicted, the surrogate correlations were significantly lower than the actual correlations (median$\leq$0.1; $p<0.001$ for all comparisons, two-sided Wilcoxon's rank sum test; compare the black and coloured distributions in Fig. 2.3B). Thus, the difference in LFP-latent dynamics correlations across LFP frequencies could not be trivially explained by similar spectral characteristics between the signals. We obtained similar LFP-latent dynamics correlation profiles for different manifold dimensionalities (Fig. S3), and when pooling all putative single neurons on an electrode into 'multi-units' (Fig. S4A). Our results also held when modifying the specific frequency bands used for LFP pre-processing within reasonable intervals (Fig. S4B), and could not be explained by differences in LFP variance across bands (Fig. S4C), or by differences in the LFP-LFP relationship across bands (Fig. S5). To confirm our results were not biased by the dimensionality reduction applied to the single neuron firing rates, we computed the 'latent LFP signals' that explained most of the variance of the multi-channel LFP recordings within each frequency band (Methods). These latent LFP signals had the same frequency-dependent relationship with the neural population latent dynamics as the full-dimensional LFP power in each channel (Fig. S6).

M1 latent dynamics [17, 100] and some LFP bands [31, 59, 72, 112] allow accurate prediction or 'decoding' of movement kinematics. We thus asked whether the strength of the LFP-latent dynamics correlation anticipates how well an LFP band predicts movement kinematics. We used standard linear decoders to predict hand velocity either from all electrodes within each LFP band or from the latent dynamics [17, 31] (Fig. 2.4A) (Methods). Decoding accuracy varied considerably across LFP bands, with some (e.g., the LMP and 200–400 Hz) being nearly as predictive as the latent dynamics (Fig. 2.4B and Fig. S7A). The accuracy of each LFP band was strongly correlated with its similarity to the latent dynamics (R^2=0.78; $p<0.001$, F-test; see Fig. 2.4D), and all of the accurate decoders made similar predictions (pairwise correlation: >0.6; Fig. S7B). However, the predictive accuracy of an LFP band did not depend on its variance, since it was similar across all bands (Fig. S4C). Predicting hand velocity from individual LFP channels was notably worse than predicting from all the channels (compare Fig. 2.4C and Fig. S7C; also see

Fig. S7D). Thus, for M1, the relationship between LFP power and latent dynamics varies across frequency bands and predicts how strongly each band relates to the ongoing motor output.

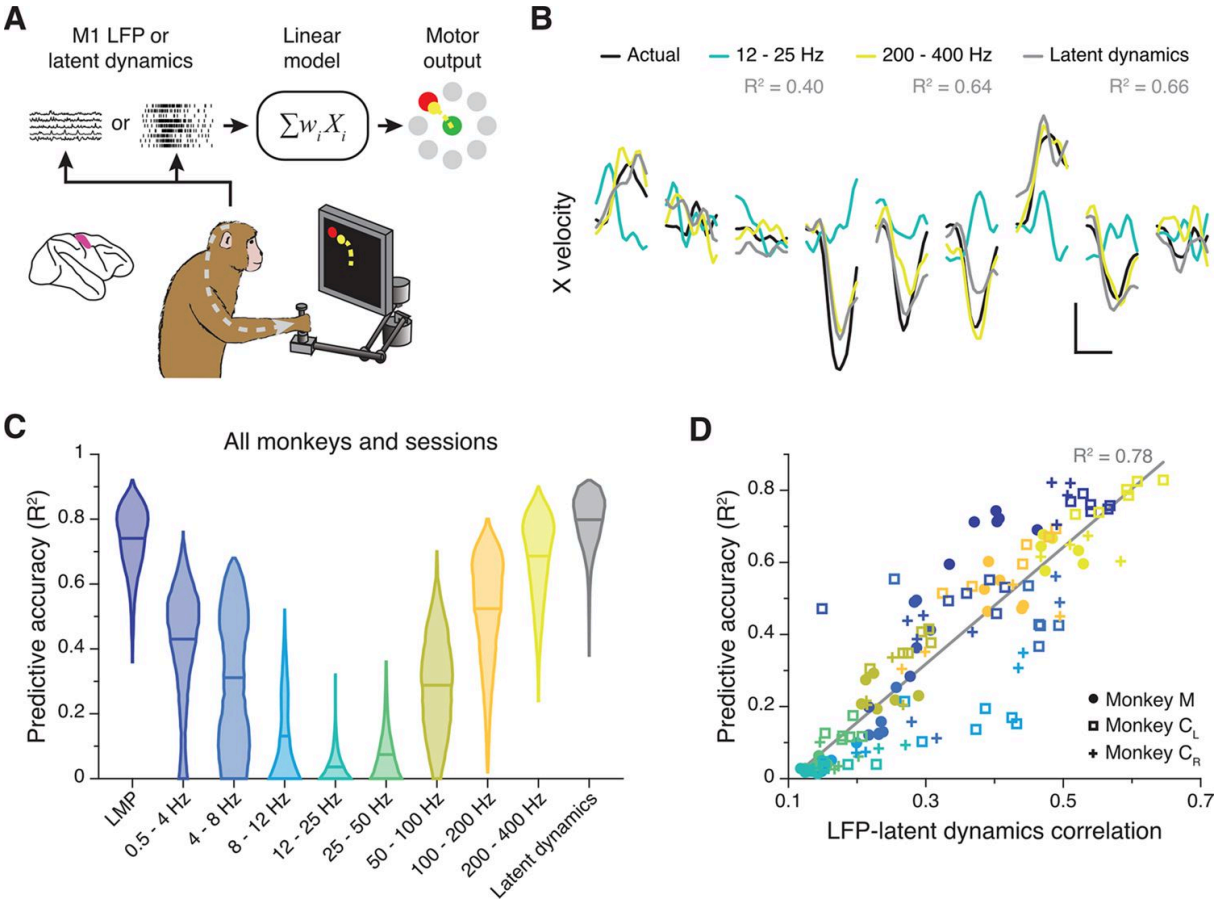

Figure 2.4: Decoding movement kinematics from primary motor cortex (M1) local field potentials (LFPs). **(A)** We trained linear decoders to predict hand velocity in the x and y axes from the either the latent dynamics or the LFP power in different frequency bands. **(B)** Example x-axis velocity predictions during nine randomly selected trials from one session from Monkey C_L. Note the clear difference in predictive accuracy between the example LFP bands (values for each input signal during the example trials shown are indicated in the legend). Scale bars: horizontal, 300 ms; vertical, 10 cm $\cdot$ s^{-1}. **(C)** Predictive accuracy of each LFP band (colours) and the latent dynamics (grey) pooled across all M1 monkeys and sessions. Violin, probability density for one frequency band (n=1600); horizontal bars, median. **(D)** Linear relationship between the M1 LFP-latent dynamics correlation coefficients and LFP predictive accuracy. Each marker denotes one LFP band from one session; different monkeys are shown using different markers (legend). Markers are colour-coded as in **C**. Grey line, linear fit to pooled data and goodness of fit (R^2).

2.3.3 The LFP-latent dynamics correlation profiles are stable between movement planning and execution

We identified a robust frequency-dependent relationship between M1 LFPs and latent dynamics during movement execution. Since M1 is also involved in movement planning [73, 127, 136], we asked whether this relationship remains stable between these two processes underlying behaviour. Repeating the previous CCA alignment procedure for the instructed delay epoch showed that the LFP-latent dynamics correlation profile was virtually identical between movement planning and execution (Fig. 2.5A): a correlation analysis between all pairs of LFP-latent dynamics correlations across all frequency bands, sessions, and monkeys revealed a strong linear association (R^2=0.87; p<0.001, F-test; Fig. 2.5B).

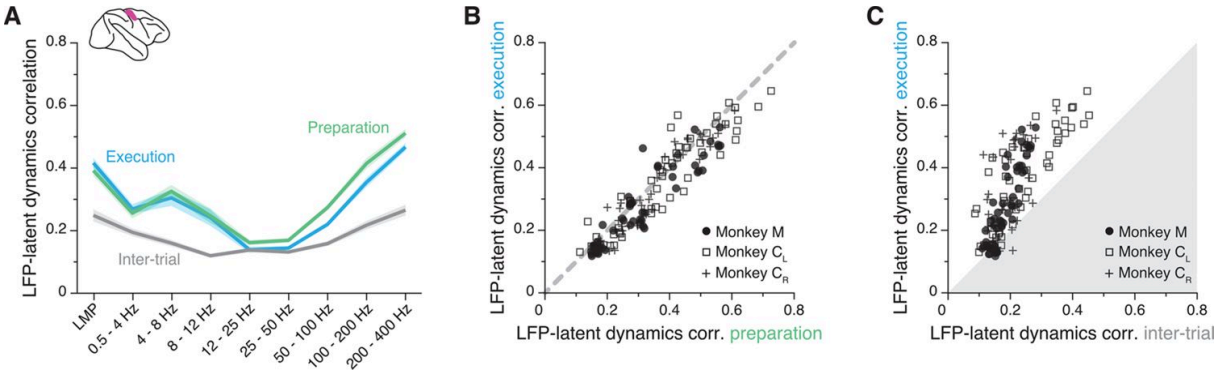

Figure 2.5: The primary motor cortex (M1) local field potential (LFP)-latent dynamics correlation profile is preserved between movement planning and execution. **(A)** LFP-latent dynamics correlation profiles across all sessions from all three M1 monkeys. Line and shaded areas, mean ± s.e.m. across all session medians (n=16). Each epoch is shown in a different colour. **(B)** Comparison between the LFP-latent dynamics correlations during movement preparation and execution. Each marker shows one frequency band for one session; each monkey is represented using a different marker (legend). Note the very strong similarity between epochs. **(C)** Comparison between the LFP-latent dynamics correlations during movement execution and the inter-trial period. Data formatted as in **B**. Note the marked decrease in correlations during the inter-trial period.

To provide a reference for the observed correlations, we repeated the alignment procedure during the inter-trial interval, when monkeys were not actively planning or executing the behaviour. In this period, M1 is not actively engaged in movement generation so neurons typically fire less (Fig. S8A,B) and become less correlated (Fig. S8C). This is due in large part to the lack of strong synaptic currents driving the population. Since

synaptic currents are the main contributors to LFP generation [4], we predicted that the LFP-latent dynamics correlations should decrease during the inter-trial interval when M1 is less actively engaged in producing behaviour. As expected, the correlations indeed became much lower during the inter-trial intervals than during movement execution (Fig. 2.5A and C).

The observation that the LFP-latent dynamics correlation profile remains preserved between movement planning and execution extends our previous findings to a more abstract process occurring in the absence of movement. It also indicates that the observed association is not a trivial epiphenomenal consequence of the frequency content of movement, but instead likely reflects underlying physiological processes related to the production of behaviourally relevant neural activity. The decrease in LFP-latent dynamics correlations during the inter-trial intervals provides additional support for this interpretation. Even though the exact circuit architecture is likely stable between these periods, M1 is not actively engaged in generating behaviour. As expected, the absence of synaptic inputs and the resultant reduction in LFP and neural population activity during these intervals in which monkeys were not engaged in the task greatly reduced the LFP-latent dynamics correlations in M1.

2.3.4 The LFP-latent dynamics correlation profiles change between primary motor and premotor cortices

The previous sections demonstrated a frequency-dependent relationship between the LFPs and latent dynamics in M1. We next asked whether similar associations could be found in other sensorimotor cortical regions. We studied PMd, a 'higher' motor region with a different cytoarchitecture from M1 [70] thought to be key for planning an action [73, 129, 130, 137, 138]. We repeated the comparison between latent dynamics and LFP bands first focusing on the movement planning epoch (Fig. 2.6A). As for M1, the LFP-latent dynamics correlations were markedly frequency-dependent (Fig. 2.6B), much stronger in the low and high frequencies, for which they greatly exceeded the surrogate control (two-sided Wilcoxon's rank sum test, $p<0.001$ for all comparisons). Interestingly, virtually

all the datasets (10 of 12; pooled results in Fig. S9A) exhibited a strong correlation at 12–25 Hz, which was notably absent in M1 despite the two regions having similar power spectra (Fig. S10A), as well as power spectral density (Fig. S10B) and variance (Fig. S10C) in each LFP band.

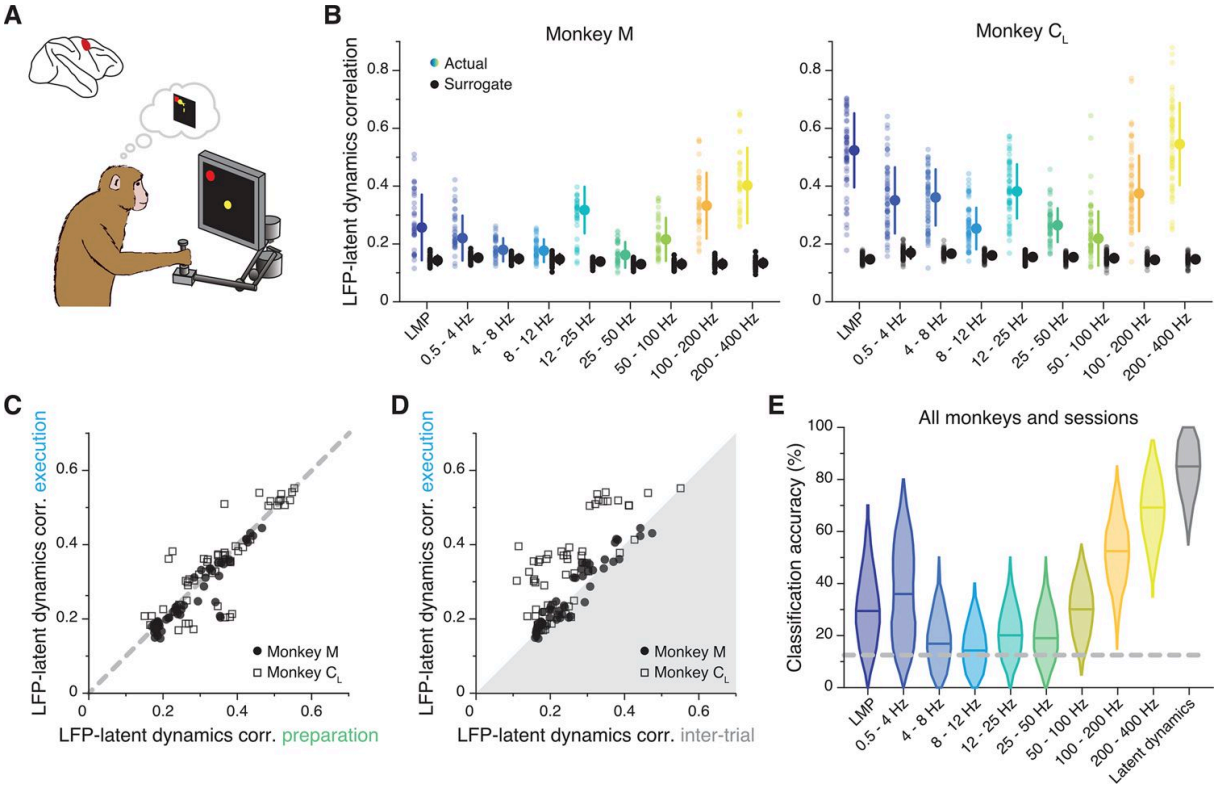

Figure 2.6: Similarity between dorsal premotor cortex (PMd) latent dynamics and each local field potential (LFP) band during movement preparation. (A) We focus on the instructed delay period following target presentation and preceding the subsequent go cue. **(B)** Correlation between each LFP band and the latent dynamics (coloured markers) during one representative session from each PMd monkey. Black markers show the control correlation values obtained after generating surrogate neural activity using tensor maximum entropy (TME). Error bars, median ± s.d. (n=40, 63 for Monkeys M and C_L, respectively). **(C)** Comparison between the LFP-latent dynamics correlations during movement preparation and execution. Each marker shows one frequency band for one session; each monkey is represented using a different marker (legend). Note again the very strong similarity between epochs. **(D)** Comparison between the LFP-latent dynamics correlations during movement execution and the inter-trial period. Data formatted as in **C**. **(E)** Accuracy of classifiers that predict the reach direction based on each LFP band (colours) and the latent dynamics (grey); data pooled across all PMd monkeys and sessions. Violin, probability density for one frequency band (n=600); horizontal bars, median; dashed horizontal line, chance level.

Similar to M1, PMd is involved in both movement preparation and execution. Thus, we tested whether the LFP-latent dynamics correlation profile was preserved as in M1, and

found it to be stable between movement planning and execution (R^2=0.79; p<0.001, F-test; Fig. 2.6C; Fig. S10D, E show that LFP power remains constant across these two epochs too). During the inter-trial period when PMd neurons become less active, the correlations dropped, though less markedly than for M1 (Fig. 2.6D). Therefore, the frequency-dependent relationship between LFPs and latent dynamics changes across cortical regions but, within a region, remains stable across different behaviour-related processes those regions are involved in.

Last, we studied whether the LFP's ability to predict the direction of the upcoming movement [139] was also frequency-dependent. The performance of classifiers that predicted the upcoming target from the instructed delay activity [17, 128] (Methods) varied dramatically across LFP bands (Fig. 2.6E; individual examples in Fig. S9B): high-frequency bands were almost as accurate as the latent dynamics, while classifier accuracy at low frequencies was worse and quite variable across monkeys (50% accuracy for Monkey C; close to chance for Monkey M). Interestingly, not all bands that were significantly correlated with the latent dynamics also predicted the upcoming target (Fig. S9C); most notably, despite being among the most correlated bands, 12–25 Hz prediction accuracy was barely above chance.

2.3.5 The LFP-latent dynamics correlation profiles are different for area 2 of primary somatosensory cortex

While LFPs in PMd and M1 primarily relate to preparing and generating movement, oscillations in area 2 of somatosensory cortex likely reflect somatosensory feedback processing [34, 140], perhaps in combination with efference copy signals [115]. We thus asked whether the LFP-latent dynamics correlation profile in area 2 is different from that of motor regions (Fig. 2.7A). Both the low- and high-frequency LFP bands of area 2 were strongly correlated with the latent dynamics (Fig. 2.7B; two-sided Wilcoxon's rank sum test, p<0.001 for all comparisons; pooled results in Fig. S11A), similar to the profile shown above in both M1 and PMd. Intriguingly, we observed a strong, significant correlation at 12–25 Hz (and for Monkey H also at 25–50 Hz), which we previously observed in PMd

but not in M1. Again, all LFP-latent dynamics correlations dropped dramatically during the inter-trial period (Fig. 2.7C), when monkeys did not move, and area 2 neurons go practically silent [75].

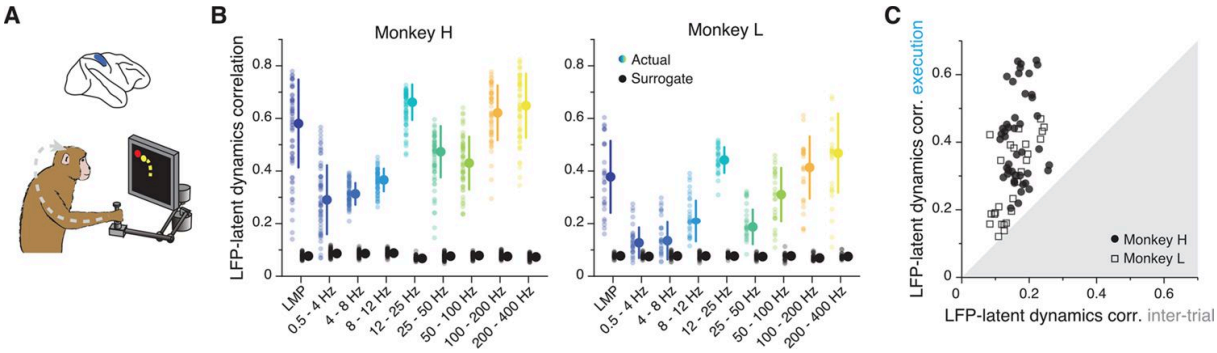

Figure 2.7: Similarity between area 2 latent dynamics and each local field potential (LFP) band during somatosensory feedback processing. **(A)** We focused on movement execution, when area 2 receives proprioceptive input about the state of the limb. **(B)** Correlation between each LFP band and the latent dynamics (coloured markers) during one representative session from each area 2 monkey; same format as Fig. 2.3B and Fig. 2.6B (n=60, 28 for Monkeys H and L, respectively). Note the clear frequency-dependent relationship between LFP and latent dynamics. Error bars, median ± s.d. **(C)** Comparison between the LFP-latent dynamics correlations during movement execution and the inter-trial period; same format as Fig. 2.5C and Fig. 2.6D. Note the marked decrease in LFP-latent dynamics correlation during the inter-trial period.

As in the other regions, we examined whether the ability to make predictions of the monkey's movement kinematics from each LFP band was correlated with the strength of its association with the latent dynamics. We again used linear decoders, in this case to predict the past state of the limb, consistent with the anticipated causality with area 2 [7, 74] (Methods). As was the case for M1, the low- and high-frequency bands provided the best predictions (Fig. S11B, C). Yet, contrary to M1 (though similar to PMd), the strength of the correlation between an LFP band and the latent dynamics did not predict how well that band could be used to decode movement (Fig. S11D). Therefore, we have found robust region-specific LFP-latent dynamics correlation profiles that remain stable for different processes underlying behaviour (Fig. 2.3B, Fig. 2.6B and Fig. 2.7B; Fig. S12), and whose correspondence to behavioural decoding (movement kinematics in M1 and area 2; reach direction in PMd) is strongest for M1.

2.3.6 LFP-latent dynamics correlations are not predicted from single neuron features

Our results indicate stable, region-specific, and frequency-dependent relationships between LFP bands and latent dynamics. However, the observed correlations could simply have arisen due to trivial relationships between LFP and single neuron spiking activity [61, 44, 62, 113, 116, 118, 119], rather than being dependent on population-wide latent dynamics. To directly address this, we asked whether LFP bands that were well correlated with the latent dynamics mostly captured the firing rates of the neurons recorded on that electrode. If that were the case, the correlations between LFPs and the firing rate of neurons on the same electrode ('within-electrode correlation') should exceed the analogous correlations between LFPs and the firing rates of neurons on different electrodes ('across-electrode correlation') (Methods). In fact, most within-electrode correlations for M1, PMd, and area 2 were very low, statistically no different from the across-electrode correlations (two-sided Wilcoxon's rank sum test, $p \geq 0.1$, except for M1 LMP ($p=0.03$) and 0.5–4 Hz ($p=0.05$); Fig. 2.8A), suggesting that the observed similarities between LFPs and latent dynamics cannot be predicted from the single neuron activity.

As a second control, we compared the LFP-latent dynamics correlation directly to the correlation between the LFP and the single neuron activity on that electrode. Notably, the LFP-single neuron correlations were always much lower than the LFP-latent dynamics correlations (Fig. 2.8B). Although there was a linear association between the two ($R^2=0.54$; $p<0.001$, F-test) the slope was very low ($\beta 1=0.15$), which indicates that even for high LFP-latent dynamics correlations, the LFP-single neuron correlations were still rather weak. The LFP-single neuron activity correlations remained low after we 'denoised' the single neuron firing rates by subtracting their projection outside the neural manifold, seeking to eliminate the components of the single neuron activity that are not shared across the population (Methods) (Fig. S13A–C). Combining all the neurons on the same electrode as 'multi-units' also did not change the results (Fig. S13D,E). These controls show that specific LFP bands capture population-wide features of the latent dynamics

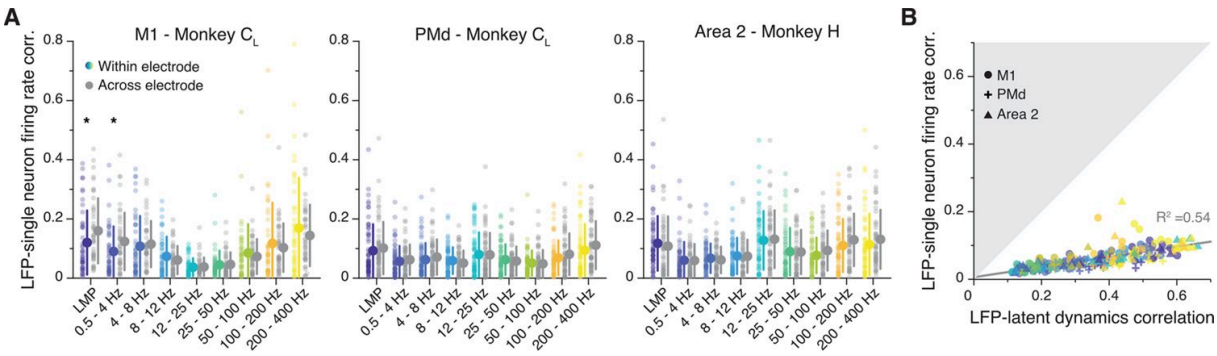

Figure 2.8: Single neuron features do not account for the local field potential (LFP)-latent dynamics correlations. (A) Comparison between the correlation of each LFP band with the firing rate of single neurons on the same electrode (coloured), and the activity of single neurons on a different electrode (grey). Each plot shows a representative session from one monkey implanted in each of the three studied regions. Error bars, median ± s.d. (n=50, 63, 60 for the primary motor cortex (M1), dorsal premotor cortex (PMd), and area 2 session, respectively). *p<0.001 two-sided Wilcoxon's rank sum test. **(B)** Comparison between the LFP-latent dynamics correlations and the correlation between the LFP and the firing rate of neurons on the same electrode. Data from all sessions and monkeys, shown using region-specific markers that indicate the session median (legend). Markers are colour-coded according to frequency as in **A**. Grey line, linear fit to pooled data and goodness of fit (R^2). Note how the LFP-latent dynamics correlations are much stronger than the correlations between LFPs and the activity of neurons on the same electrode.

that contribute to their generation in a region-specific, stable manner.

Although we have demonstrated that the reported correlation between LFPs and neural population dynamics is not a trivial consequence of the LFP reflecting single-neuron activity, characterising single-neuron properties and their contribution to the manifold directly could give further insight into the region-specific results. Across sessions, we recorded 60 ± 20 (mean ± s.d.) putative single neurons in M1, 128 ± 52 in PMd, and 47 ± 25 in area 2 (Fig. S14A), with mean firing rates of 10 ± 2 Hz, 5 ± 3 Hz, and 6 ± 1 Hz, respectively (Fig. S14B). PMd had significantly more neurons than the other two regions, while firing rates were generally higher in M1 ($p < 0.001$, two-sided Wilcoxon's rank-sum test).

The neural manifold dimensionality, that is, the number of dimensions required to explain a high percentage of the data variance, increased non-linearly with the number of sampled neurons (Fig. S14C). The dimensionality chosen for each region (10 for M1, 15 for PMd,

and 8 for area 2) captured the population activity well, although in some PMd sessions with many recorded neurons, up to 60% of the variance remained unexplained. In area 2, by contrast, where neurons were fewer and sparser, very few dimensions were needed to capture most of the variance, with 8 dimensions explaining more than 90% of the variance. Importantly, the relationship between LFPs and neural population dynamics held for as few as 5 dimensions (Fig. S3), indicating that it reflects a robust relationship already present in the top dimensions of the neural manifold.

2.4 Discussion

LFPs are intriguing neural signals: changes in power within specific LFP frequency bands correlate well with behavioural and cognitive processes such as initiating a movement [32, 33, 34, 113, 139] or choosing among different stimulus-response rules [107], yet their relationship with the activity of the neurons driving these processes has long been elusive. LFP biophysics are determined by circuit architecture [5] and the correlation among synaptic inputs to neural populations [46]. Since both of these also help define the neural covariance structure and shape the latent dynamics, we hypothesised that there should be a fundamental relationship between LFPs and latent dynamics. Here, we report, for each of three different regions of monkey sensorimotor cortex, that associations between LFPs and latent dynamics are stable and frequency-dependent (Fig. 2.3B, Fig. 2.6B and Fig. 2.7B; Fig. S12). They are also region-specific, as expected if main function and circuit biophysics play a role in determining them.

2.4.1 Relation to previous single neuron studies

Previous studies investigated the relationship between single neuron activity and LFP power within different frequency bands. Spanning a broad range of regions, from primary visual to motor cortices, they present a puzzling story in which only the medium-high frequency LFPs (typically, 40–80 Hz) are associated with the activity of single neurons or multi-units [61, 44, 62, 113, 116, 118, 119]. The lack of single neuron-LFP correlations at

low frequencies seems at odds with biophysical models predicting shared synaptic inputs – expected to drive single neuron activity – to be most strongly represented in the low-frequency LFPs [47]. Studies of M1 multielectrode LFP recordings align better with this prediction, since both the high-frequency and low-frequency LFPs allow muscle activity [31] and movement kinematics [59, 72] to be decoded accurately (Fig. 2.4). Assuming that a large component of the firing of M1 neural populations relates to the motor commands, the ability to predict behaviour from the low-frequency LFPs hints at a relationship between the two. Such a relationship has been recently modelled as a 'mode' that captures behaviour-related dynamics that are shared between the low- and high-frequency motor cortical LFPs and the neural population activity [49].

Here, we present a novel approach that explicitly quantifies the relationship between each LFP band and the population-wide latent dynamics. This allows us to show, directly, that low-frequency LFP bands are indeed correlated with the latent dynamics for each of primary motor (Fig. 2.3B), dorsal premotor (Fig. 2.6B), and somatosensory cortices (Fig. 2.7B). We also extend our investigation beyond movement execution by studying motor planning, to demonstrate that the LFP-latent dynamics correlation profiles remain stable throughout different processes of behaviour (Fig. 2.5A, B and Fig. 2.6C). This result constitutes an important control. Since there is no overt movement during the planning epoch, the robust LFP-latent dynamics correlations during the instructed delay period indicate that their association is not only driven by large common inputs reflecting motor commands: it extends to more cognitive processes such as planning an action.

