

Collective Neural Dynamics in Primate Motor Cortex

Michael Rule

Thesis

Submitted in partial fulfillment of the requirements for the Degree of
Doctor of Philosophy in the Department of Neuroscience at Brown
University

Providence, Rhode Island

May 2016

(C) Copyright 2016 by Michael Everett Rule

This dissertation by Michael E. Rule is accepted in its present form by the Department of Neuroscience as satisfying the dissertation requirement for the degree of Doctor of Philosophy.

Date _____
Wilson Truccolo, Ph.D., Advisor

Recommended to the Graduate Council

Date _____
John Donoghue, Ph.D., Reader

Date _____
Christopher Moore, Ph.D., Reader

Date _____
Michael Paradiso, Ph.D., Reader

Date _____
George Bard Ermentrout, Ph.D., External Reader

Approved by the Graduate Council

Date _____
Peter M. Weber, Ph.D., Dean of the Graduate School

Curriculum vitae

Michael Everett Rule

Ph.D. Candidate, Department of Neuroscience, Brown University.

Michael_Rule@brown.edu

Education

- Brown University**, Providence, RI 2010-
Ph.D. Candidate, Neuroscience
- Marine Biological Laboratory**, Woods Hole, MA 2014
Methods in Computational Neuroscience summer course
- Carnegie Mellon University School of Computer Science**, Pittsburgh, PA 2005-2009
B.S. Computer Science
- Center for the Neural Basis of Cognition**, Pittsburgh, PA 2008-2009

Research experience

- Ph.D. candidate** 2012-
Department of Neuroscience, Brown University
Supervised by: Dr. Wilson Truccolo
Topic: Collective dynamics in motor cortex
- Graduate rotation project** 2011
Department of Neuroscience, Brown University
Supervised by: Dr. John Donoghue
Topic: Regularization in Kalman filter decoding of kinematics from primate motor cortex
- Graduate rotation project** 2010
Department of Neuroscience, Brown University
Supervised by: Dr. David Sheinberg
Topic: Animated stimuli for biological motion perception
- Research assistant** Fall 2009
Department of Neuroscience, Carnegie Mellon University
Supervised by: Dr. Nathan Urban
Topic: A parallel spiking model of the olfactory bulb
- Undergraduate research assistant** Summer 2009
Department of Neuroscience, Carnegie Mellon University
Supervised by: Dr. Rick Gerkin, Dr. Nathan Urban
Topic: Spiking network simulations for the NVIDIA Tesla 'desktop supercomputer'

Undergraduate research assistant 2008
Center for the Neural Basis of Cognition, Carnegie Mellon University
Supervised by: Dr. Bard Ermentrout
Topic: A model for geometric visual hallucinations induced by flickering light

Undergraduate research assistant 2007
Institute for Complex Engineered Systems, Carnegie Mellon University
Supervised by: Dr. Phil G. Campbell
Topic: Computational modeling of stem cell populations

Undergraduate research assistant 2006
Institute for Complex Engineered Systems, Carnegie Mellon University
Supervised by: Dr. Julie Jadlowiec, Dr. Phil G. Campbell
Topic: Biocompatibility of micro-robotically fabricated biological scaffolds

Teaching and service

Carnegie Mellon University 2009
Teaching assistant, "Human information processing and artificial intelligence"
Instructor: Dr. John R. Anderson

Brown University 2012
Teaching assistant, "Experimental Neurobiology"
Instructor: Dr. John J. Stein

Lesson material development in K-12 STEM education: Guler, S. D. and Rule, M. E. (2013).
Invent-abling: enabling inventiveness through craft. In *Proceedings of the 12th International Conference on Interaction Design and Children*, pages 368–371. ACM

Peer-reviewed publications

Nain, A. S., Chung, F., Rule, M., Jadlowiec, J., Campbell, P. G., Amon, C., Sitti, M., et al. (2007). Microrobotically fabricated biological scaffolds for tissue engineering. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 1918–1923. IEEE

Rule, M., Stoffregen, M., and Ermentrout, B. (2011). A model for the origin and properties of flicker-induced geometric phosphenes. *PLoS Comput. Biol.* 7(9):e1002

Rule, M. E., Vargas-Irwin, C., Donoghue, J. P., and Truccolo, W. (2015). Contribution of LFP dynamics to single neuron spiking variability in motor cortex during movement execution. *Frontiers in Systems Neuroscience*, 9:89

Scientific presentations

Rule, M., Vargas-Irwin, C., Donoghue, J.P., Truccolo, W. (2015). Spatiotemporal dynamics in primate motor cortex local field potentials. Program No. 427.22. 45th Annual Meeting of the Society for Neuroscience meeting, Chicago, IL. October 17-21 (Poster)

Rule, M., Stoffregen, M., Ermentrout, G.B. (2012) A model for the origin and properties of flicker-induced geometric phosphenes. Invited talk at the Society for Industrial and Applied Mathematics (SIAM) minisymposium "Spatiotemporal dynamics in networks of the brain."

Academic fellowships and awards

NSF Graduate Research Fellowship	2012-2015
SIAM Student Travel Award	2012
Levy Fellowship, Brown University	2010-2011
CNBC Undergraduate Research Fellowship in Computational Neuroscience	2008-2009

Acknowledgements

This section is usually written quite formally, but my acknowledgement of the individuals here is not a formality. In theory one focused on those individuals instrumental to shaping the Ph.D. candidate's life, culminating in the production of this document. In practice, these acknowledgements will be read by relatively few people: my thesis committee, my friends, family, colleagues to whom I might send this document, and perhaps a half-dozen students who may wish to pick up this research where I leave off. Therefore, it seems appropriate to me to write something more intimate.

I will, in a moment, get to thanking the numerous mentors and colleagues that were instrumental in shaping this thesis. But first, I would like to, as a matter of ritual, acknowledge some random twists of fate that have lead me here. It was through an entirely chance interaction with a classmate that I was informed of neuroscience outreach activities organized by Dr. Edda Thiels at the University of Pittsburgh, which initiated my serious studies in neuroscience. I must credit my interest in programming to an old grade-school friend, who introduced me to it all those years ago. I thank especially those teachers who permitted me to do experiments that were, so to speak, against the rules. I thank, incidentally, whomever designed that electronics set that I got all those years ago, and whomever threw away the computer on which I first learn to program.

An honest history would be incomplete without mentioning a subject that I discuss with difficulty. I have been compelled toward research in the motor system by someone who was close to me, my sister, Elizabeth, who suffered a stroke and generalized hypoxia at birth. Her perceptual and emotional faculties were intact, but the damage to areas of frontal cortex and subcortical structures necessary for movement was irreparable. Our family spent many hours at the local rehabilitation hospital, and I took to hanging around the rehabilitation engineering shop. The resident engineer John Cotellero would improvise and fabricate assistive devices, customizing solutions for each individual. It was there that my enjoyment of electronics and programming met my interest in neuroscience, and

I chose to pursue computational neuroscience for a career.

I briefly studied at the California Institute of Technology, and later transferred to Carnegie Mellon, for reasons that would take far more than the length of this thesis to relate. I mention it now because several individuals that I met there have been supportive during my doctoral studies. Two of these individuals, Ashok Kumar and Natasha Cayco-Gajic, have gone on to become colleagues in neuroscience. I thank Christopher Beck for his patience in educating me on topics of theoretical computer science and discrete mathematics over the years, and for the challenging games of go. I thank Keegan McAllester for conversations both technical and philosophical, and for teaching me digital fabrication techniques and microcontroller programming. I also owe a fair bit of my basic understanding of mathematics and statistics to Marena Lin and Eric Stansifer. While on sentimental topics, words are not enough to thank my sister Hannah, and my friends Adrian Bondy, Annika Roseninge, Austin Redwood, Carolina Ramos, Emily Matteson, Hannah Sherril, and Michael Burnam-Fink for their support through difficult times. I would not have completed this program without them. We should also acknowledge ‘the turks’, Ali Arslan, Altar Sorkaç, and Mustafa Talay, my housemates-in-neuroscience for much of my time here, and de-facto housemates Jackie Hynes, James Niemeyer, and Aslı Sahin, for tolerating the negative externalities of some of my hobbies.

I remember fondly, and would like to thank, the faculty members that advised me while at Carnegie Mellon. I am grateful to Drs. Phil Campbell and Julie Jadowiec for providing me with my first research internship in bioengineering. Dr. Nathan Urban provided some of my first experiences in computational neuroscience, and was instrumental in connecting me with the broader field. It was Nathan who introduced me to Dr. Bard Ermentrout, under whom I completed my first neuroscience paper. If you are familiar with the field of mathematical neuroscience, then Bard needs no introduction. (The topic of that paper, the mathematical modeling of visual hallucinations, reflects on Bard both personally and professionally.) Bard has been a supportive mentor over the past seven years,

and has taught me everything I know about neural field theory and the applications of dynamical systems to neuroscience. I would not have had these educational opportunities if not for him. It was also during a meeting in Nathan's lab that I first encountered a paper written by Dr. Wilson Truccolo at Brown University, on statistical modeling of functional relationships between neurons in motor cortex. It was then that I resolved to apply to Brown University, with the aim of acquiring the statistical skills to apply to questions of dynamics and computation in the motor system.

If he permits me to say this in the thesis, at Brown, I had the fate of being Dr. Wilson Truccolo's first graduate student. When I arrived there was no such thing as the Truccolo lab. We needed special permission from the graduate school for him to serve as an advisor. Because the Truccolo lab did not as yet exist when I arrived, most of my close interactions were with students and post-doctoral researchers affiliated with the lab of Dr. John Donoghue and the Brain Gate research group. I remember many engaging, if not entirely scientific, conversations, with Lachlan Franquemont, Naveen Rao, Kaivon Paroo, Carlos Vargas-Irwin, and Jonas Zimmerman, of the Donoghue lab, as well as Mark Homer of Brain Gate. Additionally, none of the work in this thesis would have been possible but for the generous sharing of datasets collected in the Donoghue lab by Lachlan Franquemont, Carlos Vargas-Irwin, and Jonas Zimmerman. The past few years the Truccolo lab has grown and I have enjoyed the company and conversation of post-doctoral researchers Mehdi Aghagolzadeh, Yun Park, Fabien Wagner, Ernest Ho, Burak Yildiz, and Felipe Gerhard. I also thank Dr. Matthew Harrison for exposing me to a more systematic understanding of statistics and machine learning, as well as Drs. David Sheinberg and John Donoghue for supervising me on rotation projects at Brown, and providing professional advice over the years. Drs. Christopher Moore, Michael Paradiso, David Berson, and Gilad Barnea, have also been valued sources of career advice. I thank Drs. Moore, Paradiso, and Donoghue, also for serving on my thesis committee, and Dr. Bard Ermentrout, for serving as external reader. Dr. Wilson Truccolo has been a highly invested advisor, as-

sisting me in drafting applications and preparing presentations far above and beyond what is expected. We've had many engaging conversations, not just about statistical methods in neuroscience but also about broader scientific and cultural topics.

I can think of no way to say what follows without clichés, so I will use clichés, and perhaps irregular punctuation. My parents. Above all I admire their strength to continue, every day, through the past thirty years. My mother worked to give my sister the best life she could, and made tremendous sacrifices for it. My father has held his own in the highly competitive reality that is academia, all the while supporting and caring for the family. My parents did not steer me toward science, they encouraged a broad education and provided every opportunity. They were, however, extremely supportive of my interest in science, taking me to the public library, financially supporting projects, and providing technical assistance. So I thank my parents, for hanging in there, and also my grandma, for sending me love and tea.

Table of Contents

1	Collective neural dynamics in primate motor cortex	7
1.1	Primate motor cortex	7
1.1.1	Primary motor and premotor cortices	8
1.1.2	Connections with parietal cortex	8
1.1.3	Connections with subcortical structures	9
1.1.4	Laminar organization	9
1.1.5	Types of neurons	13
1.1.6	Lateral connectivity	15
1.1.7	Encoding in and decoding	17
1.1.8	Motor cortex as a pattern generator	18
1.1.9	Low-dimensional manifold trajectories	19
1.2	Collective dynamics and variability during movement	22
1.2.1	Variability in the nervous system	22
1.2.2	Implications of variability and redundancy for decoding	24
1.2.3	The local field potential reflects collective dynamics	25
1.2.4	Decoding from local field potentials	26
1.2.5	The LFP can reflect intrinsic sources of variability	26
1.2.6	A note on the output-null space	28
1.3	β -LFP oscillations in steady-state motor control	30
1.3.1	The functions of beta in motor control	32
1.3.2	Cognitive functions of beta oscillations	34
1.3.3	Computational roles and mechanisms of oscillations	37
1.3.4	The Origins of Cortical Beta Oscillations	43
1.3.5	The relationship between β -LFP oscillations and spiking	44
1.3.6	Reconciling cognitive and motor roles of beta oscillations	46
1.4	Waves in motor cortex	49
1.4.1	Origins of cortical waves	50
1.4.2	Waves as an indicator of collective dynamics	53
1.4.3	A computational role for beta waves in motor cortex	54
1.4.4	Optogenetics	54
1.5	Statistical methods	57
1.5.1	Extraction of LFP oscillations	58

1.5.1.1	Spike contamination artifacts	60
1.5.1.2	The analytic signal: instantaneous phase and amplitude	61
1.5.2	Stochastic point process modeling of spiking activity	62
1.5.2.1	Point process GLMs	63
1.5.2.2	Incorporating spiking history	66
1.5.2.3	Assessing predictive power	66
1.5.2.4	Cross-validation	68
1.5.2.5	Regularization	68
1.5.3	Characterizing spike-LFP phase coupling	69
1.5.3.1	Point process GLMs for spike-LFP phase coupling	70
1.5.3.2	Pairwise phase consistency	70
1.5.3.3	Addressing remaining source of bias	72
1.5.4	Classification of oscillatory LFP spatiotemporal patterns	73
1.5.4.1	Characterizing population activity	73
1.5.4.2	Spatiotemporal pattern analysis	75
2	Contribution of LFP dynamics to spiking variability in motor cortex during movement execution	83
2.1	Introduction	85
2.2	Methods	87
2.3	Results	94
2.3.1	Features of LFP oscillations predict single-neuron spiking with substantial power	95
2.3.2	LFP features contributing to prediction of single-neuron spiking	100
2.3.3	Predictive power of kinematics during naturalistic reach and grasp movements	101
2.3.4	LFP features are mostly redundant to kinematics when explaining single-neuron spiking variability	104
2.3.5	Intrinsic spiking history adds substantial power to kinematics in the prediction of single-neuron spiking variability	106
2.3.6	Conditioned on spiking history, contribution of LFP features to kinematic models is further reduced	108
2.4	Discussion	112
3	Dissociation between single-neuron spiking β-rhythmicity and transient β-LFP oscillations during movement preparation in primate motor cortex	117
3.1	Introduction	120
3.2	Methods	122
3.3	Results	131
3.3.1	Single units show sustained β -rhythmic spiking during steady-state movement preparation periods	132
3.3.2	β -LFP during steady-state movement preparation periods can be dissociated from β -rhythmic spiking	140
3.3.3	β -LFP appears as transient bursts	147

3.3.4	Population spiking activity also shows only weak coupling to on-going β -LFP oscillations	151
3.3.5	Single units show weak coupling to measures of population activity	152
3.3.6	Thin- and thick-spike cells show similar weak phase-coupling to β -LFP oscillations during preparatory steady-states	154
3.4	Discussion	156
4	Phase diversity and spatiotemporal wave dynamics in primate motor cortex local field potentials	164
4.1	Introduction	167
4.2	Methods	169
4.2.1	The CGID task	169
4.2.2	Neural recordings	170
4.2.3	Optogenetically induced spatiotemporal waves in motor cortex	170
4.2.4	Signal processing and LFP phase extraction	171
4.2.5	Assessing spatial synchrony and coherence	171
4.2.6	Spatial gradient of the Hilbert phase	173
4.2.7	Critical point analysis	174
4.2.8	Averaged phase maps for optogenetically induced spatiotemporal waves	175
4.3	Results	175
4.3.1	Motor evoked potentials and beta oscillations in the CGID task	176
4.3.1.1	Beta wave categorization	178
4.3.1.2	Task-related modulations in beta spatiotemporal activity	181
4.3.1.3	Beta wave statistics correlate with amplitude	184
4.3.1.4	Beta wavelength correlates with amplitude	185
4.3.1.5	Spatial variations in both the phase and amplitude of β -LFP may be important for understanding beta wave events	188
4.3.2	Optogenetically evoked traveling waves	191
4.3.2.1	Correlations in the summary statistics of optogenetically induce 50 Hz gamma waves	192
4.3.2.2	Comparison of spontaneous beta waves with optogenetically induced gamma	199
4.4	Discussion	202
5	Conclusion	209

List of Figures

1.1	Anatomy of visually-guided reaching and grasping	10
1.2	Laminar organization of rodent motor cortex	11
1.3	Neuronal subtypes in motor cortex.	12
1.4	Neural representations in motor cortex.	21
1.5	Noise in the nervous system	23
1.6	Local field potentials explain neural variability	27
1.7	Jasper and Andrews (1938): Recordings of sensorimotor beta oscillations in humans	31
1.8	Harvey et al. (2009): Phase coding in the hippocampus	38
1.9	Spike-LFP phase locking at beta frequencies.	45
1.10	Traveling waves in motor cortex	48
1.11	Three origins of traveling waves in cortex (Figure 1 from Ermentrout and Kleinfeld 2001.)	52
1.12	A model for spatiotemporal waves in beta oscillations (Figure 1 from Heit- mann et al. 2012)	55
1.13	Signal processing for reaching and grasping.	59
1.14	Point-process modeling of spiking activity	67
1.15	Detection and categorization of spatiotemporal structure in LFP oscillations	77
2.1	Behavioral and neural signals during the free-reach and grasp task	96
2.2	Features of ongoing and evoked multiband LFP oscillations predict single- neuron spiking: examples	98
2.3	Features of ongoing and evoked multiband LFP oscillations predict single- neuron spiking: summary across animals and areas	99
2.4	Breakdown of LFP predictive power by frequency band and LFP feature . .	102
2.5	Kinematic features predictive power for single-neuron spiking: summary across animals and areas	103
2.6	Kinematics and features of ongoing and evoked multiband LFP oscillations achieve similar predictive power on a neuron by neuron basis	105
2.7	LFP features add little predictive power to a kinematics model	107
2.8	Intrinsic spiking history carries complementary information to kinemat- ics features	110

2.9	Conditioned on intrinsic spiking history, the contribution of LFP features to kinematic models is redundant	111
3.1	The Cued Grasp with Instructed Delay (CGID) task	123
3.2	A subset of units fire rhythmically during steady-state movement preparation periods of the CGID task	133
3.3	Single-unit ISI statistics change across different stages of the task	134
3.4	Well-isolated single units fire rhythmically at beta frequency, and firing rates are stable across trials	135
3.5	The preferred firing frequency of rhythmic units varies, but typically falls within the beta band	136
3.6	Unimodal and bimodal units display similar statistics in the slow rhythmic component of the inter-spike interval distribution	137
3.7	~20 Hz beta oscillations are enhanced during the steady-state movement preparation periods of the CGID task	141
3.8	Single units exhibit β -rhythmicity that appears dissociated from the phase of the β -LFP oscillations	142
3.9	Single units exhibit β -rhythmicity that is dissociated from the phase of the β -LFP oscillations: Additional examples.	143
3.10	Spike-LFP phase coupling at the peak beta frequency is typically small during the steady-state movement preparation periods	146
3.11	β -LFP oscillations typically occur in transients and exhibit high trial-to-trial variability	148
3.12	Pairwise phase consistency increases only marginally for some units when the analysis is restricted to high beta transient events	149
3.13	Single unit firing rates during steady-state movement preparation periods are not affected by beta transients	150
3.14	Contrasting collective neural dynamics between steady-state movement preparation and movement execution periods	153
3.15	Single units cluster into thin- and thick-spike groups, but these groups show no consistent differences in PPC values	155
4.1	Evoked potentials and modulations in multi-unit spiking activity and β -LFP in the CGID task	177
4.2	Transient beta spatiotemporal wave patterns are diverse	179
4.3	Summary of the prevalence of different beta wave patterns the steady movement preparation period	180
4.4	Changes in beta spatial synchrony correlate with the CGID task stages	182
4.5	Modulations in beta spatiotemporal activity are task-related	183
4.6	Amplitude, phase gradient alignment, and spatial synchrony of spontaneous β -LFP oscillations are related	186
4.7	The amplitude and wavelength of traveling β -LFP plane waves are correlated	187
4.8	Evolution of a wave event in the analytic signal phase space	189
4.9	Limitations of phase statistics in summarizing LFP population dynamics	190
4.10	Optogenetically induced spatiotemporal gamma waves in motor cortex	193

4.11	Constant 1-second optogenetic stimulation of motor cortex induces complex oscillatory states	194
4.12	Spatial distribution of LFP spectral response to optogenetic stimulation . .	195
4.13	Spatial synchrony decreases and coherence increases during optogenetic stimulation	196
4.14	Induced gamma oscillations exhibit spatiotemporal wave structure	197
4.15	Phase-triggered averages of multi-band γ -LFP waves	198
4.16	Amplitude, wavelength, and synchrony of optogenetically induced γ -LFP spatiotemporal oscillations	200
4.17	Comparison of statistics between spontaneous β -LFP and induced γ -LFP spatiotemporal waves	201

Glossary of abbreviations

Neuroscience terms

ADP	After-Depolarization Potential (also DAP)	M1	Primary motor cortex
AHP	After-Hyperpolarization Potential	MEA	Multi-Electrode Array
AIP	Anterior Inferior Parietal lobule/gyrus	MIP	Medial Inferior Parietal lobule/gyrus
ALS	Amyotrophic Lateral Sclerosis	PMd	Dorsal Premotor cortex
CGID	Cued Grasp with Instructed Delay task	PMv	Ventral Premotor cortex
DAP	Depolarizing Afterpotential (also ADP)	PT	Pyramidal Tract neuron (synonym of PTN)
EEG	Electroencephalogram(graph)	PTN	Pyramidal Tract Neuron (synonym of PT)
EMG	Electromyogram(graph)	RS	Regular Spiking excitatory pyramidal cell
FRG	Free Reach to Grasp task	SAHP	Slow After-Hyperpolarization Potential
FS	Fast Spiking inhibitory interneuron	TMS	Transcranial Magnetic Stimulation
LFP	Local Field Potential	β-LFP	Beta oscillations (~20 Hz oscillations)
LIP	Lateral Intraparietal Area	γ-LFP	Gamma oscillations (~50 Hz oscillations)

Statistical terms

ACF	Auto-Correlation Function	PGD	Phase Gradient Directionality
AUC	Area Under the ROC Curve	PP	Predictive Power
CIF	Conditional Intensity Function	PPC	Pairwise Phase Consistency
CV	Coefficient of Variation	ROC	Receiver Operating Characteristic
GLM	Generalized Linear Model	SNR	Signal-to-Noise Ratio
ISI	Inter-Spike Interval distribution	STA	Spike-Triggered Average
PCA	Principal Component Analysis		

Preface

Understanding the neural computations involved in motor control is of fundamental scientific and clinical importance. From the standpoint of basic science, we wish to understand the mechanisms of voluntary action as one facet of the larger problem of understanding consciousness and cognition. Motor, premotor, and other areas of frontal cortex exhibit similarities in their local architecture and in their connectivity with other brain systems, and insights about the functioning of motor cortex may generalize to insights into how the frontal cortex plans and executes complex sequences of behavior. From a clinical standpoint, understanding the processes of motor control is essential for treatment of disorders of the motor system. A clear understanding of how actions are selected, timed, gated, executed, and controlled, is necessary for optimal decoding of action signals for neuroprosthetics applications for individuals with amyotrophic lateral sclerosis and upper cervical cord injury. A clear understanding of the regulation of movement is also necessary for the treatment of movement disorders, including Parkinson's, and the numerous movement disorders that can arise from injury to the motor system such as dystonia and chorea-ballism.