Reassuringly, the LFP-latent dynamic correlations drop during the inter-trial intervals, when monkeys are not engaged in the task and the synaptic currents across the neural population presumably become much weaker, as suggested by the decrease in single neuron firing rates (Fig. S8A, B) and the disappearance of large population-wide covariation patterns (Fig. S8C). This drop suggests that the LFP-latent dynamics correlations are a consequence of strong, coordinated synaptic currents across the circuit. Interestingly, such a decrease in correlation during the inter-trial period in PMd was not as marked

as in M1, perhaps because PMd may retain some aspects of the task during this period, as recently shown for the prefrontal cortex (PFC) [141, 142]. Indeed, PMd activity but not M1 activity reflects the probability of possible upcoming reach directions [143], and it also has stronger reward-related signals than M1 [144]. Under this interpretation, we would expect 'higher' brain regions such as PFC to exhibit large LFP-latent dynamics correlations in the absence of overt behaviour, provided that the animal is deliberating or reminiscing.

2.4.2 A neural population view is necessary

A critical aspect of our results is that the stable LFP-latent dynamics correlation profiles cannot be predicted based on the relationship between LFPs and single neurons. We interpret this based on the biophysics underlying LFP generation [4, 5, 30]. Biophysical models suggest that strong LFP signals may require correlated synaptic inputs to the neurons generating the electric fields, with the spatial extent of such correlated inputs determining the magnitude of the LFP [46]. These correlated inputs will lead to correlated neural firing patterns that will be captured in the latent dynamics computed with PCA, but much less so in the discharge of any given neuron. Thus, our finding of robust LFP-latent dynamics correlations is likely driven by the biophysical mechanisms of LFP generation and requires explicitly computing the dominant neural covariance patterns, since these seem to be largely determined by shared synaptic inputs across the population. In contrast, the activity of individual neurons is determined by a high-dimensional set of inputs affecting any given neuron through unknown weights, blurring their relationship with the LFP.

For a given level of correlated synaptic input, the LFP propagates in a frequency-dependent manner, with its spatial reach decreasing quite markedly at higher frequencies [47]. Yet, high-frequency components, especially broadband signals at frequencies >50 Hz, seem also to reflect local neural spiking [44]. It is thus possible that the robust associations defining our roughly V-shaped LFP-latent dynamics correlation profiles (Fig. 2.3B, Fig. 2.6B and Fig. 2.7B) have two different sources: at low frequencies, they may be dominated by the correlated synaptic inputs leading to the emergence of low-dimensional latent dynamics

in the sensorimotor cortices during behaviour; at high frequencies, they may reflect a combination of correlated synaptic input across the population and more local neural spiking. In agreement with this, 'multi-unit' firing rates at 100–200 Hz and 200–400 Hz are more strongly correlated with the LFPs than single neuron firing rates [72, 113] (compare Fig. 2.8 and Fig. S13D), although still less so than the LFP-latent dynamics (Fig. S13E).

Unlike the narrowband rhythms typically described at lower frequencies, such as beta, activity above 80 Hz is largely broadband, lacking a clear spectral peak, and is thought to arise predominantly from the summed extracellular signature of local action potentials rather than from rhythmic synaptic input [44]. This raises a question for our results: could the strong correlation between these bands and the latent dynamics simply reflect the LFP picking up the same spikes that constitute the latent dynamics themselves? Our controls suggest not. Although within-electrode correlation between LFPs and multi-unit firing rate was significantly higher than across-electrode for 100–200 Hz and 200–400 Hz bands (Fig. S13), it was always smaller than the correlation with the neural population latent dynamics. Even the high-frequency LFP bands thus capture a population-wide signature of spiking activity, extending beyond the specific neurons recorded on that electrode, rather than a trivial read-out of nearby single-neuron activity.

2.4.3 Function as well as circuit biophysics influences the differences between cortical areas

Together with functional specialisation, the importance of circuit biophysics in determining the LFP properties may also explain the differences in LFP-latent dynamics correlations across regions. Likely due to their varying evolutionary past [66], each of the three cortical regions we studied has different neuron types, layer organisation, and local and long-range connectivity patterns [3], all of which influence the biophysics of LFP generation [4, 5, 46, 47]. Because of this, it is reasonable for the 'mapping' between latent dynamics and LFPs captured by our correlation profiles to be region-dependent.

This interpretation assumes that our recordings sample a comparable mix of cortical

layers across regions. However, all arrays used the same fixed electrode length, while cortical thickness and layer proportions differ substantially between M1, PMd, and area 2. The same electrode length can therefore sample different layers in each region. Since layer-specific circuits differ in their local connectivity and synaptic input, some of the regional differences we report could reflect this systematic difference in sampled cortical layer, rather than a difference in the region's overall circuit properties. Many of our results do not depend on across-region comparison, such as the stability of the correlation profiles across days and animals, and their preservation across movement planning and execution; these are unaffected by this concern. Comparisons between cortical areas, however, should be interpreted with this caveat in mind.

The most striking difference in LFP-latent dynamics correlation profiles across regions appears at the mid-range 12–25 Hz band. Surprisingly, there is not a clear rostro-caudal gradient in their relationship; instead, the correlation is high in PMd (Fig. 2.6B) and even higher in area 2 (Fig. 2.7B), but extremely low in the intermediately located M1 (Fig. 2.3B). What is common across all three regions is that the 12–25 Hz band is never a good predictor of the movement kinematics or directionality of the upcoming reach, even when the LFP-latent dynamics correlations reach 0.5–0.6 (Fig. 2.4C and Fig. 2.6E; Fig. S11B,C). This lack of correlation with movement parameters is consistent with studies of M1 LFPs during movement [32, 120], which describe a marked decrease in LFP 12–25 Hz power prior to movement initiation [32, 33, 34, 113], with some transient oscillations only during specific tasks. These oscillations often appear when monkeys pick up treats, explore their surroundings, or perform relatively challenging finger tasks, but are very rare during more 'automatic' wrist tasks [32, 114, 120] – the latter being most similar to the centre-out reaching task we studied. Combined, these observations suggest that the 12–25 Hz LFP band may be a by-product of population activity related to attention [120] or other high-level cognitive features such as keeping track of task structure. As such, 12–25 Hz LFP power could potentially be driven by the activity of neural populations in regions projecting to the sensorimotor cortices, rather than by the activity of local populations [45]. Indeed, LFP 12–25 Hz power in PFC modulates as monkeys select

between stimulus-response rules [107], suggesting that PMd 12–25 Hz oscillations could be correlated with task structure. Analogously, parietal cortical LFP 12–25 Hz power changes as monkeys choose between movement types based on visual cues [60], again suggesting a potential origin of area 2 12–25 Hz oscillations related to task structure.

Taken together, could the relationships we describe between LFPs, latent dynamics, and behaviour be trivial? Some relationship between these three signals was always expected, since LFPs, spiking activity, and behaviour ultimately arise from the same underlying circuit; a complete absence of any association would itself have been surprising. The real question is not whether LFPs and latent dynamics are related at all, but whether the specific, frequency-dependent, region-specific pattern reflects a genuine physiological relationship beyond what was previously reported by single-unit comparison. Several control results argue against a trivial origin. First, surrogate data that preserved the spectral content of the latent dynamics but not their temporal structure produced markedly weaker correlations with the LFPs than the actual latent dynamics did, showing that shared spectral characteristics alone cannot explain our results (Fig. 2.3B). Second, LFP power was significantly more related to the latent dynamics than to the firing of neurons recorded on the same electrode, showing that the relationship is not a simple consequence of local spike contamination of the LFP signal (Fig. 2.8A). Third, the LFP-latent dynamics correlation profile weakened substantially during inter-trial intervals, when the same circuitry is present but not actively engaged in producing behaviour, indicating that the relationship depends on ongoing circuit engagement rather than being a fixed property of the recording. Finally, the relationship between LFPs and some aspects of behaviour is not a trivial consequence of the relationship between LFPs and latent dynamics: in PMd and area 2, the 12–25 Hz band is consistently among the most strongly correlated with the latent dynamics, yet it is a poor predictor of reach direction or movement kinematics, respectively. If the LFP-latent dynamics relationship were simply a trivial reflection of shared behavioural information, the high predictive accuracy of the latent dynamics would imply an equally good prediction from the most correlated LFP bands. This dissociation instead shows that our correlation captures genuine structure in the data, rather than an

automatic mathematical consequence of computing LFPs and latent dynamics from the same neural population.

2.5 Conclusion

A wealth of studies addressing regions spanning the entire brain have reported on frequency-dependent changes in LFP power that correlate with features of behaviour. Yet, the relationship between the LFPs and the activity of neurons driving behaviour has remained elusive. Here, we show that, likely due to the biophysics of their generation, LFPs are fundamentally related to the shared patterns that dominate the activity of a neural population, rather than to the activity of the neurons themselves. While this frequency-dependent association varies across the sensorimotor cortex, for a given cortical region and frequency band it is stable throughout the behaviour.

We anticipate that our approach will uncover stable LFP-latent dynamics correlations not only within the sensorimotor cortices but beyond, as low-dimensional latent dynamics have been found across many cortical and subcortical brain regions [7, 19]. We further expect to uncover stable LFP-latent dynamic correlations during cognitive tasks, since the two are correlated during 'abstract' movement planning. Identifying these relationships offers an exciting opportunity to integrate studies based on either type of signals, including a large body of field recordings in neurological patients [145, 146, 147, 148, 149, 150, 38, 151].

2.6 Methods

2.6.1 Behavioural task

Four monkeys (males, *Macaca mulatta*) were trained to sit in a primate chair and perform a centre-out reaching task using a planar manipulandum. All monkeys completed the task with the hand contralateral to the implanted hemisphere. During the task, the monkey started a trial by bringing the cursor to a target in the centre of the workspace. After a variable waiting period, the monkey was presented with one of eight outer targets (four for

Monkey L), which were equally spaced along a circle of 6–8 cm radius. Monkeys C and M were trained to wait for the auditory go cue in the centre during a variable delay period of 0.5–1.5 s in which the target remained visible. Monkeys H and L were not subjected to this delay period. To receive a liquid reward, the monkey had to move the cursor into the outer target within 1 s; Monkeys M and C were required to hold the cursor there for 0.5 s, whereas Monkeys H and L were only required to hold for 0.1 s, to ensure they had decelerated and ended the trial within the target. To start a new trial, the monkeys had to return the cursor to the central target. During the task, the endpoint position of the manipulandum was recorded at a sampling frequency of 1 kHz; the timing of the task events was digitally logged. Hand velocity was computed as the derivative of the hand position. We analysed the following number of sessions for each monkey: Monkey C_L: 6; Monkey M: 6; Monkey C_R: 4; Monkey H: 5; Monkey L: 3. In all the analyses, we only considered successful trials (an average of 307±221 trials per session; mean ± s.d.).

2.6.2 Neural implants and recordings

All surgical and experimental procedures were approved by the Institutional Animal Care and Use Committee (IACUC) of Northwestern University under protocol #IS00000367. We recorded the spiking activity of neural populations and intracortical LFPs using chronically implanted 96-channel Utah electrode arrays in different regions of cortex. Monkey M was implanted in the M1 and PMd of the right hemisphere. Monkey C was implanted two times: first, he received an array in the right M1 (denoted by C_R throughout the text) and then, in a second procedure, he received implants in the left M1 and PMd (denoted by C_L). Monkeys H and L were implanted in area 2 of primary somatosensory cortex of the left hemisphere.

Neural activity was acquired using a Cerebus system (Blackrock Microsystems, Salt Lake City, UT) at 30 kHz sampling frequency. The recordings on each channel were band-pass filtered (250–5000 Hz), and then converted to spike times based on threshold crossings. The threshold was selected according to the root mean squared (RMS) activity in each channel (Monkeys C and M: 5.5×RMS; Monkeys H and L: 5×RMS). The spike times

were later sorted to identify putative neurons using specialised software (Offline Sorter v3, Plexon, Inc, Dallas, TX). The number of identified neurons varied across cortical area (M1: 60±20; PMd: 128±52; S1: 47±25; all mean ± s.d.). We simultaneously recorded the LFP on each channel at a sampling frequency of 1–2 kHz.

2.6.3 Behavioural epochs

For each trial, we isolated different epochs in the task to identify different 'aspects' of the behaviour, namely movement preparation and execution. These windows were adjusted to the behavioural idiosyncrasies of each monkey; however, our results did not change qualitatively within reasonable changes to these windows. For Monkeys C and M, we used a movement execution window starting 120 ms before movement onset and ending 420 ms after onset. For Monkeys H and L, the movement execution window started at the go cue, and finished 600 ms after it. For the monkeys subjected to an instructed delay period (Monkeys C and M), we defined a movement preparation window starting 390 ms before movement onset and finishing 60 ms after movement onset. Finally, we also examined neural activity during the inter-trial period – when monkeys did not move the manipulandum – to study the relationship between neural population activity and LFP when monkeys are not engaged in a behaviour using a window starting 540 ms before target presentation and finishing at target presentation.

2.6.4 LFP processing

LFP signals were processed following a fixed pipeline, applied identically to all channels and sessions. The LFP signals were first de-referenced by replaced subtracting computing their common average across channels, and then centred at zero by subtracting their average over time for the entire session. Power line interference was removed using a notch filter (zero-phase Butterworth, second order, $f_c = 60$ Hz). These pre-processed LFPs were then band-pass filtered (zero-phase Butterworth, third order) at the following frequency bands: 0.5–4 Hz, 4–8 Hz, 8–12 Hz, 12–25 Hz, 25–50 Hz, 50–100 Hz, 100–200 Hz, and 200–400 Hz. The band power of the LFP filtered at each frequency band was

computed as the mean squared amplitude using an overlapping 50 ms window and 10 ms step. Since a 50 ms window does not cover the whole period of the frequency bands with minimum frequency less than 20 Hz, the band power at those frequency bands was computed using a window size of $1/f_{min}$ seconds, maintaining the 10 ms step. All LFP pre-processing procedures follow [59]. We also computed the so-called LMP by calculating the moving average of the LFP (de-referenced, detrended, and after removing the power line interference) using the same 50 ms window and 10 ms step as for band power. Note that our results held when we varied the LFP frequency bands within a reasonable range (Fig. S4B; details in 'Additional analyses including controls' below).

Although we recorded LFPs from all 96 array channels in the arrays, we only considered in the analyses the LFPs from channels containing at least one putative neuron clearly identified through manual spike sorting, resulting in 46±16 (mean ±s.d.; ranging from 15 to 77) LFP channels per session. We subsampled the post-processed LFP signals to 30 ms bins, to match the bin size used for the single neuron activity, and smoothed them using a Gaussian kernel (s.d.: 50 ms) as we did for the single unit firing rates (see below).

2.6.5 Neural population latent dynamics

To characterise the dynamics of the neural population activity, we first computed the smoothed firing rate for each putative single neuron by first binning the spike counts into 30 ms bins. The binned spikes were square-root transformed and then smoothed by applying a Gaussian kernel (s.d.: 50 ms). Although we excluded neurons with low firing rate (<1 Hz mean firing rate across all bins), we did not perform any additional preselection, for example, based on directional tuning. For each behavioural epoch [152] (execution or planning, when present) in each session, this produced a neural data matrix $\mathbf{X}$ of dimension n by T, where n is the number of neurons and T is the total number of time points from all concatenated trials; thus, T is the number of trials per session $\times$ the number of points per trial.

We represented the activity of the n identified neurons during an epoch in a session in

an n-dimensional neural state space, in which each axis represents the smoothed firing rate of one neuron. Within this space, we computed the low-dimensional manifold that spans the dominant neural covariation patterns [7, 9] by applying PCA on $\mathbf{X}$. We defined an m-dimensional neural manifold by keeping only the leading m principal components. Based on previous work by our group and others [14, 17, 152], the number of components m considered for each area where: m=10 for M1, m=15 for PMd, and m=8 for area 2; for those sessions with an instructed delay period, we used the same manifold dimensionality during the planning and execution epochs. Note that our results held across a wide range of dimensionality values m (Fig. S3). We computed the neural population latent dynamics by projecting the smoothed firing rates of the n neurons onto the neural manifold. This produced a data matrix $\mathbf{L}$ of dimension m by T, where m is the dimensionality of the manifold.

2.6.6 Alignment of latent dynamics and LFPs

We hypothesised that there should be an association between LFPs and latent dynamics, and that such an association would be region-specific, frequency-dependent, and stable during movement planning and execution. To address these hypotheses, we computed the similarity between the latent dynamics and the time-varying changes in LFP power in each band for each behavioural epoch, monkey, session, and cortical area. We used CCA, a method that quantifies the similarity between two sets of multidimensional signals by finding the linear combinations that, applied to each set, maximally correlates the resulting projections [133]; we refer to this process as 'aligning' an LFP band and the latent dynamics.

We started the alignment procedure by considering the concatenated latent dynamics, $\mathbf{L}$, and the concatenated LFP power in one electrode at one frequency band b, $\mathbf{F}_{b,i}$, where b is one of the nine frequency bands listed above (including the LMP), and i the electrode number; the vector $\mathbf{F}_{b,i}$ thus has dimension 1 by T. For each LFP band, we align separately the latent dynamics and the LFP power at one frequency band in each of the k electrodes with at least one putative neuron, which yields k canonical correlation coefficients, CCs, per

LFP band. Throughout the paper, we summarise these distributions based on their mean ± s.d. (see, e.g., Fig. 2.3B, Fig. 2.6B and Fig. 2.7B). For details on the implementation of CCA, see our previous work [17].

To test whether the CCs between each LFP band and the latent dynamics capture a significant association between these two types of signals, we devised a control using TME, a method that generates surrogate neural population data preserving selected features of the original data [135]. Here, we used TME to generate surrogate neural firing rates that preserved the covariance across neurons but that had different covariance across trials and over time (i.e., different dynamics) than the actual data. We then smoothed these surrogate firing rates with a Gaussian kernel (s.d.: 50 ms), and applied PCA to obtain a matrix of surrogate latent dynamics, which were very similar to the actual latent dynamics based on their spectral content (Fig. S2). To assess whether the experimentally measured CCs captured a significant association between an LFP band and the latent dynamics, we compared them to the CCs obtained after aligning the same LFP signals to the surrogate latent dynamics (shown in colour and in black, respectively, in Fig. 2.3B, Fig. 2.6B and Fig. 2.7B).

2.6.7 Decoding hand velocity from M1 and area 2 of somatosensory cortex movement activity

After quantifying to what extent each M1 and area 2 LFP band correlates with the latent dynamics, we asked how well it predicts behaviour. We computed Wiener filters [153] to predict hand velocity from each of the nine LFP bands, including the LMP, as well as from latent dynamics. The filters included three bins of recent LFP history, for a total of 90 ms. These additional inputs seek to account for the intrinsic LFP dynamics and transmission delays and were only applied in this analysis. Given that M1 activity causes movement, for this region we included additional bins preceding the hand velocity signals. In contrast, since area 2 activity is largely driven by responses to the ongoing movement, for this region we included additional bins lagging the hand velocity signals. We further

added a third-order polynomial at the output of the filter to compensate for nonlinearities [154]. Finally, to obtain accuracy values against which to compare the LFP decoders, we computed similar Wiener filters using the M1 or area 2 latent dynamics as inputs. Decoder prediction quality was measured based on the squared correlation coefficient (R^2) between the actual and predicted velocities. We built separate decoders for the X and Y hand velocities; as they had similar accuracy, we report their mean. The same approach was followed when predicting hand velocity from single channel LFPs (Fig. S7C,D).

To avoid overfitting, we used 90% of the trials in the session to train the decoder and the remaining 10% to test it. We repeated this procedure 50 times, randomly selecting the non-overlapping training and testing trials in each iteration. We averaged the R^2 value across all the test blocks to obtain the final reported value.

2.6.8 Predicting reach direction from PMd planning activity

We trained naïve Bayes classifiers [128] to predict the direction of the upcoming reach based on the pre-movement LFP activity in PMd. Similar to the velocity decoders, we built nine sets of LFP-based classifiers, each taking the LFP signals at one frequency band as inputs, as well as latent dynamics-based classifiers whose performance we used as reference. In all cases, we used PMd activity within a 450 ms window as input to the classifier. Within this window, we averaged all LFP activity (or latent dynamics) to obtain a single bin per trial, and then standardised it using a z-transform. Prediction accuracy was measured as the percentage of targets predicted correctly.

To protect against overfitting, we trained the classifier in 90% of the trials and tested the performance on the remaining 10%. As for velocity decoders, we repeated the procedure 50 times randomly selecting the non-overlapping training and testing trials in each iteration. We averaged the percentage of correct predictions to obtain the final reported value.

2.6.9 Additional analyses including controls

We performed all the analyses in the paper using the standard LFP bands described above. In addition, we also verified that our alignment results held even if the LFPs were filtered in different bands by repeating the analysis on partly overlapping bands that swept the entire LFP spectrum (Fig. S4B). We chose pre-processing frequency bands that had width similar to that of the 'classic' LFP bands, which resulted in 53 bands with increasing width, ranging from 4 to 100 Hz. Filtering windows had a 50% overlap, and the following widths: 4 Hz between 0.5 and 14.5 Hz, 12 Hz between 10.5 and 26 Hz, 25 Hz between 12.5 and 57.5 Hz, 50 Hz between 25 and 120 Hz, and 100 Hz between 50 and 440 Hz.

To compute the LFP variance at each frequency band (Fig. S4C, Fig. S10C), we first normalised the post-processed LFP signals during the epoch of interest (preparation or execution) between 0 and 1 for each band and trial separately; this ensured a fair comparison of LFP variability across frequency bands. Then, we computed the variance of these post-processed, normalised LFP signals, and calculated the LFP variance for each frequency band by pooling all the channels together.

We calculated the similarity between LFP bands by computing the pairwise correlation between the post-processed signals after concatenating all the trials from each session (Fig. S5A). We followed the same approach to compare the frequency-dependent movement decoder predictions (Fig. S7B).

In order to verify that the identified frequency-dependent associations between single LFP channel and neural population latent dynamics were not biased by applying PCA on the single neuron firing rates, we repeated the CCA on the 'latent LFP signals', which we identified by applying PCA on the post-processed multi-channel LFP recordings at each frequency band (Fig. S6).

To study the LFP frequency content during different behavioural epochs, we computed the power spectral density of the post-processed LFP signals as the power of its direct Fourier transform divided by the frequency of each point (Fig. S10A). To avoid any potential

distortions due to power line interference (60 Hz) and its harmonics that could remain after pre-processing, we interpolated linearly the calculated power spectral density between –2 Hz and +2 Hz with respect to each harmonic. We followed the same approach when comparing the frequency content of M1 and PMd signals between movement preparation and execution (Fig. S10D). To directly compare the power spectral density across regions (Fig. S10B) or behavioural epochs (Fig. S10E), we computed the mean of the power spectral density for each band.

To visually compare the LFP-latent dynamic associations across regions and behavioural epochs, we calculated for each session the mean canonical correlation per frequency band, and then averaged them per region (Fig. 2.5A shows all three epochs for M1, and Fig. S12 summarises our results for all three sensorimotor cortical regions).

We devised two control analyses to test whether the similarities between LFP bands and latent dynamics could be explained by features of the single neuron activity. In the first control, we verified that the LFP on an electrode does not strongly reflect the activity of single neurons in that electrode. Thus, we verified that the relationship between LFP bands and single neurons on the same electrode did not exceed that between LFP bands and single neurons on different electrodes. We computed the Pearson correlation between the activity of each putative neuron and the power of each LFP band on that same electrode; this yielded nine distributions of n correlation values. We then computed the Pearson correlation between each LFP band and the activity of a randomly selected neuron on a different electrode; this again yielded nine distributions of n correlation values. Note that neurons detected on neighbouring electrodes were excluded to make sure the distance in the across-electrode group was not small. Direct comparisons between these paired sets of distributions allowed us to show that single neuron activity does not explain away our results (Fig. 2.8A). As a second control, we compared directly, for each LFP band, the correlation between the LFP and the single neuron activity to the CC between the LFP and the latent dynamics (Fig. 2.8B).

We repeated these controls after summing up the spikes of all putative single neurons on

an electrode to form a 'multi-unit' (Fig. S13D,E), and after 'denoising' the single neuron firing rates based on PCA (Fig. S13A–C). This denoising consisted in subtracting from the smoothed firing rate of each neuron its projection outside the neural manifold, that is, the projections onto principal component (dimensions) m+1 to n. Notably, none of these manipulations changed the results qualitatively.

2.6.10 Statistics

We applied statistical tests to compare results for the actual LFP-latent dynamics correlations and the surrogate LFP-latent dynamics correlations. In addition, we tested whether each of the LFP bands was more similar to the firing rates of each sorted single neuron or to the multi-unit activity on each electrode. For all these comparisons, we used a two-sided Wilcoxon's rank-sum test to compare the distributions, which were not necessarily assumed to be normal. When examining linear relationships between variables, model fits were compared to a null model in which only the intercept was different from zero using an F-test. Throughout all analyses we used a significance threshold of $p<0.001$. Sample sizes are reported in the corresponding figure captions; when a single number is reported, it refers to all of the distributions plotted in that panel. No statistical methods were used to predetermine sample sizes, although our dataset included a large number of sessions and subjects, many more than typical studies in the field. The data analysis supporting our conclusions required no experimental intervention, so no randomisation into groups was needed and the experimenters were not blinded to the nature and goals of the experiment.

Supplementary figures

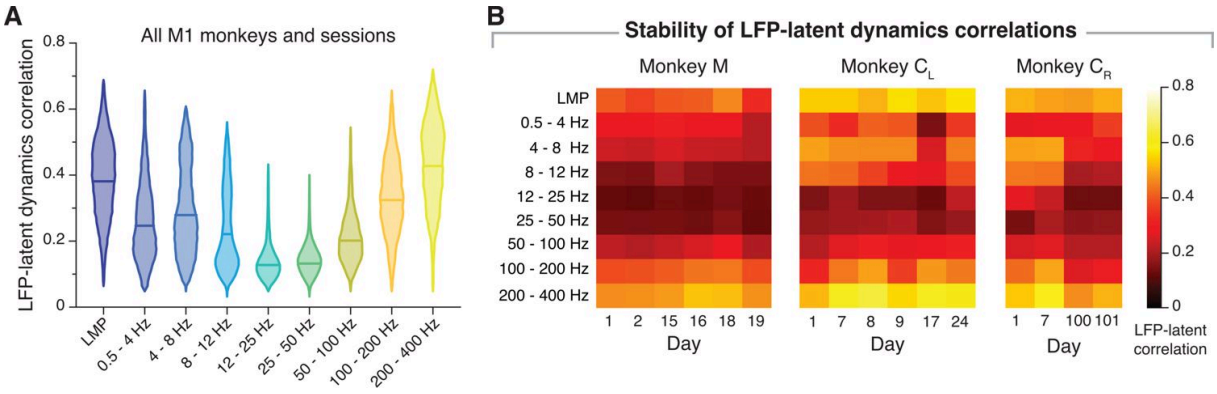

Figure S1: Similarity between primary motor cortex (M1) latent dynamics and each local field potential (LFP) band. **(A)** Correlation between each LFP band and the latent dynamics; data pooled across all M1 sessions and monkeys. Note that the correlation profile matches well the example sessions in Fig. 2.3. Violin, probability density for one frequency band (n=606); horizontal line, median. **(B)** Stability of the LFP-latent dynamics correlation profiles across different sessions from each monkey. Coloured squares represent the mean value across all LFP channels for each frequency band. Note that the frequency-dependent correlations are very similar across sessions; the only noticeable change is a shift in magnitude between days 7 and 100 for Monkey C_R.

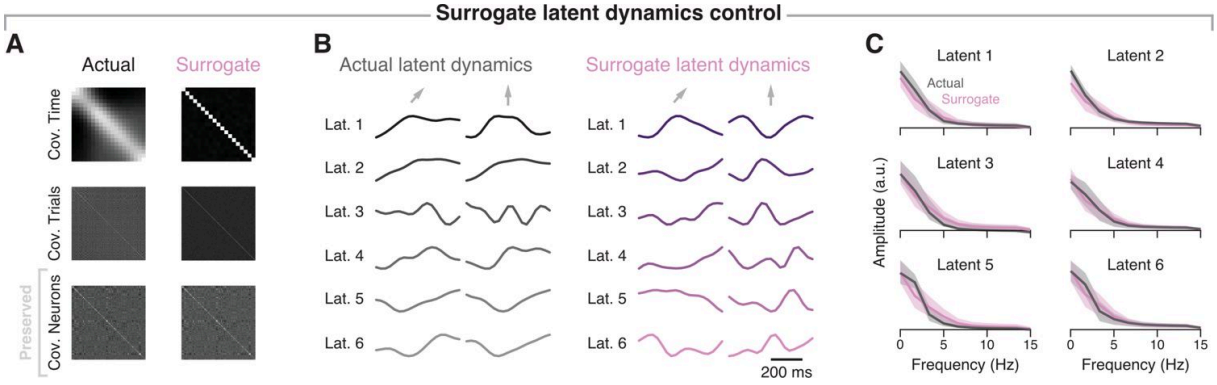

Figure S2: Surrogate latent dynamics control. **(A)** Covariance over time, across trials and across neurons for the example session from Monkey C_L shown in Fig. 2.3 ('Actual'), compared to the same covariance matrices from surrogate data that preserves the covariance across neurons ('Surrogate') that is, the neural manifold. **(B)** Example actual latent dynamics and surrogate latent dynamics generated using tensor maximum entropy (TME) for two randomly selected trials from the same example session. Arrows, target direction. Only six of ten dimensions are shown for visualisation purposes. **(C)** Amplitude spectrum of the actual and surrogate latent dynamics for the same session. Line and shaded areas: mean ± s.d. across all the trials in the session (n=286). Note the strong similarity in spectral content between actual and surrogate latent dynamics.