Historically, the brain has often been likened to a very big version of whatever most complex technology man had available at the time, the telephone switchboard, the super-computer. These analogies fail dramatically, the failure of a single transistor can cause a computer to crash, but the loss of a few neurons is hardly noticeable. This robustness is possible because single neurons rarely exert overwhelming influence over the trajectory of the brain state as a whole. Instead, the dynamics emerge from average behaviors and (pairwise and higher-order) interactions among the neurons. The planning and execution of movements involves coordination between recurrently connected populations of billions¹ of single neurons distributed across multiple areas of the brain. This number remains large even for the smaller neocortical patches examined in this thesis: a rough

¹Order of magnitude estimate. See Braendgaard et al. 1990; Gredal et al. 2000

estimate indicates that each $4 \times 4 \text{ mm}^2$ recorded patch in primary, ventral and dorsal motor areas contains about 2 million neurons. Neurons in motor cortex operate collectively to prepare and to generate movements, forming a distributed system that can maintain normal motor functioning in the face of noise, uncertainty, and unreliability in single neuron activity. Because the neural correlates of action preparation and execution are distributed, we can decode actions with substantial accuracy from a random subsample of single neurons picked up by chance in a multi-electrode array recording. However, the distributed and collective nature of the neural dynamics means that we cannot understand the operation of the motor system without considering the simultaneous operation of all its components. Simultaneous recordings from all neural structures involved in motor control is currently infeasible, and thus we are faced with the task of piecing together a theory of motor control based on observations of small sections of the motor system.

Current research in the motor system is limited by the technologies for recording and manipulating neural activity. One common recording approach, extended recently also to brain-machine interface research for paralyzed individuals, is microelectrode array (MEA) technology, which permits simultaneous recording from up to a few hundred microelectrodes covering a few square milliliters of cortex. These MEAs provide information about neural activity at multiple spatial and temporal scales, and can simultaneously record both dozens of isolated single neurons as well as indicators of population activity like multi-unit spiking activity (MUA) and ongoing local field potentials (LFPs) in multiple frequency bands. Previous work in the Truccolo lab has focused on examining collective dynamics at the level of spiking activity in ensembles of single neurons (e.g. Aghagolzadeh and Truccolo 2015; Truccolo et al. 2010). In this thesis, I focus on the collective neural dynamics reflected in LFPs. Although still a matter of intense debate (even the dipole models commonly used as LFP generators may be wrong; Bedard and Destexhe 2013; Destexhe and Bedard 2012), LFPs are thought to reflect primarily spatial averages of postsynaptic potential over large neuronal populations. The aim here is to better understand how these

signals at varying spatial and temporal scales relate, and to characterize the features of collective neural activity that may be observed. We will also address data coming from the emerging technology of optogenetic stimulation in primate motor cortex.

Specifically, the thesis examines three main topics in collective dynamics in motor cortex. First, what processes influence neural spiking in motor cortex during movement? Second, what is the relationship between neural firing and local population dynamics during movement planning and preparation? Third, how do local populations of neurons interact and coordinate across motor cortex during the preparation of movement? Chapter 1 covers relevant background, including neuroanatomy and our current understanding of neural coding, variability, oscillations, and spatiotemporal wave dynamics in primate motor cortex. Chapter 1 section 5 introduces the statistical methods used to model single neuron spiking, to examine its relationship to ongoing LFP oscillations, and to characterize spatiotemporal wave patterns.

Chapter 2 of this thesis explores variability in single neuron spiking in primate motor cortex during movement execution, and addresses the question of what accounts for the excess trial-to-trial variability observed in neural spiking during movement. As stated above, neuronal networks in neocortex are highly recurrent. Recurrent connectivity is known to lead to complex spontaneous activity, which by itself can contribute to single neuron variability. Understanding the nature of such neural variability is an important step toward a theory of how motor cortex robustly generates movement. The question is also important for improving neural decoding for brain-machine interface applications. The main finding of Chapter 2 is that collective dynamics reflected in LFPs, even though highly predictive of single neuron spiking, are mostly redundant to the information available in movement parameters (e.g. kinematics) and, therefore, cannot explain the excess variability. Specifically, information related to spiking variability in local field potentials is composed of two components: a slow, motor-evoked component that is related to generating the dynamics of reaching and grasping actions, and a fast-timescale component

that is related to stochastic spiking history effects. These results come in contrast to similar studies in sensory cortex under anesthesia, which found that local field potentials reflect intrinsic neural dynamics unrelated to stimulus encoding, and add weight to the conjecture that the spontaneous collective dynamics reflected in motor cortex LFPs are used for movement generation.

Chapter 3 of this thesis investigates motor cortex ~ 20 Hz rhythmic beta activity, a salient feature of steady-state movement preparation periods. Although many have proposed that these oscillations play some role in stabilizing motor steady-states, the computational and functional significance of this rhythmic activity is unresolved. The main finding of Chapter 3 is that the activity of rhythmically spiking single neurons can be dissociated from population oscillations evidenced in local field potentials. This result provides an important constraint on collective beta-frequency dynamics for models of motor steady-states, and raises questions about the potential computational significance of heterogeneous and variable relationships between rhythmic neural firing and population oscillations at the same frequency as reflected in beta LFPs.

Chapter 4 of this thesis investigates spatiotemporal waves observed in local field potential beta oscillations during movement preparation steady states, as well as complex multi-band spatiotemporal gamma oscillations induced by optogenetic stimulation in primate motor cortex. Spatiotemporal wave phenomena have been extensively studied in, for example, the retina and visual cortex, but remain relatively unexplored in the motor system. Studies to-date have focused on planar traveling wave activity in beta, and we find that the true activity is in fact much more variable and complex, especially in area PMv, which has not been examined before. The results of Chapter 4 support the theory that beta spatiotemporal dynamics represent spatial organization of local ongoing beta oscillation phases, but also show that amplitude may be an important feature of beta spatiotemporal dynamics. The analysis of optogenetically induced spatiotemporal gamma oscillations provides an important contrast or reference for the interpretation of

the examined spatiotemporal beta waves. The induced gamma waves offer a clear example of waves generated in an excitable medium. Chapter 4 also provides a thorough description of the response of motor cortex to optogenetic stimulation, and demonstrates that the gamma spatiotemporal dynamics in cortex can be comprised of multiple interacting narrow-band oscillations. Clarifying the mechanisms behind this phenomenon is an important first step for developing closed-loop stimulation methods for research and neuroprosthetic applications.

-

Chapter 1

Collective neural dynamics in primate motor cortex

MICHAEL RULE

This section reviews relevant background information for Chapters 2 through 4 of this thesis. The work presented in this thesis has been done in motor cortex of rhesus macaques, and so we begin with a brief review of the primate motor system. Then, we discuss collective dynamics in motor cortex during movement, including what is known about population activity and sources of variability in neural firing. Third, we discuss oscillatory states in motor cortex that can be observed during movement preparation. Fourth, we cover spatiotemporally organized population activity in the brain, including wave dynamics in motor cortex. The chapter concludes with an introduction to the statistical methods that we use to investigate these topics of collective dynamics in motor cortex.

1.1 Primate motor cortex

This thesis focuses on collective dynamics in motor and premotor cortices during both naturalistic and cued visually guided reaching and grasping. This section briefly reviews major inputs and outputs of motor and premotor cortex involved in these tasks, as well

as the laminar organization and what is known about neural cell types in these areas of the brain. Additionally, we review current understanding of how motor cortex represents and generates the neural activity required for movement.

1.1.1 Primary motor and premotor cortices

Rizzolatti et al. (1998) provide a review of the anatomy of sensorimotor cortex in rhesus macaques. Primary motor cortex (M1) in rhesus macaques lies on the anterior bank of the central sulcus and the precentral gyrus, and constitutes Brodmann areas 4 and 6 (Rizzolatti et al., 1998). A number of cortical areas connect directly to motor cortex, including the supplementary motor area (SMA), the cingulate motor areas, and dorsal and ventral premotor cortex (Dum and Strick, 2002). The dorsal and ventral premotor (PMd and PMv) areas are located anterior to M1 (Dum and Strick, 2002), and are on the outer surface of the brain and therefore accessible for multi-electrode array (MEA) recording approaches (Figure 1.1). PMd is adjacent to M1, in an anterior and dorsal position in the precentral gyrus, and dorsal to the spur of the arcuate sulcus. In macaques, PMd is located ventrally and slightly anterior to an anatomical feature called the superior precentral dimple. PMv is located below the spur of the arcuate sulcus (Dum and Strick 2002, see Figure 1.1).

1.1.2 Connections with parietal cortex

Areas PMv and PMd are involved in visually guided reaching and grasping, and connect to several parietal areas that represent the sensory information necessary for these behaviors. Tanné-Gariépy et al. (2002) found that PMd and PMv project to separate sets of parietal areas, corresponding roughly to information about spatial geometry necessary for reaching, and information about object shape important for grasping, respectively. Area PMd receives inputs from areas MIP, PEc, PGm (also called MDP or 7m; Shipp et al. 1998), PO, 7b and PE (Tanné-Gariépy et al., 2002). Area MIP neurons encode goal-directed hand motion information (Grefkes and Fink, 2005). Area PEc neurons are tuned to hand and eye

position (Ferraina et al., 2001), as are neurons in area PGm (Ferraina et al., 1997). Area PO plays a role in encoding visual space (Boussaoud et al., 1990). Taken together, the parietal areas that project to PMd convey information about visual space, spatial representations for reaching, and the conjunction between the two for hand-eye coordination. Area PMv receives inputs from areas AIP, 7b, PE, and somatosensory cortex (Tanné-Gariépy et al., 2002). Area AIP encodes object shape information relevant for grasp shaping (Grefkes and Fink, 2005). Area 7b encodes somatosensory information as well as visual information (Hyva et al., 1981), and area PE is thought to relate to somatosensory functions (Johnson et al., 1996). Taken together, the parietal areas that project to area PMv convey information important for grasping, including visual information about object geometry as well as tactile feedback important for grasp shaping.

1.1.3 Connections with subcortical structures

Motor cortex is involved in a number of recurrent connectivity loops with subcortical structures, including interactions with the basal ganglia, extrapyramidal motor structures of the brainstem, and the cerebellum. Inputs to motor cortex from these structures arrive from the ventral lateral and ventral anterior thalamic nuclei¹, and synapse primarily in the superficial cortical layers (I/II), and also in layer V (Bosch-Bouju et al., 2015). These other structures are beyond the scope of the thesis, but are relevant to the interpretation of local field potentials (LFPs) observed in motor cortex, which arise in part from synaptic inputs (Buzsáki et al., 2012).² Thalamic inputs may form part of a system that gives rise to intrinsic neural dynamics for movement pattern generation (§ 1.2, Chapter 1), and may also play a role in gating beta oscillations in the motor system (§ 1.3, Chapters 3, 4).

¹Nomenclature is not standardized in primates (Bosch-Bouju et al., 2015), see Krack et al. (2002).

²Several studies have called into question the assumption that LFPs reflect synaptic inputs. See § 1.2.3.

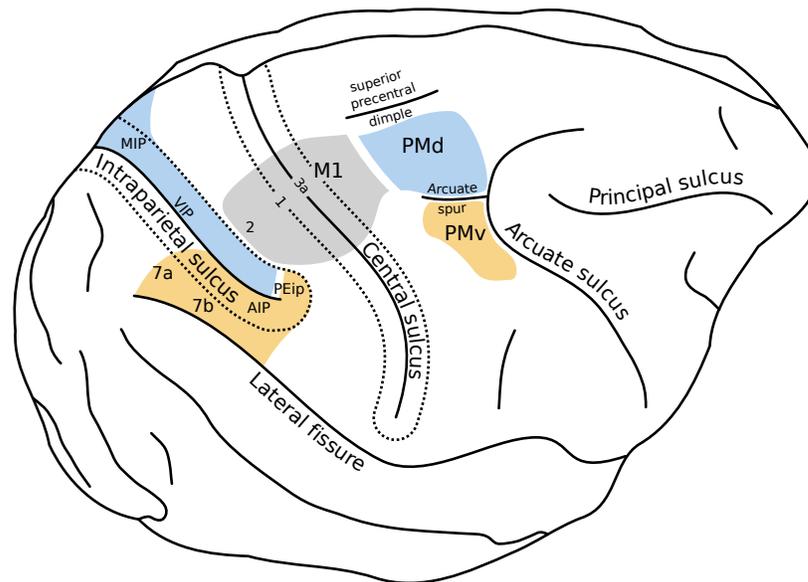


Figure 1.1: *Anatomy of visually-guided reaching and grasping.* During visually guided reaching and grasping, the arm and hand area of M1 coordinates with the dorsal and ventral premotor areas PMd and PMv. In this illustration, reciprocally connected motor and parietal areas are shaded in common colors. Premotor areas receive segregated streams of visual information from parietal cortex. Area PMd receives information about spatial geometry important for reaching from dorsal parietal areas (shaded in blue). Area PMv receives information about object geometry important for grasping from the parietal areas shaded in orange. Area M1 also receives feedback from somatosensory cortex (areas 3a,1,2, shaded in grey). Connections between parietal and premotor cortex are taken from Tanné-Gariépy et al. (2002), and anatomical boundaries of premotor areas are taken from Dum and Strick (2002).

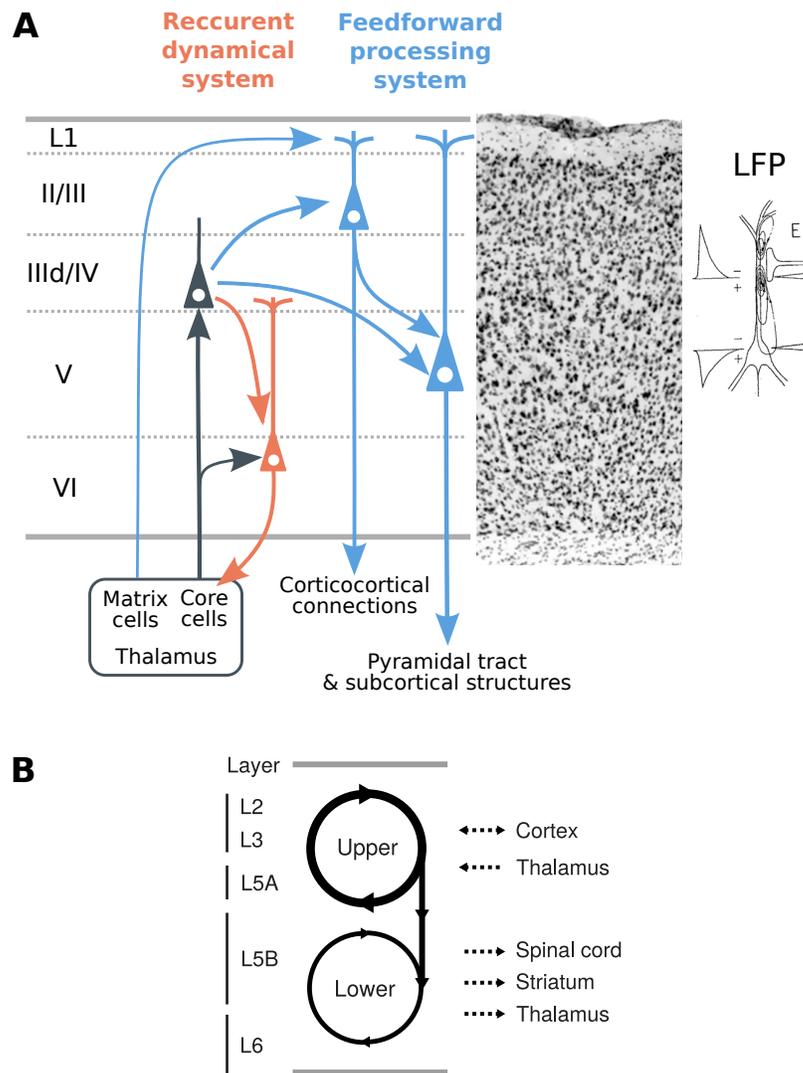


Figure 1.2: Laminar organization of rodent motor cortex. (A) Figure adapted from Kaneko (2013). Layer V of motor cortex gives rise to output to spinal cord and subcortical structures. Layer II/III participates in corticocortical and thalamocortical connections. Recurrent dynamics between motor thalamus and layers VI and III d may play a pattern-generating role in motor output. Feed-forward transformation of corticocortical and thalamic inputs in layers II/III and 1 may play a role in converting sensory and contextual information into motor output. Importantly, this study was conducted in rodents, whose motor cortex exhibits an identifiable layer IV not present in the motor cortex in primates and may exhibit a different pattern of thalamocortical projections. (A, **Illustration inset**) Nissl-stained forelimb representation area of M1 in a mouse (Yamawaki et al., 2015). (A, **LFP illustration**) Inset detailing generation of LFP adapted from Niedermeyer and da Silva (2005). Dendritic currents arising from synaptic conductances contribute to local field potentials (LFPs). The LFP provides complementary information to spiking output, and the relationship between spikes and LFP in motor cortex may reflect aspects of ongoing neural computation. (B) Recurrent connectivity in motor cortex, adapted from Weiler et al. (2008) figure 4. Recurrent intra- and inter-laminar connections form upper and lower processing loops. Inputs to motor cortex arrive primarily in the upper layers, and outputs arise from the lower layers.

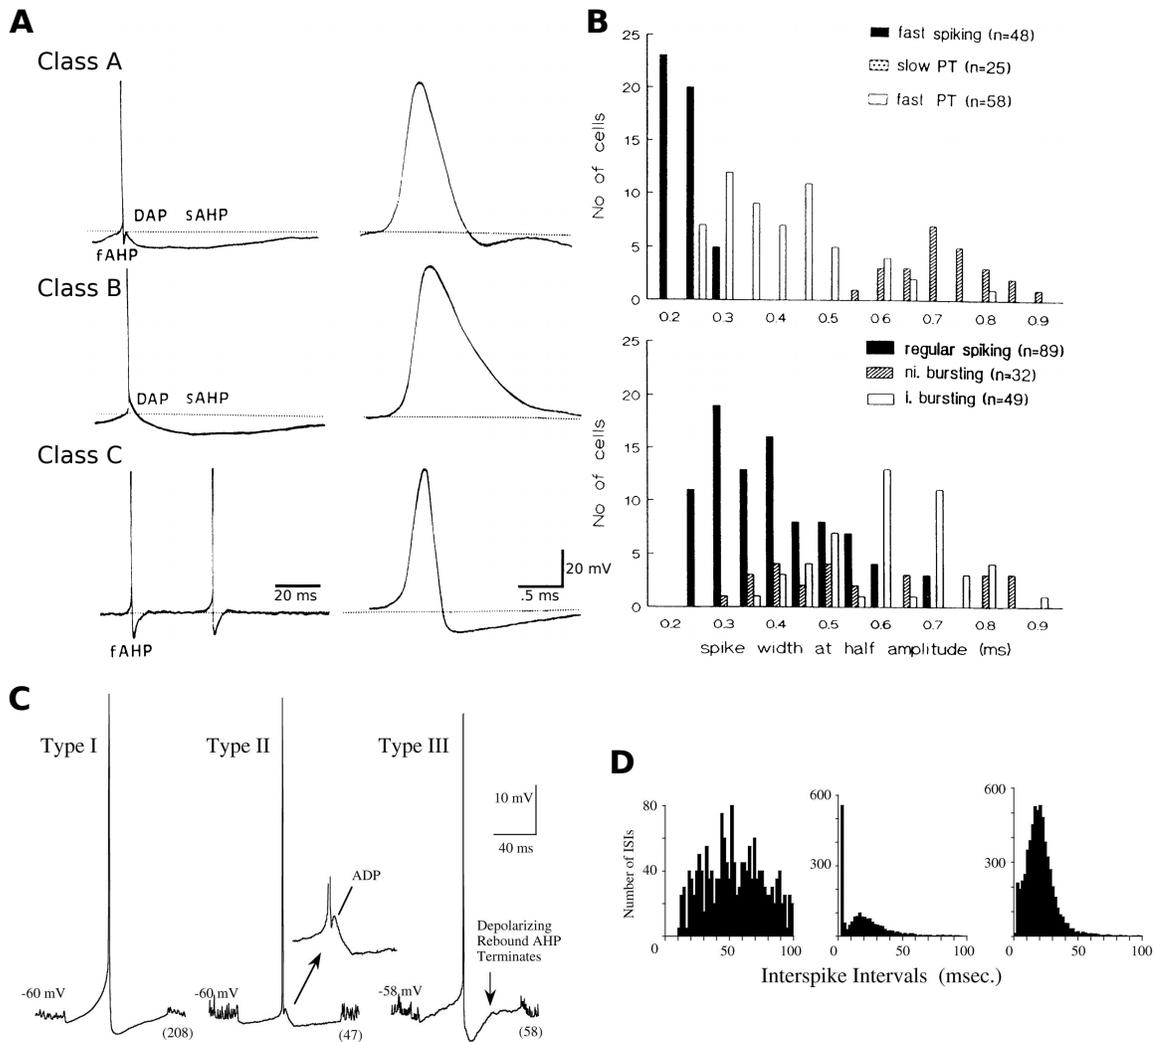


Figure 1.3: Neuronal subtypes in motor cortex. Subfigures A and B are adapted from Figure 1 of Baranyi et al. (1993b). Subfigures C and D are adapted from Figure 2 of Chen and Fetz (2005). **(A)** Three different types of action potentials were identified by Baranyi et al. (1993b) in cat motor cortex. *Class A*: Fast spiking pyramidal tract neurons (PTNs), and some types of non-PTN neurons, exhibit depolarizing afterpotentials (DAPs) that enable firing of spike doublets and bursts. These cells also exhibit slow afterhyperpolarization potentials (SAHPs). *Class B*: Slow spiking PTNs and some types of non-PTN neurons exhibit wide spikes and SAHPs. *Class C*: Fast spiking (FS) cells exhibit short spike widths, lack DAPs, and have very short afterhyperpolarization potentials. **(B)** Distribution of spike widths for various cell types. FS cells exhibit the narrowest spikes, but overlap to some extent with fast-spiking PTNs. Slow PTNs exhibit the widest spikes on average. Non PTN neurons including regular spiking cells, and inactivating (i) and noninactivating (ni) bursting neurons, exhibit varied spike widths. **(C)** Motor neuron classes in primates are similar to those observed in cats (Chen and Fetz, 2005). *Type I* is similar to type B from Baranyi et al. (1993b), exhibiting only a SAHP. *Type II* is similar to type A from Baranyi et al. (1993b), exhibiting both a DAP (or ADP as abbreviated by Chen and Fetz (2005) and a SAHP. *Type III* is similar to type C from Baranyi et al. (1993b), exhibiting narrow spikes and rapid repolarization after spiking. **(D)** Spike waveform properties correlate with differences in inter-spike interval (ISI) histogram distributions. Type I cells fire irregularly at low rates, consistent with the slow PTNs and regular spiking cells observed by Baranyi et al. (1993b). Type II cells exhibit a bursting component related to the ADP, and a rhythmic component related to the SAHP. Type III cells fire rhythmically at high rates, corresponding to the Type C FS cells from Baranyi et al. (1993b).