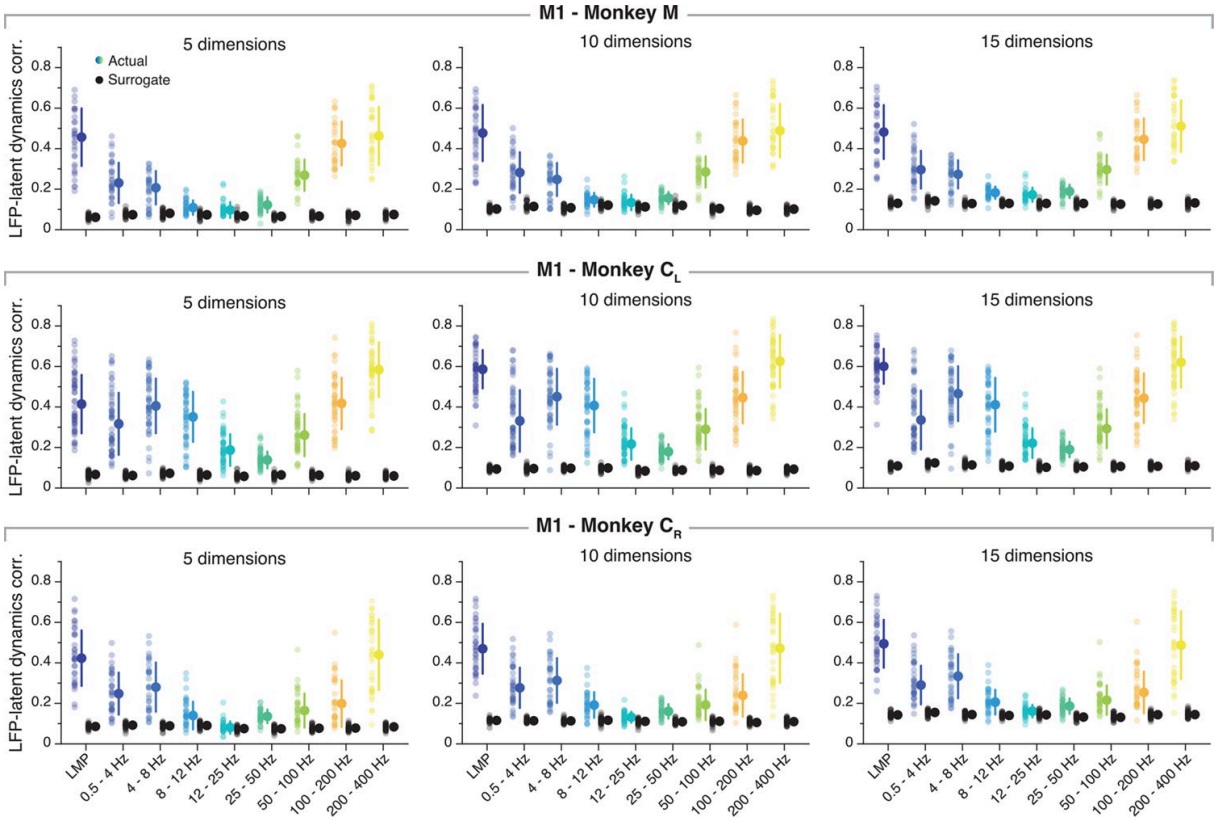

Figure S3: The local field potential (LFP)-latent dynamics correlation profiles are preserved when changing neural manifold dimensionality. Similarity between primary motor cortex (M1) latent dynamics and each LFP band for three different manifold dimensionalities. The plots show the same representative session from each M1 monkey shown in Fig. 2.3 using the same format. Note that the trend in the LFP-latent dynamics correlation profile does not depend on the assumed dimensionality of the neural manifold. For all comparisons, two-sided Wilcoxon's rank sum test, $p<0.001$ except when comparing the 12–25 Hz LFP band to the latent dynamics in a five-dimensional manifold for Monkey C_R ($p=0.037$).

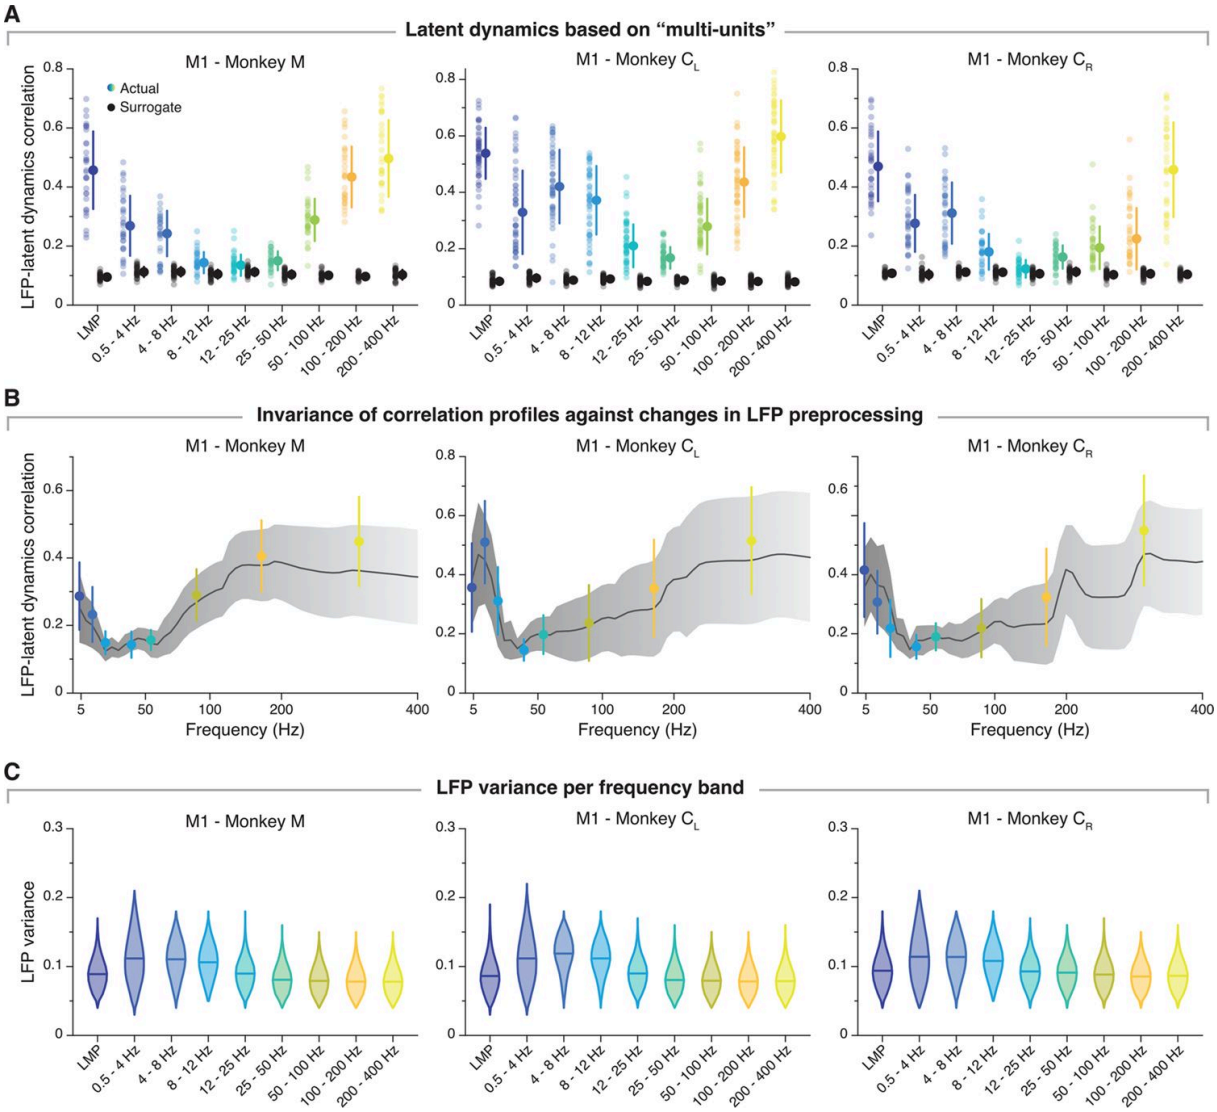
A
Latent dynamics based on "multi-units"
M1 - Monkey M
M1 - Monkey C_L
M1 - Monkey C_R
Actual
Surrogate
LFP-latent dynamics correlation
LMP
0.5 - 4 Hz
4 - 8 Hz
8 - 12 Hz
12 - 25 Hz
25 - 50 Hz
50 - 100 Hz
100 - 200 Hz
200 - 400 Hz
B
Invariance of correlation profiles against changes in LFP preprocessing
Frequency (Hz)
C
LFP variance per frequency band
LFP variance

Figure S4

Figure S4: Further control analyses to illustrate the frequency-dependent relationship between primary motor cortex (M1) latent dynamics and each local field potential (LFP) band. (A) Correlation between each LFP band and the latent dynamics (coloured markers) after pooling all single neurons on each electrode as a 'multi-unit'. Data for the same representative sessions from each M1 monkey shown in Fig. 2.3 and Fig. S3. Black markers show the control correlation values obtained after generating surrogate neural activity using tensor maximum entropy (TME). Note that pooling single neurons into 'multi-units' does not alter the results: all the trends are the same as in Fig. 2.3B. Error bars, median ± s.d. (n=31, 50, 34 for Monkeys M, C_L, and C_R, respectively). For all comparisons, two-sided Wilcoxon's rank sum test, $p<0.001$ except when comparing the 12–25 Hz LFP band to the latent dynamics for Monkey C_R ($p=0.039$). **(B)** The latent dynamics-LFP correlation profile does not depend critically on the frequency bands used for LFP pre-processing; it follows the same trend if the LFP power is computed on partly overlapping bands that sweep the entire signal spectrum (see Methods). The example sessions from each M1 monkey are the same as in Fig. 2.3B (the coloured error bars are reproduced from that same figure). Line and shaded areas, mean ± s.d. (n=31, 50, 34 for Monkeys M, C_L, and C_R, respectively). **(C)** LFP power variance in each frequency band during the same example session from each M1 monkey shown throughout the paper. LFP variance is quite similar across bands, and the bands with highest correlations with the latent dynamics (local motor potential [LMP], 100–200 Hz, 200–400 Hz) are among those with lowest variance. Violin, probability density for one frequency band (n=677, 14,300, 6086) for Monkeys M, C_L, and C_R, respectively; horizontal line, median.

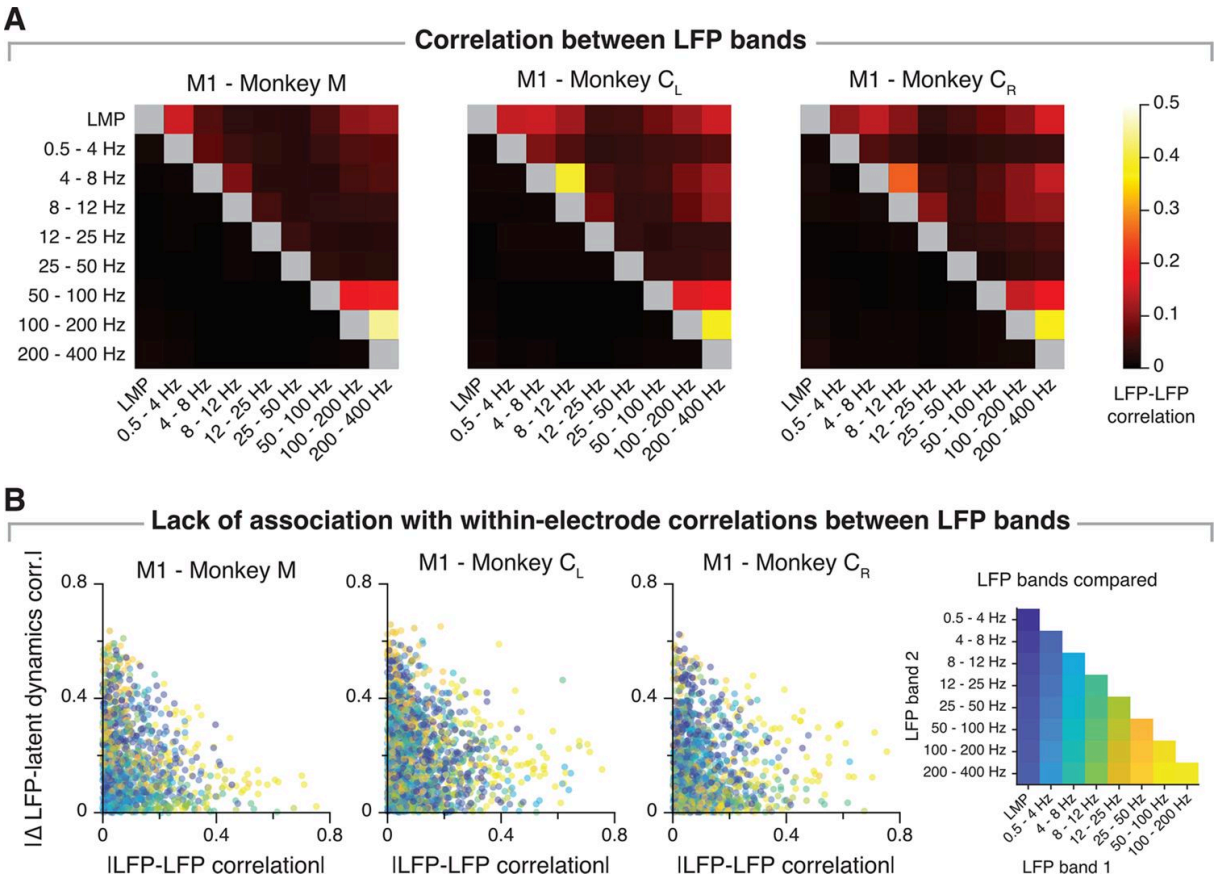

Figure S5: The relationship between local field potential (LFP) bands does not predict their association with the latent dynamics. (A) Correlation between LFP bands for each of the M1 monkeys. Colour bar, correlation magnitude; upper triangle, mean correlation; bottom triangle, s.e.m. (n=31, 50, 34). Note that the cross-band correlations are quite low, even for bands that are strongly correlated with the latent dynamics such as the local motor potential (LMP) and 200–400 Hz. **(B)** Left: The correlation between LFP bands within the same electrode does not predict the correlation of each of these bands with the neural population latent dynamics. Data for all sessions from each of the three M1 monkeys. Right: Legend, pair of bands being compared.

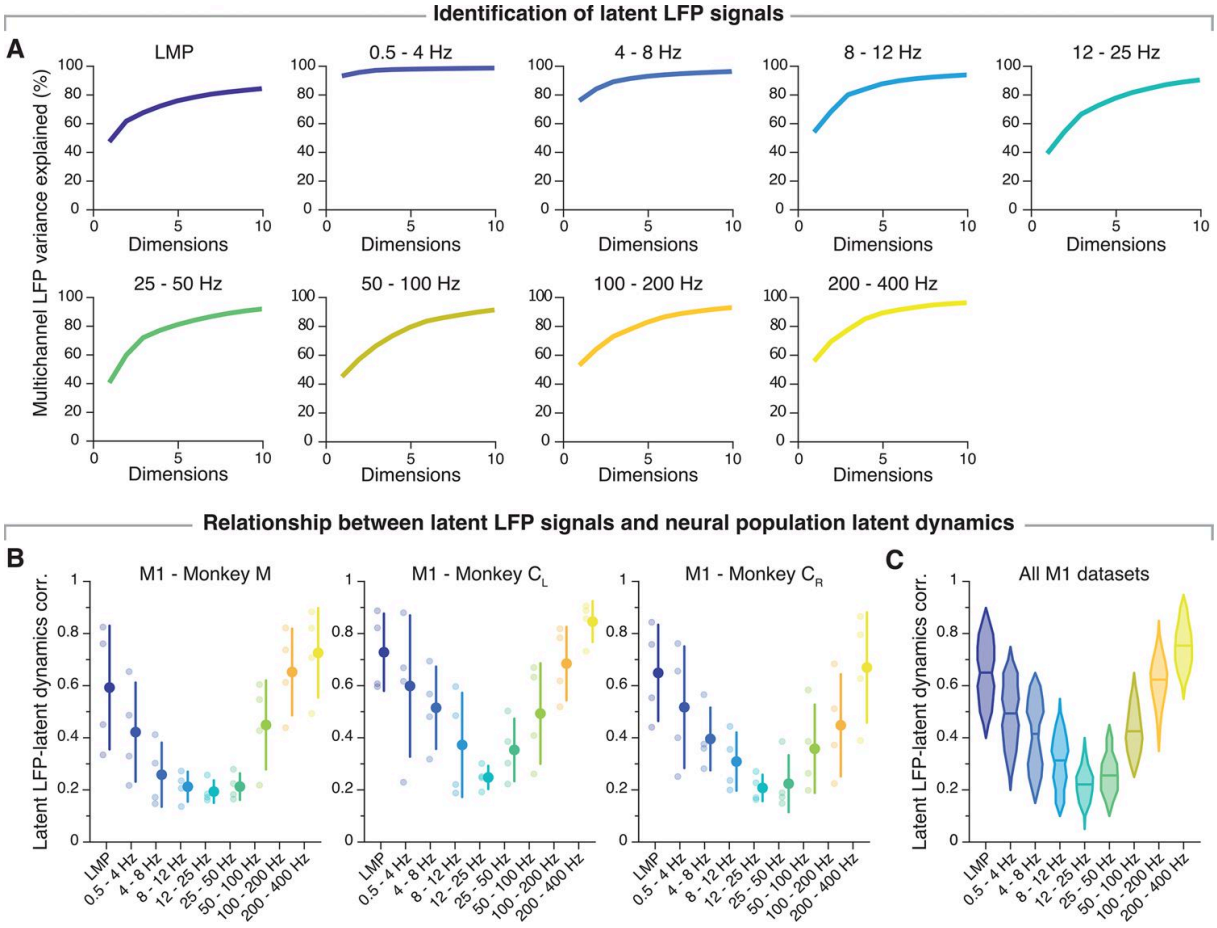

Figure S6: The latent local field potential (LFP) signals have a similar association with the neural population latent dynamics as the individual LFP signals do. (A) Percentage multichannel LFP variance explained as number of LFP dimensions (principal components) for each frequency band. Data from the same example session from Monkey C_L shown throughout the paper. Note that the apparent dimensionality does not predict the correlation with the neural population latent dynamics: frequency bands with similar eigenvalue distributions, such as the 25–50 Hz and 200–400 Hz bands, have very different correlations with the latent dynamics. **(B)** Frequency-dependent association between the latent LFP signals and the neural population latent dynamics. Markers, individual sessions; error bars, median ± s.d. (n=4 for each monkey). **(C)** Correlation between the frequency-specific latent LFP signals and the neural population latent dynamics; data pooled across all primary motor cortex (M1) sessions and monkeys. Note that the correlation profile matches well that obtained when using the individual LFP signals rather than the latent LFP signals (compare with Fig. S1). Violin, probability density for one frequency band (n=64); horizontal line, median.

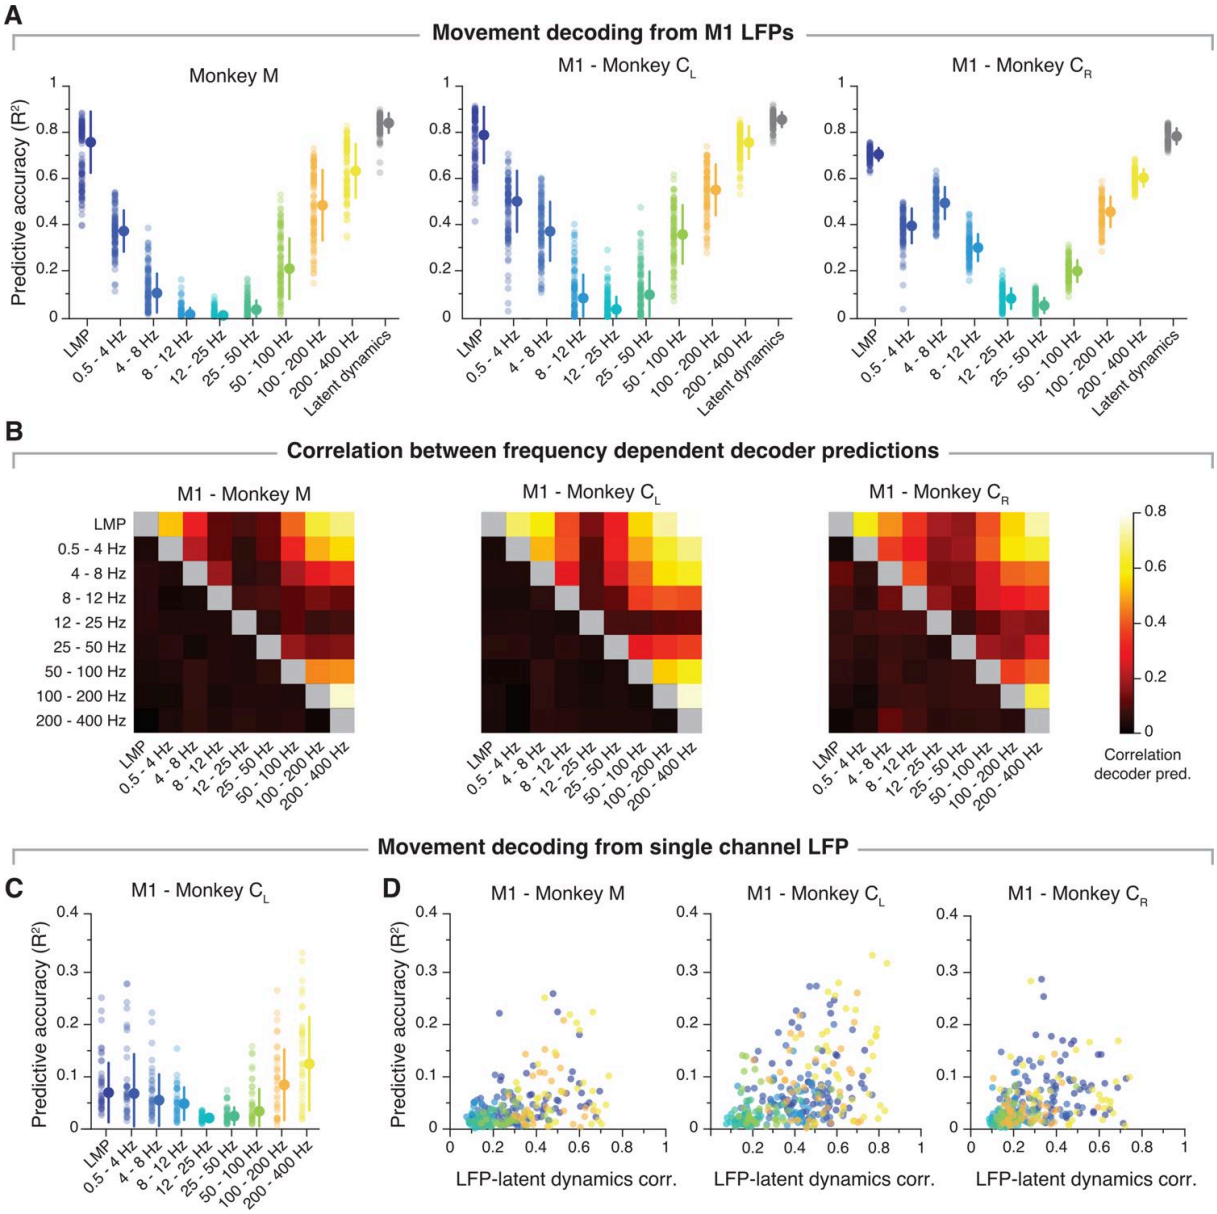

Figure S7: Additional analyses to illustrate movement decoding from primary motor cortex (M1) local field potential (LFP). (A) Predictive accuracy of each LFP band (colours) and the latent dynamics (grey). Error bars, median ± s.d. (n=100). **(B)** Correlation between hand velocity predictions obtained from decoders trained on each LFP band. Colour bar, correlation magnitude; upper triangle, mean correlation; bottom triangle, s.e.m. (n=100). **(C)** Movement decoding from single channel LFP for the same example session from Monkey C_L used throughout the paper. Markers, individual channels; error bars, median ± s.d. (n=50). **(D)** Comparison between the predictive accuracy of band-specific velocity decoders trained on single channel LFP (as in Panel **C**) and that band's correlation with the neural population latent dynamics. Note the lack of a very strong association between the two. Data for the same representative sessions from each M1 monkey shown in Fig. 2.3.

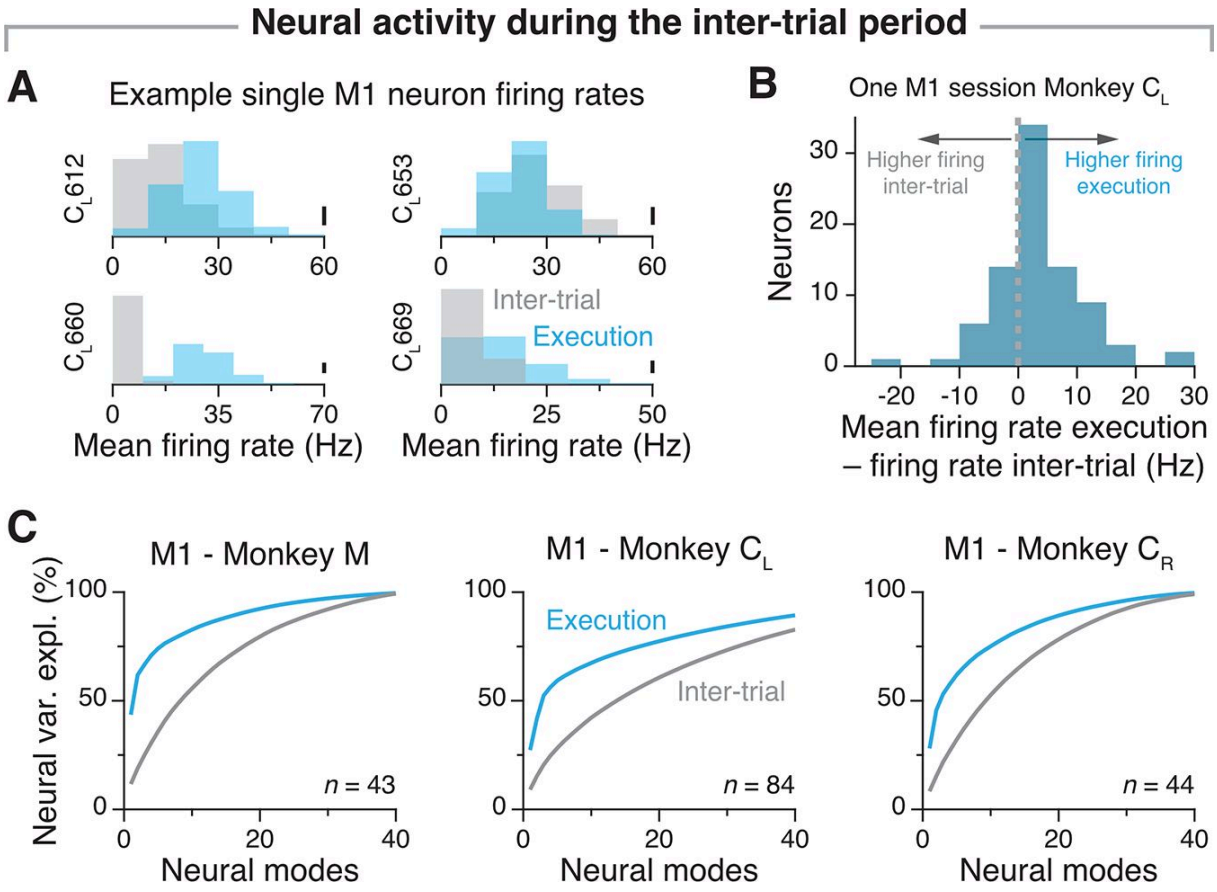

Figure S8: Decrease in primary motor cortex (M1) firing rates and neural correlations when monkeys are not engaged in the task. (A) Change in mean firing rate between the movement epoch and the inter-trial epoch – when monkeys were not engaged in the task – for four example neurons. Data from the example session from Monkey C_L shown throughout the paper. Each data point represents one trial. Scale bar: 20 trials. **(B)** Difference in average single neuron firing rates between the movement execution and inter-trial epochs for all 84 single neurons recorded during the same session as in **A**. As expected, the firing rate of M1 neurons tends to increase when animals are moving. **(C)** Percentage neural variance explained by manifolds with increased dimensionality (a 'neural mode' is a principal component defining the neural manifold; see [7]). During the inter-trial period, neural manifolds with a given dimensionality m explain a considerably lower percentage of total variance than manifolds with the same dimensionality m do during movement. Here, n indicates the number of single neurons used to compute the manifold.