1.1.4 Laminar organization

Recent studies in rodents have suggested that motor cortex is functionally divided into two interacting subsystems (Figure 1.2). Weiler et al. (2008) found that recurrent connectivity in motor cortex clusters into a superficial system in layers II/III and the upper portion of layer V, and a deep system in layer VI and the lower portion of layer V (Figure 1.2B). Kaneko (2013) hypothesizes that the superficial system is related to feed-forward transformation of inputs into motor output, and that the deep layers form part of a recurrent dynamical system in interaction with motor thalamus (Figure 1.2B). Therefore, one expects neural signals originating from deep versus superficial layers to reflect different aspects of neural computation related to motor control. This is particularly relevant for interpreting local field potentials, which can reflect both summed synaptic inputs into the local spiking population as well as volume-conducted potentials from other cortical layers. We note, however, that the laminar organization of primate motor cortex remains controversial, with layer IV initially described in adult humans by Cajal, but reported as absent in subsequent studies. More recent work (García-Cabezas and Barbas, 2014) supports the existence of layer IV cells in motor cortex of adult primates, although these cells may not form a distinctive layer.

In this work we analyze recordings from multi-electrode arrays (MEAs; Utah Array, Blackrock Microsystems) with an electrode length of 1.5 mm. It is thought that the single units isolated in these experiments are in layer V. Some amount of compression and cortical thinning through cell loss may occur in the tissue under the MEA implant, making the laminar origin of MEA recordings uncertain, and layer III is also a possibility. Histological assessments, which were not available at the time of this writing, would be necessary to verify the targeted layers.

1.1.5 Types of neurons

The diversity and number of distinct neuronal subtypes in neocortex remains an open question (for review, see Markram et al. 2004). Determining the functional roles of neuron types is important for understanding the dynamics of motor cortex, as each subtype may play a different role in network dynamics. In the simplest possible schema, one can identify two major types of neurons (interneurons vs. principal cells) based on extracellular spikes recorded on MEAs (Barthó et al., 2004; McCormick et al., 1985). In sensory cortex single units can be classified into putative fast spiking (FS) inhibitory and regular spiking (RS) pyramidal cells based on the width of the extracellularly recorded spike waveform (e.g. Cardin et al. 2007, Woloszyn and Sheinberg 2012), and these classifications fit reasonably well with major morphological differences (Foehring et al., 1991). Some studies have categorized cells in motor cortex based on spike width, but surprisingly found little difference between the dynamics of cells with thin and thick spikes (Kaufman et al., 2010). Vigneswaran et al. (2011) showed that in motor cortex, pyramidal tract neurons can exhibit narrow spikes and high firing rates, making them difficult to distinguish from putative inhibitory FS interneurons. Therefore, the lack of electrophysiological differences between thin and thick spiking cells recorded on MEAs from motor cortex may result from both being subtypes of excitatory pyramidal tract neurons. Additionally, the relative sparseness of the thin-spike, aspiny neurons (~20%) and their restricted electrical field compared to pyramidal neurons, means that they may be less likely to be sampled³.

Studies of intracellular action potentials and spiking statistics have found a diversity of neurons in motor cortex. Baranyi et al. (1993a) and Baranyi et al. (1993b) identified a number of different types of neurons in cat motor cortex, including inhibitory FS interneurons, slow and fast spiking PTNs, regular spiking non-PTNs, and bursting non-PTN types with and without spike-rate adaptation. These units had overlapping distributions of spike widths (Figure 1.3B). Baranyi et al. (1993b) observed three main types of action

³Dr. John Donoghue, personal communication.

potentials in primate motor cortex: thin spikes with fast repolarization, thick spikes with slow afterhyperpolarization potentials (SAHPs), and thin spikes with a depolarizing afterpotential (DAP) followed by a SAHP (Figure 1.3A). Chen and Fetz (2005) also found three classes of neurons in primate motor cortex that can be distinguished by spike waveforms: regular spiking cells with thick spikes, and two types of fast spiking cell: one that exhibits bursting and rhythmicity as a result of a short DAP and a SAHP, and another that fires regularly and at high rates, exhibiting neither bursting nor SAHP (Figure 1.3C,D). Overall, the results of Baranyi et al. 1993a and Chen and Fetz (2005) suggest that units in motor cortex can be classified into at least three groups based on extracellular recordings. Thick spikes correspond to regular spiking excitatory pyramidal cells, a subset of which are PTNs. Thin spikes correspond to two categories: inhibitory FS interneurons, and fast-spiking PTNs. Both exhibit rhythmicity, but only the fast-spiking PTNs exhibit bursting and spike-doublets, as a consequence of the DAP.

Vigneswaran et al. (2011) note that the fast spiking PTNs in motor cortex likely correspond to the fast-conducting pyramidal tract fibers (Calvin and Sypert, 1976; Koike et al., 1970, 1968; Takahashi, 1965), which were found to be associated with fast, transient "kinetic" motor output, as opposed to slow-conducting fibers which relate to sustained or "tonic" motor output (Koike et al., 1970). This is interesting from a standpoint of motor decoding, as fast spiking PTNs should be interpreted as having more of a phasic relationship to motor output compared to slow-spiking PTNs. Fast spiking PTNs activate at high rates earlier in movements and display spike-rate adaptation. For the purposes of understanding computational and neural circuit mechanisms in motor cortex, the ambiguity of cell types recorded on MEAs frustrates attempts to analyze excitatory and inhibitory neurons separately. Recent work by Best et al. (2016) found that motor cortex thin-spike cells convey more information than thick-spike cells for movement decoding, supporting the theory that some thin-spike cells isolated on MEA recordings of primate motor cortex are a subtype of PTNs.

1.1.6 Lateral connectivity

Chapter 4 of this thesis investigates two spatiotemporal wave phenomena in motor cortex. To clarify interpretations of functional connectivity and collective wave activity in motor cortex, it is necessary to review what is known about the underlying anatomical connectivity in motor cortex. Takahashi et al. (2015) analyze functional connectivity, i.e. pairwise time-causal statistical relationships between isolated single units, in motor cortex during beta (10-45 Hz, see ave events. They found that (1) *"The circular distribution of excitatory connection directions became bimodal and closely oriented to the beta wave axis"* and (2) *"The circular distributions of inhibitory connection directions were also bimodal but oriented almost orthogonal to the beta wave axis"*. They also found that the beta state modulated spiking functional connectivity, and it is conceivable that this connectivity reflects the impact of beta waves on single-unit activity.

Measurements of conduction velocities in motor cortex are relevant for understanding which effects in Chapter 4 might be mediated by direct spread of neural activity through excitatory connections. Takahashi et al. (2015) found that the conduction velocity between functionally connected pairs of units is approximately 10 cm/s, and noted that this is comparable to both the observed propagation velocity of motor cortex beta waves, and the experimentally measured conduction speeds of 10 cm/s observed in electrical stimulation in rat motor cortex (Aroniadou and Keller, 1993). Another study in motor cortex, Capaday et al. (2011), found that multi-unit activity in cat motor cortex spreads from a stimulation site at about 14 cm/s. These measured propagation speeds relate to polysynaptic transmission through excitatory connections, since the distances over which the observed phenomena propagate are larger than the range of typical monosynaptic connections. Although some activity may spread via direct long-range connections, the gradual radiation of activity from the stimulation site observed in the above studies suggests that these conduction velocities reflect propagating activity mediated by lateral excito-

excitatory connections.

Aroniadou and Keller (1993) found that typical monosynaptic excitatory lateral connections extend for up to 450 μm . Similarly, Asanuma and Rosén (1973) found that monosynaptic connections from layers II, III, and IV spread up to 500 μm . Although their work did not explicitly separate excitatory and inhibitory populations, they suggested that the spread of putative excitatory connections from superficial and intermediate layers is approximately 450 μm and that the spread of putative inhibitory connections from deep layers is approximately 600 μm . Aroniadou and Keller (1993) also observed that electrical stimulation elicits an outward traveling LFP response that decays exponentially with a space constant of approximately $\lambda = 450 \mu\text{m}$.

Huntley and Jones (1991) found complex lateral connectivity in motor cortex over a variety of spatial scales that integrates areas of motor cortex needed for reaching and grasping (e.g. finger, wrist, and elbow representations). This pattern of connectivity is very different from the almost crystalline arrangement of cortical columns with center-surround inhibition (“Mexican hat”) often considered in studies of waves and pattern formation (e.g. Ermentrout and Cowan 1979; Rule et al. 2011). Some models of spatiotemporal wave phenomena in motor cortex (Heitmann et al., 2012) require center-surround connectivity, however there is little evidence that the connections in motor cortex reflect this arrangement. Understanding how spatiotemporally organized waves emerge in motor cortex is therefore an open problem.

1.1.7 Encoding in and decoding

Georgopoulos et al. (1982) found that single-unit firing rates in motor cortex can be tuned to the direction of reaching movements (Figure 1.4A). Georgopoulos et al. (1986) subsequently showed that the direction of movement can be decoded as a population average of the preferred directions of individual neurons, weighted by each neuron’s firing rate (Figure 1.4B). This provided a good first approximation of motor encoding that has proven

sufficient for accurately decoding movement based on populations of neurons in motor cortex (Schwartz et al., 2001). However, there is no reason to assume that neurons encode a particular physical quantity such as position, velocity, or force. Shoham et al. (2005) found that neurons can exhibit tuning up to the 3rd derivative of position, indicating that neurons in motor cortex can encode a combination of position, velocity, acceleration, and jerk.

Modern brain machine interfaces use a Kalman-filtered population velocity vector average (Wu et al., 2006), combined with a collection of heuristics to address common sources of decoding errors, such as slow firing rate nonstationarities (e.g. Bishop et al. 2014). Chhatbar and Francis (2013) found that treating neurons as if they encode a combination of torque and position can improve brain machine interface control in some circumstances. Optimizing models for brain-machine interface decoding remains a subject of active research. Curiously, motor cortex neurons can also show visual responses related to action observation and planning (Rizzolatti et al., 1996; Vargas-Irwin et al., 2015), indicating that neurons in motor cortex can exhibit complex tuning and variations in firing rate not necessarily related to motor output.

1.1.8 Motor cortex as a pattern generator

This subsection briefly covers what is known about the role of motor cortex in action execution, including theories about how motor cortex encodes movement information, and how motor cortex might prepare and generate dynamic output. We begin with a discussion of what motor cortex is thought to encode. We cover the history of decoding from motor cortex, and discuss the implications of this work for understanding motor cortex. We will also cover neural population dynamics during reaching and grasping, including low-dimensional manifold trajectories, and discuss some of the implications of these models for observing motor cortex.

Churchland et al. (2010b) has noted that the function of motor cortex is not to encode anything per se, but to generate the spatiotemporal patterns required to drive the spinal cord and musculoskeletal plant in the proper way to generate coordinated movements. If motor cortex serves as a pattern generator for movement, we would expect motor cortex to contain correlates of the dynamical variables required for generating movement patterns. Accordingly, it has been shown that motor cortical neurons are tuned to extended velocity trajectories, representing the recent and upcoming dynamical evolution of motor output (Figure 1.4C; Hatsopoulos et al. 2007). These “pathlets” form an additive basis of motor primitives for composing dynamic movements (Hatsopoulos and Amit, 2012).

Churchland and colleagues (Churchland and Cunningham, 2015; Churchland et al., 2012; Shenoy et al., 2013) have proposed that pattern generation in motor cortex is fundamentally oscillatory. The essence of this theory is that mammalian motor control can be viewed as a highly evolved version of primitive oscillatory pattern generators. Multiple damped oscillations can be combined to generate the sequences of muscle activations needed to perform complex actions. The theory that dynamic movements are composed of a collection of transient oscillatory components is attractive because such components constitute a flexible basis for approximating arbitrary smooth functions. However, this theory leaves unresolved exactly how populations of single-neurons might act collectively to form such a basis.

1.1.9 Low-dimensional manifold trajectories

Using machine learning, it is possible to predict complex movements with substantial accuracy using a small, sparse, random subsample of the firing of motor cortical neurons. This implies that information about movement in motor cortex is distributed. This is to be expected because the ultimate targets of motor cortex output (muscles), generate forces that are low-pass filtered summations of the activity of many hundreds to thousands of alpha motor neurons (McComas, 1991). Accordingly, the features of spiking

population outputs that relate to movement tend to be slowly varying population trajectories distributed over many neurons. It seems that motor cortex contains far more neurons than are strictly required for representing motor output. Reconciling the apparent low-dimensional and low-frequency output of the motor system with the potentially high-dimensional and fast-timescale spiking dynamics remains a major open problem.

Work in Aghagolzadeh and Truccolo (2014, 2015); Byron et al. (2009); Vargas-Irwin et al. (2015) has shown that the intrinsic dimensionality of neural variability in motor cortex can be low-dimensional (Figure 1.4D). In these studies, the observed neural variability in motor cortex was concentrated in a sub-manifold much smaller than what would be expected if neurons were independent. During movement execution, neural variability may be confined to a low-dimensional manifold containing no more degrees of freedom than strictly required for motor output. However, it could also be the case that the low-dimensional manifold activity exhibits more degrees of freedom than the motor output. The extent to which such low-dimensional manifolds might remain similar over time or across contexts is unknown. Movement decoders cannot assume that the low-dimensional manifolds representing motor output will be stable across time and context, and efficiently estimating and tracking changes in the low-dimensional structure of motor cortex remains an open problem. In the case that low-dimensional manifolds change, the larger subspace over which these manifolds vary constitutes an additional source of variability. Because short-timescale variability remains confined to a lower-dimensional subspace, statistically sampling this larger, slowly-varying space is especially challenging. Continuous recalibration of decoders could address this issue (Bishop et al., 2014).

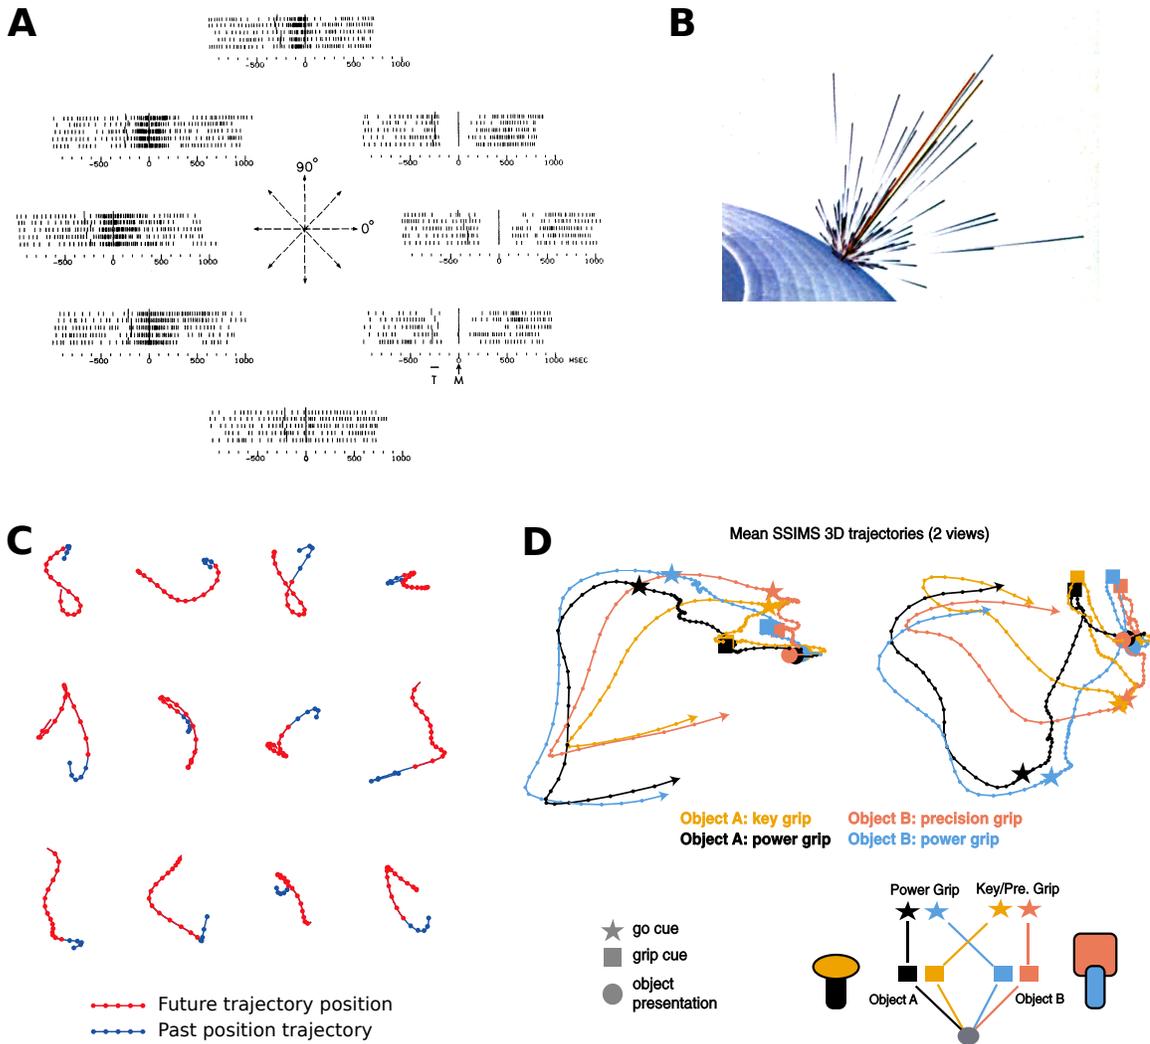


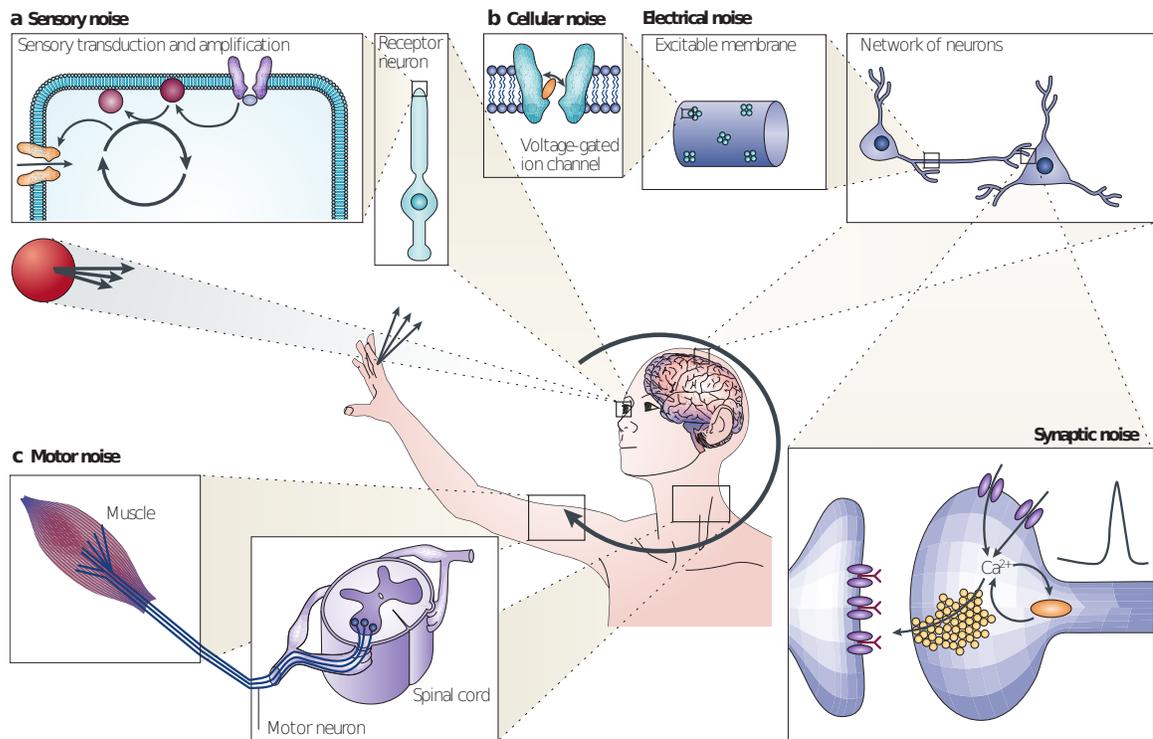
Figure 1.4: Neural representations in motor cortex. (A) Georgopoulos et al. (1982) found that, during reaching movements, neurons in motor cortex are approximately cosine tuned to a single reaching direction. (B) Georgopoulos et al. (1986) showed that the direction of reaching can be decoded from populations in motor cortex by taking a population average of neurons' preferred directions, weighted by their firing rates. Therefore, movement representations in motor cortex are distributed over many neurons. This figure has been denoised and value-inverted from the original version to improve readability. (C) Adapted from Figure 3 from Hatsopoulos et al. (2007). Hatsopoulos et al. (2007) show that neurons in motor cortex encode not “instantaneous” motor output, but rather extended trajectories reflecting recent and upcoming motor dynamics. (D) Motor cortex is now thought to encode dynamical trajectories related to motor processing and pattern generation. For example, Figure 7 in Vargas-Irwin et al. (2015) shows that reaching and grasping trajectories lie in a low-dimensional state space. Preparatory activity in response to visual cues correlates with trajectories in motor state space, which configure the system to generate upcoming reaching and grasping movements.

1.2 Collective dynamics and variability during movement

Motor cortex is the primary source of cortical innervation to the spinal motor neurons that drive muscle contraction, and is fundamentally involved in the generation of movement. However, it is still not known exactly what computational problems are solved by motor cortex, and how neurons in motor cortex implement solutions to these problems. We investigate this question by considering the statistical relationships between the neural signals that we observe in motor cortex and behavioral output.

1.2.1 Variability in the nervous system

Understanding what information is contained in the firing output of motor cortical neurons is a crucial step for understanding the function of motor cortex. Decoding information from spiking units in motor cortex is complicated by the fact that neurons exhibit a large amount of variability that cannot be related directly to movement. Cortical responses recorded in vivo can be highly variable (Azouz and Gray, 1999): a single neuron can fire very different sequences of spikes in response to the same stimulus. Spontaneous collective dynamics, in the form of spatiotemporal patterns of activity, are likely to arise in the highly recurrent neuronal networks found in cortex. A major contributor to the variability of single neurons in vivo may be the intrinsic collective dynamics (Arieli et al., 1996; Stevens and Zador, 1998) of the brain. Therefore, variability in motor cortex unrelated to movement may relate to other cognitive processes. Motor output is conditioned on behavioral, motivational, and cognitive state, as well as environmental and task context and sensory feedback. We expect to find correlates of these processes in motor cortex, which could explain the apparent variability of spiking activity. However, spiking variability may also arise purely from noise and stochasticity, and it is important to understand what components of neural variability are noise, and what components are signatures of perceptual and cognitive processes.



Nature Reviews | Neuroscience

Figure 1.5: *Noise in the nervous system.* Noise in the nervous system gives rise to neural variability. All stages of sensorimotor processing exhibit noise to some extent. Sensory transduction (a) and muscle activation (c) exhibit noise. Noise may arise during sensorimotor processing from cellular and electrical noise (b) within the brain, as well as synaptic release failure. Distinguishing neural variability arising from noise from that arising from intrinsic neural processes is important for understanding neural computation. Additionally, understanding how the motor system solves sensorimotor control problems robustly in the face of noise is important for clarifying the computational mechanisms of motor cortex. (Figure 1 from Faisal et al. 2008)

The contributions of various processes to neural variability has been widely studied. It is known that stimulus onset reduces neural variability in a variety of brain areas (Churchland et al., 2010a), and that naturalistic stimuli give rise to less variable neural responses than artificial stimuli (Hasson et al., 2010; van Steveninck et al., 1997). These results suggest that neural variability depends on stimulus and behavioral context. Intrinsic sources of noise in the nervous system include synaptic release failure, thermal and electrical noise, and chaotic and stochastic effects (Figure 1.5; Faisal et al. 2008). Neural populations exhibit noise correlations: correlated variability that is seemingly unrelated to sensory input (Abbott and Dayan, 1999; Averbeck et al., 2006; Averbeck and Lee, 2006; Maynard et al., 1999). Correlated variability in motor cortex that is unrelated to motor output can be thought of as a dual phenomenon to noise correlations in sensory systems. Variability in neural systems can arise spontaneously as a feature of the highly recurrent nature of networks in the brain (Arieli et al., 1996). For example, Litwin-Kumar and Doiron (2012) demonstrated that neuronal spiking networks can exhibit two timescales of variability: slow dynamics, and fast spiking irregularity (see also Churchland and Abbott (2012)). In this work, the slow components of variability related to persistent state attractor dynamics, and reflected shared variability across the population. The fast components of variability arose from the stochastic nature of spiking and intrinsic and ensemble history effects.

1.2.2 Implications of variability and redundancy for decoding

The fact that motor cortex exhibits variability apparently unrelated to motor output has implications for our ability to build stable decoding models based on random, sparse subsamples of the neural population. In the case where motor-cortical representations are distributed, i.e., the number of units demonstrating movement modulation is much larger than the total degrees of freedom associated with the task, one can in theory exactly decode movements based on a sample of only a few neurons. As long as each neuron encodes

a more or less random combination of motor degrees of freedom, a sufficiently large random sample will form a complete basis for encoding motor output. However, decoding is more challenging if there are a large number of unobserved degrees of freedom also represented in neural firing rates (e.g. context, planning, attention). A dramatic example of this can be seen in hippocampal place cells. For a given environment, place cells in the hippocampus fire maximally when an animal is in a location corresponding to the place cell's receptive field. One can accurately decode position based on place cell firing rates (Brown et al., 1998; Jensen and Lisman, 2000). However, place cells re-map when an animal enters a new environment, or when aspects of an environment or task goal are changed (Moser et al., 2008), making across-condition decoding from place cells challenging or impossible. Slowly varying contextual effects encoded in motor cortex firing rates may appear as slow firing rate nonstationarities, and impair generalization of decoding models over time and across conditions.

1.2.3 The local field potential reflects collective dynamics

The local field potential (LFP) is an important signal for understanding information processing in motor cortex. In order to understand the computational processes of motor cortex, it is useful to examine the relationship between inputs and outputs of motor cortex. Microelectrode array recording approaches record both spiking output and LFPs. LFPs are believed to contain a signature of average synaptic activation (Buzsáki et al., 2012) and therefore reflect information about the inputs to a neural population. Determining the relationship between spikes and the LFPs observed in motor cortex remains an open problem, and the relationship between average electrical potentials and single-neuron activity has been an open problem since the early days of electroencephalography. For example, in 1934 Adrian and Matthews wrote, in relation to EEG recordings:

It has been shown that the potential waves are generated in the cortex... there is no doubt that they are due to the activity of the cortical neurones. Record-

ing them would seem to offer the most direct method of investigating cortical activity, but for the difficulty that they are certainly summated effects compounded out of the potential changes in many neurones. It is most unlikely that the change in each neuron is an exact copy, on a smaller scale, of the massed effect. (Adrian and Matthews, 1934)

LFPs are believed to reflect in part synaptic input to the local neural population (Buzsáki et al., 2012), but also contain significant contributions from local spiking activity (Bedard and Destexhe, 2013; Denker et al., 2011; Destexhe and Bedard, 2012; Reimann et al., 2013; Waldert et al., 2013). Therefore, LFPs contain signatures of both the local population activity and neural signals arriving from other brain structures. In MEA recordings, the population of isolated single spiking units reflects the output of the computation of motor cortex. The LFP, containing in it signatures of synaptic inputs from local and remote sources, reflects in part the input to this computation. Observing the relationship between these signals may shed light on the information transformations occurring in motor cortex.

1.2.4 Decoding from local field potentials

Perge et al. (2014) found that LFP ensembles recorded on MEAs can provide comparable decoding accuracy to unsorted multi-unit spiking activity for a cursor-based brain-machine interface, and Bansal et al. (2011) found that reaching and grasping movements could be decoded from motor cortex LFPs. Bansal et al. (2012) demonstrated that motor cortex LFP has, overall, less movement-related information than the spiking population. In regards to relating LFP and spiking signals, Rasch et al. (2008) successfully inferred low-frequency structure in spike trains from low-frequency LFPs and power modulation in gamma-frequency (γ -) LFP, but noted that fast-timescale spiking variability is poorly predicted by LFP. Rasch et al. (2009) also showed that LFP can be inferred from spiking, confirming that spiking and LFP reflect common sources of shared variability to some extent. Although Bansal et al. (2012) and Perge et al. (2014) showed that motor cortex LFP

contains information about kinematics, they do not explicitly address whether kinematics explain all of the shared variability between spikes and LFP.

1.2.5 The LFP can reflect intrinsic sources of variability

Neural variability arising from intrinsic collective dynamics can interfere with decoding. This problem can be mitigated by finding ways to sample variability along “noise” dimensions and explicitly modeling effects of intrinsic sources of variability on the neural population. Kelly et al. (2010) provide a good example of this in estimating the tuning functions of neurons in visual cortex. In the anesthetized state, the neural population responses to natural videos exhibit trial-to-trial variability due to slow population rate fluctuations (Figure 1.6A). Kelly et al. (2010) found that slow fluctuations in the local field potential correlate with the intrinsic fluctuations in population firing rate (Figure 1.6B). By using the LFP as an indicator of collective dynamics, and incorporating it as a co-regressor in their tuning models, they were able to identify sharper (i.e. more informative) models of neural tuning to the natural video stimulus (Figure 1.6C). Banerjee et al. (2012) have also demonstrated that ongoing LFP activity in area LIP can be used to explain variability in neural responses to sensory input, and Haslinger et al. (2006) demonstrated that the phase of slow oscillations affects the response of barrel cortex to sensory input.

1.2.6 A note on the output-null space

This discussion of neural variability is related to another concept that has been promoted in recent publications: the so-called “output-null” subspace of neural population dynamics in motor cortex (Churchland et al., 2010b; Shenoy et al., 2011, 2013). If the dimensionality of neural activity in motor cortex is larger than that required for movement, many neural states will map to the same motor output. In this case, dimensions of neural variability in motor cortex can be factored into “output-potent” and “output-null” subspaces, with the output-null (i.e. movement-null) subspace representing degrees of free-

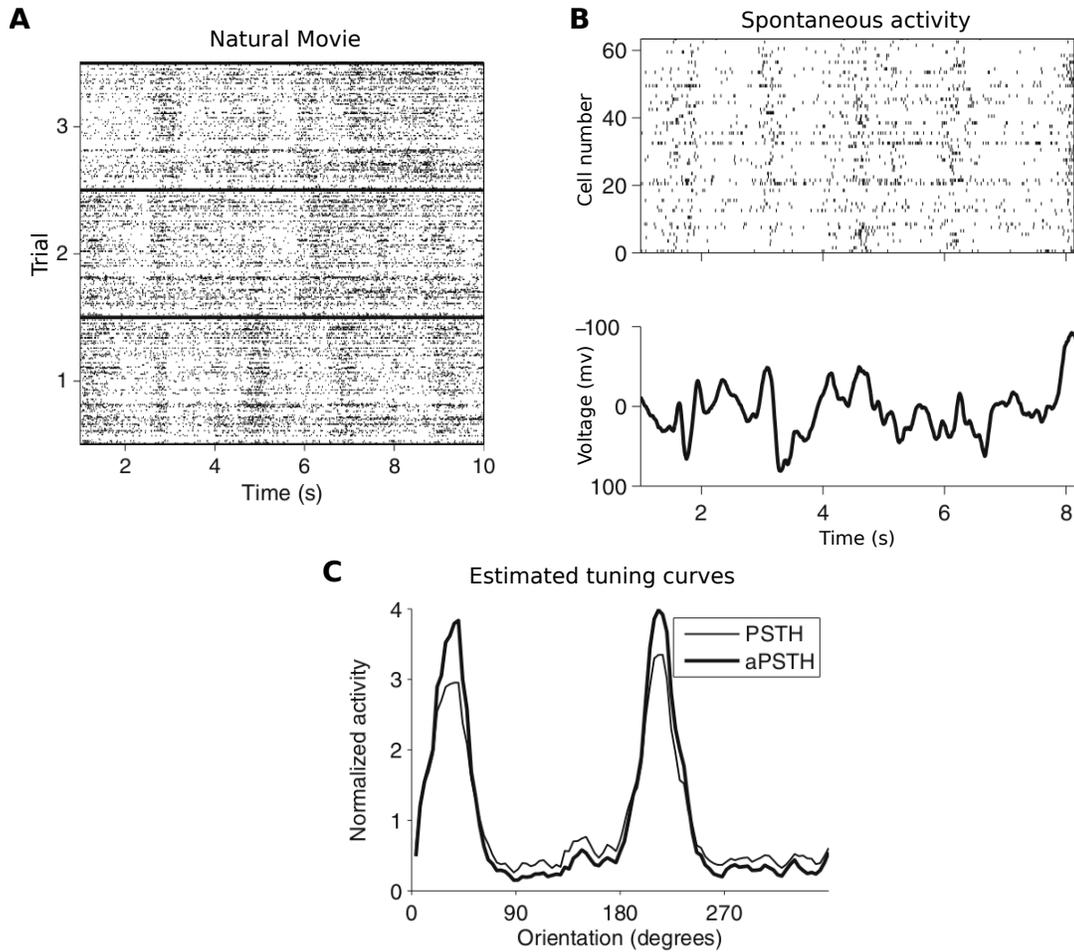


Figure 1.6: *Local field potentials explain neural variability.* In primary visual cortex of macaques under anesthesia, neural firing exhibits considerable trial-to-trial variability arising from intrinsic collective dynamics. **(A)** Spike rasters of population responses of neurons in primary visual cortex to three repeated presentations of the same natural movie. Slow fluctuations in the population rate arise from intrinsic neural dynamics, and are mixed with the visually evoked activity. **(B)** The low-frequency local field potential (bottom) correlates with fluctuations in the population firing rate (top). Thus, the low-frequency local field potential can be used as a proxy for the intrinsic dynamics that give rise to neural variability unrelated to visual input. **(C)** By accounting for the intrinsic variability reflected in the LFP and population firing (thick curve, adjusted PSTH), it is possible to estimate more precise orientation tuning properties of single neurons in visual cortex. (Selected figures from Kelly et al. 2010)

dom in cortical activity that have no effect on muscle output (Kaufman et al., 2014). The output-null space proposed in Kaufman et al. (2014) refers to dimensions of neural activity that are unrelated to motor output, but still serve an important cognitive function. This is to be distinguished from dimensions of neural variability that are simply “noise”. The output-null space was proposed to explain how motor cortex might prepare, observe, and imagine movements without generating the corresponding motor output. If motor output were simply gated by inhibition, one would expect inhibitory interneurons to fire more during preparatory states. However, putative FS inhibitory interneurons in motor cortex do not show such elevated firing (Kaufman et al., 2010). It is important to note that Vigneswaran et al. (2011) found that inhibitory FS cells cannot be unambiguously identified from extracellular recordings in motor cortex, and so the results in Kaufman et al. (2010) may be confounded by fast-spiking PTNs. Despite this, the output-null space is an important concept for reasoning about movement-related and movement-independent neural variability in motor cortex.

There are two related senses of output-null activity which must be clarified. On one hand, motor cortex is able to prepare and simulate movements without giving rise to motor output, and this activity occurs in an output-null subspace. On the other hand, if motor cortex is a pattern generator, it must encode dynamical states, and these additional degrees of freedom also constitute an output-null space. For example, if motor cortex implements an ensemble of oscillatory pattern generators, the “real-valued” component of which drives motor output, and the “imaginary-valued” component of which relates to the dynamics of pattern generation, we would expect the output-null subspace of neural variability to contain the “imaginary” components (Churchland and Cunningham, 2015; Shenoy et al., 2013). Preparatory and “dynamical” null-spaces can be reconciled if movement preparation is accomplished by setting the initial conditions in a “dynamical” null space (Churchland et al., 2010b). The extent of output-null variability in motor cortex during movement execution remains an open question. Output-null dimensions during

movement execution are expected to contain at minimum the correlates of dynamical state for pattern generation, and could reflect additional cognitive variables. Both true noise and variability in movement-null subspaces would appear as “noise” to a decoder.

1.3 β -LFP oscillations in steady-state motor control

Two modes of motor control may be distinguished: a kinetic mode associated with movement and large fluctuations in firing rates, and a steady-state mode associated with maintenance of steady posture and isometric force, which is different from the relaxed state of quiet wakefulness. This steady-state mode is associated with a prominent increase in beta (β -) LFP power, which ranges from 15-30 Hz but is more typically observed around 20 Hz (Baker, 2007; Chen et al., 2007; Jenkinson and Brown, 2011; ?). This section briefly covers the history of research into beta oscillations, including issues that remain to be clarified. It covers potential roles for beta oscillations in motor control, as well as proposed cognitive roles, and potential computational roles for oscillations in the brain. Finally, it covers what is known about the origins of beta oscillations and the relationship between single-unit activity and β -LFP.

More thorough discussions of motor system β -LFP are available in several recent reviews. Salenius and Hari (2003) and van Wijk et al. (2012) provide reviews of oscillations and synchrony in the motor system, including beta oscillations. Baker (2007) focuses on the relationship of beta oscillations to the peripheral motor system. Jenkinson and Brown (2011) review previous literature and hypothesize that beta suppression is mediated by dopamine signaling, and is a correlate of the likelihood that a new voluntary action will occur. Engel and Fries (2010) suggest that beta oscillations “signal the status quo” in the motor system. Beta oscillations have received considerable attention for their role in Parkinson’s disease (Brittain and Brown, 2014; Brown, 2007; Eusebio and Brown, 2009; Hammond et al., 2007; Stein and Bar-Gad, 2013; Weinberger et al., 2009). Kilavik et al. (2013) provides a thorough review of modulation of beta oscillations during visually cued tasks with instructed delays, and is especially relevant to Chapters 3 and 4 of this thesis. Khanna and Carmena (2015) review beta oscillations and propose new experimental directions.

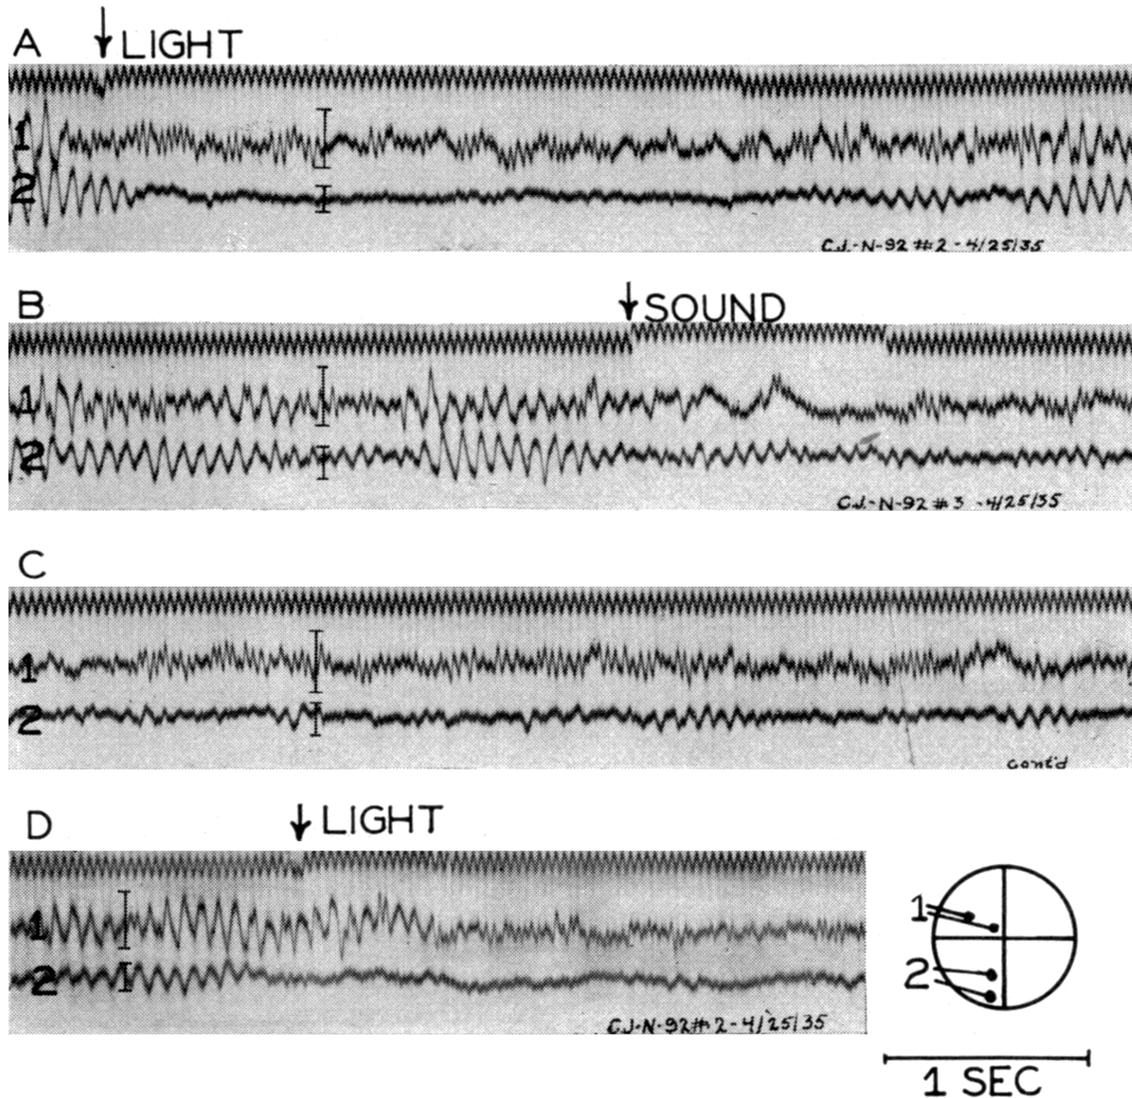


Figure 1.7: Figure 3 from Jasper and Andrews (1938) "Electro-encephalography: III. Normal differentiation of occipital and precentral regions in man". Beta oscillations are identified as a characteristic oscillation recorded on EEG over sensorimotor cortex (trace 1), to be distinguished from lower-frequency alpha oscillations seen over occipital cortex (trace 2). These are the first recordings showing beta oscillations in sensorimotor cortex in humans. (A) It was observed that, unlike alpha oscillations, sensorimotor beta oscillations are not strongly modulated by light. (B,C) Frontal cortex beta could, however, be suppressed by a sound stimulus. (D) Sensorimotor cortex also exhibited alpha oscillations, which modulated independently from the occipital alpha oscillations.

History of research into beta oscillations β -LFP oscillations were first described in the original electroencephalography (EEG) experiments of Hans Berger (Berger, 1929), and subsequently studied in greater detail by Herbert Jasper and colleagues. It was noted by Rheinberger and Jasper (1937) that β -frequency oscillations in EEG recordings from awake behaving cats were “correlated with increased restlessness or alertness, preparation for movement, or movement itself.” Subsequently, Jasper and Andrews (1938) demonstrated that 25-50 Hz oscillations could be identified over human sensorimotor cortex, and were dissociated from 10 Hz alpha oscillations (Figure 1.7). Jasper and Penfield (1949) found that sensorimotor cortex beta was suppressed during voluntary movement, but was elevated during sustained contraction. Thus, it has been known since the early 20th century that the beta rhythm associated with sensorimotor cortex is absent during quiet wakefulness, is elevated during attentiveness, postural maintenance, isometric force, and movement preparation, and is suppressed surrounding the initiation of voluntary movements. Despite this long history of research into β -LFP, the origin and function of beta oscillations in motor cortex has not been fully clarified.

1.3.1 The functions of beta in motor control

The fact that β -LFP power is enhanced during postural maintenance and isometric force may lead one to presume that beta serves primarily to inhibit movement. Modulations in beta power outside of isometric force or steady posture would then be explained by variations in the extent to which planned movements would need to be inhibited. The observation that beta oscillations are enhanced not only during isometric force and postural maintenance, but also during attentive waiting, has led to the hypothesis that beta is a correlate of the maintenance of motor steady-states, representing an efficient way to drive downstream targets, enhancing proprioception, or facilitating sensorimotor computations required for corrective movements. Studies of Parkinson’s disease, which displays abnormalities in beta synchronization, have also shed some light on the role of beta os-

cillations in steady-state motor control.

Recruitment of downstream units Baker et al. (1999) suggested that one function of beta synchrony in motor cortex is the efficient recruitment of muscle activity. Spikes synchronized to arrive within a short time window can act synergistically to recruit more muscle activity at lower firing rates. Supporting this theory is evidence that beta oscillations are coherent between motor cortex and downstream targets (Baker et al., 1999; Kristeva et al., 2007; Riddle and Baker, 2006; van Wijk et al., 2009; Witham et al., 2011; Witte et al., 2007), and that even weakly synchronized units sum up to yield measurable coherence between cortex and the periphery (Baker et al., 2003). Baker (2007) posit that the reason beta oscillations are observed only during motor steady-states is that the oscillations impair neural processing associated with movement, and thus are only feasible when the motor system is idle.

Improved proprioception and somatosensation The observation that transient beta oscillations are seen during difficult manual tasks requiring somatosensory feedback (e.g. Murthy and Fetz 1996a), and that the power of beta oscillations appears to be proportional to lever compliance in isometric force tasks (Kilner et al., 2000), raises the possibility that beta oscillations play a role in facilitating somatosensation. Schaefer et al. (2006) found that membrane potential oscillations can have an important role in reducing spike-train variability by improving the precision of action potentials. Oscillatory drive creates spike-timing precision by reducing stochastic history effects that give rise to spiking variability. Beta oscillations may render the effects of cortical output on muscles predictable, such that changes in muscle stretch due to motor output can be distinguished from external perturbations (Baker, 2007; Khanna and Carmena, 2015).

Facilitating feedback control The motor system is not, strictly speaking, "idle" during motor steady-states. These positions are activity stabilized against gravity and

other outside forces by continually applying corrective movements. The mapping between sensory signals and the appropriate corrective movements varies with the position stabilized and the forces at work, and therefore requires rapid context-dependent sensorimotor transformations. Beta synchronization in the motor system might support stabilization of motor positions through rapid negative feedback mechanisms. Androurlidakis et al. (2006) find that visually guided corrective movements are potentiated during beta oscillations, suggesting that beta oscillations have a facilitatory role not just in somatosensory-motor stabilization, but in generating corrective movements based on sensory feedback in general. Androurlidakis et al. (2007) find that beta power increases when anticipating a postural challenge and that this increase improves performance, indicating that beta oscillations can be an internally gated mode of motor cortex and that they facilitate processing needed to maintain motor steady-states against perturbation. Solving how (or whether) beta oscillations mediate corrective feedback computations would be a major breakthrough in our understanding of motor control.