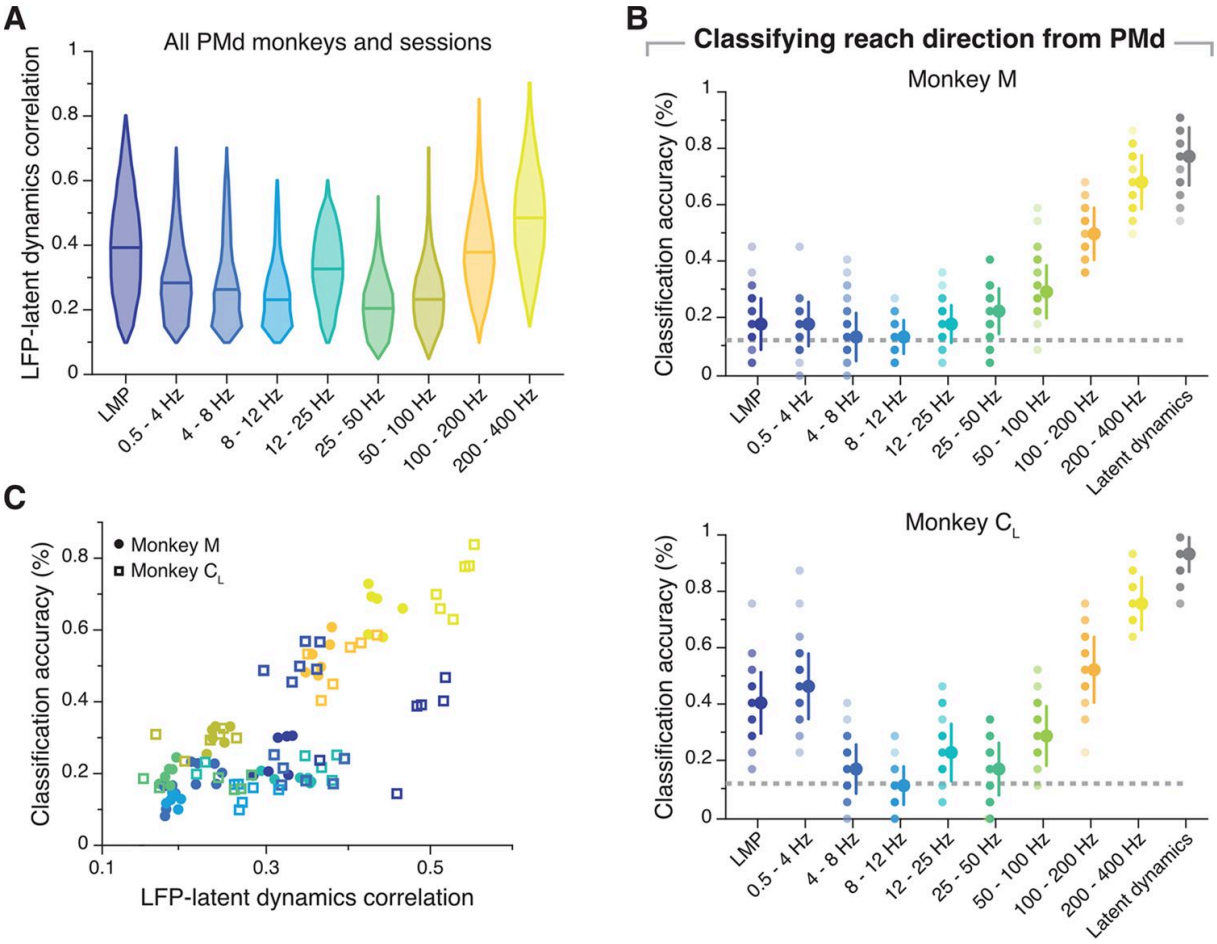

Figure S9: Similarity between dorsal premotor cortex (PMd) latent dynamics and each local field potential (LFP) band, and predicting reach direction from PMd LFPs. **(A)** Correlation between each LFP band and the latent dynamics; data pooled across all PMd sessions and monkeys. Note that the correlation profile matches well the example sessions in Fig. 2.6. Violin plot, probability density for one frequency band (n=723); horizontal line, median. **(B)** Accuracy of classifiers that predict the reach direction based on each LFP band (colours) and the latent dynamics (grey) during one representative session from each PMd monkey (same sessions shown in Fig. 2.6). Error bars, median ± s.d. (n=50). **(C)** The relationship between the PMd LFP-latent dynamics correlation coefficients and LFP classification accuracy is not linear, in contrast to primary motor cortex (M1). Each marker denotes one LFP band from one session; different monkeys are shown using different markers (legend). Markers are colour-coded as in **B**.

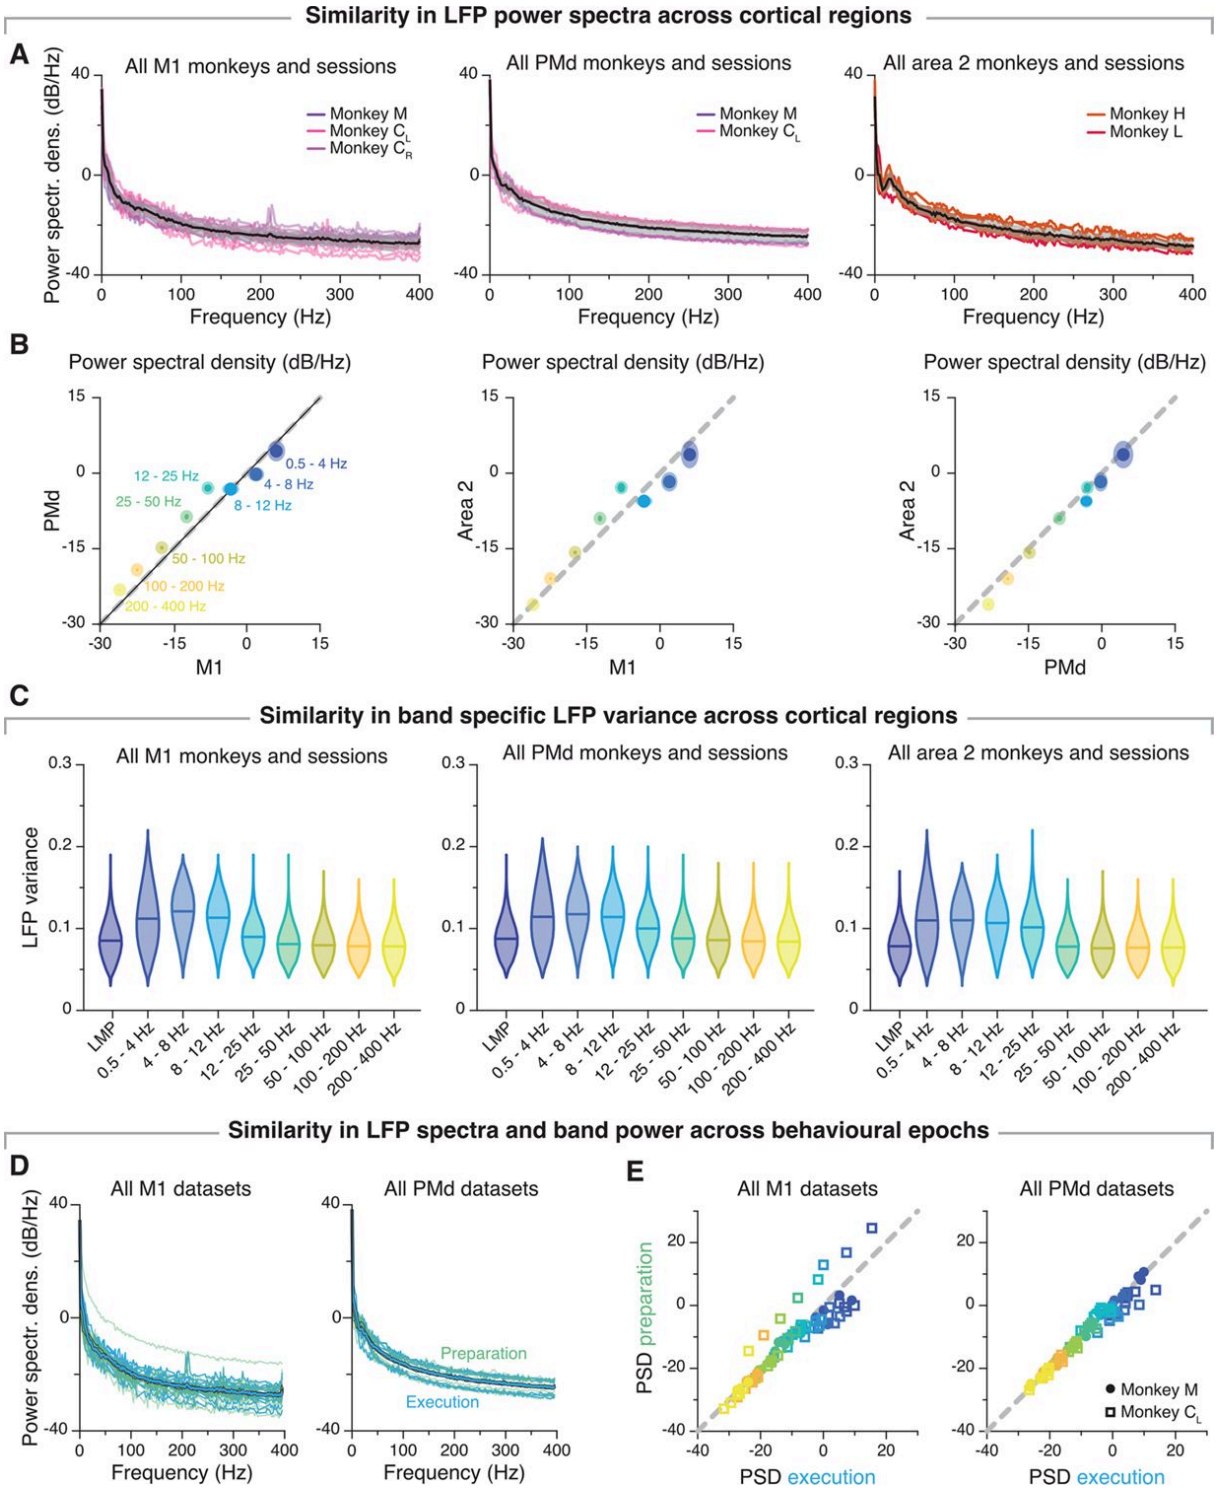
Similarity in LFP power spectra across cortical regions
A
All M1 monkeys and sessions
All PMd monkeys and sessions
All area 2 monkeys and sessions
Power spectr. dens. (dB/Hz)
Frequency (Hz)
Monkey M
Monkey C_L
Monkey C_R
Monkey H
Monkey L
B
Power spectral density (dB/Hz)
PMd
M1
Area 2
0.5 - 4 Hz
4 - 8 Hz
8 - 12 Hz
12 - 25 Hz
25 - 50 Hz
50 - 100 Hz
100 - 200 Hz
200 - 400 Hz
C
Similarity in band specific LFP variance across cortical regions
LFP variance
LMP
0.5 - 4 Hz
4 - 8 Hz
8 - 12 Hz
12 - 25 Hz
25 - 50 Hz
50 - 100 Hz
100 - 200 Hz
200 - 400 Hz
Similarity in LFP spectra and band power across behavioural epochs
D
All M1 datasets
All PMd datasets
Preparation
Execution
E
PSD preparation
PSD execution
Monkey M
Monkey C_L

Figure S10

Figure S10: Similarity of local field potential (LFP) spectral properties across frequency bands and behavioural epochs. **(A)** LFP power spectral density for primary motor cortex (M1) (left), dorsal premotor cortex (PMd) (middle), and area 2 (right). Coloured traces, individual sessions colour-coded by monkey (legend); line and shaded areas, mean ± s.e.m. across all monkeys and sessions (n=16, 12, 8 for M1, PMd, and area 2, respectively). **(B)** Comparison between the band-specific power spectral density across all pairs of regions (X and Y axes). Data pooled across all sessions. Line and shaded areas, mean ± s.e.m. (n=16, 12, 8 for M1, PMd, and area 2, respectively). Note the strong similarity in LFP spectral content across regions. **(C)** LFP power variance in each frequency band for M1 (left), PMd (middle), and area 2 (right). Note that the LFP variance is quite similar across regions, and that the bands with highest correlations with the latent dynamics (LMP, 100–200 Hz, 200–400 Hz) are among those with lowest variance. Violin, probability density for one frequency band; horizontal line, median. **(D)** LFP power spectral density for M1 (left) and PMd (right) during movement preparation and execution (legend). Coloured traces, single sessions from all M1 (left) and PMd (right) sessions from Monkey M and Monkey C_L ; black traces, mean. **(E)** Comparison between the band-specific power spectral density across behavioural epochs. Markers, single comparisons colour-coded by frequency band as in Panel B; different markers indicate different monkeys (legend); grey dashed line indicates identity. Note that the LFP spectral properties are very well preserved between movement preparation and execution except for an outlier M1 session from Monkey C_L.

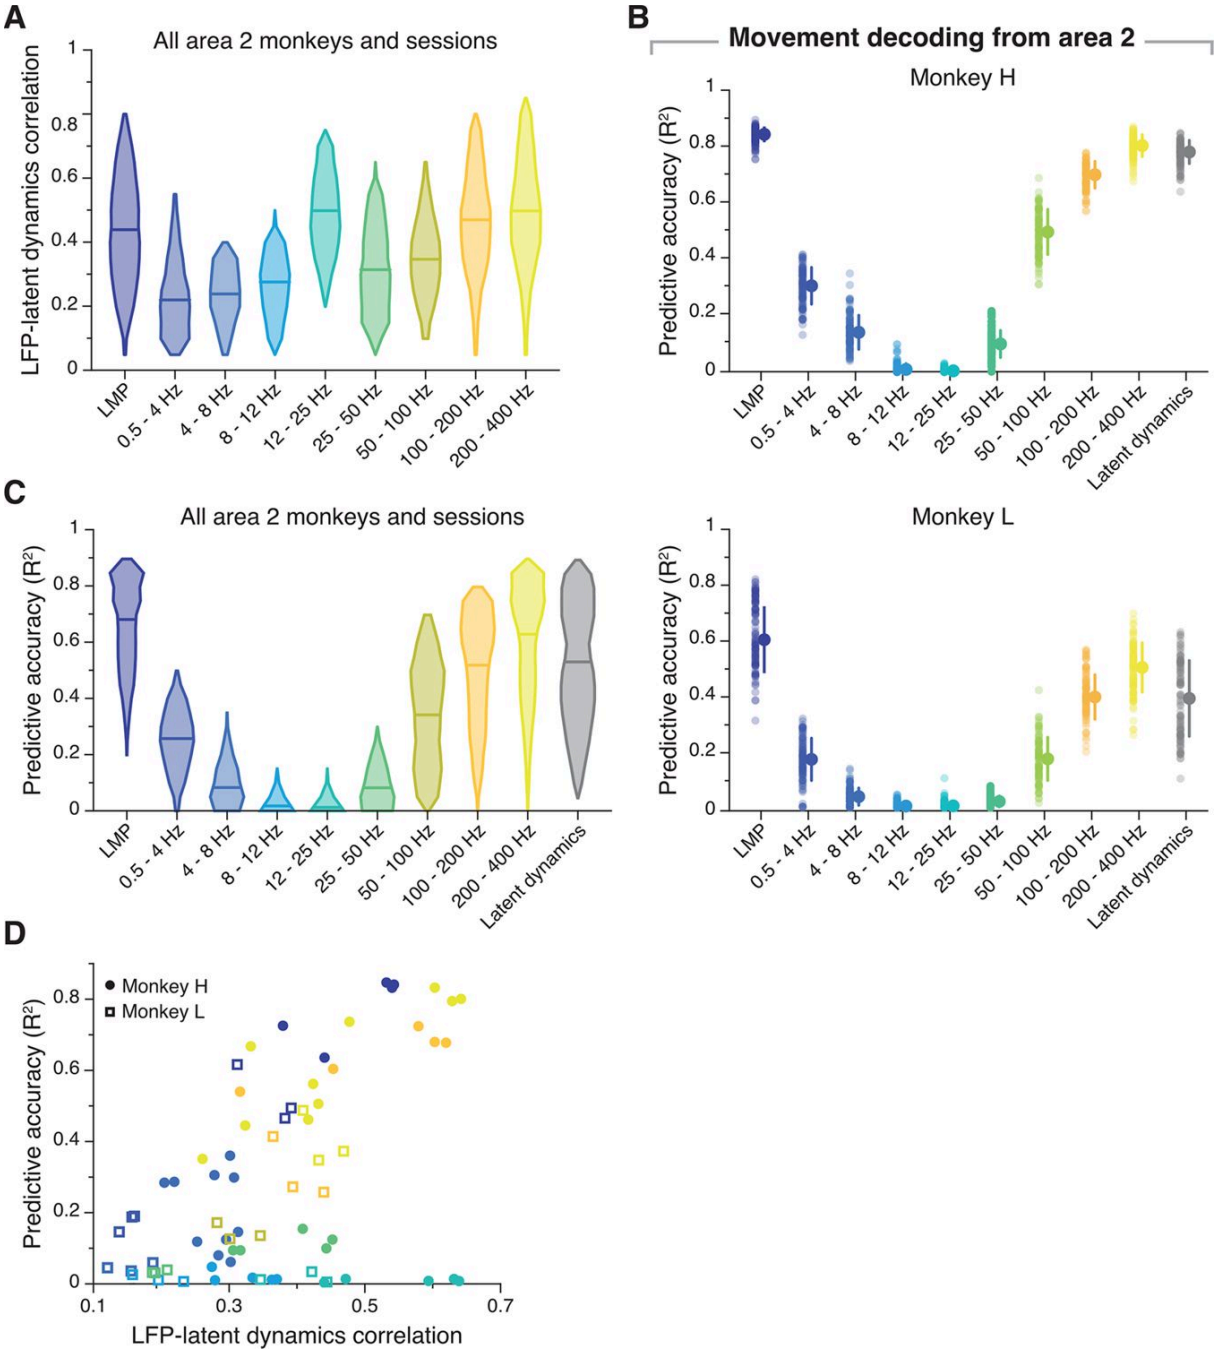

Figure S11: Similarity between area 2 latent dynamics and each local field potential (LFP) band, and decoding hand kinematics from area 2 LFPs. (A) Correlation between each LFP band and the latent dynamics; data pooled across all area 2 sessions and monkeys. Note that the correlation profile matches well the example sessions in Fig. 2.7. Violin plot, probability density for one frequency band (n=305); horizontal line, median. **(B)** Predictive accuracy of each LFP band (colours) and the latent dynamics (grey) for one representative session from each area 2 monkey (sessions are the same used in Fig. 2.7). Error bars, median ± s.d. (n=100). **(C)** Predictive accuracy of each LFP band (colours) and the latent dynamics (grey) pooled across all area 2 monkeys and sessions. Same format as **A**. **(D)** The relationship between the area 2 LFP-latent dynamics correlation coefficients and LFP predictive accuracy is not linear, in contrast to M1, but similar to PMd. Each marker (colour-coded as in **A**) denotes one LFP band from one session; different monkeys are shown using different markers (legend).

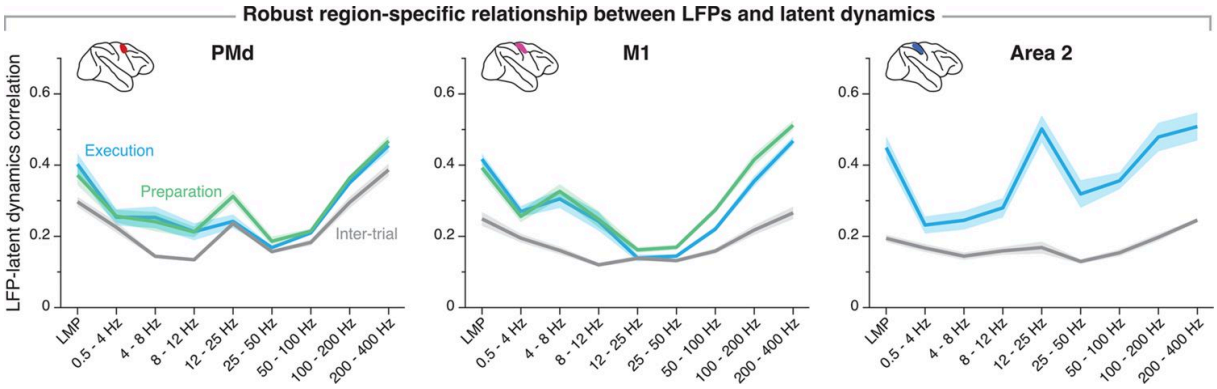

Figure S12: Region-specificity of the local field potential (LFP)-latent dynamics correlation profile. Comparison of the correlation between each LFP band and the latent dynamics across three sensorimotor cortical regions. For both dorsal premotor cortex (PMd) and primary motor cortex (M1), the correlation profiles remain virtually identical between movement preparation and movement execution. For all three regions, the LFP-latent dynamics correlation profiles decrease drastically when monkeys are not engaged in the task. Line and shaded areas, mean ± s.e.m. across all session medians from all monkeys for each area (n=16, 16, 8 for PMd, M1, and area 2, respectively). The M1 panel is reproduced from Fig. 2.5A.

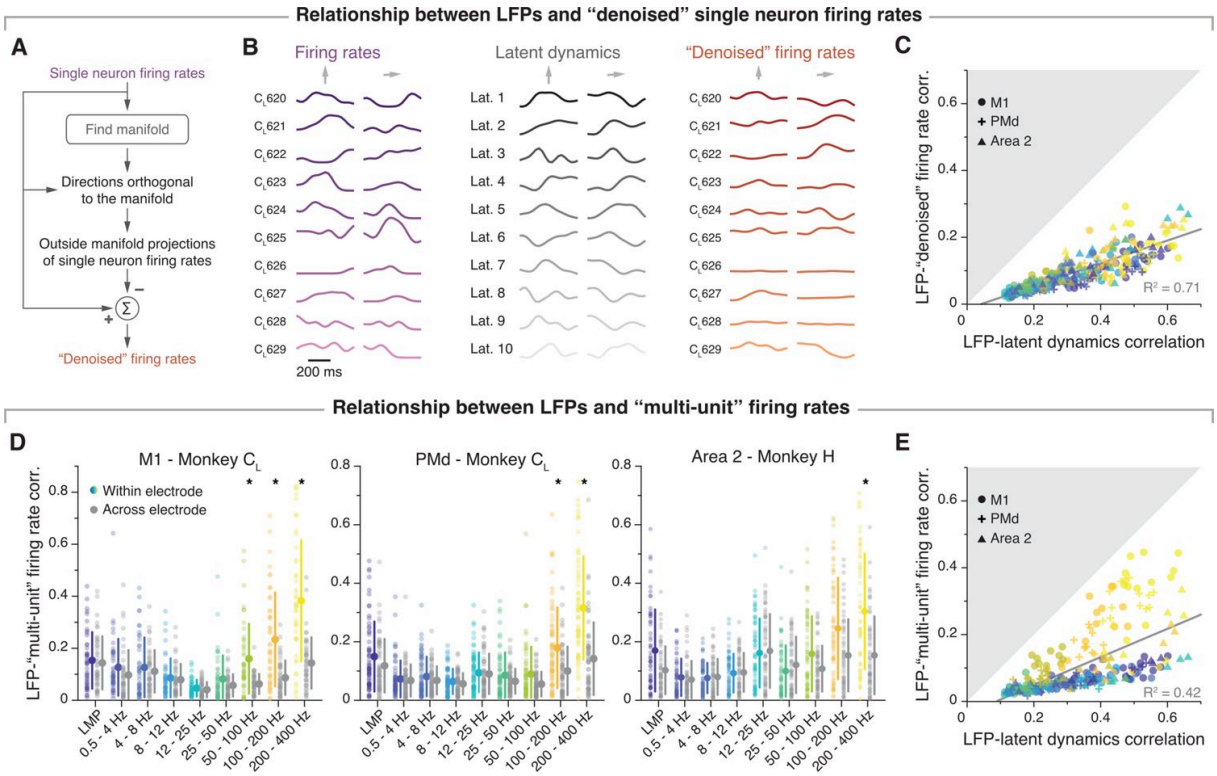

Figure S13: Additional data: Single neuron features do not account for the local field potential (LFP)-latent dynamics correlations. **(A)** To match the smoothness of the single neuron firing rates to that of the latent dynamics, we 'denoised' the former by subtracting their projections outside the neural manifold (see Methods). **(B)** Examples of firing rate 'denoising' for two randomly selected trials from one example session from Monkey C_L. Each of the three columns shows, from left to right: the single neuron firing rates for 10 randomly selected neurons, the neural population latent dynamics along each of the 10 manifold dimensions, and the denoised firing rates from the same neurons shown on the left. **(C)** Comparison between the LFP-latent dynamics correlations and the correlation between the LFP and the 'denoised' firing rate of neurons on the same electrode. Data formatted as in Fig. 2.8B. **(D)** Comparison between the correlation of each LFP band with the 'multi-unit' activity on the same electrode (coloured), and the 'multi-unit' activity on a different electrode (grey). As in previous analysis, 'multi-units' were obtained by pooled together all single neurons detected on each electrode. Each plot shows a representative session from one monkey implanted in each of the three studied regions. Error bars, median ± s.d. (n=50, 63, 60 for Monkeys M, C_L, and C_R, respectively). *p<0.001 two-sided Wilcoxon's rank sum test. **(E)** Comparison between the LFP-latent dynamics correlations and the correlation between the LFP and the 'multi-unit' activity on the same electrode. Data formatted as in **C**. Note that neither denoising the firing rates nor pooling the activity of single neurons into multi-units affects the conclusion that the neural population latent dynamics are more strongly related to the LFPs than single neurons (or 'multi-units') are.

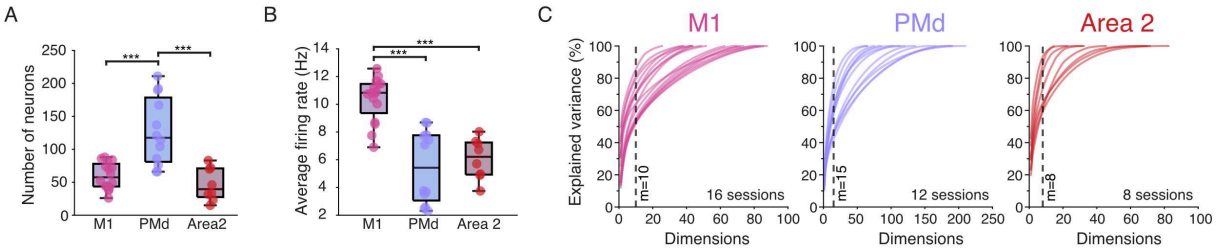

Figure S14: Single-neuron population statistics differ across regions, but manifold dimensionality reliably captures population activity. (A) Number of putative single neurons recorded per session in the primary motor cortex (M1), dorsal premotor cortex (PMd), and area 2. Each dot represents one session; box plots show the median, interquartile range, and range. ***p<0.001, two-sided Wilcoxon's rank-sum test. **(B)** Average firing rate per session for the same regions as in A. Each dot represents the session mean across all recorded single neurons. ***p<0.001, two-sided Wilcoxon's rank-sum test. **(C)** Explained variance as a function of the number of manifold dimensions, for all sessions in M1 (left), PMd (middle), and area 2 (right). Each line represents one session (n=16, 12, and 8 sessions, respectively). Dashed line indicates the manifold dimensionality (m) used throughout the chapter for each region.

Chapter 3

A brain-computer interface paradigm reveals different constraints on cortical and striatal activity

In the previous chapter, we showed that the distinct circuit architectures in the sensorimotor cortex shape the relationship between LFPs and neural population dynamics. We found a frequency-dependent relationship between the two signals that varied across cortical regions. However, for a given region, it remained stable across the distinct neural processes underlying movement planning and execution.

In this chapter, we investigate the circuit constraints between cortical and subcortical structures in the sensorimotor system. Among the regions of the sensorimotor system, M1 and DLS are among the most directly connected. Furthermore, they present well-characterised and markedly distinct cytoarchitectural properties. Hence, using an operant BCI task on head-fixed mice, we studied the volitional modulation of neurons in both

regions to characterise the constraints imposed by their distinct cytoarchitecture and connectivity.

The work presented in this chapter is an ongoing study. My contributions to this work included conceptualisation, experiment design, experimental set-up assembly, animal handling, surgery and training, data collection and preprocessing, analysis and validation of the data and custom software development.

3.1 Abstract

The primary motor cortex (M1) and striatum are two highly connected regions whose functional contributions to behaviour remain elusive. Some studies suggest that striatal activity primarily reflects cortical inputs, whereas others propose a more sophisticated contribution. Here, we used a BCI task where head-fixed mice were required to increase the firing rate of a set of Neuropixel channels located in M1 or DLS. We extensively compared the volitional modulation of the two areas, including neurons that were not targeted by the BCI.

We found that mice could intentionally increase the firing rate in both M1 and DLS to similar extents. Furthermore, we showed that the effect of the volitional modulation in the surrounding neural populations was very different depending on the area, consistent with different cytoarchitectural properties that impose constraints on the neural population activity. Finally, we showed that animals on the M1 or DLS BCIs use fundamentally different strategies to solve the BCI task. Altogether, these results demonstrated that striatal population activity is more than a reflection of ongoing M1 activity; instead, although volitionally controllable, these two areas have fundamental differences in the constraints of their neural population activity.

3.2 Introduction

The primary motor cortex (M1) and striatum are two of the key nodes for motor control in the brain. Although anatomically well characterised, the exact contribution of each of these regions to behaviour remains elusive. The striatum is the main input structure of the basal ganglia, a subcortical cluster of nuclei that integrates cortical signals and influences motor output through downstream projections to the brainstem, and back-projections to the cortex [76, 77, 81]. Within this circuit, the striatum serves as an integrative node, creating a convergence zone for motor-related information from many cortical regions, including the primary motor (M1), premotor and somatosensory cortices [71, 82, 83]. This integration provides the area with a comprehensive set of signals related to ongoing behaviour, from motor and sensory descriptive information to higher-level cognitive information related to reward or task context [71, 78, 155].

The cortical projections to the striatum are topographically organised, with the ventral and dorsomedial striatum traditionally associated with cognitive information, and the dorsolateral striatum (DLS) mostly related to sensorimotor information, predominantly from the primary motor cortex (M1) [71, 156]. Given their strong functional coupling, disentangling their individual contributions to behaviour remains a central challenge in motor control neuroscience. Some recent studies suggest that striatal activity primarily reflects cortical inputs [91], whereas others propose more sophisticated functional contributions [79, 92].

Although M1 and DLS neural populations are modulated by similar parameters of the ongoing movement during some tasks, such as direction or vigor [6, 157, 158], these two brain areas are anatomically very distinct, which imposes distinct constraints on the neural activity patterns they can produce. In M1, the neural population activity has been well described by a relatively small number of covariance patterns relevant to the behaviour, the "neural manifold" [7, 8]. In contrast, DLS neural population activity tends to be sparser and less correlated, consistent with the presence of fast-spiking inhibitory interneurons [87,

159]. In a recent study directly comparing both neural populations, the authors showed that the striatal neural population activity is higher-dimensional than M1 activity during the same reaching and grasping task, although both were well described by non-linear manifolds [22].

To directly probe the distinct constraints on the neural population in M1 and DLS, a controlled manipulation of their activity would be necessary. Brain-computer interfaces (BCIs) have been traditionally designed to decode the neural activity of a specific brain region to control external devices, frequently inducing modifications to the underlying neural activity [15, 16, 94]. The ability to produce such modifications in the brain has recently expanded the perception of BCIs not only as a mapping tool, but also as a perturbation tool, to shift the neural activity to a specific state [94, 95, 96]. Unlike other perturbation tools such as optogenetics or electrical stimulation, BCIs exploit the subject's own volition to generate naturalistic changes that follow intrinsic constraints of the neural population [96].

Recent studies have exploited this property by designing operant conditioning BCI tasks where the subjects are rewarded for a certain brain activity, such as increasing the firing rate for some neurons [94, 160, 161, 162], or pushing a certain state on their neural population activity [163, 164], hence making them likely to repeat the same activity pattern in the future. However, these studies have largely focused on cortical activity, mainly from M1 [15, 16, 161, 162, 164, 165, 166], leaving subcortical structures like striatum largely unexplored [95]. Here, we address this gap by exploring the ability to intentionally modulate striatal activity. We chose DLS as our target region given its close relationship with M1. We designed an operant BCI task in head-fixed mice to compare the volitional modulation of M1 and DLS neural populations, investigating the extent to which each of their distinct population-level properties imposes different constraints on their neural activity.

3.3 Results

3.3.1 Operant Brain Computer Interface task

To explore the volitional control and neural constraints on different brain regions, we designed a brain-computer interface (BCI) task for head-fixed mice on a spherical treadmill (Fig. 3.1A). In this setup, the animals could freely modulate the running speed and direction, allowing for a wide variety of movements. While the animals were head-fixed, we acutely inserted a Neuropixel 1.0 probe targeting the primary motor cortex (M1) and the dorsolateral striatum (DLS) region, key nodes for motor control. Then, we selected five consecutive neuropixel channels from either layer V of M1 or DLS, and calculated the combined rate of theshold crossings based on a RMS threshold for each channel, the "BCI threshold crossings rate", as the control signal for the BCI. Neuropixel channels are very close to each other, often causing the same neurons to be detected in multiple channels [167]. By selecting five channels to control the BCI instead of a single one, we gave our algorithm a better estimation of the firing rate happening around the BCI channels and made it more robust to possible drifts on the probe.

Mice performed the BCI task for 60 to 90-minute sessions, depending on the engagement of the animals. They received a reward anytime they increased the BCI threshold crossings rate above a threshold for a 200ms holding time (Fig. 3.1B). The threshold was set based on previous BCI threshold crossings rate values and was updated every 5 minutes according to previous activity [Methods Section 3.6.7]. This adaptive threshold method ensured that the BCI was challenging but fair for the long recording times, independently of the overall BCI threshold crossings rate.

The BCI threshold crossings rate was mapped to a bidirectional visual cursor on an LED strip located in a semi-circle in front of the animal (Fig. 3.1A, C). The visual cursor received values from 0 to 100, corresponding to the sides and the centre of the LED strip, respectively. The BCI map assigned a cursor value of 0 to having no firing in the BCI channels, and a value of 100 to any BCI threshold crossings rate value above the adaptive

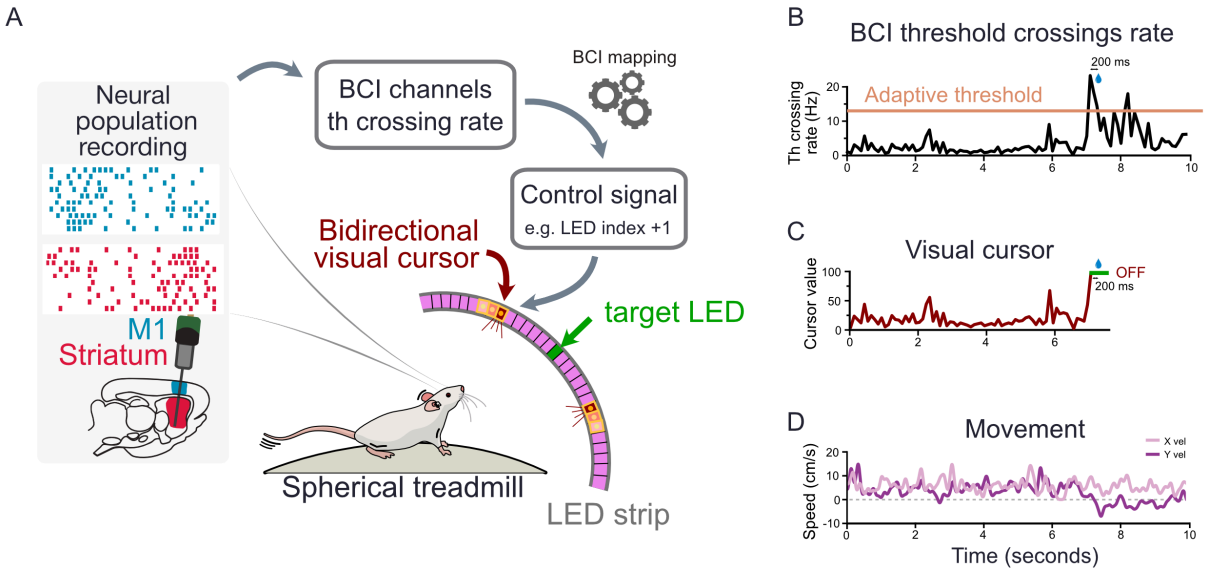

Figure 3.1: Operant brain-computer interface (BCI) task schematic. **(A)** We recorded the spiking activity of neural populations from primary motor cortex (M1) and striatum from head-fixed mice on a spherical treadmill. From the neural recordings, we selected five consecutive channels from either M1 or the dorsolateral striatum (DLS) to be the BCI channels. The BCI threshold crossings rate was mapped to a 1-dimensional control signal, translated to a visual cursor on an LED strip in front of the animal for BCI feedback. **(B)** The BCI threshold crossings rate was compared against an adaptive threshold [Methods Section 3.6.7]. If the BCI threshold crossings rate remained above the threshold for more than 200ms, a reward was triggered. **(C)** The bidirectional visual cursor was mirrored to both sides of the animal, and changed proportionally to the BCI threshold crossings rate. When the BCI threshold crossings rate was above the threshold, the cursor was maintained at the centre/target LED. After a reward was triggered, the cursor remained in the target LED for 0.5s before turning off to restart the task. **(D)** We simultaneously recorded the displacement of the treadmill in the X and Y directions corresponding to longitudinal and lateral movement, respectively.

threshold. That is, when the animals reached the firing rate threshold, the visual cursor remained on the centre until the BCI threshold crossings rate decreased again or the holding time of 200ms was achieved, resulting in a successful trial. In that case, in addition to receiving a liquid reward, the LED cursor remained in the centre for an extra 500ms to mark the success, and it was followed by an intertrial period of 5 seconds, during which the LED strip was off. After that, the BCI was engaged again and the task resumed.

We recorded data from 5 mice on the BCI task. In each session, animals were asked to increase the firing rate of neurons located either in M1 (12 sessions) or DLS (30 sessions). The BCI channels were manually selected on the real-time streaming of the electrophysiological data (Fig. 3.2A) from the channels that showed a high firing rate,

to ensure a large number of neurons available to contribute to the BCI control. We preprocessed the signals using a common average reference (CAR) and a high-pass filter (HPF), and detected putative spikes by thresholding each channel based on the unique root-mean square (RMS) value to get the threshold crossing counts per 100ms bins (Fig. 3.2B, C). Finally, we applied a gaussian weighted average on the five channels to get the BCI threshold crossings rate (Fig. 3.2C, D), the control signal for the BCI [additional information in methods Section 3.6.7].

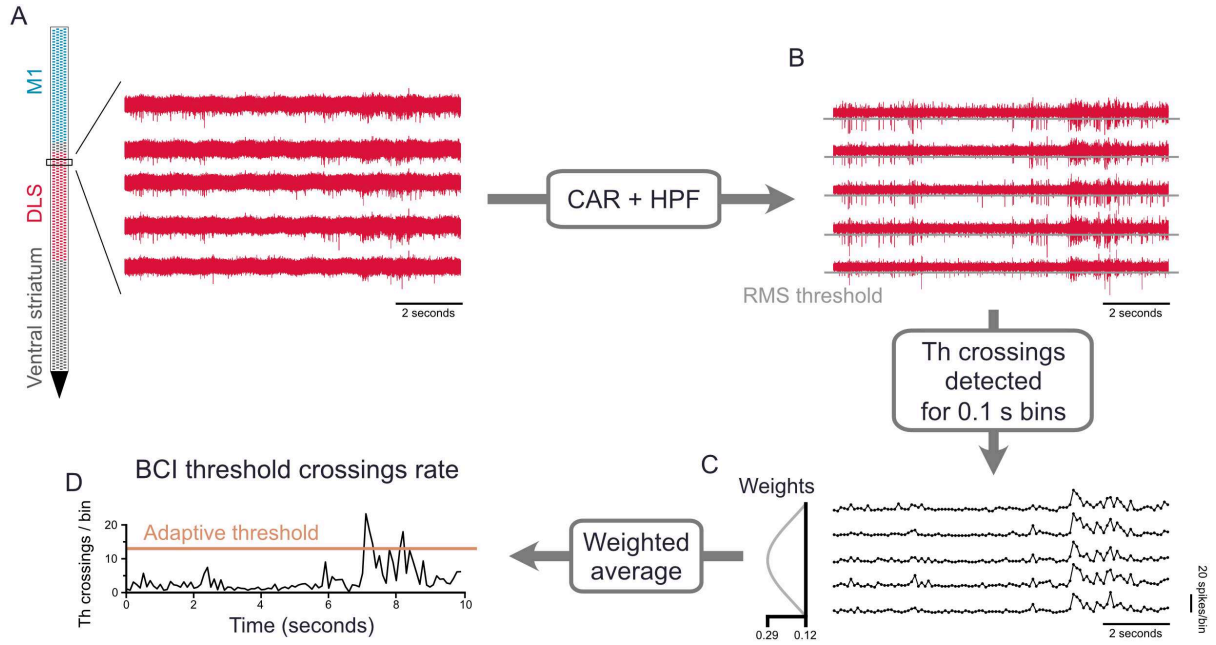

Figure 3.2: Real-time brain-computer interface (BCI) firing rate calculation from electrophysiological data. **(A)** Same example as in Fig. 3.1B. From the raw electrophysiological real-time data, we manually selected five consecutive channels, depending on the region of interest (DLS in this example). **(B)** For each channel, we applied a common average reference (CAR) and high-pass filter (HPF) to clean the signal. **(C)** Using prerecorded data, we established an RMS value specific to each channel. The RMS values were used to calculate a spike detection threshold to be applied in real time [Methods Section 3.6.7]. We estimated the number of spikes per time bin by counting the number of times the channel signal crossed the threshold, using 100ms bins. **(D)** We used a weighted average on the spike counts across the channels to obtain the final BCI threshold crossings rate. The BCI threshold crossings rate was then compared in real-time to the adaptive threshold and mapped into the visual cursor.

We recorded the neural activity from hundreds of neurons in M1 and DLS. We also recorded the neural activity from the ventral striatum, although this data was outside the scope of the project. In addition to neural activity, we simultaneously recorded the animal movement by using two displacement sensors on the treadmill, to capture each

perpendicular direction (Fig. 3.1D). Finally, we also recorded the licking events during the task. Altogether, these neural and behavioural recordings provide a comprehensive dataset to explore the BCI control of neural populations in M1 and DLS.

3.3.2 Volitional increase of the firing rate in the BCI channels.

To analyse the changes in neural activity, we first used Kilosort 4 [168] to identify putative neuronal activity. As part of the preprocessing, the spiking of each neuron was square-root-trasnformed and smoothed using a gaussian filter to get the firing rate for each neuron [Methods Section 3.6.8]. We then studied the activity of the BCI neurons, that is, the neurons identified in the BCI channels, and noticed that some BCI neurons showed a firing rate increase throughout the sessions (Fig. 3.3A). To study specifically the change of the BCI control signal, we estimated the BCI threshold crossings rate as the average firing rate of all the BCI neurons called "BCI firing rate". By fitting a linear model and measuring the slope, we calculated the global increase or decrease of the BCI firing rate throughout the session in Hz per hour (Fig. 3.3B).

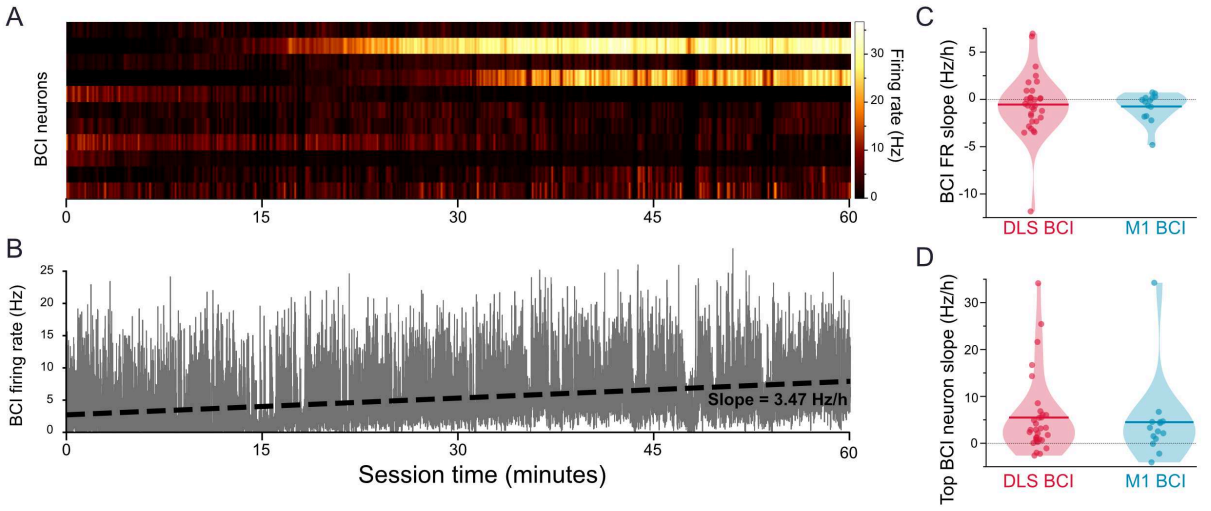

Figure 3.3: BCI firing rate increased throughout the session. **(A)** Raster plot of the smooth firing rate of the BCI neurons for an example 60-minute session. **(B)** Average firing rate of the neurons in (A) overlaid by the fitted linear model as a dashed line. The slope value is indicated under the line. **(C)** Summary of all the BCI firing rate slope values for all the recorded sessions, colour-coded by BCI area. Each dot is a single session. The horizontal line corresponds to the mean value of each distribution. **(D)** Same as (C) but only including the BCI neuron with the highest slope value per session.

We initially expected to see a clear increase in the BCI firing rate during the session.

However, we found that, in most sessions, the slope of the BCI firing rate was around 0 or even negative (Fig. 3.3C). One possibility is that the BCI control relied only on a few neurons, which we miss when averaging all the BCI neurons together, especially if other neurons had a higher overall firing rate. Hence, we repeated the same analysis but using only the neuron that increased its firing rate the most (Fig. 3.3D). We found that most sessions included at least one neuron whose activity increased over the session, with a mean change of 5Hz per hour. Interestingly, the effect was similar regardless of the area, M1 or DLS, suggesting that DLS has similar capabilities for modulation of its activity as M1.

Despite this promising result, the analysis also showed that there was a lot of variability in the change of the BCI firing rate across sessions, obscuring the significance of the result. This could be due to several reasons. First, we performed acute recording, so the BCI neurons recorded each day were not the same, and their capability to increase their firing rate may be different. Second, as explained in the previous section, our selection of BCI channels was biased towards significantly active channels, which means that the BCI firing rate usually started already high at the beginning of the session. Even using an adaptive threshold, animals may find it difficult to increase the BCI firing rate beyond the high starting point. Lastly, we were not truly rewarding a long-term increase in the BCI firing rate, but the fast increase above the threshold in a given moment. Furthermore, an overall increase in the BCI firing rate complicated the task by increasing the adaptive threshold. Altogether, this analysis revealed that the animals were likely not producing a long-lasting increase in the activity of BCI neurons; they were increasing the BCI firing rate only momentarily to achieve rewards without an effect over the whole session.

To measure the BCI performance and volitional control in this task was very challenging, mainly because the use of an adaptive threshold meant that the reward condition was constantly changing. Hence, we could not use traditional metrics such as the number of rewards or the time to acquire the next reward. We therefore looked at the behaviour of the animals for evidence of volitional control. Interestingly, we observed that once

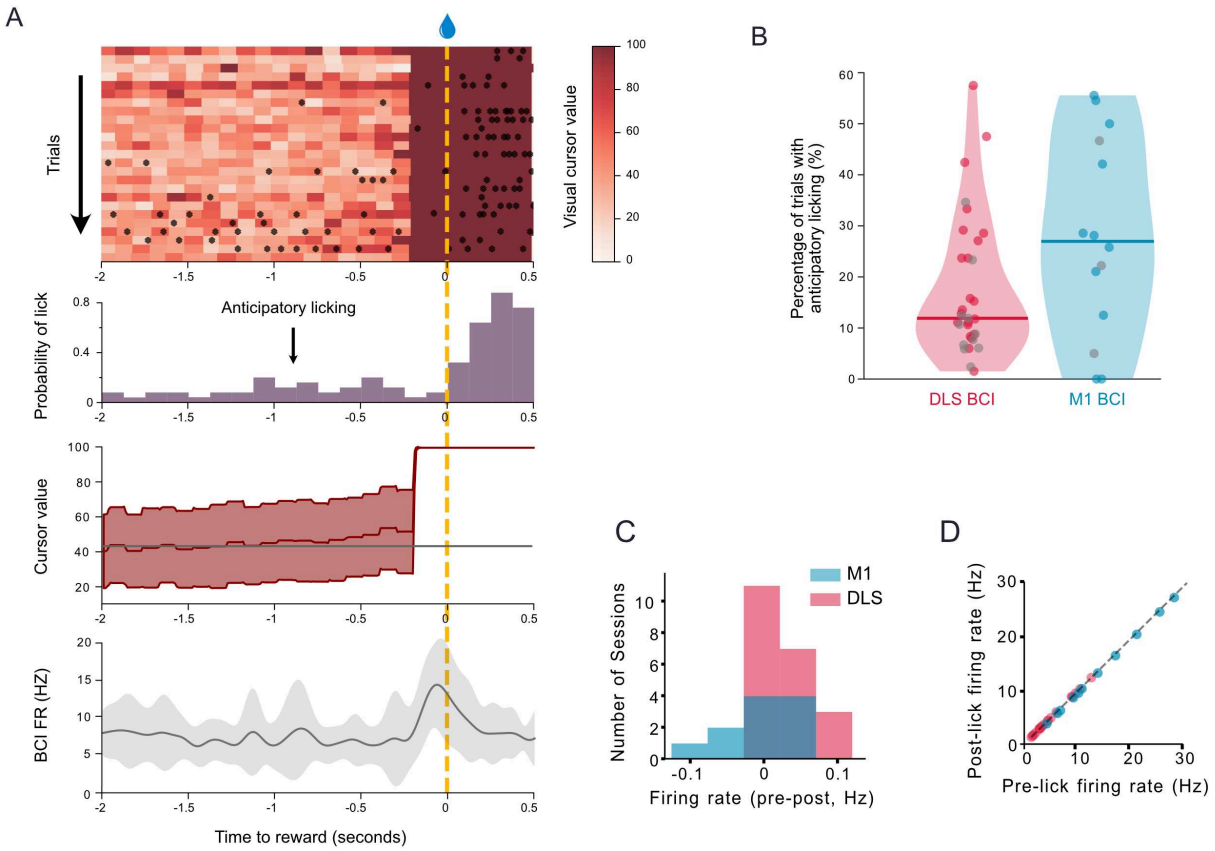

Figure 3.4: Anticipatory licking shows intentional control of the BCI. (A) Example DLS BCI session showing clear anticipatory licking not linked to visual cursor value or BCI firing rate. All the plots share the same x-axis, time, from -2 seconds before reward to 0.5 seconds after reward delivery. From top to bottom: (1) Heatmap of the visual cursor values for all the rewarded trials in the session, ordered increasing by trial time. Each dot represents a lick event. (2) The animal's probability of licking around the reward for the above trials. (3) Mean (line) and standard deviation (shaded area) of the cursor value for the above trials. (4) Same as (3) for BCI firing rate. **(B)** Summary of percentage of rewarded trials with anticipatory licking, colour-coded by BCI area. Each dot is a single session. The horizontal line corresponds to the median value of each distribution. Grey dots represent sessions where the measure of the licks was not accurate. **(C)** Change in BCI firing rate before and after random licking events in the session for all sessions, colour-coded by BCI area. **(D)** BCI firing rate before vs after random lick events for all the sessions, colour-coded by BCI area. Each dot represents a session.

the animals understood the task, they licked very soon after the reward was delivered. Furthermore, in some sessions, animals started licking even before the target was achieved. This anticipatory licking occurred mostly in the late trials in the session (Fig. 3.4A), happening between 1 and 0.5 seconds before reaching the target.

Critically, this anticipatory licking could not be explained by the visual cursor value or the BCI firing rate, where the upcoming reward could only be predicted around 0.5

seconds before the target acquisition (Fig. 3.4A). In other words, the animal knew a reward was coming even before the BCI feedback indicated it. The anticipatory licking happened in most of the sessions, independently of the BCI area, with some sessions showing anticipatory licking in up to 60% of the rewarded trials (Fig. 3.4B). Importantly, licks themselves are not sufficient to affect BCI firing rate; in all the sessions, random licks throughout the experiment did not significantly change the BCI firing rate (Fig. 3.4C, D). Unfortunately, we could not reliably track this effect in many sessions due to a hardware issue in the lickometer, the device to measure the lick events; those sessions are coloured in grey in Fig. 3.4B.

The anticipatory licking showed that the animals could predict that the reward was coming before there were any signs of it. Furthermore, the ability to predict the reward was learned by experience on the BCI task, since the anticipatory licking happened mostly in the later trials of the session. The only mechanism by which they could predict the reward with such anticipation is if they were intentionally deciding to increase the firing rate of the BCI neurons, provoking a reward delivery, hence volitionally controlling the BCI. Crucially, this was the case for both M1 and DLS BCIs, demonstrating that animals could volitionally modulate neural activity in subcortical regions as previously reported for cortical areas [15, 94, 161, 165].

3.3.3 BCI-driven firing rate increases reveal distinct constraints on neuronal population activity

We have shown that mice have the capacity to volitionally increase the firing rate of neurons in M1 or DLS for a short period of time, but it is still unknown what changes are induced in the neural populations due to these volitional perturbations. We first addressed this at the level of individual neurons, asking which specific neurons changed their firing rate around rewarded events, and where they were located relative to the BCI neurons. In this section, we examined the modulation of neighbouring neurons around the rewarded events, that is, when the BCI firing rate increased above the adaptive threshold

for more than 200ms. As described above, we recorded hundreds of units from M1, DLS and ventral striatum. As illustrated in Fig. 3.5A, trial-averaged raster plots revealed that, independently of the BCI region, the increase in firing rate of the BCI neurons also affected other neurons in the recordings.

To quantify this effect, for each rewarded trial, we calculated the percentage of change of the firing rate for each neuron with respect to their baseline firing rate. We calculated the average firing rate of each neuron around the reward using a window from 0.1 seconds before to 0.1 seconds after the reward, and compared it to their baseline firing rate defined as the session-long average firing rate of that neuron (Fig. 3.5B). We calculated the percentage of change for each rewarded trial individually [Methods Section 3.6.9] (Fig. 3.5C). Based on visual inspection of the distributions, we identified the "modulated neurons" as those whose average percentage of change was higher than 300%.

The number of detected neurons per area was very variable per session, with values ranging from 100 to 400 per area (Fig. 3.5D). When looking just at the number of modulated neurons per region, we found similar results independently of the location of the neurons controlling the BCI (Fig. 3.5E). In both M1 and DLS BCIs, there were non-BCI neurons modulated by the reward, with a median value of 10 neurons. Interestingly, focusing on the number of modulated neurons across regions, that is, modulated DLS neurons for M1 BCI and modulated M1 neurons for DLS BCI (grey distributions in Fig. 3.5E), we found that M1 BCIs were able to modulate DLS neurons in some sessions, but the inverse was not true in any session. This is consistent with the known connectivity between the two areas, since M1 directly projects to DLS [71, 82], driving the modulation downstream. Intriguingly, the absence of modulated neurons in M1 for a DLS BCI points towards a volitional DLS modulation that is not directly related to M1 activity, but likely had a different origin.

Finally, we quantified the distance between the BCI neurons and the same-area modulated neurons. Despite the number of same-area modulated neurons being similar for both M1 and DLS BCIs, we found a striking difference in the location of those neurons. When the

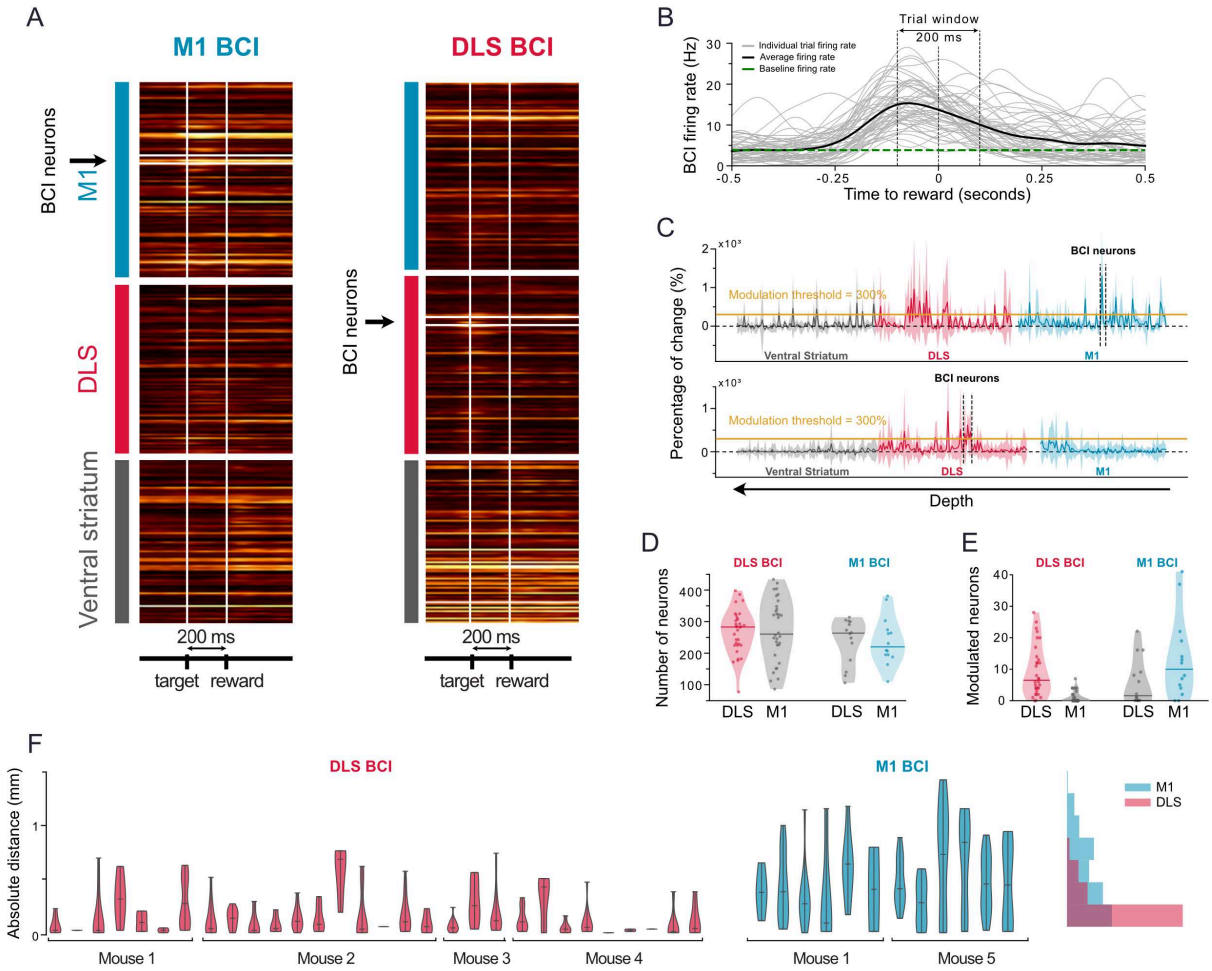

Figure 3.5: Firing rate increase revealed different neural population constraints. **(A)** Raster plot of the average neural population activity at rewarded trials for two example sessions, of M1 and DLS BCI, respectively. The left label indicates the brain region. White horizontal lines mark the BCI neurons. White vertical lines mark the target acquisition and reward times. **(B)** BCI firing rate at rewarded trials for the example DLS BCI session. Each grey line represents a single trial. The black thick line represents the average firing rate for all the individual trials. The green dashed line represents the baseline firing rate. The vertical dashed lines represent the time window used to calculate the percentage of change. **(C)** Percentage of change during reward events for all the neurons recorded in the two example sessions in (A). Mean and standard deviation of the individual trials are shown as a line and shaded area, respectively. The yellow line shows the 300% threshold to consider a neuron "modulated". **(D)** Total number of neurons detected in M1 and DLS for all the recorded sessions, separated by BCI area. The horizontal line shows the median values. **(E)** Same as (D), only for the number of modulated neurons. **(F)** The absolute distance of the modulated neurons from the BCI channels for each individual session, colour-coded by BCI area. Only neurons from the same BCI region. The summary histogram on the right shows all the distributions combined.

BCI neurons were in DLS, the modulated neurons were located close to those BCI neurons, extending no more than 0.2mm in most sessions. Conversely, when the BCI neurons were

located in M1, the modulated neurons were distributed across the entire recorded area (Fig. 3.5F).

This finding suggests fundamental differences in the constraints on cortical and striatal circuits. The widespread modulation in M1 neural populations is consistent with its activity being more correlated and explained in a few covariance patterns [7]. In contrast, the localised DLS modulation, also consistent with the sparse activity previously reported [87], potentially allows for a more targeted control of specific neural subpopulations.

3.3.4 Animals adopt distinct control strategies for cortical and striatal BCIs

These fundamental differences in the M1 and DLS BCI may not be limited to modulation constraints. In this section, we explored whether mice adopted different strategies to modulate the activity of these motor-related areas.