Parkinson's disease: a case of over-stabilization? Related to the theory that beta oscillations stabilize corticospinal output is the conjecture that beta oscillations have a functional role in inhibiting movement. For example, beta oscillations could impair information transmission through the basal ganglia (Cruz, 2012; Hammond et al., 2007) and therefore prevent the dynamics that would normally cause the system to evolve along the dynamical trajectories associated with movement. In Parkinson's disease, hyper-synchronous beta is associated with rigidity and bradykinesia (Chen et al., 2007; Kühn et al., 2009; Weinberger et al., 2009), and the disruption of hyper-synchronous beta oscillations is thought to mediate the therapeutic effects of deep brain stimulation (Brittain and Brown, 2014; Mallet et al., 2008; McIntyre and Hahn, 2010; Mehanna and Lai, 2013). Beta oscillations have a causal role in bradykinesia, as boosting beta power over motor cortex via transcranial magnetic stimulation (TMS) leads to movement slowing (Pogosyan et al., 2009). Although movement is impaired in Parkinson's, reflexive responses of motor cortex

to muscle stretch are preserved, and in fact exhibit a reduced latency compared to controls (Pasquereau and Turner, 2013). Consistent with this, Gilbertson et al. (2005) found that beta oscillations in healthy subjects correlate with slowed movement and potentiated long-latency stretch reflexes which stabilize motor steady-states. Leventhal et al. (2012) conjecture that beta oscillations stabilize the neural representation of planned motor actions, and that motor rigidity in Parkinson's is an instance of 'over-stabilization'.

1.3.2 Cognitive functions of beta oscillations

Maintaining motor steady-states may be a special case of more general neural mechanisms for stabilizing neural states in the brain. This conjecture is motivated by a collection of observations. First, the variation in motor cortex beta power is more complex than what can be explained by motor-control demands alone (Kilavik et al., 2013), and appears to depend on cognitive variables. Second, beta oscillations are enhanced and task-modulated during instructed delay periods in visually cued tasks. Third, beta oscillations can be observed in brain areas not directly involved in motor output, where they appear to be a correlate of attention or working memory. Although there is a general consensus that beta has something to do with "maintaining the status quo" (Engel and Fries, 2010), the relationship between motor-control and cognitive functions of beta needs to be clarified. With the exception of Jenkinson and Brown (2011), current models remain somewhat abstract.

Sources of variation in beta power A variety of cognitive processes appear to influence beta power. Observation and imagination of movements evoke similar firing rate modulations as movement in single units in motor cortex ('mirror neurons', Rizzolatti and Craighero 2004). Similarly, beta oscillations are suppressed both during movement and imagined movement (Chatrian et al., 1959; Kühn et al., 2006; Pfurtscheller et al., 2005). van Ede et al. (2010) found that elevated beta power in anticipation of a tactile stimulus was greater when attending to the stimulus, implying that beta power is modulated by

attention. Beta oscillations are also enhanced in anticipation of visual cues (Saleh et al., 2010), and van Wijk et al. (2009) found that beta desynchronization accompanies selection of a behavioral response. Fujioka et al. (2012) and Fujioka et al. (2015) showed beta reflects the timing of auditory beats. None of these conditions are directly associated with changes in motor output, suggesting that stabilizing motor steady-states is not the only function of beta.

Beta oscillations in cued motor tasks with instructed delays Chapters 3 and 4 of this thesis explore beta oscillations during a cued reaching and grasping task with instructed delays, and several studies have examined how beta oscillations modulate in similar tasks (e.g. Alegre et al. 2004, 2006; Doyle et al. 2005; Haegens et al. 2011; Kilavik et al. 2012; Saleh et al. 2010; Sanes and Donoghue 1993; van Wijk et al. 2009). Kilavik et al. (2013) report a number of common trends across studies. They note that beta oscillations are enhanced in anticipation of cues, even before an explicit motor plan forms (Kilavik et al., 2013; Saleh et al., 2010). Sensory cues are often followed by a period of beta suppression, putatively corresponding to cue processing. Kilavik et al. (2013) note that β -rebound commonly occurs following cue-related beta suppression, and that beta power is often sustained across delay periods. The frequency of beta oscillations seen during cue-anticipation (before movement planning) is lower than that seen during instructed delays immediately preceding the go cue (Kilavik et al., 2012). Beta oscillations may play slightly different roles in the pre-cued vs post-cued delay periods, with the latter reflecting stabilization of preparatory state. Overall, Kilavik et al. (2013) propose that beta reflects “anticipatory up-regulation of sensorimotor processing beyond somatosensation.” It is not clear whether the beta states seen during motor preparation, which may subserve a working memory-like role, are functionally similar to those seen during isometric force: the need for motor cortex to sustain a cued preparatory state may result in very different beta dynamics than the maintenance of steady isometric force or posture.

Beta oscillations outside primary motor cortex Ohara et al. (2000) found that beta suppression occurs also in the supplementary motor cortex during self-initiated movements, preceding motor cortex beta suppression, and indicating that more cognitive areas of frontal cortex can also exhibit beta dynamics reminiscent of primary motor cortex. Beta oscillations are observed in prefrontal cortex (Van Aerde et al., 2009), where they are correlated with rule (context) selection (Jensen and Bonnefond, 2013) and inhibition (Hwang et al., 2014). Beta power and synchrony in association areas of frontal cortex are correlated with attention (Sacchet et al., 2015). Abnormally low beta power relative to theta LFP power is present in subtypes of attention deficit hyperactivity disorder (ADHD), in which attentional control is impaired (Dockstader et al., 2008). In the visual system, beta has been observed to behave like a “carrier of visual attention”, reflecting top-down influences of frontal areas on sensory cortices (Wróbel, 2000). Theoretical work by (Kopell et al., 2011) suggests that β_1 (~15 Hz) oscillatory dynamics might be critical to coordinating cell assemblies during working memory. These results point to the possibility that the beta is a correlate of behavioral or attentional stabilization. It remains unclear whether the beta oscillations observed in other frontal areas share common mechanisms with beta observed in motor cortex, and how proposed roles for beta in attention and working memory can be reconciled with the sensorimotor functions of beta oscillations.

1.3.3 Computational roles and mechanisms of oscillations

Several roles for beta oscillations in motor control and cognition have been proposed. Some of these, such as recruitment of downstream areas via improved synchrony, rely on seemingly straightforward mechanisms. However, several more sophisticated functions for oscillations in neocortex have also been suggested, including phase coding and gating communication between cell assemblies. Oscillations also have implications for the capacity of neural populations to encode and transmit information, and beta oscillations may also be part of dynamics that support stabilization and long-timescale activity in the

brain. It is not known which, if any, of these theories might apply to motor cortex beta oscillations. This subsection reviews computational and encoding principles related to oscillations that may be important for unraveling both the motor and cognitive roles of beta oscillations.

Phase coding

Neurons can encode information in the timing of spikes relative to a common background oscillation. This phase coding has been proposed as a mechanism for conveying information in the brain (Fries et al., 2007; Kayser et al., 2009; MacLeod et al., 1998; Montemurro et al., 2008). Rate codes are inherently limited because neural firing must be averaged over time or over a population in order to extract an estimate of the instantaneous firing rate (Quiroga and Panzeri, 2009). Mazurek and Shadlen (2002) show that correlations within a population of neurons can interfere with rate coding rapidly varying signals. Phase coding might circumvent these limitations because signal intensity can be encoded in the phase of firing relative to an ongoing shared oscillation (Buzsáki and Draguhn, 2004).

Evidence for phase coding has been found in a number of brain systems. Place cells of the hippocampus (O'Keefe and Recce, 1993; Skaggs et al., 1995) and grid cells in the entorhinal cortex (Hafting et al., 2008) of rats encode aspects of an animal's location in terms of the phase of firing relative to ongoing theta (6-10 Hz) oscillations. A relevant phenomenon for Chapter 3 is the phase drift between place cells' theta membrane potentials and the population theta LFP (Figure 1.8). Additionally, there is evidence that information is encoded in the phase at which neurons fire relative to stimuli-entrained low-frequency (delta) oscillations in both visual (Montemurro et al., 2008) and auditory (Kayser et al., 2009) cortex. In visual cortex, the sequence in which neurons fire relative to 20-60 Hz gamma oscillations also encodes stimulus information (Havenith et al., 2011). It is not known whether single units exhibit phase coding relative to ongoing LFPs in motor cortex, although the phase of ongoing beta oscillations can affect the timing of ar-

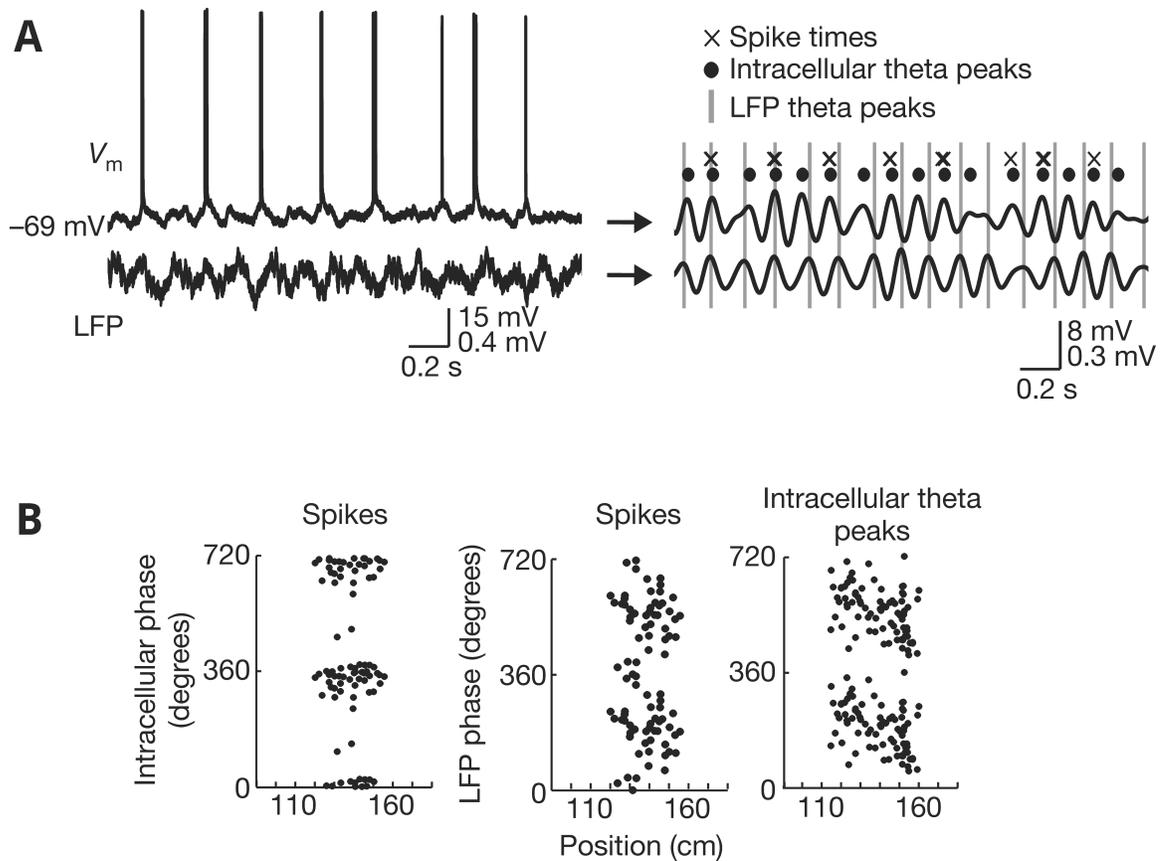


Figure 1.8: *Phase coding in the hippocampus, adapted from Harvey et al. (2009).* Place cells in the hippocampus fire periodic bursts at theta frequency (6-10 Hz). The phase of place cell firing relative to ongoing theta LFP oscillation varies and encodes information about the animal's position relative to the cell's place field. **(A)** Intracellularly recorded potentials (top trace, left) reveal that place cells fire bursts of action potentials phase-locked to the peak of an intracellular theta oscillation. This intracellular oscillation can occur at a different frequency from the LFP theta oscillation (bottom trace, left). When extracellular spike times are viewed relative to the LFP oscillation, their phase of firing appears to drift. **(B)** Hippocampal place cell firing is tightly locked to the phase of intracellular theta oscillations (left). Relatively little stable phase locking is observed between spiking and the LFP theta oscillation, as the preferred phase of firing varies in a way that depends on the animal's position (center). Phase precession of place field firing relative to LFP theta oscillations is mediated by shifts in the phase of intracellular theta oscillations (right). This indicates that ensembles of neurons responsible for the input to place cells can shift their preferred firing frequencies relative to the ongoing theta oscillations, providing an example in which the activity of subpopulations can be dissociated from the collective activity recorded in the local field potential.

rival of sensory information in motor areas (Reimer and Hatsopoulos, 2010), and Canolty et al. (2012) find that preferred phase of firing of single neurons relative to β -LFP can shift across task conditions.

Whether spike-LFP phase relationships are instrumental in neural computation, or merely an epiphenomenon, remains unclear in some cases. For example, Ray and Maunsell (2010) note that the frequency of gamma oscillations in visual cortex depends on stimulus contrast, and differences in contrast across the visual field give rise to gamma oscillations at different frequencies, limiting their utility in communication. Gamma phase-of-firing coding may be an epiphenomenon of gamma-timescale neural processing. Nevertheless Besserve et al. (2015) found that gamma-band phase shifts organized themselves along the direction of local, lateral spiking information transfer in V1, indicating that the relationship between spiking activity and LFP phase may be an important signal for understanding neural information processing, even if it is epiphenomenal. It is therefore useful to explore spike-LFP interactions in motor cortex, both for the potential of extracting additional information to aid in decoding, and from the standpoint of understanding the relationship between spiking and local field potentials.

Gating communication between neural assemblies

Another potential role for oscillations in the nervous system is to bind groups of neurons together into functional assemblies and control communication between such assemblies (Akam and Kullmann, 2014; Maris et al., 2016; Singer, 1993; Singer and Gray, 1995; Varela et al., 2001). Population oscillations can create regular time windows in which neurons are more or less sensitive to external inputs. For example, Dugué et al. (2011) found that the likelihood of perceiving visual phosphenes induced by transcranial magnetic stimulation (TMS) varied with the phase of ongoing alpha (~ 10 Hz) EEG oscillations. Neurons may be less sensitive to inputs during phases where neural firing is suppressed, either due to recurrent inhibition or refractory and recovery effects. By this mechanism, Fries (2005)

propose that “only coherently oscillating neuronal groups can interact effectively, because their communication windows for input and for output are open at the same times.” Simultaneous beta oscillations at different frequencies have been observed in different somatotopic regions of motor cortex, suggesting that beta oscillations may involve synchronous sub-assemblies of neurons (Neuper and Pfurtscheller, 2001; Pfurtscheller et al., 2000). Aumann and Prut (2015) propose that cell assemblies in motor cortex oscillate at slightly different beta frequencies that depend on which muscle groups they innervate, and that different characteristic frequencies represent functional cell assemblies associated with different muscle synergies. In addition to frequency diversity, phase diversity may also play a role in binding together cell assemblies (Maris et al., 2016).

Information coding capacity of synchronized networks

If neurons are constrained to fire synchronously and at the same frequency, the capacity for the population to encode information is reduced. Rate coding requires diversity in firing rates, which requires that units either fire at different frequencies (leading to asynchronous population activity), or fire only during some oscillation cycles. This has been related to motor system beta oscillations in two ways. First, Baker (2007) propose that beta synchronization is a metabolically efficient way to drive muscles during motor steady-states, but that this synchrony limits the capacity of the system to plan and execute new movements. Second, in Parkinson’s disease hypersynchrony in beta oscillations has been proposed to limit the amount of information that can propagate through deep brain structures, impairing motor function (Cruz, 2012; Hammond et al., 2007).

Steady-states and long timescales

Maintenance of a static posture or of isometric force requires sustained output from the motor system. If beta oscillations have a causal role in maintaining firing rates at constant levels, they may correspond a general mechanism for sustaining information in cor-

tex. Understanding the origin of the long timescales and persistent information of neural steady-states remains an open problem. For review of these topics, see Brunel (2003) and Vogels et al. (2005), as well as the “persistent activity” issue of *Cerebral Cortex* (Wang, 2003). Network oscillations and neural steady-states are often treated as distinct phenomena in neural networks (Vogels et al., 2005), but limit cycle attractors can combine features of steady-states with oscillations. Neural steady-states are assumed to require some degree of self-excitation to maintain neural firing, but balancing this self-excitation in a way that leads to stable states that can store information can be challenging. In order for a neural system that exhibits steady-states to encode information, it must exhibit multiple stable configurations, and it must not drift from one stable state to another. Beta oscillatory activity correlates with maintenance of not only a continuum of postural configurations, but also dimensions of cognitive and preparatory information. Models of motor steady-state must therefore reflect this structure and also exhibit a continuum of stabilized states, as well as be compatible with oscillatory dynamics. Neural network models that exhibit a continuum of steady states have been demonstrated by fine-tuning the balance of excitation and inhibition to yield manifold attractors (e.g. Brunel (2003); Camperi and Wang (1998); Compte et al. (2000); Goldman et al. (2003); Seung (1996); Vogels et al. (2005)). Bump attractors, a special case of a mutually-inhibitory network topology, can yield a continuum of stable states (Kilpatrick and Ermentrout, 2013; Laing and Chow, 2001; Lim and Goldman, 2014; Wimmer et al., 2014; Wu et al., 2008b), but appear to be very different from the neural dynamics observed during motor steady-states.

Rule et al. (2011) demonstrated a neural field model which exhibits both oscillatory dynamics and a continuum of limit-cycle solutions as a consequence of rotational and translational symmetries in the solution space. Rule (2014) (unpublished, Woods Hole Methods in Computational Neuroscience course project) showed that this model can, with minor modification, also apply to neural networks that display a large number of disjoint limit cycle attractors, and that information can be written into and stabilized in such a

network by controlling the level of extrinsic oscillatory drive. These models may offer a way of reconciling long-timescale stabilization of neural states with oscillatory dynamics. However, Rule et al. (2011) studied neural dynamics in the mean-field case, in which it is assumed that the individual actions of spiking units are asynchronous and “average out” to yield population firing rate fluctuations. If similar dynamics are at play in motor steady-states, we would expect highly synchronous neural activity at a single common frequency. In models like Rule et al. (2011), different oscillatory solutions are distinguished not by variations in the mean rates in units, but in how different subpopulations of units organize their relative phase of oscillations. This is incompatible with observation that firing rates can exhibit some diversity in motor steady-state (Chapter 3) and that planned information can be decoded from firing rates during preparatory steady-states (Vargas-Irwin et al., 2015).

A tantalizing clue may lie in models that achieve neural steady states using negative-derivative feedback control (Lim and Goldman, 2013, 2014). These models stabilize arbitrary network states by monitoring neural firing and applying compensatory feedback when this firing changes. Gating mechanisms that turn on and off the negative feedback can allow the system to be more or less sensitive to input at different times, and activation of feedback stabilizes firing rates at their current level. This model is a neural analogue of the electronic sample-and-hold circuit, which uses negative feedback from operational amplifiers to stabilize an arbitrary voltage. If the negative feedback mechanisms can be generalized to respond not only to deviations in local firing rates, but also to visual and somatosensory feedback carrying information about errors from external perturbations, negative-derivative feedback control may offer a promising model to unify the motor control and the cognitive dimensions of motor cortex steady-states. However, the model in Lim and Goldman (2013) is a firing rate model that does not exhibit oscillations during steady-states, and is therefore inconsistent with the beta oscillations observed in motor steady-states. It seems that the regular β -timescale firing observed during motor steady-

states lends naturally to population oscillations, and it may be that spike timing effects are critical to the mechanisms for stabilizing motor steady-states. It may be that beta oscillations in motor steady states are not limit-cycle attractors in the conventional sense, but rather an emergent phenomenon that results from highly regular spiking activity within a system that is otherwise using firing-rate coding.

Kopell et al. (2011) addresses a mechanism for oscillatory working memory at β_1 (~15 Hz) timescales in a spiking model. This mechanism allows for multiple cell assemblies to store information without competition. Kopell et al. (2011) note that they investigate association cortex beta, and that it is not known whether this mechanism applies to motor steady-states. Clarifying the structure of spiking and network oscillations during motor steady states may help resolve this uncertainty. The mechanism of Kopell et al. (2011) encodes information in the presence or absence of firing relative to a common network oscillation, and further work is needed to verify whether this mechanism can be reconciled with the diverse firing rates observed during motor steady-state beta oscillations seen in Chapter 3.

1.3.4 The Origins of Cortical Beta Oscillations

The origins of beta oscillations in the motor system remain unclear, and the oscillations appear to arise in a number of brain structures simultaneously. Beta oscillations can be synchronized between many brain structures, including motor cortex, parietal cortex, deep brain structures, and the cerebellum (Cassidy et al., 2002; Engel and Fries, 2010; Schnitzler and Gross, 2005). How the level of beta oscillatory power and synchrony is regulated or gated is not known.

Studies of the laminar origin of beta oscillations in motor cortex point to layer V PTNs as being the source of motor cortex β -LFP, with the highest beta power being found between a depth 1 and 2 mm (Murthy and Fetz, 1996a; Witham and Baker, 2007). Consistent with this, the motor cortex pyramidal tract neurons that send their efferents to the spinal

cord can exhibit intrinsic β -rhythmicity, and contain subthreshold membrane currents that support β -periodic firing without requiring collective network reverberation (Chen and Fetz, 2005; Wetmore and Baker, 2004). Kopell et al. (2011) also proposed a model in which beta oscillations arise from intrinsic currents. In slice preparations, Roopun et al. (2006) found an intrinsic β_2 (20-30 Hz) oscillation in layer V, mediated by subthreshold currents and synchronized by gap junctions. Roopun et al. (2008) have proposed that one period of this β_2 oscillation concatenates with the gamma oscillations in superficial layers to generate a slower β_1 (~15 Hz) rhythm. Any of these mechanisms may confer a tendency to fire rhythmically in the beta range. However, the level of β -LFP power and β -rhythmic spiking varies considerably, and it is not the case that intrinsic β -rhythmicity necessarily gives rise to population beta oscillations. Other mechanisms must be at play to control the level of overall beta oscillations at the level of the population.

Beta oscillations appear also in somatosensory cortex. They may be conveyed from the spinal cord (Baker, 2007) or arise locally as a result of ~10 Hz thalamic drive (Jones et al., 2009). The Granger causality analysis of Brovelli et al. (2004) suggests that somatosensory beta oscillations entrain motor cortex β -LFP, although interpretation issues with Granger causality mean that this cannot be taken as a definitive causal connection. Reconciling the observation that motor cortex β -LFP signals appear to be related to somatosensory cortex inputs with the observation that layer V PTNs in motor cortex exhibit their own intrinsic β -rhythmicity remains an open problem in understanding the origins of motor cortex β -LFP. It will be important to clarify the relationship between PTN β -rhythmicity and the β -LFP recorded in the same population in order to understand the significance of these neural signals.