Many BCIs controlled by M1 neurons were solved by movement strategies [161, 169], that is, the subjects learned to move certain ways to elicit the desired neural activity. Here, we explored whether this was the case for our M1 BCIs and, more importantly, whether the animals would also use the same strategy for BCIs controlled by DLS neurons. We evaluated the prediction of the BCI firing rate, i.e. the control signal, from the motion data recorded with the treadmill sensors (Fig. 3.1D). If the animals were using the motion to control the BCI, the velocity profiles should be able to predict the BCI firing rate.

We used the recorded longitudinal and lateral velocity profiles to train a linear model to "encode" the BCI firing rate (Fig. 3.6A) [Methods Section 3.6.10]. We measured the performance of the encoders in terms of how much variance of the original BCI firing rate was explained by the predicted BCI firing rate, the explained variance (EV). The EV of the encoders ranged from 0 to 0.3, with 0.3 corresponding to a close prediction of the BCI firing rate despite its low value (Fig. 3.6B). Comparing both BCI areas, there was a clear distinction where most BCIs based on M1 were, as expected, modulated by the running

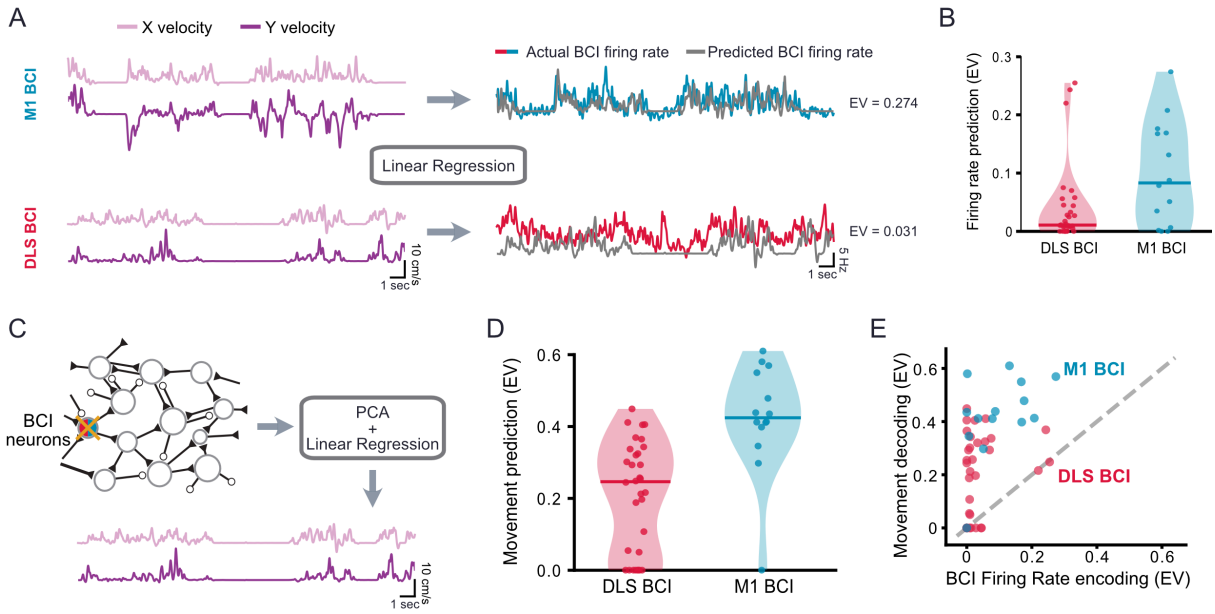

Figure 3.6: Motion can predict BCI firing rate on M1 but not DLS. (A) Example sessions for DLS and M1 BCIs. On the left, an example of ten seconds of longitudinal (X) and lateral (Y) running speeds during the session, which were used as predictor variables in the linear model trained to generate the BCI firing rate. On the right, the predicted and actual BCI firing rate for both examples, together with the explained variance (EV) for that session. **(B)** Summary plot of the linear model explained variance for all the sessions. The horizontal line shows the mean of the distribution. **(C)** Locomotion decoding from population dynamics. M1 or DLS population dynamics obtained by PCA are used as input for a linear regression to predict the motion variables. The BCI neurons were excluded from this population analysis. **(D)** Summary plot of the linear model scores for all the sessions. The horizontal line shows the mean of the distribution. **(E)** Locomotion decoding vs BCI firing rate encoding scores show no relationship between the two.

speed of the animals. In contrast, when the BCI was controlled by activity in the DLS, the locomotion of the animal was not related to the BCI control signal in most sessions.

This result was surprising since we expected the locomotion to modulate the BCI firing rate in both cases, as both M1 and DLS are predominantly movement-related areas in the brain. One possibility could be that the DLS population that we were targeting was not as related to locomotion as we expected. To address that possibility, we did a second analysis assessing the movement-related information present in each of these areas. We trained a linear decoder to predict the mouse running speed from the neural population activity of each area obtained by PCA (Fig. 3.6C) [Methods Section 3.6.11], excluding the BCI neurons.

Comparing the prediction scores, we found that, as expected, both areas were able to predict the animal locomotion, although M1 predictions were slightly more accurate than DLS ones (Fig. 3.6D). Interestingly, these results show that, although the DLS neural population contains movement-related information for many sessions, the animals were not using locomotion as a strategy to solve the BCI task. In fact, if we compare the BCI firing rate encoding and the movement decoding per session (Fig. 3.6E), many DLS BCI sessions had a high explained variance for movement prediction but close to 0 explained variance for the BCI firing rate encoding.

This result suggests that animals used different control strategies depending on whether the BCI is controlled by M1 neurons or DLS neurons. When the BCI was controlled by M1 neurons, mice tended to adapt their running speed to control the BCI and achieve the rewards. However, this was not the case for DLS BCIs. The striatum is a region that integrates information from many cortical regions, and, even though we were specifically targeting the main projection area from M1, it still may have information from other areas, such as the premotor or somatosensory cortices. It may be difficult for the animals to associate the DLS neuronal activity directly with locomotion, or, more intriguingly, they may have been able to find other cost-effective strategies to modulate their DLS activity that did not intrinsically relate to locomotion. Combined with the localised modulation effect, these results suggest that DLS can be volitionally modulated through mechanisms that are independent of M1.

3.3.5 Neural manifolds in DLS reorganise over the course of striatal BCI learning, but not M1 manifolds in cortical BCI

Having characterised the BCI effects at the level of individual neurons, we next asked whether BCI-driven modulation was also reflected at the population level, by examining its relationship to the region's neural manifold. Consistent with previous studies [22], DLS manifolds were generally higher-dimensional than M1 manifolds across BCI sessions

(Fig. 3.7A, B), although this difference did not reach significance ($p<0.1$). This is in line with the expectation that DLS's non-layered architecture supports a richer repertoire of activity patterns than M1's more structured, correlated circuitry, while suggesting this baseline difference in dimensionality is more modest in our specific task than previously reported.

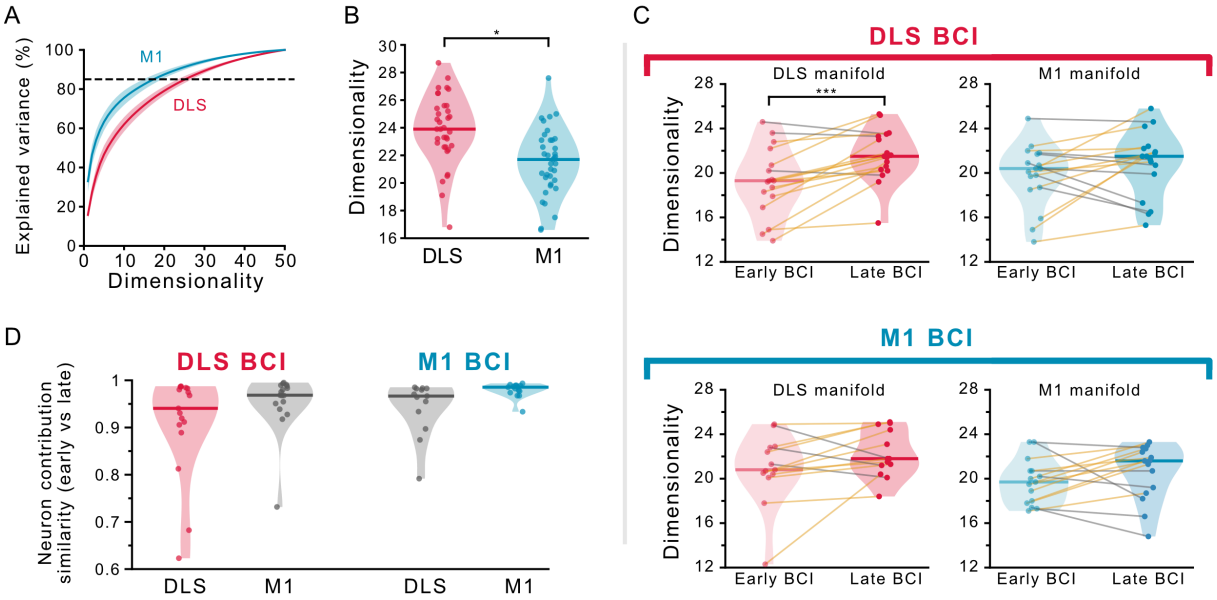

Figure 3.7: DLS manifolds reorganise over the course of BCI learning, while M1 manifolds remain comparatively stable. **(A)** Explained variance as a function of manifold dimensionality for M1 (blue) and DLS (red). Shaded area represents [mean ± s.d.] across 50-neuron iterations. Dashed line indicates the variance threshold (85%) used to define manifold dimensionality throughout this analysis. **(B)** Manifold dimensionality for DLS and M1, pooled across all sessions and BCI types. Each dot represents one session; horizontal lines indicate the median. *$p<0.1$, two-sided Wilcoxon's rank-sum test. **(C)** Manifold dimensionality during the first (Early BCI) and last (Late BCI) 15 minutes of each session, shown separately for DLS BCI sessions (top) and M1 BCI sessions (bottom), and for the manifold of the region being controlled by the BCI and of the other, simultaneously recorded region. Each pair of connected dots represents one session; lines are coloured by the direction of change (increase/decrease). ***$p<0.001$, two-sided Wilcoxon's signed-rank test. **(D)** Similarity between the neurons contributing most to the manifold (top 10 principal components) in the early and late windows of each session, for the DLS and M1 manifolds during DLS BCI sessions (left) and M1 BCI sessions (right). Each dot represents one session; horizontal lines indicate the median.

We then asked whether manifold dimensionality changed over the course of a BCI session, comparing the first and last 15 minutes of each recording (Fig. 3.7C). In DLS BCI sessions, DLS manifold dimensionality increased significantly by the end of the session in 80% of sessions. In M1 BCI sessions, by contrast, M1 manifold dimensionality increased in only

60% of sessions, and this change was not statistically significant. The manifold of the region not being controlled by the BCI, that is, DLS during M1 BCI sessions and M1 during DLS BCI sessions, also tended to increase in dimensionality over the session in most cases, but this increase did not reach significance either.

Finally, we examined how stable the specific neurons contributing to each manifold were between the first and last 15 minutes of the session (Fig. 3.7D). The neurons contributing most to the M1 manifold during M1 BCI sessions were highly preserved across the session. In contrast, the neurons contributing to the DLS manifold during DLS BCI sessions changed substantially, with different neurons contributing at the start and end of the session.

Together, these results suggest a link between the spatial extent of BCI-driven modulation and its effect on the underlying manifold. In M1, the widespread modulation across the population appears to operate within a stable manifold: many neurons are recruited, but the manifold's dimensionality and its contributing neurons remain largely unchanged. In DLS, by contrast, the more localised modulation restricted to a smaller set of neurons is accompanied by a reorganising manifold, one that gains dimensions and recruits different contributing neurons over the course of learning. This raises the possibility that, in DLS, achieving the required change in a small set of neurons is not simply a matter of recruiting existing activity patterns, as in M1, but instead requires the manifold itself to change to accommodate it.

3.4 Discussion

The differences and similarities between M1 and DLS are still one of the main open questions in motor control. Despite them showing similar activity in some contexts [88, 91], many studies suggest specialised, complementary functionality [79, 92]. Furthermore, the distinct anatomy and cytoarchitecture of the areas likely impose constraints on the population activity they can generate. Here, we tested that hypothesis using a causal perturbation method, an operant BCI. During the BCI sessions, animals were required to

increase the firing rate of neurons from either M1 or DLS. While imposing that change, we simultaneously recorded hundreds of neurons from both areas to explore the effect of the perturbation in the surrounding neural populations. We found that animals could modulate DLS neurons to the same extent as M1 neurons. Furthermore, this modulation had a distinct effect on the whole neural population, with M1 showing a wider effect compared to DLS.

3.4.1 An Operant BCI task as a tool to study neural constraints on volitional modulation

BCI tasks are a tool to probe the capacity of the brain to volitionally control its neural activity. They are arguably one of the best state of the art manipulation techniques to create perturbations on the brain to establish causal relationships. One of their main advantages is that the perturbations are always internally generated, which ensures that they lie within the natural constraints in the neural populations. By measuring the performance of the BCI task, one can estimate the extent and efficacy of said perturbations.

Our specific task design offers some further advantages, beyond those shared by BCI tasks in general. The self-paced spherical treadmill allowed animals to move freely rather than being restrained, giving a more naturalistic behavioural repertoire than tasks requiring a trained, stereotyped movement, and requiring comparatively little pre-training before recording could begin. Furthermore, because the BCI firing rate is a single, continuous control signal derived directly from a small set of neurons, it provides a common metric that can be compared across brain regions. This type of control signal also carries limitations, however. Its binary, threshold-based reward is a considerable simplification compared to the continuous kinematic control required in naturalistic reaching, which may limit how directly our findings generalise to more complex volitional control. Additionally, the reward condition does not account for differences in firing statistics between M1 and DLS neurons: if the two regions differ systematically in their baseline firing rate or variability, the threshold could be effectively easier or harder to reach in one region than the other,

independent of any true difference in volitional control. The adaptive threshold, based on each region's own recent activity, may alleviate this limitation, but it is still worth bearing in mind when comparing overall BCI performance between M1 and DLS.

In our specific case, tracking performance was challenging for several reasons. First, the BCI had an adaptive threshold that changed every 5 mins, to ensure mice engagement and a working BCI over an extended period of time. Hence, usual performance metrics, such as the number of rewards or time to achieve the reward, could not be used as a learning proxy. Second, the neural recordings used to control the BCI were acute, that is, we inserted and retracted the probe every day, inevitably recording different neurons each day. This made it impossible to track long-term learning and performance of the BCI.

Furthermore, we did not see a significant increase in the BCI firing rate throughout the session, although we did see an increase in some individual neurons. Importantly, a session-long increase in firing rate was neither needed to solve the BCI task nor explicitly rewarded, so the lack of a strong effect does not contradict volitional control. Additionally, every session, the selection of the BCI channel was based on visual inspection of the firing rate before the BCI, which biased the channels toward those containing already highly active neurons and could explain why the further increase in firing rate may have been challenging.

Remarkably, the animals still showed evidence of intentional control of the BCI by anticipatory licking. The anticipatory licking effect could not be explained by changes in the BCI firing rate or the visual feedback. The animals were able to predict the reward without any external clues, which implies that there were internal cues predicting the reward delivery. That is, volitionally controlling the BCI to get the reward. The anticipatory licking effect showed significant variability across sessions but mostly occurred during the later trials in the recordings, consistent with BCI learning within the session.

This result is informative beyond simply confirming volitional control. Licking itself does not drive the BCI firing rate: random licks throughout the session did not change it

(Fig. 3.4C, D), so licking is not part of the control strategy itself, but rather a read-out of an internal state that anticipates the outcome of that strategy. The anticipatory licking indicates that, whichever strategy was used, the animal had some internal representation of its own likely success before any external feedback confirmed it, since visual cursor value only began to reliably predict the reward around 0.5 seconds before target acquisition. Finally, this effect was equally present in M1 and DLS BCIs, indicating that both regions supported genuinely volitional, learned control, rather than BCI success arising by chance or through some qualitatively different, non-volitional process.

From these results, we can infer that the animals learned to volitionally control the BCI by momentarily increasing the BCI firing rate enough to get the reward, although long-term changes in the BCI firing rate were not achieved. By observing these fast periods of increased BCI firing rate, we can study how the volitional control of the BCI is not only creating perturbations on the BCI neurons, but also on other neuronal population activity, thus providing insight into neural constraints that may limit their population interactions.

3.4.2 Neural population constraints on BCI-induced neuronal changes

We have found that the volitional increase of neuronal activity in both M1 and DLS influenced the activity of other neurons. We specifically looked at neurons that consistently increased their firing rate together with the BCI neurons during individual trials. The number of modulated neurons was very variable across sessions and not significantly different between the two brain regions, although the spatial distribution of the modulated neurons revealed key differences between the regions.

Previous studies have shown that DLS neural population activity is more sparse and less correlated compared to M1, where the activity is usually well described by a few covariance patterns [7, 22]. Consistent with those studies, we found that, when the BCI neurons were located in DLS, the additionally modulated DLS neurons were very close to the BCI channels. The complete opposite effect was found in M1 BCIs, where the modulated M1

neurons were evenly spread throughout the recorded M1 area.

Interestingly, we found that for some sessions of M1 BCIs, there was also significant modulation in DLS neurons, but this was never the case for DLS BCIs. This was expected given that we were targeting the DLS region where M1 is directly projecting [71, 82]. However, this result not only pointed to that unidirectional anatomical constraint, but it also confirmed our expectations that volitional control of M1 neurons can influence DLS. This carries important implications for the DLS BCIs, where the animals could use the volitional increase of M1 firing rate to provoke an equal increase in DLS neural modulation, but that was never the case. The DLS volitional modulation was always independent of M1 activity, proving these two areas have a great degree of functional independence.

During each recording session, we used a single neuropixel probe, that is, a single-shank long probe. When talking about population-level effects, the electrode dimensions always determine the neural sample acquired, in this case, a single vertical column. This may have favoured areas like the cortex, where the neurons are usually organised by functional columns [3, 170, 171]. The lack of large-spanning modulation in DLS, with modulated neurons located only around the BCI channels, may just mean that the other modulated neurons were located a few micrometres to the side, and they were not sampled. Using other electrode types, such as multiple-shank electrodes, could help address this limitation in the future.

A related consideration is whether the neurons we recorded and selected are representative of the broader cell-type composition of each region, particularly in DLS. Unlike M1, where the recorded population is dominated by pyramidal projection neurons, DLS contains a mix of cell types: medium spiny neurons (MSNs), which make up the vast majority of striatal neurons but fire at very low rates, and fast-spiking interneurons, which are comparatively rare but fire at much higher rates [87, 159]. Because we manually selected BCI channels based on visibly high firing rates, our DLS recordings may have been biased toward fast-spiking interneurons rather than MSNs. We did not classify recorded neurons by cell type so we cannot rule out that a disproportionate number of our DLS BCI and

nearby modulated neurons were interneurons rather than MSNs. This distinction matters for interpretation: MSNs are the principal projection neurons and output of the striatum, whereas interneurons act locally and do not project outside DLS. If our DLS results were disproportionately driven by interneurons, the localised modulation we observed may reflect local inhibitory circuitry. Future work classifying recorded units by waveform [172] would help establish how representative our DLS sample is of the region's principal output neurons, MSNs, and would clarify whether the differences we report between M1 and DLS generalise across cell types or are specific to the more excitable neurons.

A separate, more general concern applies to high-firing neurons in either region, irrespective of cell type. Because the BCI reward was based on an increase in firing rate relative to an adaptive threshold, neurons that already fired at a high baseline rate may have found this increase easier to achieve than slower-firing neurons, simply due to greater natural trial-to-trial variability in their firing. Future work relating each neuron's baseline firing statistics to its individual contribution to solving the BCI task would help establish how much of our results depend on this selection bias, in addition to the cell-type composition discussed above.

It is also important to note that we have only considered as "modulated" neurons those that increased the firing rate together with the BCI neurons. Although visual inspection of the data did not suggest other types of modulation, further analysis may be beneficial to include other types of modulation, such as neurons decreasing their firing rate.

3.4.3 BCI control strategies reveal differences in M1 and DLS

One remaining question regarding the volitional control of neuronal firing rate in M1 and DLS is the underlying process to achieve such modulation. The strategy that animals adopt to solve the BCI task also provides insights into the most efficient methods to intentionally increase the neural activity. In previous studies, when the BCI control signal came from the M1, animals tended to solve the task using movement-based strategies [161, 169]. Although many studies purposely suppress the movement to probe the motor cortex

capacity to generate neural activity beyond motor output [15, 161, 165], here we did not limit the animals to specific behaviours during the task.

As expected, when the BCI neurons were located in M1, the running speed predicted the BCI firing rate in most sessions, since the animals were using running as a means to increase the firing rate of the M1 neurons. In contrast, when the BCI neurons were located in DLS, the running speed could not predict the BCI firing rate in the vast majority of sessions, even if the DLS neural population had the required movement-related information. The lack of M1 modulation during DLS BCIs, together with this result, proves that the animals use fundamentally distinct strategies to solve the BCI.

The striatum is the main input of the basal ganglia; specifically, DLS integrates inputs not only from multiple cortical regions like somatosensory or premotor cortices [71, 82], but also from subcortical regions like other basal ganglia or the thalamus [81, 173]. This provides the animals with a rich set of possible strategies to modulate their DLS activity not related to motion. Although we know that animals on a DLS BCI do not use locomotion as a control strategy, the specific mechanisms remain elusive, and it is unclear how many alternative strategies can actually be volitionally used in a BCI context.

Broadly, animals could modulate the BCI neurons through at least three, non-mutually exclusive mechanisms. The first is overt behaviour: generating an actual movement or postural change that happens to co-activate the BCI neurons, as we found for M1 BCIs, where running speed predicted the BCI firing rate [161, 165, 169]. The second is an internal strategy that does not require any overt movement, relying instead on other processes, such as changes in arousal, attention, or motivation, that can modulate cortical and striatal firing rates without a corresponding behaviour change, as proven by classic work where movement was explicitly discouraged or physically prevented [94]. The third is an imagined action, that is, internally simulating a movement without executing it. This strategy is well documented in BCI studies with paralysed human patients, who can achieve fine BCI control purely through attempted or imagined movement in the absence of any overt behaviour [174]. Our results indicate that DLS BCIs did not rely

on the first mechanism, since locomotion could not predict the BCI firing rate. However, our task cannot distinguish between other types of movement, or the second and third mechanisms, as they would produce a similar pattern of results; future work measuring additional physiological or behavioural correlates would be needed to disentangle them.

3.4.4 Neural manifolds as a static versus dynamic constraint on BCI learning

While physiological or behavioural correlates of these strategies remain to be measured directly, the neural manifold itself offers a complementary window onto how each region achieved BCI control. In M1, achieving the BCI target through widespread modulation did not require the manifold to change. The same stable, correlated population and the same fixed dimensionality supported BCI control throughout learning, allowing the BCI target to be reached simply by recruiting existing activity patterns more widely. In DLS, by contrast, the more localised modulation was accompanied by a manifold that reorganised over the course of learning, gaining dimensionality and recruiting a changing set of contributing neurons. This suggests that DLS's architecture does not offer the same repertoire of readily-recruitable, widely-shared activity patterns that M1 does, demanding the manifold itself to slightly change.

This analysis has some limitations worth noting. First, the baseline difference in dimensionality between M1 and DLS did not reach significance ($p<0.1$). Although the unconstrained movement could be an explanation of the lack of structure in both M1 and DLS, a larger number of sessions would be needed to establish it more firmly. Furthermore, because manifold dimensionality estimated through PCA is sensitive to the number of neurons recorded, we repeated this analysis by randomly subsampling 50 neurons per region, over 10 iterations, confirming that our dimensionality estimates were not simply driven by differences in neuron yield between regions or recording windows. Second, the apparent change in which neurons contributed most to the DLS manifold could partly reflect our choice to examine only the top principal components. A lack of neuron contribution could

reflect a priority change in which its component simply shifts to a later position in the manifold. Reassuringly, the striking stability we observed in neural contributions to the M1 manifold, measured with the same approach, suggests that our method can detect genuine stability when it is present, and provides an upper bound against which the comparatively large change we observed in DLS can be judged. Despite these limitations, this analysis is, to our knowledge, the first to directly link the spatial extent of BCI-driven modulation to changes in the underlying neural manifold, offering a promising path for understanding how different circuits support learning.

Overall, these results support our view that the striatal population activity is more than a reflection of the ongoing activity in upstream M1. Furthermore, it points toward these areas having key differences in their neural population activity, with striatal activity being more locally constrained than cortical activity, possibly allowing for finer BCI control.

3.5 Conclusion

In this study, we used a BCI task on head-fixed mice to probe the volitional control of individual neurons in primary motor cortex (M1) and dorsolateral striatum (DLS). We showed that, similar to what was previously reported for M1, animals could also intentionally increase the activity of striatal neurons. Furthermore, this volitional control induced changes in the neural populations that were distinct for each region and constrained by the different cytoarchitecture of the neural populations. Intriguingly, we found that the modulation induced by DLS BCIs did not affect M1 population activity and that the strategy to produce such modulation is different for both regions, proving these two areas have a great degree of functional independence.

These findings support the view that DLS makes independent functional contributions to behaviour beyond simply reflecting cortical inputs. Furthermore, the ability to volitionally modulate DLS activity, as traditionally reported for other cortical areas, opens the possibility to use BCI approaches to causally probe subcortical constraints, like the tight correlation between M1 and DLS. These results provide an exciting starting point for BCI

perturbation studies in subcortical structures.

3.6 Methods

3.6.1 Subjects

This study included five adult female CD1 mice. All animal surgeries and procedures were performed in accordance with the guidelines established by the UK Home Office under the name "Animals (Scientific Procedures) Act 1986". All the procedures were performed under the project licence number PP1496969. The animals were mostly caged in groups with up to three other mice of the same litter per cage.

3.6.2 Surgeries and regulated procedures

Head implant

After an initial habituation (details in Section 3.6.5 "Habituation and training protocol"), we performed the first surgery to attach a head implant for head-fixation during the experiment. Mice were anesthetised with isoflurane, with an initial concentration of 5% for induction and a continuous concentration of 1.5-2% during the surgery. Breathing and temperature of the animals were tracked during the surgery to ensure light but effective anaesthesia. Once the animals were fully anaesthetised, we gave them subcutaneous injections of carprofen and buprenorphine. Eye ointment was applied during the surgery to protect and moisten the eyes while under anaesthesia. We prepared the scalp with a shaver and hair removal cream until the skin on the top of the head was clear of hair. We disinfected the area with iodine and then, using forceps and sharp scissors, we removed the scalp to expose the skull. Once the bregma and lambda landmarks were exposed, we measured the distance between them for accurate trajectory planning (details in Section 3.6.6 "Electrophysiological recordings"). Once the skull was dry, we attached the head implant using clear SuperBond dental cement to visually locate bregma, needed for calibration during the recordings. Tissue glue was applied to any remaining open wounds

around the implant. Animals were given water mixed with carprofen for three days after the surgery. Animals were closely monitored and weighed daily for seven recovery days before any other procedure.

Craniotomy

Once the animals were trained for the BCI, we performed a small craniotomy to expose the brain for the neural recordings. Animals were put under anaesthesia as previously described. Again, carprofen and buprenorphine were administered subcutaneously for this surgery. After disinfecting the area, we used a small 0.5mm diameter drill to perform a small craniotomy located +0.5mm AP, +1.5mm ML from bregma. Small craniotomies are less prone to infection and can be maintained for more days. After washing the craniotomy with saline, we covered it with Kwik-cast bio-compatible silicone to keep it moist and protected for the recording session. Same as in the previous surgery, animals were given water mixed with carprofen for three days after the surgery. After this surgery, mice were kept in a separate cage to avoid other cagemates damaging the Kwik-cast cover.

Perfusion

Once the experiment was finished, the animals were euthanised, and a perfusion was performed to extract the brain. First, we gave an overdose IP injection of sodium pentobarbital of around 0.1 ml. Once the animal was at the surgical plane of anaesthesia, we exposed the heart and, using a pump, we let PBS flow through the circulatory system to remove the blood. Once the blood had been cleared, we changed the PBS to a 4% PFA solution and let it flow through the animal for tissue fixation. Finally, once the animal was rigid, we used scissors and forceps to expose the brain and transfer it to a PFA vial left in the fridge for 2-24h for further fixation. When we were ready to section the tissue, we transferred the brain to a PBS solution for 24 hours before histology.

3.6.3 Histology

Histology post-experiment was performed to verify the electrode location. During the last recording session, we impregnated the electrode with Dil stain before insertion to locate the electrode placement after the experiment. Then, to prepare the tissue for imaging, we sliced the brain into 60 μm thickness coronal slices using a vibrotome. The slices were transferred to Superfrost slides and left to dry overnight before staining. When ready to image, we applied a mixture of Fluoromount-G and Fluoromount-G with DAPI, and slowly applied a coverslip, spreading the solution through the slide, minimising air bubbles. After one hour of drying time, the tissue was ready to image on a fluorescence microscope. The obtained images were later compared with a mouse brain atlas to confirm the electrode placement.

3.6.4 Experimental setup

The experiment was performed in a spherical treadmill consisting of a 16-inch-diameter hollow styrofoam ball on top of seven ping-pong balls. The ping-pong balls were set in air cannons, allowing compressed air to push the ping-pong balls, and holding the spherical treadmill floating and moving freely. The animals were head-fixed on top of the treadmill using the RIVETS system [175]. The experimental setup included multiple peripheral devices to interact with the animal controlled by a standard microcontroller called Pycontrol.

To provide rewards and measure the licking events, we used a Sanworks optical lickometer. It consists of a stainless steel hypodermic drinking tube surrounded by an infrared emitter and a phototransistor that detects when the animal's tongue touches the tube to get a reward. The drinking tube was connected to a reward syringe with a solenoid to control the flow. The reward solution consisted of an artificial sweetener (acesulfame potassium) mixed with tap water at a ratio of 800mg/L. Both the lick detection and reward delivery were controlled by Pycontrol.

To provide visual feedback to the animals, an Adafruit DotStar flexible 72-LED strip was placed as a semi-circle in front of the treadmill. The LED strip was controlled by a separate Rasberry Pico Pi 2 microcontroller running CircuitPython to use the Adafruit custom software. The Raspberry microcontroller interacted with Pycontrol via UART. LEDs could show any RGB colour independently, but we mostly used a background red colour in all the LEDs and two consecutive LEDs tuning white for the cursor position(details in Section 3.6.5 "Habituation and training protocol").

To record the rotation of the ball, two standard mouse sensors were placed on the equator of the ball, recording the longitudinal and lateral displacement with respect to the head-fixed mouse. To get a better estimation of the animal's behaviour, we also installed seven FLIR cameras around the setup for a whole-body pose estimation. We used JARVIS software to record the videos at 100 frames per second synchronous to the other Pycontrol signals. Unfortunately, the pose estimation software traditionally used (SLEAP or DeepLabCat) did not yield good results for these sessions due to the visual cursor and ball rotation confusing the models. Although we are currently working on the processing of the videos, they were not included in the thesis.

3.6.5 Habituation and training protocol

Habituation

Once the animals arrived at our facilities, they were allowed 7 days of acclimatisation before any handling. After this period, we did 15-20 minutes of handling sessions daily, where the animals were picked up and allowed to freely move on the experimenter's forearm. Once the animals were comfortable and relaxed being picked up and handled, we proceeded with the first surgery to get the head implant (details in Section 3.6.2 "Surgeries and regulated procedures"). Animals were allowed 7 days of recovery after the surgery, during which they were daily weighed and minimally handled. After that period, animals started the habituation to the experimental setup.

The first day, the animals were exposed to the experimental room, where they were placed

in a square-shaped foam that mimics the treadmill material for at least 5 minutes to freely roam. During this time, we also turned on the compressed air valve that raises the treadmill to habituate them to the ambient noise that will be present in the experiment. On the second day, the animals were carefully head-fixed in the spherical treadmill for no more than 5 minutes. On the following days, we increased the duration of head-fixation to 10 minutes, and then to 20 minutes. Any of the previous steps could be extended if the animal still showed signs of stress or discomfort.

Once the animals were comfortable with the head-fixation, didn't show any sign of distress, and they mostly ran on the treadmill, the habituation was complete, and we proceeded to water restriction for behavioural training.

Water restriction

All mice were water-restricted to increase the motivation for reward during the experiment. We removed the water bottle from the cage and carefully controlled the water intake every day. We recorded the initial weight before the water restriction and weighed the animals daily to control for dehydration. In the unlikely case that the weight dropped more than 15% of the initial value, animals were taken out of the water restriction protocol until the weight was recovered.

The animals would mostly get their daily water intake from the training or experiment sessions, although if necessary, mice received extra water in their cages after the session by giving them access to the water bottle for a fixed amount of time.

Training protocol

The training protocol was divided into two steps, with the training time varying from 30 minutes in the first step to 45 minutes in the second step. We occasionally increased the training time by 15 more minutes if the animals had not lost engagement on the task.

On the first training task, we introduced the reward delivery system and the reward sign. Mice were head-fixed and, after a few minutes, we moved the reward spout near their

mouth. Some reward drops were left on the tip of the spout to encourage initial licking. Anytime the animals licked, they would get a reward, and the central LED light would turn on, signalling a reward. After one or two days in this task, mice were usually comfortable, engaged in licking and getting a lot of continuous rewards. At this point, we moved to the second task.

On the second task, we introduced the visual feedback, that is, the whole LED strip. We showed a bilateral visual cursor, similar to the BCI feedback, continuously moving toward the centre position. Once the centre LED was reached by the moving cursor, it was held there for 3 seconds waiting for the animal to lick. If the animals licked during this holding time, the reward was delivered. Whether the animal licked or the holding time passed, the trial restarted, and the bilateral cursor went back to the initial side positions. Animals took between three and five days to learn this task, with around 1/3 of the animals losing interest and not learning the task at all. Once the animals licked consistently at the correct times, we considered the training completed and the animal ready for the BCI task.

3.6.6 Electrophysiological recordings

We used Pinpoint [176] to plan a trajectory for the Neuropixel probe that hit our main targets, the forelimb area of M1 [170] and the projecting region in dorsal striatum [156]. Pinpoint provided the angles and coordinates of insertion to accurately hit both targets with the same Neuropixel probe. For this experiment, the planned insertion coordinates were +0.5AP, +1.5ML from bregma with 115 and 75 degrees of yaw and pitch angles, respectively.

Once the craniotomy was performed (details in Section 3.6.2 "Surgeries and regulated procedures") and the exposed brain was covered in Kwik-cast, we proceeded with the electrophysiological recordings. We head-fixed the animal and removed the Kwik-cast. We added saline to the craniotomy to keep it moist and protected from external agents before inserting the Neuropixel probe.

The insertion was performed using the New Scale Tech Multi-Probe Micro-Manipulator

System, which allowed us to move the probe in stereotactical coordinates to the desired location and perform automatic insertions at a controlled speed. To avoid tissue damage, we inserted the probe at 150 μm/minute for 4100 μm, until all the recording sites were inside the brain.

All the recordings were made with SpikeGLX, an open-source software for extracellular, high-density probes. SpikeGLX simultaneously recorded the data of the 384 Neuropixel channels at a sampling frequency of 30 kHz, and also stream the data to a server that will be read by the BCI code in Python.

After the recording session, we extracted the Neuropixel probe at 150 μm/minute to minimise the tissue damage again. We soaked the Neuropixel probe with a Tergazyme solution of 150g/L for cleaning and then in water for rinsing until the next recording session. The craniotomy was also cleaned a couple of times with saline before reapplying the Kwik-cast to protect the craniotomy for the next recording session.

We recorded at the same craniotomy site for up to eight insertions until the tissue was too damaged to insert safely or the number of neurons detected was very low.

3.6.7 Real-time Brain Computer Interface

Once the Neuropixel was inserted, we started the session with a prerecording in which the animal was freely moving on the treadmill. We recorded the neural activity for 2 minutes. Using this prerecording data, we detected and discarded the channels, if any, with a high noise value. We first applied a Common Average Reference (CAR) of the data by subtracting the average signal across all channels for each time point to reduce common noise across the channels. The result was filtered using a 300-6000Hz Butterworth band-pass filter for each channel, optimal for spike sorting. Then the root mean square (RMS) of each channel is calculated as:

$$\mathrm{RMS}_i = \sqrt{\frac{1}{N}\sum_{j=1}^{N} x_{ij}^2}$$

where i corresponds to each channel and j to each time point. These RMS values were used throughout the BCI for spike detection. Finally, we discarded any channel where $\text{RMS}_i > \text{median}(RMS_i) + 2.5\ \text{std}(RMS_i)$.

Using the prerecording data, we also calculated the BCI threshold crossing rate threshold to initialise the BCI. To do so, we first separated the filtered data from each BCI channel into 100ms bins, the same size that was used during the BCI. Then, for each time bin, we calculated the number of putative spikes as the number of times the signal went under $-3.5 \times RMS_i$. This way, we obtained the threshold crossing rate of the BCI channels for a 100ms step. Then we removed the outliers by discarding the time bins where the threshold crossing rate was higher than the 99th percentile of the distribution. Finally, we got the threshold for the BCI as the 99th percentile of the remaining samples.

After that, the BCI started. During the BCI, we were collecting and processing the neural data in real time. Our custom Python pipeline collected the data and accumulated it for 100 ms. After that, we applied the CAR, filter and threshold crossing detection to the data for that 100 ms bin, like for the preprocessing data. The BCI threshold crossings rate was calculated as the weighted average of the BCI channels' threshold crossing rate in that time bin. The weights were defined by a gaussian window of s.d: 1.5 points. Once the BCI threshold crossings rate for that time bin was calculated, it was mapped into the visual cursor, that is, a number between 1 and 100, where 1 would be the side LEDs, and 100 would be the centre LEDS. The cursor value for the time bin j was calculated as:

$$\text{cursor value} = \max\left(1, \min\left(\text{round}\left(\frac{BCI\ FR_j}{BCI\ FR_{\text{th}}} \times 100\right), 100\right)\right)$$

The BCI threshold was updated every 5 minutes following the same steps as for the preprocessing threshold. This made the BCI robust and engaging for the animals, even if the BCI threshold crossings rate changed during the session.

3.6.8 Data processing

For neural data processing after the recordings, we used a similar approach to that in the previous study. We did a spike sorting on the electrophysiological recordings using Kilosort4 [168] to identify the individual neuron spiking. Kilosort4 identifies individual neurons using clustering algorithms, accounting for possible drift in the channels during the recordings. Then, we converted the spike times of the identified neurons into a smooth firing rate per neuron by applying a Gaussian kernel (s.d: 50 ms) to the binned square-root-transformed spikes, using 30 ms bins. We only excluded units with a mean firing rate smaller than 1Hz.

We identified the BCI neurons as the neurons whose action potentials were mainly detected on the BCI channels, as informed by Kilosort4. We calculated the BCI firing rate as simply the standard average between all the firing rates of the BCI neurons.

During the recording sessions, the neural data were synchronised with the behavioural recordings, that is, the movement data and the licking events. Both signals were also binned at 30ms to match the neural data. Furthermore, the movement recordings were similarly smoothed by applying a Gaussian kernel of 50ms s.d.

3.6.9 Firing rate percentage of change

We measured the modulation of each neuron to the reward to evaluate how much the BCI changes are affecting the neighbouring neural populations. This modulation is measured as a percentage of change during the rewarded trials compared to the baseline activity of that neuron. We defined the baseline activity as the all-time average firing rate of that neuron.

For each rewarded trial, we calculated the mean firing rate in a window of -0.1 seconds to 0.1 seconds after the reward events. Hence, we calculated the percentage of change individually for each trial, i, and neuron, n, as follows:

$$Percentage\ of\ change_{n,i} = \frac{|reward\ firing\ rate_{n,i} - baseline\ firing\ rate_n|}{baseline\ firing\ rate_n} \times 100$$

The mean and standard deviation of all the trials per neuron are shown in Fig. 3.5C. To classify a neuron as "modulated", we established a threshold of 300% mean percentage of change based on data observation.

3.6.10 Encoding BCI firing rate from movement

To determine how much the motion of the animal was affecting the BCI control, we trained a linear model to predict the BCI firing rate from the longitudinal (x) and lateral (y) motion of the animal.

Following standard methods, we first standardise both predictor variables as follows:

$$[x, y]_s = \left[\frac{x - \mu_x}{\sigma_x}, \frac{y - \mu_y}{\sigma_y}\right]$$

then, we used the following linear model:

$$\hat{FR} = \beta_0 + \beta_x x + \beta_y y$$

where $\hat{FR}$ is the predicted BCI firing rate, and β_0, β_x and β_y are the model coefficients estimated by minimising the sum of squared errors:

$$\arg\min_{\beta_0,\beta_x,\beta_y} \sum_{i=1}^{n} \left(FR_i - \beta_0 - \beta_x x_i - \beta_y y_i\right)^2$$

where FR is the actual BCI firing rate and i is each time point.

The model was fitted using 10-fold cross-validation, that is, the session was divided into 10

equally sized folds. Then, 9 folds were used for training and 1 for testing, sequentially for all the folds. We estimated the model performance of each fold as the explained variance:

$$\mathrm{EV} = 1 - \frac{\mathrm{Var}\left(FR - \hat{FR}\right)}{\mathrm{Var}\left(FR\right)}$$

The final performance for each session was the average explained variance for all the folds.

3.6.11 Decoding movement from M1 or DLS population activity

To predict the movement from the neural population, we first applied PCA to the population firing rate to get the neural population latent dynamics as explained in Section 2.6.5. Similar to previous studies, we kept 10 components, hence reducing the data dimensionality.

We then repeated the standardisation and linear model fit explained in Section 3.6.10 with two main differences. First, in this case, we had 10 predictor variables instead of 2. Second, we fitted two different models, one to predict longitudinal velocity and another to predict lateral velocity. The final performance for each session was the average between the two models' explained variance.

3.6.12 Neural manifold analysis

Manifold dimensionality was defined operationally as the minimum number of principal components required for the cumulative explained variance to reach 85%. Because this estimate depends on the number of neurons entering the PCA, a fixed population size of 50 neurons was used throughout: for every session and area, 50 neurons were drawn at random from all available neurons in that area (sessions/areas with fewer than 50 neurons were excluded), and the dimensionality estimate was computed on that subsample. To reduce sensitivity to any single random draw, this procedure was repeated 10 times per session and area, and the resulting dimensionality values (and cumulative-variance curves) were averaged across repetitions to give one value per session and area.

To ask whether prolonged engagement with the task reorganises either region's manifold, each session's neural data were further split into an early window (the first 15 minutes of task) and a late window (the last 15 minutes). The same fixed-population-size, 10-repetition dimensionality procedure described above was applied to the early and late windows, separately for DLS-BCI and M1-BCI sessions and separately for the DLS and M1 populations. They were statistically compared using a two-sided Wilcoxon signed-rank test on the paired per-session differences.

Finally, to test whether it is the same neurons that define the manifold before and after prolonged task engagement rather than simply asking how many dimensions are needed, PCA was fit separately on the early- and late-window data (same 50-neuron subsamples), fixed to 10 principal components in each window, yielding two neuron-by-component loading (weight) matrices. These loading matrices were compared via canonical correlation analysis (CCA), treating neurons as observations and the 10 PC weights as features, and the mean of the top four canonical correlations was taken as a single per-session index of neuron-contribution similarity between the early and late manifolds (1 = an identical set of neurons and weightings defines the manifold in both windows; lower values indicate that different neurons come to dominate the manifold over the course of the session). As above, this was computed per 50-neuron subsample and averaged across 10 repetitions.

Chapter 4

Conclusion

4.1 Summary

Different brain regions present a wide variety of circuit architectures that constrain the neural signals they produce. However, the way these constraints shape neural signals and circuit function remains poorly understood. In this thesis, we addressed this question from two complementary angles: how the relationship between neural signals generated by the same cortical area depends on the specific cytoarchitecture of the region, and how causal manipulation of neural dynamics provokes different effects in cortical and subcortical populations, consistent with their cytoarchitectural constraints.

In Chapter 2, we characterised the relationship between the local field potentials (LFPs) and the neural population dynamics of three regions of the primate sensorimotor cortex during a delayed-reaching task. The correlation between the latent dynamics and the LFPs was frequency-dependent and varied across regions: low- and high-frequency bands were correlated in all three cortical areas, but mid-range frequencies were near zero in

M1 but elevated in PMd and area 2. The relationship found was stable not only during movement execution, but also throughout planning, extending the finding to the more abstract cognitive process of movement planning. This correlation is mainly driven by synaptic currents within the neural population, evidenced by the marked decrease in LFP-latent dynamics correlations during inter-trial intervals, when there was minimal activity in the region. Furthermore, these stable correlations cannot be predicted by single neuron activity, requiring a population point of view to fully capture the synaptic current effects. Together, these results showed that the cytoarchitectures of the circuits, which shape the synaptic currents, impose stable constraints on the neural activity of the regions. Furthermore, these robust associations help bridge the studies using either type of recording individually [7, 30, 64].

In Chapter 3, we causally tested the cytoarchitectural constraints in two highly connected cortical and subcortical regions of the sensorimotor system, the primary motor cortex (M1) and the dorsolateral striatum (DLS). We designed an operant brain-computer interface (BCI) task where head-fixed mice running on a spherical treadmill increased the firing rate of a set of neurons in M1 or DLS. The animals demonstrated volitional control of both M1 and DLS BCIs, evidenced by anticipatory licking that could not be predicted by task parameters. The volitional increase in firing rate of specific neurons affected the surrounding neural population to a different spatial extent for each region. M1 BCI modulation was widespread across the M1 region, whereas DLS BCI modulation was localised around the BCI neurons. Additionally, M1 BCIs also modulated some neurons in DLS but never the other way around. These modulation patterns are consistent with the different cytoarchitectural properties of either region and show how these properties constrain the neural activity. The animals also adopted different control strategies to solve the BCI, from an expected movement-based strategy for M1, to a non-movement-based strategy in DLS despite the area also containing movement information. Together with the lack of M1 modulation on DLS BCIs, these results point toward DLS activity being more than a passive reflection of M1 inputs. It operates under distinct constraints that allow for targeted, independent modulation.

Together, both chapters show that circuit architecture in the sensorimotor system imposes distinct and measurable constraints on the neural population signals that are reflected in LFPs, population latent dynamics, and volitionally modulated firing rates. These findings show that circuit architecture shapes the neural signals stably and distinctly, which in turn constrains the circuit function of each brain region that collectively produces a behaviour.

4.2 Impact

Although extensive research has focused on the contributions of each brain region to behaviour, the specific circuitry of each region imposes constraints on the generated neural activity, limiting their contribution to behaviour. This thesis explores the effect of those constraints in multiple sensorimotor regions, aiming to build a bridge between circuitry and function.

Here, we have presented a stable relationship between the neural population dynamics and the LFP at traditional frequency bands. In M1, the strength of these relationships also predicted how strongly each frequency band related to ongoing motor output, which shows how these circuit-level constraints are informative of the functionality of the region. More importantly, these results provide a crucial bridge between two largely separate bodies of literature, focused on the contributions of either LFPs or neural population dynamics. This systematic characterisation of the LFP frequency bands with respect to the neural population dynamics provides a framework for translating the findings across recording modalities. This has particular relevance for regions highly studied through LFPs, such as the hippocampus [117, 132]. Furthermore, it also provides a methodological framework to characterise LFP-population dynamics relationships using canonical correlation analysis (CCA) that could be applied in any brain region where simultaneous LFP and spiking recording is available [9, 17], and captures population-level relationships that traditional single neuron analyses overlook. These relationships are not only limited to the sensorimotor system and correlational studies, but could have a higher impact in clinical settings, where the use of LFP is common as a biomarker for diagnosis

and closed-loop stimulation [54, 56, 57]. For example, the finding that the 12-25 Hz band is strongly correlated in PMd and area 2 despite carrying very little movement information [107, 120] suggests that it may reflect circuit states related to cognitive processes such as attention or task structure, which could be useful for different purposes.

Furthermore, in this work we have shown that exploring these cytoarchitectural constraints could contribute to long-standing debates in the field such as the functional independence of DLS from M1. Here, we provided causal evidence that DLS can be modulated independently of M1, despite the strong anatomical connectivity [91, 92, 95]. Additionally, we have shown that M1 and DLS have distinct constraints on the spatial extent of their neural population modulation consistent with their cytoarchitecture [22, 87], providing further evidence on how the cortical and subcortical circuits are organised fundamentally differently. This again could be particularly useful for clinical applications where a focal modulation of DLS may be desirable.

Here, we also presented a new use of operant BCI tasks as tools for probing circuit constraints on head-fixed mice, using the BCI to answer fundamental questions of the neural circuits rather than the traditional use of BCIs for decoding and rehabilitation. This approach can be extended to other brain regions and neural perturbations, limited only by the natural constraints of the circuit, providing a powerful tool for studying circuit constraints at the population level [15, 16, 94, 95].

4.3 Limitations

While this thesis provides valuable insights into the constraints on population dynamics across cortical and subcortical sensorimotor circuits, a few limitations should be acknowledged.

In this work, we used non-human primates and mice to test our hypotheses. Although the sensorimotor system is well characterised in both, as for any work using animal models, it is unclear how well the findings generalise to other species, including humans. Furthermore,

we specifically characterised the relationship between LFPs and neural population dynamics only in three cortical regions (M1, PMd, and area 2 of the somatosensory cortex); whether the same relationships hold in other cortical areas or in subcortical regions remains to be tested.

Throughout the thesis, we have used a range of linear methods to test our hypotheses: principal component analysis (PCA) for data dimensionality reduction, CCA for LFP-latent dynamics alignment, linear regression for decoding of movement and encoding BCI firing rate and naïve Bayes classifiers for target classification from planning activity [9, 17, 133, 153]. Linear methods are powerful tools that are computationally efficient and capture prominent effects in the data, also providing high interpretability. However, by using PCA to identify the neural manifold and CCA to quantify its relationship with the LFPs, we assume that neurons and LFP bands relate to each other through linear combinations. Similarly, the naïve Bayes classifiers used to decode reach direction rely on largely linear decision boundaries, and the Wiener filters used for velocity decoding mapped LFP power or latent dynamics to velocity linearly.

Nonlinear methods could reveal additional relationships, especially between LFPs and neural population activity, that our linear approach may have missed. This could be particularly relevant for the intermediate frequency bands (8–50 Hz), where we found near-zero linear correlations between LFP power and latent dynamics in M1. A nonlinear relationship in these bands, if present, would suggest that these frequencies are not functionally disconnected from the population dynamics, but rather related to them in a more complex manner than the low- and high-frequency bands. Importantly, the identification of further non-linear relationships would not contradict our main findings, since the surrogate-based controls confirmed that the linear associations we report are unlikely to arise by chance, and they were robust across sessions and animals, even using different numbers of dimensions. Rather, it would suggest that the linear component we characterised here represents a lower bound on the true relationship between LFPs and neural population activity, and that the full picture of how circuit architecture constrains

these two signals may be richer than we were able to capture here.

The operant BCI task presented several limitations mentioned in Chapter 3 discussion (Section 3.4). Although the acute recordings were sufficient to see volitional control of the BCI, the fact that we recorded different neurons each day made it impossible to track long-term learning on the BCI, even if we recorded daily for more than a week in some animals. Furthermore, due to the adaptive threshold, the better the animals performed, the harder the task became, preventing us from using standard performance metrics such as number of rewards or time to get a reward. Additionally, we chose the BCI control signal, that is, the five channels that will map to the BCI cursor, visually inspecting the electrophysiology data, which may have induced inconsistencies across sessions. Reassuringly, most sessions showed some anticipatory licking, proving volitional control, and all of them showed a significant modulation increase around the reward times for the BCI neurons and also the neighbouring population.

Regarding the neural recordings, using a single-shank Neuropixel limited the neural sample to a single vertical column, significantly biasing the neighbouring neurons recorded for the modulation analysis [167]. Other electrode shapes may show a wider modulation spread in DLS than the one we found. Furthermore, we recorded the BCI task only in five animals, with most of the sessions focused on DLS BCIs. Although this was originally intentional, as we wanted to prove the volitional control of DLS, it reduced the statistical power of the cross-area comparisons. Increasing the sample size, especially in M1 BCI, could show more robust differences between the two areas and increase the statistical power of our comparisons.

4.4 Future work

The region-specific relationship between LFPs and neural population dynamics shows a fundamental connection between the two neural signals that could be extended to other brain regions where the LFP oscillations have been extensively studied, such as the basal ganglia or hippocampus [63, 132]. Characterising how traditional LFP signals, such as the

theta oscillations in the hippocampus [117] or the beta LFP in the subthalamic nucleus (STN) in Parkinson's disease [54], relate to the underlying neural population activity could contribute to the understanding of these neural phenomena. However, validating the LFP-latent dynamics framework in human patients, possibly extending to ECoG or DBS recordings, is necessary to evaluate its potential for clinical translation [50, 51].

Regarding the operant BCI, the most imminent technical improvement is the use of chronic recordings instead of acute insertions. This would allow tracking of individual neurons across days and enable direct measurement of BCI learning and long-term modulation in DLS and M1 population activity [17]. Other improvements, such as multiple shanks or multiple probes, could also provide a wider neural sample to study the extent of the BCI-driven modulation, as well as provide insight into whether DLS volitional modulation involves inputs beyond M1.

An operant BCI task provides flexibility to causally prove natural neural constraints. Extending the study to other brain regions such as STN could provide insight into neural limitations and the effects of therapies such as DBS [39, 56]. Furthermore, extending the BCI control to a more complex signal, such as the correlation between two areas [15, 164], could help disentangle the activity between highly connected regions such as M1 and DLS.

Overall, the work presented in this thesis represents a step towards a systematic understanding of how circuit architecture constrains the neural population signals across the sensorimotor system. Understanding these constraints is not only a basic neuroscience question but also a practical one, providing a better understanding of the neural signals for designing better tools to interact with the brain.

Bibliography

[1] Stephen H. Scott. "The computational and neural basis of voluntary motor control and planning". In: *Trends in Cognitive Sciences* 16.11 (Nov. 1, 2012), pp. 541–549.

[2] Paul Cisek and John F. Kalaska. "Neural Mechanisms for Interacting with a World Full of Action Choices". In: *Annual Review of Neuroscience* 33 (Volume 33, 2010 July 21, 2010), pp. 269–298.

[3] Kenneth D. Harris and Gordon M. G. Shepherd. "The neocortical circuit: themes and variations". In: *Nature Neuroscience* 18.2 (Feb. 2015), pp. 170–181.

[4] György Buzsáki, Costas A. Anastassiou, and Christof Koch. "The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes". In: *Nature Reviews Neuroscience* 13.6 (June 2012), pp. 407–420.

[5] Gaute T. Einevoll et al. "Modelling and analysis of local field potentials for studying the function of cortical circuits". In: *Nature Reviews Neuroscience* 14.11 (Nov. 2013), pp. 770–785.

[6] A. P. Georgopoulos, A. B. Schwartz, and R. E. Kettner. "Neuronal population coding of movement direction". In: *Science* 233.4771 (Sept. 26, 1986), pp. 1416–1419.

[7] Juan A. Gallego et al. "Neural Manifolds for the Control of Movement". In: *Neuron* 94.5 (June 7, 2017), pp. 978–984.

[8] Saurabh Vyas et al. "Computation Through Neural Population Dynamics". In: *Annual Review of Neuroscience* 43.1 (2020), pp. 249–275.

[9] John P. Cunningham and Byron M. Yu. "Dimensionality reduction for large-scale neural recordings". In: *Nature Neuroscience* 17.11 (Nov. 2014), pp. 1500–1509.

[10] Juan A. Gallego et al. "Cortical population activity within a preserved neural manifold underlies multiple motor behaviors". In: *Nature Communications* 9.1 (Oct. 2018), p. 4233.

[11] Krishna V. Shenoy, Maneesh Sahani, and Mark M. Churchland. "Cortical control of arm movements: a dynamical systems perspective". In: *Annual Review of Neuroscience* 36 (July 8, 2013), pp. 337–359.

[12] Peiran Gao and Surya Ganguli. “On simplicity and complexity in the brave new world of large-scale neuroscience”. In: *Current Opinion in Neurobiology*. Large-Scale Recording Technology (32) 32 (June 1, 2015), pp. 148–155.

[13] Matthew T. Kaufman et al. “Cortical activity in the null space: permitting preparation without movement”. In: *Nature neuroscience* 17.3 (Mar. 2014), pp. 440–448.

[14] MM Churchland et al. “Cortical preparatory activity: representation of movement or first cog in a dynamical machine?” In: *Neuron* 68.3 (Nov. 2010), pp. 387–400.

[15] Patrick T. Sadtler et al. “Neural constraints on learning”. In: *Nature* 512.7515 (Aug. 2014), pp. 423–426.

[16] Emily R. Oby et al. “New neural activity patterns emerge with long-term learning”. In: *Proceedings of the National Academy of Sciences* 116.30 (July 23, 2019), pp. 15210–15215.

[17] Juan A. Gallego et al. “Long-term stability of cortical population dynamics underlying consistent behavior”. In: *Nature Neuroscience* 23.2 (Feb. 2020), pp. 260–270.

[18] Mostafa Safaie et al. “Preserved neural dynamics across animals performing similar behaviour”. In: *Nature* 623.7988 (Nov. 2023), pp. 765–771.

[19] Sander W Keemink and Christian K Machens. “Decoding and encoding (de)mixed population responses”. In: *Current Opinion in Neurobiology*. Computational Neuroscience 58 (Oct. 1, 2019), pp. 112–121.

[20] Valerio Mante et al. “Context-dependent computation by recurrent dynamics in prefrontal cortex”. In: *Nature* 503.7474 (Nov. 2013), pp. 78–84.

[21] Christian K. Machens, Ranulfo Romo, and Carlos D. Brody. “Functional, But Not Anatomical, Separation of “What” and “When” in Prefrontal Cortex”. In: *Journal of Neuroscience* 30.1 (Jan. 2010), pp. 350–360.

[22] Cátia Fortunato et al. *Nonlinear manifolds underlie neural population activity during behaviour*. Pages: 2023.07.18.549575 Section: New Results. Apr. 25, 2024.

[23] Mehrdad Jazayeri and Srdjan Ostojic. “Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity”. In: *Current Opinion in Neurobiology* 70 (Oct. 2021), pp. 113–120.

[24] Daniel Durstewitz, Georgia Koppe, and Max Ingo Thurm. “Reconstructing computational system dynamics from neural data with recurrent neural networks”. In: *Nature Reviews. Neuroscience* 24.11 (Nov. 2023), pp. 693–710.

[25] Richard J. Gardner et al. “Toroidal topology of population activity in grid cells”. In: *Nature* 602.7895 (Feb. 2022), pp. 123–128.

[26] Michael Okun et al. “Diverse coupling of neurons to populations in sensory cortex”. In: *Nature* 521.7553 (May 2015), pp. 511–515.

[27] Barbara Feulner and Claudia Clopath. “Neural manifold under plasticity in a goal driven learning behaviour”. In: *PLOS Computational Biology* 17.2 (Feb. 5, 2021), e1008621.

[28] Matthew G. Perich et al. *Inferring brain-wide interactions using data-constrained recurrent neural network models*. Pages: 2020.12.18.423348 Section: New Results. Mar. 11, 2021.

[29] U. Mitzdorf. “Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena”. In: *Physiological Reviews* 65.1 (Jan. 1985), pp. 37–100.

[30] Bijan Pesaran et al. “Investigating large-scale brain dynamics using field potential recordings: analysis and interpretation”. In: *Nature Neuroscience* 21.7 (2018), pp. 903–919.

[31] Robert D. Flint et al. “Local field potentials allow accurate decoding of muscle activity”. In: *Journal of Neurophysiology* 108.1 (July 2012), pp. 18–24.