1.3.5 The relationship between β -LFP oscillations and spiking

Understanding the relationship between single-unit spiking and β -LFP oscillations may reveal how synchronized beta activity emerges in motor cortex, and provide clues as to

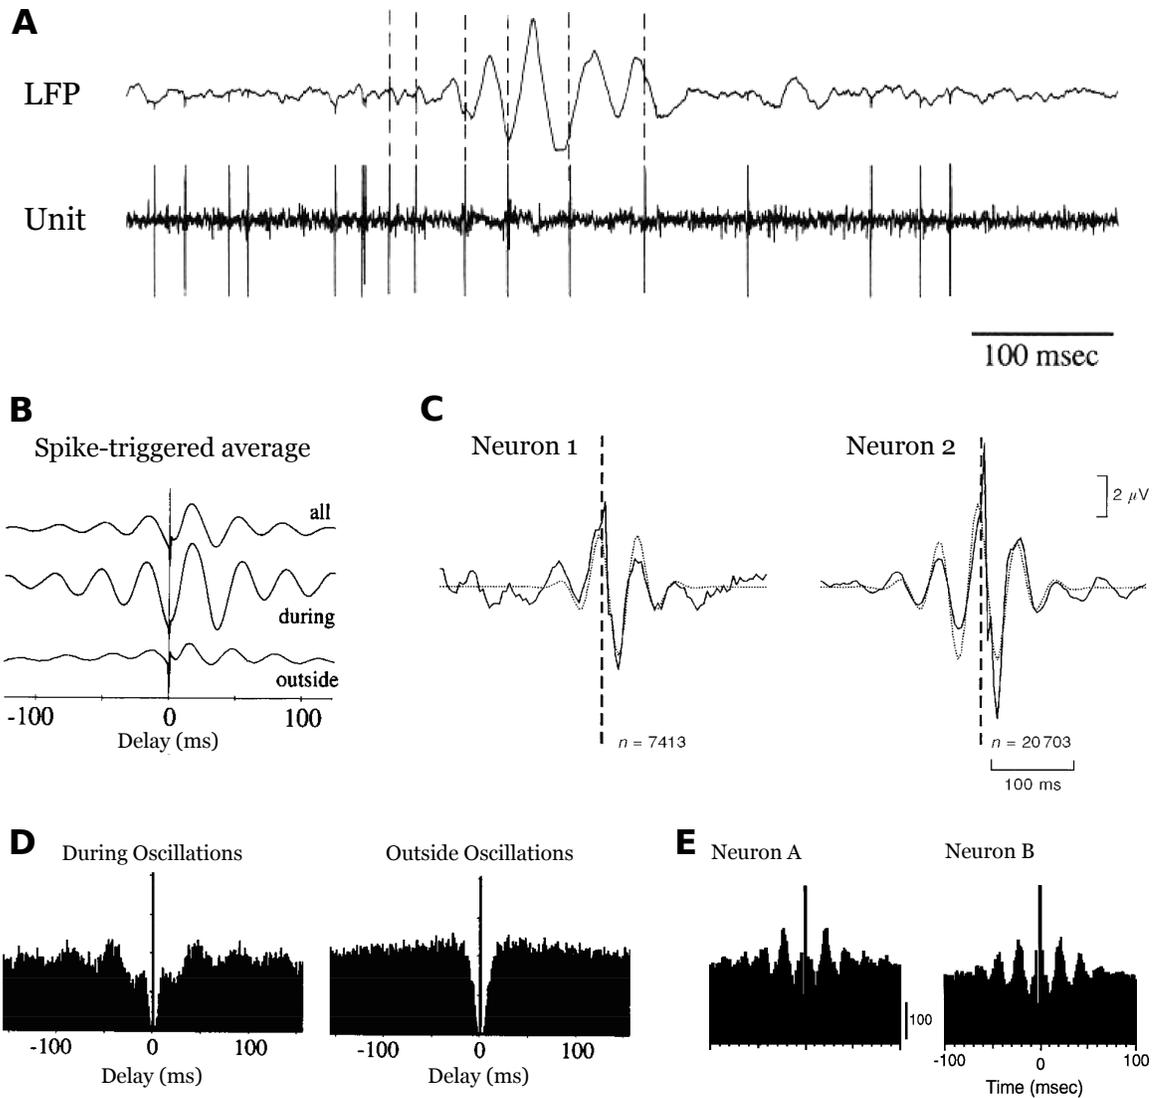


Figure 1.9: Spike-LFP beta phase locking in motor cortex based on spike-triggered averages of LFPs. **(A)** Adapted from Murthy and Fetz (1996a) and Murthy and Fetz (1996b) Figure 1. Transient beta oscillations were observed while a monkey retrieved raisins from a Kluver board, and less frequently during repetitive wrist flexion, or while at rest. **(B)** Adapted from Murthy and Fetz (1996b) Figure 2E. Spike triggered averaging of LFP reveals phase-locking to beta oscillations, both during and outside of periods where beta power is elevated. **(C)** Adapted from Baker et al. (1997) Figure 9B. **(D)** Adapted from Murthy and Fetz (1996b) Figure 2B,C. The spike trains of an example neuron show beta rhythmicity only during high-beta events. **(E)** Adapted from Donoghue et al. (1998), Figure 7. Beta rhythmicity in single-unit spike trains in motor cortex is observed primarily during a premovement delay period, coinciding with elevated β -LFP power.

how high- versus low-beta states are gated. Elevated beta oscillations are associated with reduced PTN firing rates, with units tending to fire at rates around the beta band (Baker, 2007). Although many studies report some amount of phase locking between single units and β -LFP oscillations (Figure 1.9; (Baker et al., 1997; Donoghue et al., 1998; Murthy and Fetz, 1996a,b), there is evidence that β -LFP and single-unit spiking in motor cortex may often be dissociated.

β -LFP oscillations are often observed as variable transients, rather than as sustained ongoing oscillations (Denker et al., 2007; Donoghue et al., 1998; Feingold et al., 2015). The origin of the transient nature of β -LFP has not been fully explained. Beta transients may be a correlate of brief synchronization between motor areas, or may represent bursts of oscillatory activity conveyed between different brain areas. The relationship between beta transient activity and single-unit spiking has yet to be studied in detail, and may point to important clues about the computational mechanisms at work during elevated beta states.

Donoghue et al. (1998) found that β -LFP modulations surrounding movement onset may overlap the movement-related firing rate increases, indicating that β -LFP oscillations may not necessarily be locked to the local neuronal firing. Witham and Baker (2007) found that β -LFP power need not correlate with β -rhythmic spiking. During motor steady-states, PTN firing is only weakly coherent with β -LFP oscillations, but this weak coherence may be sufficient to account for the observed coherence between cortical β -LFP and beta oscillations in the electromyograph (EMG) of effector muscles (Baker et al., 2003). Taken together, it seems that either the beta oscillations observed in LFP arise from the summation of a large number of weakly coupled neurons, which individually are much more variable, or β -LFPs in motor cortex reflect inputs from other brain areas.

1.3.6 Reconciling cognitive and motor roles of beta oscillations

To summarize, beta oscillations are prominent during motor steady-states and suppressed during movement. Beta oscillations facilitate computation of reflexive stabilizing move-

ments, including stretch reflexes, and both visual and somatosensory feedback. This indicates that the high beta state is not idle, and that it is not output firing rates that are stabilized per se but rather a particular postural or isometric force configuration. Beta oscillations appear during steady-state instructed delays, and are briefly suppressed following cues in a manner reminiscent of movement. They may play an active role in working memory and attention. A unified model that explains all functional roles of beta remains elusive, and the complexity of beta oscillations may stem from multiple beta sources serving different functions.

A weak consensus is emerging that beta oscillations are a correlate of stabilized neural states. Brittain et al. (2014) propose that “the manipulation of computational capacity by beta activity may itself present a mechanism of action selection and maintenance”. To date, no unified model has been proposed that could explain how beta could both facilitate feedback stabilization of motor steady-states, while also maintaining cognitive information locally within cortex. Although some theoretical models of working memory use beta oscillations or incorporate the concepts of negative feedback, no work has been done to evaluate whether these are consistent with the observed neural activity.

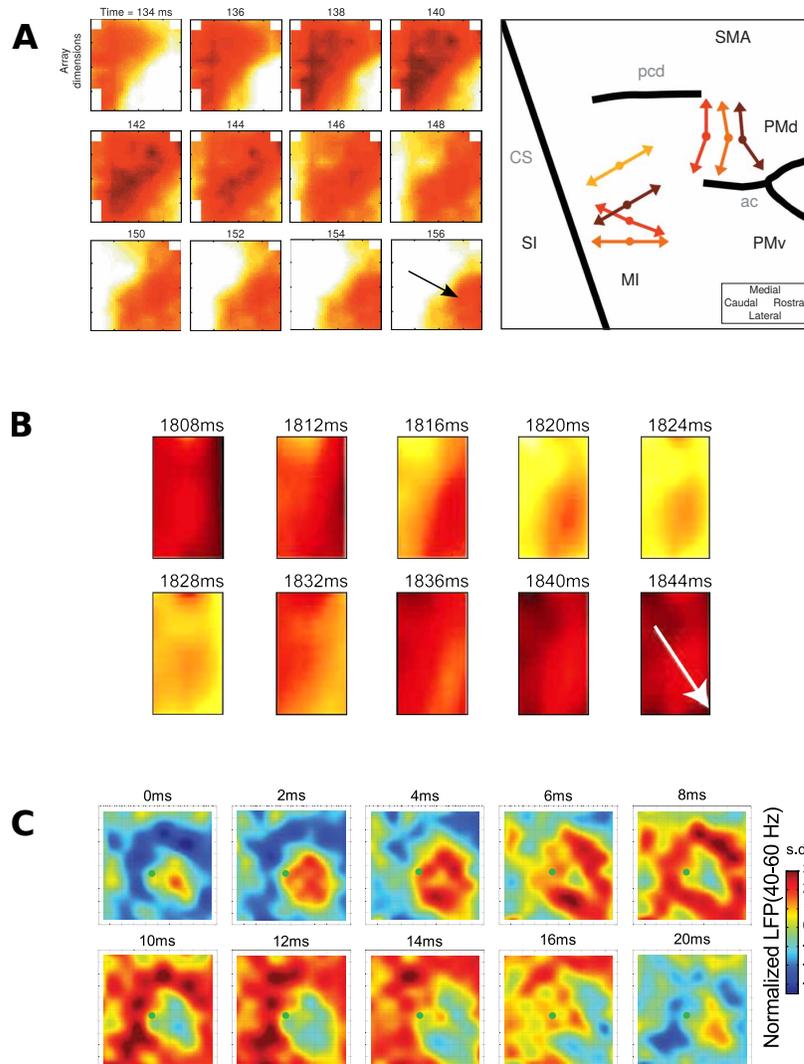


Figure 1.10: Traveling waves in motor cortex. (A) Adapted from Rubino et al. (2006). Traveling waves can be observed in beta (~ 20 Hz) oscillations in motor cortex of rhesus macaques during an instructed delay period following preparation of reaching movements (left: example wave event). The directions in which these waves travel cluster along anatomical axes (right), which may be related to the structure of lateral connections in motor cortex, or be associated with functional connectivity facilitating communication between motor areas. (B) Traveling waves in β -LFP were subsequently confirmed to occur also in human cortex (adapted from Takahashi et al. 2011). (C) Adapted from Lu et al. (2015). Constant pulse optogenetic stimulation in motor cortex induces gamma oscillations that organize as traveling waves. Shown here is an example of a radiating wave emanating from the edge of the stimulated region. Spiral waves were also observed.

1.4 Waves in motor cortex

The recurrent connectivity of the cortex, combined with the fact that cortex is organized locally as a two-dimensional sheet, gives rise to a system that can support a diversity of spatiotemporal wave phenomena. Emergent wave dynamics in the brain may play a role in neural computation and communication. Waves can spatially organize subpopulations of neurons into synchronous groups, and therefore impact phase coding (Lubenov and Siapas, 2009), detection of sensory inputs (Ermentrout and Kleinfeld, 2001; Reimer and Hatsopoulos, 2010) and communication between brain areas (Maris et al., 2016; Prechtl et al., 1997; Rubino et al., 2006). A number of the proposed functional roles for oscillations in the brain depend on phase relationships between oscillations in different areas. Phase coding requires that the downstream target experience oscillations coherent with the source. Oscillations can set up time windows where a neural population is receptive or insensitive to input, and so the relative phase lag between brain areas has an impact on their ability to communicate. Characterizing the spatiotemporal organization of oscillatory activity in motor cortex is therefore important for understanding the functional significance of oscillations across motor cortex. To date, two spatiotemporally organized wave phenomena have been reported in MEA recordings from primate motor cortex (Figure 1.10). In 2006, Rubino et al. reported that beta oscillations in motor cortex are organized as traveling waves. In 2015, Lu et al. showed that γ -LFP can also exhibit spatiotemporal organization in primate motor cortex. Lu et al. (2015) show that motor cortex exhibits waves at gamma (γ) frequency (~ 50 Hz) in response to a constant pulse of optogenetic stimulation during rest, movement preparation, and movement execution. Understanding the spatiotemporal organization of beta oscillations in motor cortex is of potential clinical relevance. Closed-loop cortical stimulation may provide a less invasive alternative to deep brain stimulation for suppressing pathological beta synchronization associated with Parkinson's disease (Beuter et al., 2014). If beta oscillations display spatial diversity in their phase across motor cortex, closed-loop feedback stimulation may benefit

from accounting for this diversity. This section provides background on spatiotemporal activity in motor cortex relevant for the study presented in Chapter 4. It covers historical background with a focus on what analysis of spatiotemporal waves can tell us about the underlying system, and what is known about waves in beta and gamma oscillations in motor cortex.

1.4.1 Origins of cortical waves

The interactions between large populations of neurons gives rise to emergent collective effects not evidenced at the single-unit level. Freeman III (1972) outlined a theory of collective “mass” action of neurons, writing:

“... when neurons strongly interact in sufficiently large numbers (on the order of 10 or more), new collective properties emerge that demand a different kind or level of conceptualization. An analogy equivalent to that given above is the notion that temperature and pressure exist only for a mass, in contrast to the thermal kinetic energy of molecules in the mass.”

This analogy is not to be taken too literally, as the spatiotemporal phenomena arising from collective neural dynamics are far richer and more complex than thermodynamic properties. Spatiotemporally patterned activity in cortex may arise through a number of different mechanisms. Ermentrout and Kleinfeld (2001) enumerate three possibilities for traveling waves in cortex: (a) delayed excitation from a common source, (b) propagation of neural activity in an excitable network, and (c) organization between coupled local oscillations (Figure 1.11). Traveling waves arising from signal propagation were first observed by Adrian (1936), who demonstrated that electrical stimulation in anesthetized cortex propagates away from the stimulation site as a radiating wave traveling approximately 25-35 cm/s. These radially propagating potentials are likely a combination of conduction delays and propagating pulses of excitation. Other mechanisms of spatiotemporal pattern formation are possible. Center surround inhibitory interactions give rise to connectivity that promotes spatial pattern formation (Ermentrout and Cowan, 1979), in

a mechanism analogous to Turing's mechanism for biological pattern formation (Turing, 1952), or the "resonance patterns" discussed in Freeman III (1972). This mechanism may play a role in abnormal states, for example drug and migraine related geometric visual hallucinations (Ermentrout and Cowan, 1979), or the Purkinje lights (Purkyně, 1819)⁴, geometric phosphene hallucinations induced by flickering light (Rule et al., 2011). In a collection of recent papers, Heitmann and colleagues have proposed that center-surround interactions between coupled phase oscillators could mediate beta synchronization and desynchronization in motor cortex (Breakspear et al., 2010; Heitmann et al., 2013; Heitmann and Ermentrout, 2015; Heitmann et al., 2012). In addition to lateral interactions and signal conduction effects, waves may also arise from emergent collective dynamics in the form of waves or wavefronts reflecting changes in neural state (e.g. cortical spreading depression observed by Leão 1944).

A great variety of traveling wave phenomena have been observed in cortex. Prechtl et al. (1997) describe stimulus evoked waves at 10-20 Hz in turtle visual cortex, and find plane waves, spiral waves, and complex waves. Huang et al. (2004) observe wave phenomenon in tangential slides of rat visual cortex spanning layers III-V, and report planar, radiating, and complex waves. In their work, higher LFP oscillation amplitudes are more typically associated with spiral waves. Like Prechtl et al. (1997), they speculate that spatially organized wave dynamics are important for organizing oscillatory activity. Huang et al. (2010) report spiral waves during sleep and drug induced oscillations. Siapas et al. (2005) observe that hippocampal theta oscillations are organized as a traveling wave. Because place cells encode information relative to the phase of theta, a spatial gradient in theta phase implies that, at each instant, a segment of space is encoded (Lubenov and Siapas, 2009). Ermentrout and Kleinfeld (2001) note that traveling waves observed in sensory cortices while "idle" could in effect scan the sensory space for novel stimuli.

⁴Purkyně is the Czech orthography, Purkinje the German. Purkinje is the more common spelling (e.g. "Purkinje cells"), but the original reference can be found under Purkyně.

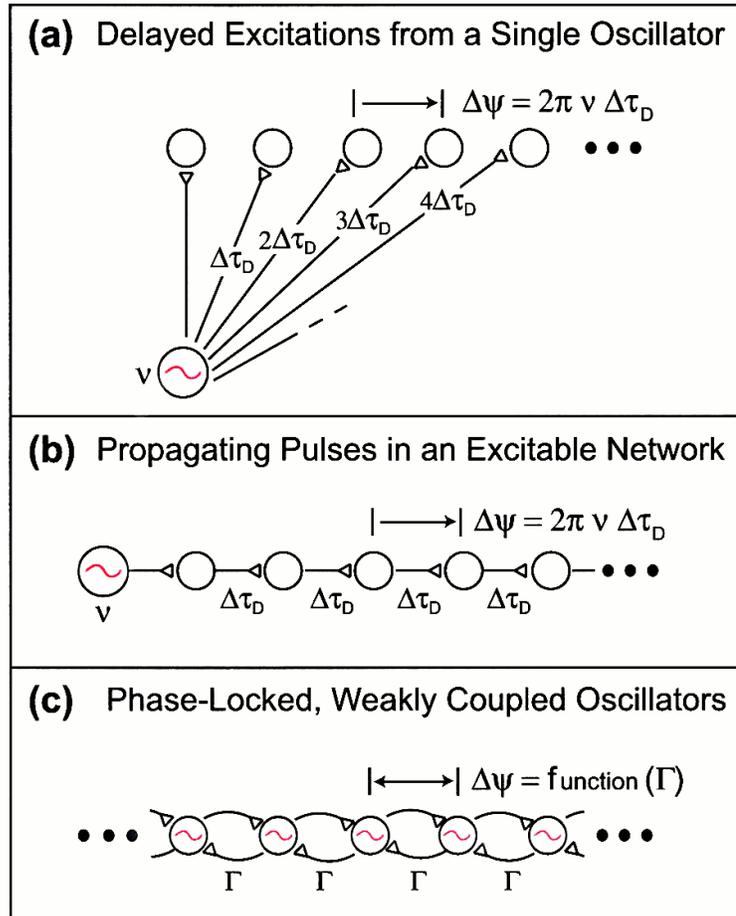


Figure 1.11: Three origins of traveling waves in cortex (Figure 1 from Ermentrout and Kleinfeld 2001.). **(A)** Apparent traveling wave phenomena arise if adjacent regions of cortex are driven by a single source with varying conduction latencies. This mechanism does not require local interactions. **(B)** Recurrent lateral connectivity in cortex can allow excitatory pulses to propagate across cortex. In this diagram, a single feed-forward chain is shown. Propagating excitatory waves are also possible in the case of bi-directional or two-dimensional connectivity. **(C)** Traveling waves arise as organized phase gradients between local, coupled oscillators. In this case, apparent wave motion does not reflect propagating excitation, but rather the spatial organization of oscillator phases. In the case where there is a gradient in frequency of coupled oscillators, oscillator coupling can entrain phase relationships such that traveling waves appear to radiate from high-frequency regions to low-frequency regions.

Rubino et al. (2006) was the first study to report traveling wave phenomena in motor cortex of macaques. They analyze traveling plane waves and find that they occur along preferred anatomical axes in M1 and PMd. Takahashi et al. (2015) find that, when beta power is elevated in motor cortex, there is directed spiking information transfer (Quinn et al., 2011) that aligns with the axis of wave propagation. Both the wave propagation axis and the direction of spiking information transfer align with the putative anatomical axis of longer-range lateral connectivity. It is interesting to note that the emergence of the directed spiking statistical relationship was conditional on the presence of beta activity, suggesting that the collective wave statistics are influencing single-unit spiking. The spatiotemporal wave phenomena may be modulating the effective, or functional, connectivity between units in motor cortex. Takahashi et al. (2011) find traveling plane waves at β frequency (15-30 Hz) in human subjects as well. These waves were found to propagate with velocities ranging from 10 to 35 cm/s. They also suggest that waves may be reflecting off the boundary between primary motor and somatosensory cortices, similar to the phenomenon at the boundary of V1 and V2 analyzed by Goulet and Ermentrout (2011).

1.4.2 Waves as an indicator of collective dynamics

Given that spatiotemporal wave activity arises from the collective dynamics of neural populations, can we infer important information about neural dynamics by observing wave activity? Schiff et al. (2007) note that changes in wave dynamics may reflect changes in underlying "biological order parameters". The velocity, frequency, spatial wavelength, and pattern of organized wave activity depend on underlying neural states, and provide another important summary of the network activity. Several papers have inferred network parameters from spatiotemporal wave dynamics. For example, Goulet and Ermentrout (2011) model the relationship between network parameters and the patterns of reflection and compression of cortical waves observed at the V1/V2 boundary by Xu et al. (2007). González-Ramírez et al. (2015) fit a neural field model to wave propagation preceding

seizure termination. Rubino et al. (2006) extract wavelengths and propagation velocities from traveling waves in motor cortex.

1.4.3 A computational role for beta waves in motor cortex

Heitmann et al. (2012) considered a population of neural oscillators coupled with center-surround or (“Mexican hat”) style lateral coupling (Figure 1.12A). They found that the system could exhibit both traveling waves and zero-lag spatial phase synchrony, depending on the strength of the surround inhibition (Figure 1.12). They also identified a regime exhibiting bistability, and found that transient perturbations could switch the system between synchronized and plane-wave solutions (Figure 1.12B). This is a potential model of transitions between high-beta motor states associated with stabilized positions, and low-beta motor states associated with movement. They noted that modulations in apparent LFP power in motor cortex may be due to changes in the synchronization dynamics. Heitmann et al. (2013) speculated that these emergent oscillatory dynamics themselves might play a role in encoding and generating motor commands. Heitmann and Ermentrout (2015) further analyzed the laterally coupled neural oscillator model and found multiple stability between wave solutions and synchronized solutions, and also observed a third state which they termed “ripple,” consisting of spatially patterned deviations from zero-lag spatial phase synchrony without bifurcation to true traveling wave dynamics (Figure 1.12C,D, fourth example). It is not known whether the dynamical transitions observed in motor cortex correspond to those in this model. For example, it is not known whether motor cortex exhibits a ripple state as a transitional mode between spatial synchrony and traveling waves. The neural oscillator model in Heitmann et al. (2012) and Heitmann and Ermentrout (2015) used center-surround connectivity, which has not been demonstrated to exist anatomically in motor cortex, but functionally equivalent inhibitory interactions between phase oscillators could still arise from network effects.