[32] J. N. Sanes and J. P. Donoghue. “Oscillations in local field potentials of the primate motor cortex during voluntary movement.” In: *Proceedings of the National Academy of Sciences* 90.10 (May 1993), pp. 4470–4474.

[33] John P. Donoghue et al. “Neural Discharge and Local Field Potential Oscillations in Primate Motor Cortex During Voluntary Movements”. In: *Journal of Neurophysiology* 79.1 (Jan. 1998), pp. 159–173.

[34] Claire L. Witham, Minyan Wang, and Stuart N. Baker. “Cells in somatosensory areas show synchrony with beta oscillations in monkey motor cortex”. In: *European Journal of Neuroscience* 26.9 (2007). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1460-9568.2007.05890.x, pp. 2677–2686.

[35] György Buzsáki and Andreas Draguhn. “Neuronal oscillations in cortical networks”. In: *Science* 304.5679 (June 25, 2004), pp. 1926–1929.

[36] Roger Traub. “Inhibition-based rhythms: experimental and mathematical observations on network dynamics”. In: *International Journal of Psychophysiology* (Jan. 1, 2000).

[37] G. Buzsaki et al. “High-frequency network oscillation in the hippocampus”. In: *Science* 256.5059 (May 1992), pp. 1025–1027.

[38] Doris D. Wang et al. “Pallidal Deep-Brain Stimulation Disrupts Pallidal Beta Oscillations and Coherence with Primary Motor Cortex in Parkinson’s Disease”. In: *Journal of Neuroscience* 38.19 (May 9, 2018), pp. 4556–4568.

[39] Hayriye Cagnan, Eugene Paul Duff, and Peter Brown. “The relative phases of basal ganglia activities dynamically shape effective connectivity in Parkinson’s disease”. In: *Brain: A Journal of Neurology* 138 (Pt 6 June 2015), pp. 1667–1678.

[40] Pascal Fries. “Neuronal gamma-band synchronization as a fundamental process in cortical computation”. In: *Annual Review of Neuroscience* 32 (2009), pp. 209–224.

[41] Pascal Fries. “Rhythms for Cognition: Communication through Coherence”. In: *Neuron* 88.1 (Oct. 7, 2015), pp. 220–235.

[42] Pascal Fries. “A mechanism for cognitive dynamics: neuronal communication through neuronal coherence”. In: *Trends in Cognitive Sciences* 9.10 (Oct. 1, 2005), pp. 474–480.

[43] Ole Jensen and Ali Mazaheri. “Shaping Functional Architecture by Oscillatory Alpha Activity: Gating by Inhibition”. In: *Frontiers in Human Neuroscience* 4 (Nov. 4, 2010).

[44] Supratim Ray and John H. R. Maunsell. “Different Origins of Gamma Rhythm and High-Gamma Activity in Macaque Visual Cortex”. In: *PLOS Biology* 9.4 (Apr. 12, 2011), e1000610.

[45] Oscar Herreras. “Local Field Potentials: Myths and Misunderstandings”. In: *Frontiers in Neural Circuits* 10 (Dec. 15, 2016).

[46] Henrik Lindén et al. “Modeling the Spatial Reach of the LFP”. In: *Neuron* 72.5 (Dec. 8, 2011), pp. 859–872.

[47] Szymon Łęski et al. “Frequency Dependence of Signal Power and Spatial Reach of the Local Field Potential”. In: *PLOS Computational Biology* 9.7 (July 2013), e1003137.

[48] Yoshinao Kajikawa and Charles E. Schroeder. “How local is the local field potential?” In: *Neuron* 72.5 (Dec. 8, 2011), pp. 847–858.

[49] Hamidreza Abbaspourazad et al. “Multiscale low-dimensional motor cortical state dynamics predict naturalistic reach-and-grasp behavior”. In: *Nature Communications* 12.1 (Jan. 27, 2021), p. 607.

[50] Francis R. Willett et al. “High-performance brain-to-text communication via handwriting”. In: *Nature* 593.7858 (May 2021), pp. 249–254.

[51] Sharlene N. Flesher et al. “A brain-computer interface that evokes tactile sensations improves robotic arm control”. In: *Science* 372.6544 (May 21, 2021), pp. 831–836.

[52] Maeike Zijlmans et al. “High-frequency oscillations as a new biomarker in epilepsy”. In: *Annals of Neurology* 71.2 (Feb. 2012), pp. 169–178.

[53] Martha J. Morrell and RNS System in Epilepsy Study Group. “Responsive cortical stimulation for the treatment of medically intractable partial epilepsy”. In: *Neurology* 77.13 (Sept. 27, 2011), pp. 1295–1304.

[54] Abbey B. Holt et al. “Phase-Dependent Suppression of Beta Oscillations in Parkinson’s Disease Patients”. In: *The Journal of Neuroscience* (Feb. 6, 2019).

[55] Valentina D’Onofrio et al. “Local field potentials survey to guide DBS programming in Parkinson’s disease: a clinical-neurophysiological longitudinal study”. In: *npj Parkinson’s Disease* 12.1 (Dec. 7, 2025), p. 1.

[56] Timothy O. West et al. “Stimulating at the right time to recover network states in a model of the cortico-basal ganglia-thalamic circuit”. In: *PLOS Computational Biology* 18.3 (Mar. 4, 2022), e1009887.

[57] Simon Little et al. “Adaptive Deep Brain Stimulation In Advanced Parkinson Disease”. In: *Annals of Neurology* 74.3 (Sept. 2013), pp. 449–457.

[58] Robert D. Flint et al. “Long term, stable brain machine interface performance using local field potentials and multiunit spikes”. In: *Journal of Neural Engineering* 10.5 (Aug. 2013), p. 056005.

[59] Sergey D Stavisky et al. “A high performing brain–machine interface driven by low-frequency local field potentials alone and together with spikes”. In: *Journal of Neural Engineering* 12.3 (June 2015), p. 036009.

[60] Hansjörg Scherberger, Murray R. Jarvis, and Richard A. Andersen. “Cortical Local Field Potential Encodes Movement Intentions in the Posterior Parietal Cortex”. In: *Neuron* 46.2 (Apr. 21, 2005), pp. 347–354.

[61] Bijan Pesaran et al. “Temporal structure in neuronal activity during working memory in macaque parietal cortex”. In: *Nature Neuroscience* 5.8 (Aug. 2002), pp. 805–811.

[62] V. N. Murthy and E. E. Fetz. “Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys”. In: *Journal of Neurophysiology* 76.6 (Dec. 1996), pp. 3968–3982.

[63] Andrew Sharott et al. “A Population of Indirect Pathway Striatal Projection Neurons Is Selectively Entrained to Parkinsonian Beta Oscillations”. In: *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 37.41 (Oct. 11, 2017), pp. 9977–9998.

[64] Cecilia Gallego-Carracedo et al. “Local field potentials reflect cortical population dynamics in a region-specific and frequency-dependent manner”. In: *eLife* 11 (Aug. 15, 2022), e73155.

[65] Jeffrey J. Hutsler, Dong-Geun Lee, and Kristin K. Porter. “Comparative analysis of cortical layering and supragranular layer enlargement in rodent carnivore and primate species”. In: *Brain Research* 1052.1 (Aug. 2, 2005), pp. 71–81.

[66] Jon H. Kaas. “Evolution of somatosensory and motor cortex in primates”. In: *The Anatomical Record Part A: Discoveries in Molecular, Cellular, and Evolutionary Biology* 281A.1 (2004). _eprint: https://anatomypubs.onlinelibrary.wiley.com/doi/pdf/10.1002/ar.a.20120, pp. 1148–1156.

[67] T. P. Pons et al. “The somatotopic organization of area 2 in macaque monkeys”. In: *Journal of Comparative Neurology* 241.4 (1985), pp. 445–466.

[68] Roger N. Lemon. “Descending pathways in motor control”. In: *Annual Review of Neuroscience* 31 (2008), pp. 195–218.

[69] Jean-Alban Rathelot and Peter L. Strick. “Subdivisions of primary motor cortex based on cortico-motoneuronal cells”. In: *Proceedings of the National Academy of Sciences of the United States of America* 106.3 (Jan. 20, 2009), pp. 918–923.

[70] R. P. Dum and P. L. Strick. “The origin of corticospinal projections from the premotor areas in the frontal lobe”. In: *Journal of Neuroscience* 11.3 (Mar. 1, 1991), pp. 667–689.

[71] Bryan M. Hooks et al. “Topographic precision in sensory and motor corticostriatal projections varies across cell type and cortical area”. In: *Nature Communications* 9.1 (Sept. 3, 2018), p. 3549.

[72] Arjun K. Bansal et al. “Decoding 3D reach and grasp from hybrid signals in motor and premotor cortices: spikes, multiunit activity, and local field potentials”. In: *Journal of Neurophysiology* 107.5 (Mar. 2012), pp. 1337–1355.

[73] Brian M. Dekleva, Konrad P. Kording, and Lee E. Miller. “Single reach plans in dorsal premotor cortex during a two-target task”. In: *Nature Communications* 9.1 (Sept. 2018), p. 3556.

[74] Raeed H Chowdhury, Joshua I Glaser, and Lee E Miller. “Area 2 of primary somatosensory cortex encodes kinematics of the whole arm”. In: *eLife* 9 (Jan. 2020), e48198.

[75] Brian M. London and Lee E. Miller. “Responses of somatosensory area 2 neurons to actively and passively generated limb movements”. In: *Journal of Neurophysiology* 109.6 (Dec. 2012), pp. 1505–1513.

[76] R. L. Albin, A. B. Young, and J. B. Penney. “The functional anatomy of basal ganglia disorders”. In: *Trends in Neurosciences* 12.10 (Oct. 1989), pp. 366–375.

[77] M. R. DeLong. “Primate models of movement disorders of basal ganglia origin”. In: *Trends in Neurosciences* 13.7 (July 1990), pp. 281–285.

[78] Ann M. Graybiel. “Habits, rituals, and the evaluative brain”. In: *Annual Review of Neuroscience* 31 (2008), pp. 359–387.

[79] Joshua T Dudman and John W Krakauer. “The basal ganglia: from motor commands to the control of vigor”. In: *Current Opinion in Neurobiology*. Neurobiology of cognitive behavior 37 (Apr. 1, 2016), pp. 158–166.

[80] Sten Grillner and Brita Robertson. “The Basal Ganglia Over 500 Million Years”. In: *Current Biology* 26.20 (Oct. 24, 2016), R1088–R1100.

[81] Charles R. Gerfen and D. James Surmeier. “Modulation of striatal projection systems by dopamine”. In: *Annual Review of Neuroscience* 34 (2011), pp. 441–466.

[82] Houri Hintiryan et al. “The mouse cortico-striatal projectome”. In: *Nature Neuroscience* 19.8 (Aug. 2016), pp. 1100–1114.

[83] Barbara J Hunnicutt et al. “A comprehensive excitatory input map of the striatum reveals novel functional organization”. In: *eLife* 5 (Nov. 28, 2016), e19103.

[84] Pieter Voorn et al. “Putting a spin on the dorsal-ventral divide of the striatum”. In: *Trends in Neurosciences* 27.8 (Aug. 2004), pp. 468–474.

[85] E. A. Stern, D. Jaeger, and C. J. Wilson. “Membrane potential synchrony of simultaneously recorded striatal spiny neurons in vivo”. In: *Nature* 394.6692 (July 30, 1998), pp. 475–478.

[86] Anatol C. Kreitzer. “Physiology and pharmacology of striatal neurons”. In: *Annual Review of Neuroscience* 32 (2009), pp. 127–147.

[87] Giovanni Barbera et al. “Spatially Compact Neural Clusters in the Dorsal Striatum Encode Locomotion Relevant Information”. In: *Neuron* 92.1 (Oct. 5, 2016), pp. 202–213.

[88] Yoshikazu Isomura et al. “Reward-Modulated Motor Information in Identified Striatum Neurons”. In: *Journal of Neuroscience* 33.25 (June 19, 2013), pp. 10209–10220.

[89] Atsushi Nambu, Hironobu Tokuno, and Masahiko Takada. “Functional significance of the cortico-subthalamo-pallidal ’hyperdirect’ pathway”. In: *Neuroscience Research* 43.2 (June 2002), pp. 111–117.

[90] Adam R. Aron and Russell A. Poldrack. “Cortical and subcortical contributions to Stop signal response inhibition: role of the subthalamic nucleus”. In: *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 26.9 (Mar. 1, 2006), pp. 2424–2433.

[91] Andrew J. Peters et al. “Striatal activity topographically reflects cortical activity”. In: *Nature* 591.7850 (Mar. 2021), pp. 420–425.

[92] Junchol Park et al. “Conjoint specification of action by neocortex and striatum”. In: *Neuron* 113.4 (Feb. 19, 2025), 620–636.e6.

[93] Risa Kawai et al. “Motor cortex is required for learning but not for executing a motor skill”. In: *Neuron* 86.3 (May 6, 2015), pp. 800–812.

[94] Eberhard E. Fetz. “Operant Conditioning of Cortical Unit Activity”. In: *Science* 163.3870 (Feb. 28, 1969), pp. 955–958.

[95] Aaron C. Koralek et al. “Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills”. In: *Nature* 483.7389 (Mar. 2012), pp. 331–335.

[96] Asma Motiwala et al. “Brain–computer interfaces as a causal probe for scientific inquiry”. In: *Trends in Cognitive Sciences* 30.1 (Jan. 1, 2026), pp. 40–53.

[97] Rodrigo Quian Quiroga. “Plugging in to Human Memory: Advantages, Challenges, and Insights from Human Single-Neuron Recordings”. In: *Cell* 179.5 (Nov. 14, 2019), pp. 1015–1032.

[98] Ueli Rutishauser et al. “The Architecture of Human Memory: Insights from Human Single-Neuron Recordings”. In: *Journal of Neuroscience* 41.5 (Feb. 3, 2021), pp. 883–890.

[99] Guosong Hong and Charles M. Lieber. “Novel electrode technologies for neural recordings”. In: *Nature Reviews Neuroscience* 20.6 (June 2019), pp. 330–345.

[100] Chethan Pandarinath et al. “Inferring single-trial neural population dynamics using sequential auto-encoders”. In: *Nature Methods* 15.10 (Oct. 2018), pp. 805–815.

[101] Matthew G. Perich, Juan A. Gallego, and Lee E. Miller. “A Neural Population Mechanism for Rapid Learning”. In: *Neuron* 100.4 (Nov. 21, 2018), 964–976.e7.

[102] Saurabh Vyas et al. “Neural Population Dynamics Underlying Motor Learning Transfer”. In: *Neuron* 97.5 (2018), 1177–1186.e3.

[103] Xulu Sun et al. “Skill-specific changes in cortical preparatory activity during motor learning”. In: *bioRxiv* (Jan. 2020), p. 2020.01.30.919894.

[104] Jing Wang et al. “Flexible timing by temporal scaling of cortical responses”. In: *Nature Neuroscience* 21.1 (Jan. 2018), pp. 102–110.

[105] Evan D. Remington et al. “Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics”. In: *Neuron* 98.5 (June 2018), 1005–1019.e5.

[106] Barbara Feulner et al. “Small, correlated changes in synaptic connectivity may facilitate rapid motor learning”. In: *Nature Communications* 13.1 (Sept. 2, 2022), p. 5163.

[107] Timothy J. Buschman et al. “Synchronous Oscillatory Neural Ensembles for Rules in the Prefrontal Cortex”. In: *Neuron* 76.4 (Nov. 21, 2012), pp. 838–846.

[108] Sébastien Tremblay et al. “Single-Trial Decoding of Visual Attention from Local Field Potentials in the Primate Lateral Prefrontal Cortex Is Frequency-Dependent”. In: *Journal of Neuroscience* 35.24 (June 17, 2015), pp. 9038–9049.

[109] Jing Liu and William T. Newsome. “Local Field Potential in Cortical Area MT: Stimulus Tuning and Behavioral Correlations”. In: *Journal of Neuroscience* 26.30 (July 26, 2006), pp. 7779–7790.

[110] Alexander Gail, Hans Joerg Brinksmeyer, and Reinhard Eckhorn. “Perception-related Modulations of Local Field Potential Power and Coherence in Primary Visual Cortex of Awake Monkey during Binocular Rivalry”. In: *Cerebral Cortex* 14.3 (Mar. 1, 2004), pp. 300–313.

[111] Charles D. Holmes, Charalampos Papadimitriou, and Lawrence H. Snyder. “Dissociation of LFP Power and Tuning in the Frontal Cortex during Memory”. In: *Journal of Neuroscience* 38.38 (Sept. 19, 2018), pp. 8177–8186.

[112] Jun Zhuang et al. “Decoding 3-D Reach and Grasp Kinematics From High-Frequency Local Field Potentials in Primate Primary Motor Cortex”. In: *IEEE Transactions on Biomedical Engineering* 57.7 (July 2010), pp. 1774–1784.

[113] Sagi Perel et al. “Single-unit activity, threshold crossings, and local field potentials in motor cortex differentially encode reach kinematics”. In: *Journal of Neurophysiology* 114.3 (Sept. 2015), pp. 1500–1512.

[114] S. N. Baker, E. Olivier, and R. N. Lemon. “Coherent oscillations in monkey motor cortex and hand muscle EMG show task-dependent modulation”. In: *The Journal of Physiology* 501.1 (1997). _eprint: https://physoc.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1469-7793.1997.225bo.x, pp. 225–241.

[115] Claire L. Witham, Minyan Wang, and Stuart Baker. “Corticomuscular Coherence Between Motor Cortex, Somatosensory Areas and Forearm Muscles in the Monkey”. In: *Frontiers in Systems Neuroscience* 4 (July 30, 2010).

[116] Claire L. Witham and Stuart N. Baker. “Coding of digit displacement by cell spiking and network oscillations in the monkey sensorimotor cortex”. In: *Journal of Neurophysiology* 108.12 (Dec. 15, 2012), pp. 3342–3352.

[117] György Buzsáki. “Theta Oscillations in the Hippocampus”. In: *Neuron* 33.3 (Jan. 2002), pp. 325–340.

[118] S. Ray et al. “Effect of Stimulus Intensity on the Spike-Local Field Potential Relationship in the Secondary Somatosensory Cortex”. In: *Journal of Neuroscience* 28.29 (July 2008), pp. 7334–7343.

[119] Michael E. Rule et al. “Dissociation between sustained single-neuron spiking and transient β-LFP oscillations in primate motor cortex”. In: *Journal of Neurophysiology* 117.4 (Apr. 2017), pp. 1524–1543.

[120] V. N. Murthy and E. E. Fetz. “Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior”. In: *Journal of Neurophysiology* 76.6 (Dec. 1996), pp. 3949–3967.

[121] Shinji Kakei, Donna S. Hoffman, and Peter L. Strick. “Direction of action is represented in the ventral premotor cortex”. In: *Nature Neuroscience* 4.10 (Oct. 2001), pp. 1020–1025.

[122] Abhranil Das and Ila R. Fiete. “Systematic errors in connectivity inferred from activity in strongly recurrent networks”. In: *Nature Neuroscience* 23.10 (Oct. 2020), pp. 1286–1296.

[123] Jean-Alban Rathelot and Peter L. Strick. “Muscle representation in the macaque motor cortex: An anatomical perspective”. In: *Proceedings of the National Academy of Sciences* 103.21 (May 2006), pp. 8257–8262.

[124] E. E. Fetz and P. D. Cheney. “Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells”. In: *Journal of Neurophysiology* 44.4 (Oct. 1980), pp. 751–772.

[125] M. M. Morrow and L. E. Miller. “Prediction of Muscle Activity by Populations of Sequentially Recorded Primary Motor Cortex Neurons”. In: *Journal of Neurophysiology* 89.4 (Apr. 2003), pp. 2279–2288.

[126] Lauren E. Sergio, Catherine Hamel-Pâquet, and John F. Kalaska. “Motor Cortex Neural Correlates of Output Kinematics and Kinetics During Isometric-Force and Arm-Reaching Tasks”. In: *Journal of Neurophysiology* 94.4 (Oct. 2005), pp. 2353–2378.

[127] Liming Shen and Garrett E. Alexander. “Neural Correlates of a Spatial Sensory-To-Motor Transformation in Primary Motor Cortex”. In: *Journal of Neurophysiology* 77.3 (Mar. 1997), pp. 1171–1194.

[128] Gopal Santhanam et al. “Factor-Analysis Methods for Higher-Performance Neural Prostheses”. In: *Journal of Neurophysiology* 102.2 (Aug. 2009), pp. 1315–1330.

[129] D. J. Crammond and J. F. Kalaska. “Modulation of preparatory neuronal activity in dorsal premotor cortex due to stimulus-response compatibility”. In: *Journal of Neurophysiology* 71.3 (Mar. 1994), pp. 1281–1284.

[130] Liming Shen and Garrett E. Alexander. “Preferential Representation of Instructed Target Location Versus Limb Trajectory in Dorsal Premotor Area”. In: *Journal of Neurophysiology* 77.3 (Mar. 1997), pp. 1195–1212.

[131] K. Cora Ames, Stephen I. Ryu, and Krishna V. Shenoy. “Simultaneous motor preparation and execution in a last-moment reach correction task”. In: *Nature Communications* 10.1 (June 2019), p. 2718.

[132] György Buzsáki and Edvard I. Moser. “Memory, navigation and theta rhythm in the hippocampal-entorhinal system”. In: *Nature Neuroscience* 16.2 (Feb. 2013), pp. 130–138.

[133] Francis R. Bach and Michael I. Jordan. “Kernel Independent Component Analysis”. In: *Journal of Machine Learning Research* 3 (Jul 2002), pp. 1–48.

[134] David Sussillo et al. “A neural network that finds a naturalistic solution for the production of muscle activity”. In: *Nature Neuroscience* 18.7 (July 2015), pp. 1025–1033.

[135] Gamaleldin F. Elsayed and John P. Cunningham. “Structure in neural population recordings: an expected byproduct of simpler phenomena?” In: *Nature Neuroscience* 20.9 (Sept. 2017), pp. 1310–1318.

[136] A. Riehle and J. Requin. “Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement”. In: *Journal of Neurophysiology* 61.3 (Mar. 1989), pp. 534–549.

[137] Machiko Ohbayashi, Nathalie Picard, and Peter L. Strick. “Inactivation of the Dorsal Premotor Area Disrupts Internally Generated, But Not Visually Guided, Sequential Movements”. In: *Journal of Neuroscience* 36.6 (Feb. 10, 2016), pp. 1971–1976.

[138] Mark M. Churchland and Krishna V. Shenoy. “Delay of Movement Caused by Disruption of Cortical Preparatory Activity”. In: *Journal of Neurophysiology* 97.1 (Jan. 2007), pp. 348–359.

[139] John G. O’Leary and Nicholas G. Hatsopoulos. “Early Visuomotor Representations Revealed From Evoked Local Field Potentials in Motor and Premotor Cortical Areas”. In: *Journal of Neurophysiology* 96.3 (Sept. 2006), pp. 1492–1506.

[140] Stuart N. Baker, Matthew Chiu, and Eberhard E. Fetz. “Afferent Encoding of Central Oscillations in the Monkey Arm”. In: *Journal of Neurophysiology* 95.6 (June 2006), pp. 3904–3910.

[141] Priyanka S. Mehta, Seng Bum Michael Yoo, and Benjamin Y. Hayden. *Signatures of processing complexity during global cognitive states in ventromedial prefrontal cortex*. Pages: 2020.10.08.331579 Section: New Results. Oct. 8, 2020.

[142] Silvia Maggi and Mark D. Humphries. “Activity Subspaces in Medial Prefrontal Cortex Distinguish States of the World”. In: *Journal of Neuroscience* 42.20 (May 18, 2022), pp. 4131–4146.

[143] Joshua I. Glaser et al. “Population coding of conditional probability distributions in dorsal premotor cortex”. In: *Nature Communications* 9.1 (May 3, 2018), p. 1788.

[144] Pavan Ramkumar et al. “Premotor and Motor Cortices Encode Reward”. In: *PLOS ONE* 11.8 (Aug. 26, 2016), e0160851.

[145] Ryan Kelley et al. “A human prefrontal-subthalamic circuit for cognitive control”. In: *Brain: A Journal of Neurology* 141.1 (Jan. 1, 2018), pp. 205–216.

[146] Elliot H. Smith et al. “Widespread temporal coding of cognitive control in the human prefrontal cortex”. In: *Nature Neuroscience* 22.11 (Nov. 2019), pp. 1883–1891.

[147] Jie Zheng et al. “Amygdala-hippocampal dynamics during salient information processing”. In: *Nature Communications* 8.1 (Feb. 8, 2017), p. 14413.

[148] Alex P. Vaz et al. “Coupled ripple oscillations between the medial temporal lobe and neocortex retrieve human memory”. In: *Science* 363.6430 (Mar. 1, 2019), pp. 975–978.

[149] Gopala K. Anumanchipalli, Josh Chartier, and Edward F. Chang. “Speech synthesis from neural decoding of spoken sentences”. In: *Nature* 568.7753 (Apr. 2019), pp. 493–498.

[150] Jonathan Miller et al. “Lateralized hippocampal oscillations underlie distinct aspects of human spatial memory and navigation”. In: *Nature Communications* 9.1 (June 21, 2018), p. 2423.

[151] Constantino Méndez-Bértolo et al. “A fast pathway for fear in human amygdala”. In: *Nature Neuroscience* 19.8 (Aug. 2016), pp. 1041–1049.

[152] Gamaleldin F. Elsayed et al. “Reorganization between preparatory and movement population responses in motor cortex”. In: *Nature Communications* 7 (Oct. 27, 2016), p. 13239.

[153] Joshua I. Glaser et al. “Machine Learning for Neural Decoding”. In: *eNeuro* 7.4 (July 2020).

[154] Eric A Pohlmeyer et al. “Prediction of upper limb muscle activity from motor cortical discharge during reaching”. In: *Journal of Neural Engineering* 4.4 (Nov. 2007), p. 369.

[155] Kazuyuki Samejima et al. “Representation of action-specific reward values in the striatum”. In: *Science* 310.5752 (Nov. 25, 2005), pp. 1337–1340.

[156] Karin Morandell and Daniel Huber. “The role of forelimb motor cortex areas in goal directed action in mice”. In: *Scientific Reports* 7.1 (Nov. 17, 2017), p. 15759.

[157] Mark M. Churchland and Krishna V. Shenoy. “Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex”. In: *Journal of Neurophysiology* 97.6 (June 2007), pp. 4235–4257.

[158] Babita Panigrahi et al. “Dopamine Is Required for the Neural Representation and Control of Movement Vigor”. In: *Cell* 162.6 (Sept. 10, 2015), pp. 1418–1430.

[159] Gregory J. Gage et al. “Selective activation of striatal fast-spiking interneurons during choice execution”. In: *Neuron* 67.3 (Aug. 12, 2010), pp. 466–479.

[160] Ryan W. Eaton, Tyler Libey, and Eberhard E. Fetz. “Operant conditioning of neural activity in freely behaving monkeys with intracranial reinforcement”. In: *Journal of Neurophysiology* 117.3 (Mar. 1, 2017), pp. 1112–1125.

[161] Mario Prsa, Gregorio L. Galiñanes, and Daniel Huber. “Rapid Integration of Artificial Sensory Feedback during Operant Conditioning of Motor Cortex Neurons”. In: *Neuron* 93.4 (Feb. 22, 2017), 929–939.e6.

[162] Kayvon Daie et al. *Functional reorganization of motor cortex connectivity during learning.* ISSN: 2692-8205. Mar. 5, 2026.

[163] Emily R. Oby et al. “Dynamical constraints on neural population activity”. In: *Nature Neuroscience* 28.2 (Feb. 2025), pp. 383–393.

[164] Matthew D. Golub et al. “Learning by neural reassociation”. In: *Nature Neuroscience* 21.4 (Apr. 2018), pp. 607–616.

[165] Karunesh Ganguly and Jose M. Carmena. “Emergence of a Stable Cortical Map for Neuroprosthetic Control”. In: *PLOS Biology* 7.7 (July 21, 2009), e1000153.

[166] Krishna V. Shenoy and Jose M. Carmena. “Combining decoder design and neural adaptation in brain-machine interfaces”. In: *Neuron* 84.4 (Nov. 19, 2014), pp. 665–680.

[167] James J. Jun et al. “Fully integrated silicon probes for high-density recording of neural activity”. In: *Nature* 551.7679 (Nov. 2017), pp. 232–236.

[168] Marius Pachitariu et al. “Spike sorting with Kilosort4”. In: *Nature Methods* 21.5 (May 2024), pp. 914–921.

[169] Eun Jung Hwang, Paul M. Bailey, and Richard A. Andersen. “Volitional control of neural activity relies on the natural motor repertoire”. In: *Current biology: CB* 23.5 (Mar. 4, 2013), pp. 353–361.

[170] Manfred J. Oswald et al. “Diversity of layer 5 projection neurons in the mouse motor cortex”. In: *Frontiers in Cellular Neuroscience* 7 (Oct. 16, 2013).

[171] V. B. Mountcastle. “Modality and topographic properties of single neurons of cat’s somatic sensory cortex”. In: *Journal of Neurophysiology* 20.4 (July 1957), pp. 408–434.

[172] Eric Kenji Lee et al. “Non-linear dimensionality reduction on extracellular waveforms reveals cell type diversity in premotor cortex”. In: *eLife* 10 (Aug. 6, 2021), e67490.

[173] Gil Mandelbaum et al. “Distinct Cortical-Thalamic-Striatal Circuits through the Parafascicular Nucleus”. In: *Neuron* 102.3 (May 8, 2019), 636–652.e7.

[174] Leigh R. Hochberg et al. “Neuronal ensemble control of prosthetic devices by a human with tetraplegia”. In: *Nature* 442.7099 (July 2006), pp. 164–171.

[175] Jason E. Osborne and Joshua T. Dudman. “RIVETS: A Mechanical System for In Vivo and In Vitro Electrophysiology and Imaging”. In: *PLOS ONE* 9.2 (Feb. 14, 2014), e89007.

[176] Daniel Birman et al. *Pinpoint: trajectory planning for multi-probe electrophysiology and injections in an interactive web-based 3D environment*. Pages: 2023.07.14.548952 Section: New Results. July 15, 2023.