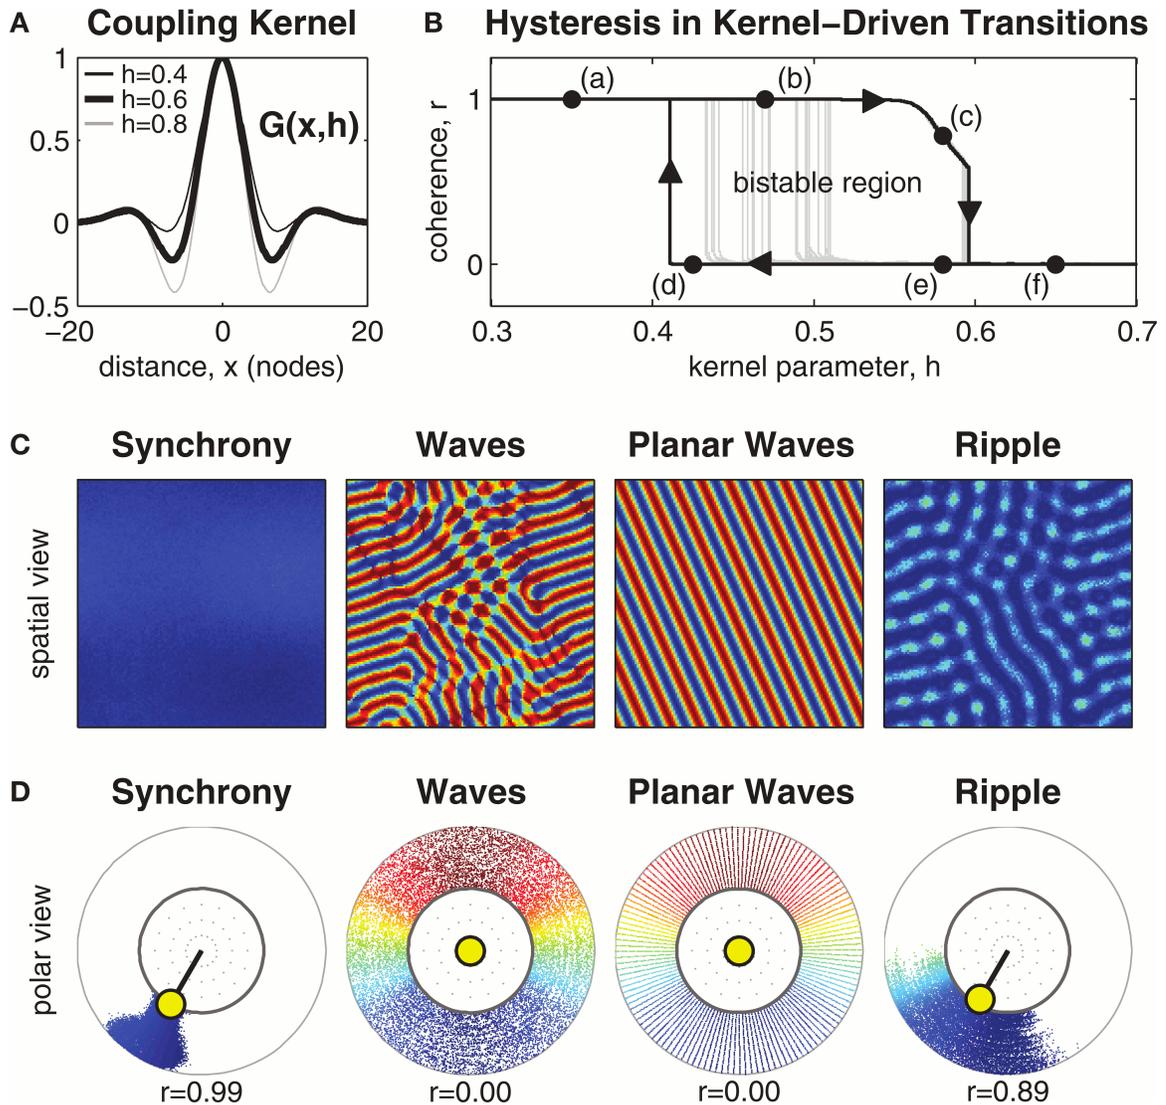


Figure 1.12: A model for spatiotemporal waves in beta oscillations (Figure 1 from Heitmann et al. 2012). Inhibitory surround interactions may give rise to spatiotemporal patterns in β -LFP oscillations. Shown here is a coupled oscillator model for spatial pattern formation in β -LFP. **(A)** Center-surround interactions can give rise to pattern formation via a Turing mechanism. Local excitatory interactions promote local synchronization between adjacent oscillators, while more distant inhibitory surround interactions push adjacent larger patches of cortex out of phase. **(B)** Depending on the strength of inhibition, different spatial patterns can be stabilized. With weak inhibition strengths, oscillations are spatially synchronous (a). With strong inhibition, spatially patterned states are stable (f). At intermediate strengths, the system exhibits multiple stability between synchronized and patterned solutions (b-f). **(C)** Example spatially pattern states for oscillations in motor cortex. In addition to spatially synchronous states and both planar and complex wave patterns, this model also exhibits a state termed ‘ripple’, which is similar to the synchronous state with the exception that subpopulations of oscillators lead or lag the average phase in a way that creates spatial patterns. **(D)** Distributions of the population of oscillator phases for the spatiotemporal states in (C). Spatially synchronized states can be identified by a highly concentrated distribution of phases. Both complex and plane waves exhibit a uniform distribution of phases. Ripple states exhibit an intermediate level of phase concentration.

1.4.4 Optogenetics

Optogenetic methods, which allow for ionic currents to be induced via optical stimulation, have recently been refined and popularized by Deisseroth and colleagues (Boyden et al., 2005). Optogenetic research in primates is an emerging field subject to ongoing investigation (Berdyeva and Reynolds, 2009; Gerits and Vanduffel, 2013; Han, 2012). Initial studies (Diester et al., 2011; Galvan et al., 2012; Han et al., 2009) attained single-unit modulation but no behavioral effects. Subsequently, behavioral effects of optogenetics in the visual system were observed. Dai et al. (2014) observed response biasing by stimulation in area LIP, and saccade generation can be modified by optogenetic stimulation in superior colliculus (Cavanaugh et al., 2012), V1 (Jazayeri et al., 2012), and arcuate sulcus (Gerits et al., 2012). Because sensory cortices are by design sensitive to inputs, these findings in the visual system may depend on the capacity of sensory systems to detect and use slight perturbations for decision making, and behavioral effects of optogenetic stimulation in motor cortex might be very different.

Lu et al. (2015) found that short constant pulse optogenetic stimulation in primary motor cortex elicits traveling gamma waves and firing rate modulations, but that these optogenetically evoked responses could be overridden by the intrinsic dynamics of motor cortex during movement. In the future, optogenetics may enable new experiments that modulate collective neural dynamics in motor cortex, for example increasing or decreasing the amplitude of different narrow-band oscillations, or altering their spatiotemporal synchronization. In the long term, we would like to develop closed-loop stimulation approaches that use optogenetic methods to stabilize neural states, for example, promoting or inhibiting β -frequency synchronization. Such closed-loop control of neural activity in motor cortex may ultimately be important for more effective neuroprosthetics for the treatment of movement disorders. In order to do this, we must build an accurate model of the response of neural tissue in motor cortex to optogenetic perturbations. Lu et al. (2015)

have done initial work analyzing the structure of optogenetic responses in motor cortex, and this thesis further explores optogenetically-evoked wave structure in Chapter 4.

1.5 Statistical methods

This section covers the statistical methods used in this thesis, including methods for analyzing LFP signals, their relationship to neuronal spiking, and their spatiotemporal structure. Figure 1.13A briefly diagrams the recording of kinematic variables during a reach and grasp task. (Chapters 2 - 4 provide detailed information on the behavioral tasks.) We begin with an overview of the signal processing steps used to extract features of LFP oscillations such as power and instantaneous phase. We then briefly review point process generalized linear models for predicting spiking, and cover methods for assessing spike-LFP phase coupling. Finally, we cover methods for analyzing LFP spatiotemporal patterns recorded via MEAs, including methods for identifying and characterizing spatiotemporally organized wave dynamics.

1.5.1 Extraction of LFP oscillations

Figure 1.13B illustrates the neural signal processing used in this thesis. Narrow-band components can be identified as peaks in the LFP power spectrum. Additionally, narrow-band LFP oscillations can modulate independently of the background power spectrum, even when power in these bands does not appear as discrete peaks. In these cases, narrow-band LFP oscillations may be identified by subtracting LFP power spectral densities across two conditions. LFP power spectra are considered a mixture of a power-law distributed background, and a collection of narrow-band oscillations. These two effects can be dissociated using a variety of methods (e.g. Yamamoto and Hughson 1993). In our data, however, theta, alpha, beta, and gamma peaks were at times sufficiently enhanced against background LFP power that they could be identified from local maxima in the power spectral density. Once the frequency band of a LFP oscillation has been identified, it is straightforward to extract the LFP oscillation using band-pass filtering. We elected to use primarily 4th-order Butterworth band-pass filters, applied both forwards and backwards so as to create zero net phase distortion or delay. The 4th order Butterworth filter provides a balance of numerical stability, low distortion and ringing, and adequate roll-off in the stop

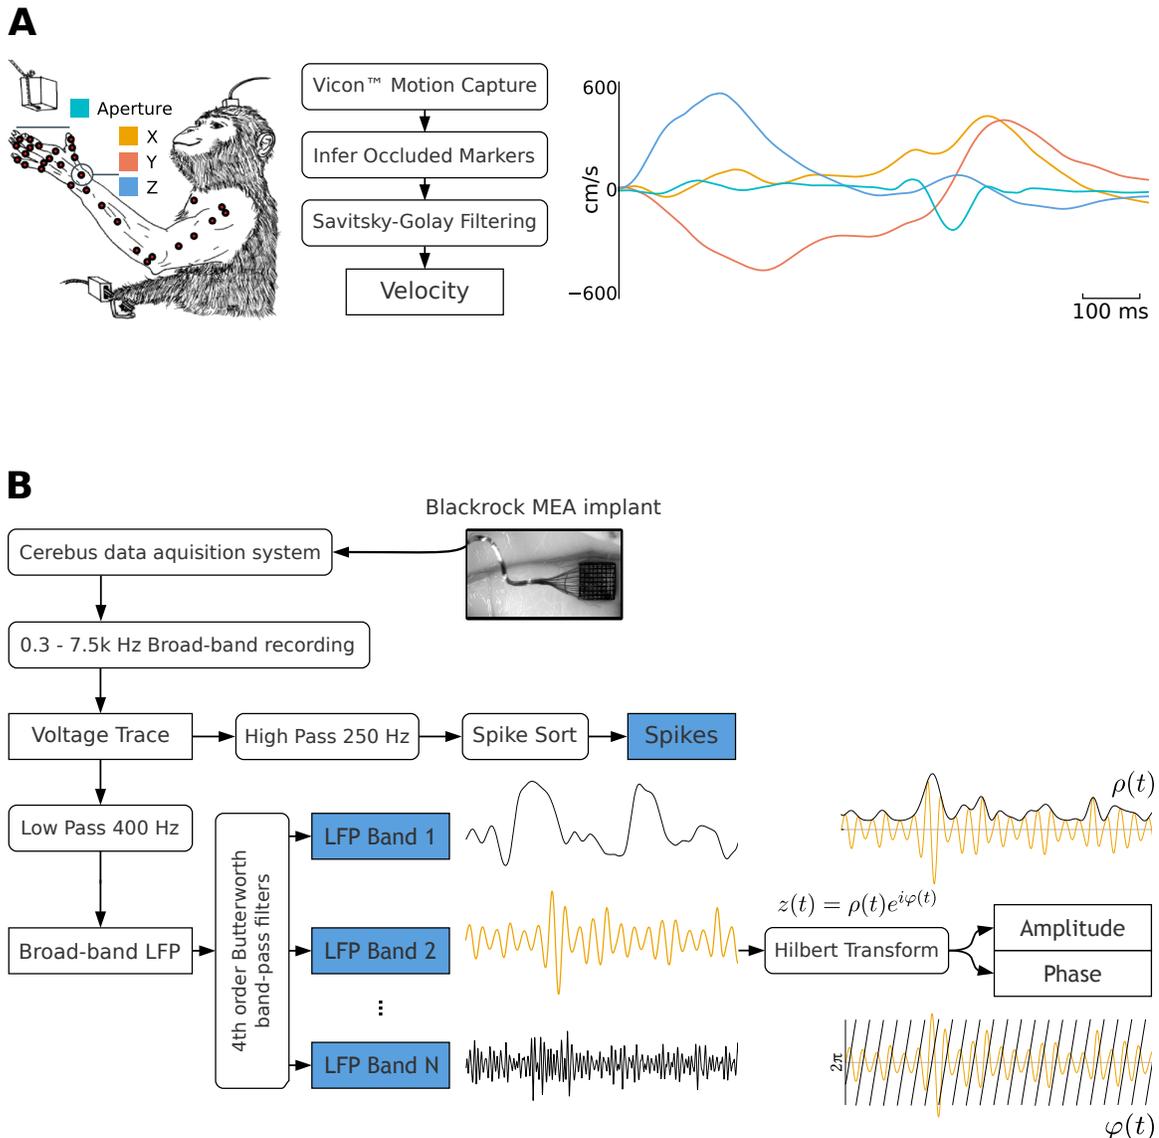


Figure 1.13: Signal processing for reaching and grasping. (A) Kinematics are recorded using a motion capture setup, which tracks the position of infrared reflective markers mounted on the forearm. Motion capture data is cleaned, and occluded markers are inferred based on a model of arm kinematics. Smoothed velocity trajectories are extracted, and normalized “pathlet” features as in Hatsopoulos et al. (2007) are generated for the grip aperture and wrist endpoint separately. The drawing of the monkey is credited to John Mislow and is from Figure 2 in Vargas-Irwin et al. (2010). (B) Neural signals are recorded on implanted microelectrode arrays (MEAs). Spikes are extracted from the high frequency components of the electrical signal, and sorted off-line. Local field potentials are extracted from the low-pass electrical signal, and separated into various bands. The Hilbert transform can be used to construct the analytic signal from narrow-band LFP oscillations, from which the instantaneous phase and amplitude envelope can be extracted.

band. If down-sampling is to be performed, it is important to note that a 4th order Butterworth low-pass filter at the Nyquist frequency is not sufficient to prevent aliasing, and one should instead low-pass filter at half the Nyquist frequency in this case.

In some cases, we employ causal (forward only) band-pass filtering, in order to enforce time-causality when assessing predictive information about spiking in LFPs. This was done both because it simplifies interpretation issues that may arise from spiking contamination of the LFP (see section 1.5.1.1 Spike contamination artifacts), and because translational applications for real-time brain machine interfaces can only make use of time-causal information. Causal filtering is not zero phase, but the 4th-order Butterworth filter exhibits nearly linear phase shift in the passband such that the signal is delayed but not distorted. In the case of ongoing narrow-band LFP oscillations, this delay has minimal impact as there is relatively little amplitude modulation from one cycle of the oscillation to the next, and phase tuning models can account for any shift in the LFP phase. We note also that phase shifts are inherent to these recordings, since LFP signals themselves may be phase-shifted compared to the neural phenomena that generate them (either as a result of the mechanism generating the LFP currents, or through phase shifts due to the impedance of the tissue and the electrode itself). Therefore, it is advantageous to use models that can account for arbitrary phase relationships, and the phase delay introduced into a narrow-band LFP feature from filtering does not affect the results here.

1.5.1.1 Spike contamination artifacts

Chapters 2 and 3 of this thesis deal with predicting single-unit spiking from local field potentials recorded on the same electrode. Although LFP reflects primarily the collective dynamics of the local population, it is also contaminated by the single-unit spiking that we wish to predict. In some cases, it may be extremely difficult to dissociate spiking contamination effects from true correlations between spiking and the population activity as reflected in the LFP. Spiking contamination of low-frequency LFPs occurs primarily

through three mechanisms: (1) single-unit extracellular spikes have a large, fast, and sharp spike waveform that contributes to broad-band LFP power, (2) slow afterpotentials following a spike contain power at lower frequency bands, and (3) spike train rhythmicity at low frequencies can contribute to LFP power at said frequencies.

It is difficult to dissociate network effects from spiking contamination based on extracellular potentials alone. Consider the case of single-unit firing phase-locked to an ongoing LFP oscillation. A spike-triggered LFP average will reveal both the spike waveform itself (due to contamination), as well as the low-frequency oscillatory component. It may be that the neuron itself contains an intrinsic resonance at that frequency, in which case the oscillatory component may correspond to a subthreshold mechanism related to spiking. Alternatively, the oscillation may be mediated entirely by network effects, and oscillatory components in the spike-triggered average emerge only because of the statistical relationship between the spike timing and the network oscillations. Similarly, it can be difficult to distinguish effects of intrinsic spiking history from network effects, as apparent rhythmicity in spiking could be generated either by an intrinsic mechanism or through a population oscillation. Therefore, to a certain extent, both spiking history and low-frequency LFP components contain combined representations of network effects and intrinsic history effects. Because of this, we address spiking contamination using causal filtering of LFP signals. This ensures that spiking contamination is confined to the past, and that any predictive information about future spiking reflects a true predictive history effect.

1.5.1.2 The analytic signal: instantaneous phase and amplitude

The instantaneous phase and amplitude envelopes of narrow-band LFP oscillations can be extracted using the Hilbert transform, which rotates each frequency component by $\pi/2$ in the Fourier domain. The complex-valued combination of the original signal $x(t)$ and its Hilbert transform $y(t)$ is termed the analytic signal $z(t) = x(t) + iy(t)$, and the instantana-

neous phase $\varphi(t)$ and amplitude envelope $\rho(t)$ can be extracted from the polar form of the analytic signal $z(t) = \rho(t)e^{i\varphi(t)}$. Numerically, the analytic signal can be generated using the function `hilbert` in MATLAB and in Python's `scipy.signal.signaltools`, which zeros out the negative frequency components of the Fourier transform and doubles the power in the positive frequency components. For real-valued data, the negative frequency components are complex conjugate to the positive frequency components, such that the imaginary terms of the positive and negative frequency components cancel to yield a real-valued signal. Removing the negative frequency components (and doubling the positive) eliminates this cancellation, and yields the $\pi/2$ phase-shifted signal as the imaginary component. The Hilbert phase can only be interpreted as the literal phase of a true oscillation if the signal consists of a single narrow-band sinusoidal component that modulates slowly in amplitude or frequency. Other procedures exist for extracting phase from LFP measurements that deviate significantly from sinusoidal in waveform shape, but in this work we band-pass LFP with narrow-band filters to ensure that the Hilbert transform is well defined. Throughout this work, we will frequently represent the instantaneous phase $\varphi(t)$ in terms of the unit-length complex phase vector $e^{i\varphi(t)}$. Representing directional information in this form makes it easy to calculate summary statistics, like the circular mean and standard deviation.

1.5.2 Stochastic point process modeling of spiking activity

The vast majority of neurons generate their output signal in the form of sequences of spikes in their membrane potentials. Relating these binary, point-like firing events to other continuous covariates, like stimuli, behavior, or local field potential features, requires statistical models explicitly designed to reflect the point-like nature of the spiking signal.

1.5.2.1 Point process GLMs

Conditionally Poisson generalized linear point-process models (Poisson GLMs or point process GLMs; Truccolo et al. 2005) can estimate the instantaneous firing rate (intensity function), $\lambda(t)$, conditioned on measured covariates such as sensory inputs, motor output, spiking history, and other neural signals. If $X(t) = [X_1(t), X_2(t), \dots, X_n(t)]^\top$ denotes a covariate vector sampled over time, then a point process GLM (Poisson, log-link function) for the conditional intensity function (CIF), $\lambda(t|X(t))$, with an offset parameter μ (related to a background spiking rate) and feature coefficient vector $\mathbf{B} = [\beta_1, \dots, \beta_n]$, can be expressed as

$$\ln[\lambda(t|X(t))] = \mu + \beta_1 X_1(t) + \beta_2 X_2(t) + \dots + \beta_n X_n(t) = \mu + \mathbf{B}X(t). \quad (1.1)$$

We note that $X(t)$ can include any nonlinear function of the original covariates. Other GLMs are also possible (e.g. binomial with logistic link function), but here we focus on the log-link (Poisson) case. The probability of a model $\{\mu, \mathbf{B}\}$, given data $\{X(t)\}$ and spiking $\{Y(t)\}$ (represented here as a sequence of Dirac delta functions at the spike times), can be expressed via Bayes' rule as

$$\Pr(\mu, \mathbf{B} | \{X(t)\}, \{Y(t)\}) = \frac{\Pr(\mu, \mathbf{B})}{\Pr(\{X(t)\}, \{Y(t)\})} \Pr(\{X(t)\}, \{Y(t)\} | \mu, \mathbf{B}). \quad (1.2)$$

As the data are fixed, the term $\Pr(\{X(t)\}, \{Y(t)\})$ is constant. Often we assume a non-informative prior such that $\Pr(\mu, \mathbf{B})$ is also constant, although regularized GLM point-process models can be formulated in terms of selecting appropriate priors for (μ, \mathbf{B}) . Without regularization, fitting the GLM point-process model reduces to maximum likelihood estimation, i.e. to finding the parameters (μ, \mathbf{B}) that maximize the likelihood of observing the data

$$\hat{\mu}, \hat{\mathbf{B}} = \underset{\mu, \mathbf{B}}{\operatorname{argmax}} [\mathcal{L}(\{Y(t)\} | \{X(t)\}, \mu, \mathbf{B})]. \quad (1.3)$$

In continuous time, the log-likelihood of observing spikes $\{Y(t)\}$, given an intensity function $\lambda(t) \triangleq \lambda(t|X(t), \mu, \mathbf{B})$, is composed of a term based on the probability of observing each spiking event, minus a term reflecting the overall expected average number of spikes

$$\ln \mathcal{L}(Y|\lambda(t)) = \int Y(t) \ln(\lambda(t)) dt - \int \lambda(t) dt. \quad (1.4)$$

In practice, these models are fit for discrete time bins of width Δ . For a given model, the probability of observing k spikes in a time interval $[t, t + \Delta)$ is Poisson distributed conditionally on the history

$$\Pr \left(\int_t^{t+\Delta} \lambda(t) dt = k \right) \sim (\Delta \lambda_t)^k \frac{e^{-\Delta \lambda_t}}{k!} + o(\Delta). \quad (1.5)$$

Conditioned on the spiking history and covariates, the likelihood of observing a sequence of (discrete-time) events factorizes, resulting in the product of the likelihoods of each event

$$\mathcal{L}(Y|X, \mu, B) = \prod_t (\Delta \lambda(t))^{Y_t} e^{-\Delta \lambda(t)} / Y_t!. \quad (1.6)$$

For numerical efficiency, we minimize the negative log-likelihood, rather than maximize the likelihood. This converts the product to a sum

$$\operatorname{argmin}_{\mu, B} \left[\sum_t -\ln \left((\Delta \lambda(t))^{Y_t} e^{-\Delta \lambda(t)} / Y_t! \right) \right]. \quad (1.7)$$

If Δ is chosen sufficiently small such that $Y_t \in \{0, 1\}$, this expression simplifies because $Y_t! = 1$. For notational simplicity, we will also let $\lambda_t = \Delta \lambda(t)$ have units of $1/\Delta$, such that the multiplication by the time window Δ is now implicit.

$$\operatorname{argmin}_{\mu, B} \left[- \sum_t Y_t \ln(\lambda_t) - \lambda_t \right]. \quad (1.8)$$

As indicated above (Equation 1.1), the point process GLM framework considers models in

which $\ln(\lambda_t)$ is a linear function in the model parameters

$$\ln(\lambda_t) = \mu + \mathbf{B}X(t), \quad i.e. \quad \lambda_t = \exp[\mu + \mathbf{B}X_t]. \quad (1.9)$$

In this case the problem takes the form

$$\operatorname{argmin}_{\mu, \mathbf{B}} \left[- \sum_t Y_t [\mu + \mathbf{B}X_t] - e^{\mu + \mathbf{B}X_t} \right]. \quad (1.10)$$

The software MATLAB (`glmfit`) and the Python package `statsmodels` contain built-in functions for solving this minimization problem using Newton's method and iteratively reweighted least squares. At the time of writing this (MATLAB R2015b, `statsmodels` 0.6.1), it is in practice faster to solve the problem using gradient descent. Either MATLAB's `minimize` function or the `newton-CG` option of the `minimize` function in SciPy (Jones et al., 2001) perform well ($\sim 3\times$ faster than iteratively reweighted least squares).

The gradient and Hessian functions to provide to the gradient solver are:

$$\begin{aligned} \frac{\partial}{\partial \mu} [-\ln \mathcal{L}(Y|X, \mu, \mathbf{B})] &= \sum_{t=1}^T [e^{\mu + \mathbf{B}X_t} - Y_t], \\ \frac{\partial}{\partial \beta_i} [-\ln \mathcal{L}(Y|X, \mu, \mathbf{B})] &= \sum_{t=1}^T [e^{\mu + \mathbf{B}X_t} - Y_t] X_{t,i}, \end{aligned} \quad (1.11)$$

$$\begin{aligned} \mathbf{H} &= \begin{bmatrix} \frac{\partial^2}{\partial \mu^2} & \frac{\partial^2}{\partial \mu \partial \beta_1} & \cdots & \frac{\partial^2}{\partial \mu \partial \beta_n} \\ \frac{\partial^2}{\partial \beta_1 \partial \mu} & \frac{\partial^2}{\partial \beta_1^2} & \cdots & \frac{\partial^2}{\partial \beta_1 \partial \beta_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2}{\partial \beta_n \partial \mu} & \frac{\partial^2}{\partial \beta_n \partial \beta_1} & \cdots & \frac{\partial^2}{\partial \beta_n^2} \end{bmatrix} [-\ln \mathcal{L}(Y|X, \mu, \mathbf{B})] \\ &= \sum_{t=1}^T e^{\mu + \mathbf{B}X_t} \begin{bmatrix} 1 \\ X_t \end{bmatrix} \begin{bmatrix} 1 \\ X_t \end{bmatrix}^\top \end{aligned} \quad (1.12)$$

where T corresponds to the total number of samples (time bins). The above corresponds to the empirical expectation (average w.r.t. the observations Y_t) of the Hessian matrix.

1.5.2.2 Incorporating spiking history

The intrinsic properties of neurons, like refractoriness and recovery times, in addition to fast-timescale variability imposed by network dynamics (e.g. network oscillations), can induce autocorrelations in spike trains. Therefore, it is important to incorporate the intrinsic spiking history of a neuron when predicting its firing pattern. Spiking history can be incorporated into the point process GLM framework by adding terms depending on the spiking activity at previous times i.e. $\beta_1 Y_{t-1} + \beta_2 Y_{t-2} + \dots + \beta_n Y_n$. Although in theory one might model the effect of the entire spiking history, in practice we consider only a finite time window and use a basis function representation in order to reduce the total number of history parameters to estimate (e.g. Figure 1.14A). These basis functions are both a computational convenience (lower dimensional models can be estimated more rapidly) and a form of regularization to enforce smoothness in the estimated filters and prevent the over-fitting that can be induced by a large number of parameters.

1.5.2.3 Assessing predictive power

Predictive power can be assessed using receiver operating characteristic (ROC) curve analysis. The area under the ROC curve (AUC) provides a summary of the ability of a model to predict the observed neural spiking (Fawcett, 2006). It is related to the probability of correctly ranking two randomly selected time bins, one that does and one that does not contain a spike. The accuracy of the AUC estimate depends on the extent to which the distributions of predicted intensity λ , conditioned on both the presence and absence of spikes, is well sampled. In general, there is negligible firing-rate dependent bias in the AUC estimate, although the variance of the estimator can increase for small sample sizes. Combined with the fact that point process GLMs explicitly model the mean firing rate

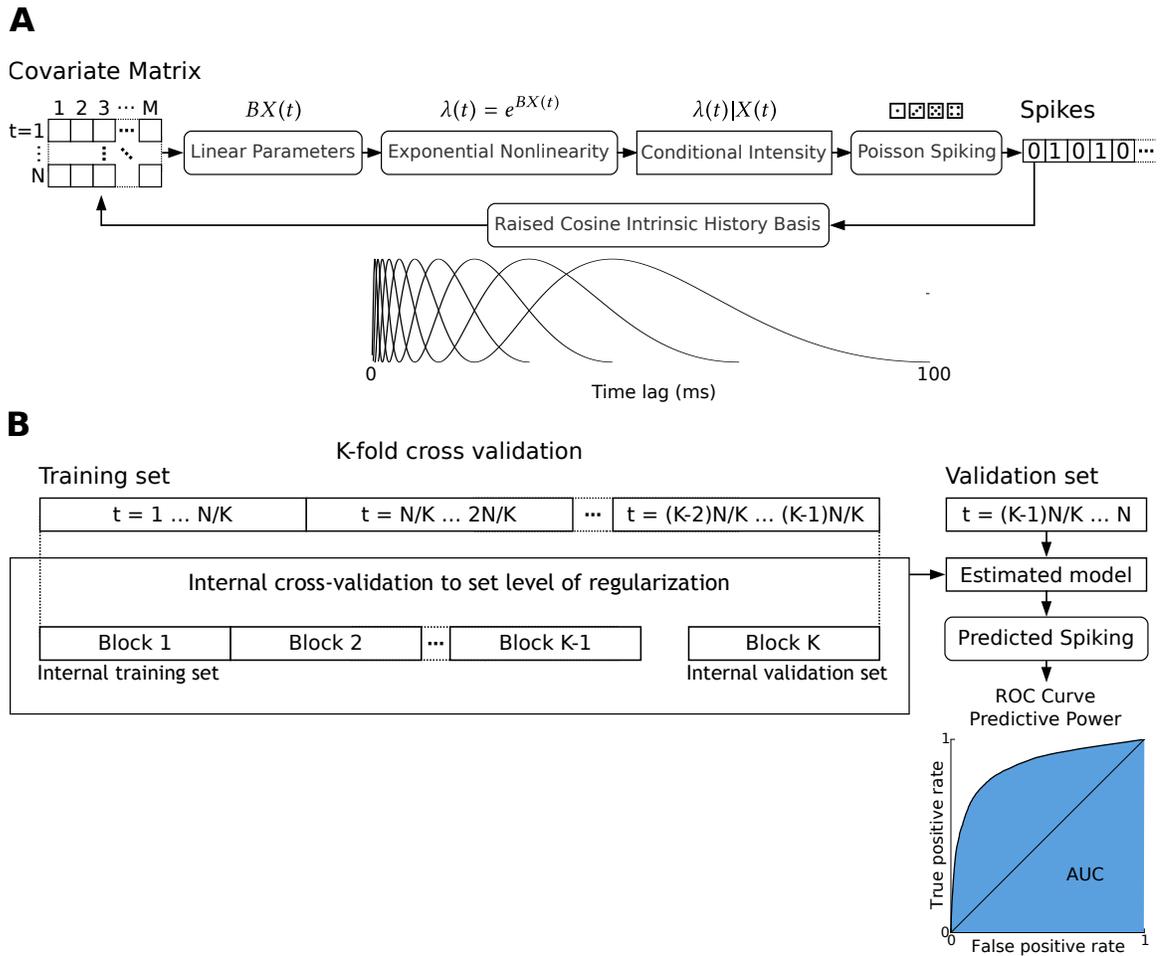


Figure 1.14: Point-process modeling of spiking activity. (A) Conditionally Poisson generalized linear models predict neural spiking as a function of extrinsic covariates and spiking history. These models predict the logarithm of the instantaneous spike rate (“conditional intensity”) as a function of the covariates that is linear in the model parameters. These models accept a matrix of features, which can relate to extrinsic (e.g. stimuli and behavior) or intrinsic (e.g. attentional and cognitive processes) features of interest. Recursive models can be generated by including intrinsic and population spiking history covariates. Typically, spiking history is convolved with a set of basis functions to reduce the dimensionality of the history feature vector and enforce smoothness. **(B)** The predictive power of GLM point-process models can be assessed using receiver operator characteristic (ROC) curve analysis. The ROC curve reflects the trade-off between false positive and false negative predictions when using the estimated conditional intensity function to predict the occurrence of single spikes. When fitting models with a large number of parameters, chance correlations between spiking and covariates can cause over-fitting and lead to spurious assessments of model predictive power. Cross-validation approaches, in which predictive power is assessed on data not used for model training, are important to avoid spuriously elevated estimates of predictive power. Regularization reduces over-fitting by defining prior assumptions about the likely distribution of model parameters. The strength of regularization is another free parameter that must be estimated. One approach to setting the strength of regularization is to test the extent of over-fitting using a second level of internal cross-validation on the training data.

of a neuron, this means that point process GLM performance as assessed by ROC curve analysis is insensitive to variations in spike rate. This is useful if one wishes to compare effects, like the degree of phase coupling or amount of spiking variability explained, between neurons with different firing rates, or in the same neuron across conditions.

1.5.2.4 Cross-validation

When using statistical models to assess the extent to which individual covariates explain neural spiking, it is important to properly address over-fitting. Models fit to a small amount of training data can often predict the training set, but this accuracy might be due to over-fitting. It is therefore necessary to assess predictive power under cross-validation. Cross-validation, in general, entails validating the statistical model on data that has not been used in model estimation. K-fold cross-validation breaks a dataset into K groups, for each group trains a model based on the remaining K-1 groups, and then tests the performance of this model on the withheld data (Figure 1.14B). As the assessment of predictive power under generalization itself exhibits some variance, repeating this procedure for all K groups can be used to reduce the variance in the estimation of model performance. In cases where predictive power is lower under cross-validation compared to the training data, regularization can often improve the accuracy of model estimation.

1.5.2.5 Regularization

In the case where model performance is worse on test data compared to training, it may be appropriate to incorporate regularization into the model estimation procedure. This can be achieved by adding a penalty term, which is a function of the parameter values, to the negative log-likelihood, such that larger parameter values are discouraged. Common regularization terms include the L2 and L1 norm of the parameter vector, corresponding to a Gaussian and Laplacian prior on the distribution of the parameters, respectively. The L0 penalty, which incorporates a constant penalty for every nonzero parameter, may be

approximated by greedy algorithms (e.g. Li et al. 2015), or by L1. Penalties that combine L1 and L2 regularization, and penalties over groups of parameters together, are also used. Determining which penalty is appropriate to use is something of an art. In general, one would use a L1 penalty when one expects most parameters to be 0, as is the case when the informative features are sparse. Greedy approximation of the L0 penalty is computationally expensive, but is sometimes used when one wishes to identify an informative subset of the features without introducing a magnitude-dependent penalty on the parameters. Approximations of the L0 penalty are often used to select an informative subset of neurons out of a larger population (e.g. Bansal et al. 2012; Menzer et al. 2014). The L2 penalty is appropriate when one expects parameters to be small, but not necessarily sparse. The L2 penalty was found to be especially appropriate in this research for cases where there was some redundancy (correlation) between features, and the parameter vector \mathbf{B} is expected to be dense, as was the case for predicting spikes based on movement trajectories.

The selection of a regularization procedure and setting the strength of regularization can also be subject to over-fitting. It is therefore important to either fix the regularization strength in advance, or identify the strength of the regularization based on the training data alone. There are several ways to do this, but the one that we employed here is to test the generalization of models under a second level of simulated cross-validation, and to select the regularization strength that led to the best generalization under this procedure (Figure 1.14B). In Chapter 2, this two-tier cross-validation approach with regularization prevented the predictive power of models from degrading even when a large number non-informative features were added.

1.5.3 Characterizing spike-LFP phase coupling

Characterizing the relationship between spikes and local field potentials is a common statistical problem in neuroscience. Traditional signal-processing methods such as coherence suffer from a number of undesirable biases when applied to point-process data. In

particular, applying the usual signal-processing definition of coherence to spike trains is biased both by the underlying firing rate of a neuron, and by the presence of variations in firing rate correlated with LFP power fluctuations (Lepage et al., 2013). Several estimators have been designed to circumvent this problem, and each has its benefits and drawbacks. In this work, we use two recently developed estimators of spike-LFP phase coupling that avoid firing rate biases: point process GLMs for phase tuning, and the Pairwise Phase Consistency (PPC).

1.5.3.1 Point process GLMs for spike-LFP phase coupling

The point process GLM framework can be used to estimate a phase-tuning model based on the instantaneous phase $\varphi(t)$ of ongoing narrow-band LFP oscillations. Lepage et al. (2013) introduce GLM phase tuning models as an alternative to spike-field coherence. Zhou et al. (2015) have confirmed that the point process GLM based measures of spike-LFP phase tuning exhibit lower bias than spike phase histogram approaches, and are especially advantageous if intrinsic history terms are added to dissociate effects of firing history from those of population oscillations. In this work, we employ GLM phase tuning models to examine the relationship between spiking and ongoing LFP oscillations. Taking $\varphi(t)$ to represent the phase of a narrow-band LFP oscillation, φ_0 to represent the preferred phase tuning, and α to represent the strength of this tuning, the form of the model for phase tuning to narrowband LFP oscillations $z(t) = \rho(t) \exp(i\varphi(t))$, with no history, is:

$$\ln(\lambda_t) = \mu + \alpha \cos(\varphi(t) - \varphi_0) = \mu + \beta_1 \sin(\varphi(t)) + \beta_2 \cos(\varphi(t)). \quad (1.13)$$

1.5.3.2 Pairwise phase consistency

In 2010, two papers addressing the firing-rate bias in spike-field coherence emerged: Grasse and Moxon (2010) proposed a bias-corrected version of the power spectrum of the spike-triggered LFP average, and Vinck et al. (2010) suggested a measure based on

the distribution of phase differences between all pairs of spike-triggered LFPs. It turns out that this pairwise phase consistency (PPC) measure is remarkably similar to the bias-corrected coherence of Grasse and Moxon (2010), and that the measure is equivalent to estimating moments of circularly distributed data (Aydore et al., 2013). Vinck et al. (2010) define Pairwise Phase Consistency (PPC) as the average dot product between all pairs $j \neq k$ of N spike-triggered phase measurements $\varphi_1, \dots, \varphi_N$:

$$\text{PPC} = \frac{2}{N(N-1)} \sum_{j=1}^{N-1} \sum_{k=j+1}^N \cos(\varphi_j - \varphi_k). \quad (1.14)$$

A more computationally efficient form of the PPC is given by Equation 11 in Aydore et al. (2013), and can be derived as follows. Let a collection of N phase measurements be denoted by the complex phase vectors $z_j = \exp(i\varphi_j) = \cos(\varphi_j) + i \sin(\varphi_j)$, $j = 1, 2, \dots, N$. If \bar{z} is the population average phase vector, then

$$N^2 |\bar{z}|^2 = \left(\sum_{j=1}^N \cos(\varphi_j) \right)^2 + \left(\sum_{j=1}^N \sin(\varphi_j) \right)^2 \quad (1.15)$$

The sum of the cosine of all pairwise phase differences (including self pairs $k = j$) is related to the squared magnitude of the sum over all phase vectors:

$$\begin{aligned} \sum_{j=1}^N \sum_{k=1}^N \cos(\varphi_j - \varphi_k) &= \sum_{j=1}^N \sum_{k=1}^N \cos(\varphi_j) \cos(\varphi_k) + \sin(\varphi_j) \sin(\varphi_k) \\ &= \left(\sum_{j=1}^N \cos(\varphi_j) \right)^2 + \left(\sum_{j=1}^N \sin(\varphi_j) \right)^2 = N^2 |\bar{z}|^2. \end{aligned} \quad (1.16)$$

Because each pair of phases enters into the above expression twice, and because the dot-product of a unit vector with itself is 1, Equation 1.16 relates to the sum in the right-hand

side of Equation 1.14:

$$\sum_{j=1}^N \sum_{k=1}^N \cos(\varphi_j - \varphi_k) = N + 2 \sum_{j=1}^{N-1} \sum_{k=j+1}^N \cos(\varphi_j - \varphi_k) \quad (1.17)$$

Equations 1.16 and 1.17 imply that (see also Aydore et al. (2013) Equation 11)

$$PPC = \frac{N|\bar{z}|^2 - 1}{N - 1}. \quad (1.18)$$

This is the form that we use to compute PPC in this thesis. In this form, the relationship between PPC and other bias-corrected coherence measurements is more obvious. Compare Equation 1.18 to Equation 10 (1.19 below) from Grasse and Moxon (2010), defining a bias-corrected measure of spike-field coherence:

$$c = \frac{N\hat{c} - 1}{N - 1}, \quad (1.19)$$

where

$$\hat{c} = \frac{Power_{STA}}{\langle Power_{LFP} \rangle}. \quad (1.20)$$

Power in this case is computed from the squared magnitudes of the Fourier coefficients from the spike-triggered ($Power_{STA}$) and the whole-session ($Power_{LFP}$) power spectrum using a multi-taper approach. The primary difference here, then, is that the PPC removes effects of power modulation by normalizing spike-triggered LFP analytic signal measurements individually, whereas Grasse's coherence normalizes the final estimate by an expected power based on the background LFP power spectrum.

1.5.3.3 Addressing remaining source of bias

Although there are a number of approaches for removing the spike-rate dependent bias from spike-LFP phase coupling estimates, some sources of residual bias remain. Tradi-

tional coherence estimators normalize the coherence by autospectral power, and therefore automatically correct the coherence magnitude for autocorrelations in the signal. The spike-field phase coupling estimators discussed here do not incorporate this, and therefore care must be taken to address biases resulting from autocorrelation in the spike trains. This can arise from highly rhythmic neural firing, in which the strong periodicity of the spike train may cause it to appear phase locked to transient LFP oscillations at the same frequency by chance. As suggested by Vinck et al. (2010), it is necessary to determine the timescale over which the spike train and LFP are autocorrelated, and exclude spikes that occur too close together (with respect to the autocorrelation decay time) from the phase locking calculation. To verify that the effects of signal autocorrelations are adequately addressed, one can perform random permutation tests using surrogate data that mimic the autocorrelation structure of the spike trains and LFP signals, but contains no phase locking. In experimental designs that include multiple trials, and in which neither spikes nor LFP are strongly locked to the trial structure, it is sufficient to compute the phase locking between the spikes and the LFPs from unrelated trials. We use such an approach in Chapter 3 to evaluate the chance level of β -frequency spike-LFP phase coupling, as the timing and phases of beta transients during the steady state periods in our experiment showed no stable phase relationships across trials. In the case where trials are not independent due to trial-locked neural activity, phase randomization can generate a surrogate LFP signal that has the same autocorrelation structure as the true data, but no stable spike-LFP phase relationship.

1.5.4 Classification of oscillatory LFP spatiotemporal patterns

Multielectrode array recording techniques provide an opportunity to observe how populations of neurons interact collectively across cortex. Methods for summarizing multi-channel activity, like population coherence and phase synchrony, provide summary statistics that may reflect features of the underlying neural dynamics. Methods borrowed from

fluid flow analysis and computer vision can let us analyze geometric relationships between LFP phases and allow us to characterize spatiotemporal wave activity.

1.5.4.1 Characterizing population activity

Population coherence The coherence spectrum is the frequency-space analogue of the cross correlation between two signals. While the cross correlation reflects the tendency for fluctuations in signal A to predict fluctuations in another signal B at some time lag, coherence reflects the extent that fluctuations in A within a particular frequency band predict (linearly) fluctuations in that same band in B . Formally, coherence is the power spectrum of the cross correlation, normalized by the power spectra of the autocorrelation functions of each signal. On K -channel multielectrode arrays, there are potentially $K(K - 1)/2$ pairwise coherence spectra. To summarize the overall coherence across a population of electrodes, one can take the m leading (largest) eigenvalue(s) from the pairwise coherence matrix at each frequency.

Measures of population zero-lag spatial phase synchrony Coherence, by design, allows for individual channels to oscillate at a relative phase lag. It therefore cannot distinguish between a large number of different spatiotemporally organized wave phenomena. Zero-lag spatial phase synchrony, traveling waves, and stable radiating, rotating, and complex waves, can all exhibit the same overall coherence. A number of statistics exist to distinguish states exhibiting zero phase lag spatial synchronization from these other cases. One such statistic, from the coupled oscillator literature, is the Kuramoto order parameter (Kuramoto, 2012). If a population of oscillations with phases φ_j are represented by their complex-valued unit phase vectors $e^{i\varphi_j}$, then let

$$R \cdot e^{i\varphi_0} = \frac{1}{N} \sum_{j=1}^N e^{i\varphi_j}, \quad (1.21)$$

where R is commonly called the Kuramoto order parameter, and φ_o is the mean phase. R is 1 when the phases of all oscillators are perfectly aligned, and tends toward 0 when the distribution of population phases approaches a uniform distribution.

The distribution of phases in a population of coupled oscillators may be described by a circular distribution. The von Mises distribution and the circularly wrapped normal distribution are two common cases. The circular standard deviation S , in particular, is a circular analogue to ordinary linear standard deviation and can be computed from the Kuramoto order parameter $S = \sqrt{\ln(1/R^2)}$. The main advantage of circular standard deviation in this work is that the familiar intuition about normally distributed data carries over to some extent. For example, approximately 95% of normally distributed data lie within two standard deviations of the mean. Similarly, circularly distributed data for which the circular standard deviation is less than $\pi/4$ will have approximately 95% of data concentrated within the same quadrant of phase space.

The complex Gaussian So far we have discussed two related measures of population zero-lag phase synchrony: Kuramoto's order parameter and the circular standard deviation. These apply to populations of phase oscillators. The narrow-band analytic signal that we extract from LFP also has both phase and magnitude components, and fluctuations in the magnitude and phase of ongoing oscillations may be of interest. At times, it may be appropriate to attempt to model the distribution of analytic signals across a population. The complex Gaussian is one such distribution, and simply treats the real and imaginary components of the analytic signal as a two-dimensional normally distributed variable.

1.5.4.2 Spatiotemporal pattern analysis

Challenges Several caveats must be addressed when examining spatiotemporal dynamics using MEA data. There are four main issues addressed here. First, the spatial

scale of cortical wave phenomena can be larger than the size of the MEA, introducing an aperture problem. Second, the Nyquist sampling theorem dictates that we should sample at at least twice the resolution as the highest spatial frequency in our data. Third, LFP signals are subject to some amount of noise, and there is currently no clear procedure for establishing ground truth as to which signals should be considered noise. Fourth, spatiotemporal LFP oscillations may exhibit complex structure, including harmonic relationships with other LFP bands, which can cause the apparent spatiotemporal dynamics to depend on band-pass filtering used in data pre-processing. The following sections will elaborate on these problems and discuss workarounds as well as interpretation caveats that remain. Figure 1.15 illustrates the main concepts and quantities used in the analyses of spatiotemporal wave patterns.

Hilbert phase gradients In analyzing spatial structure in MEA recordings, a challenge arises as the spatial scale of the wave activity can be larger than the size of the recording area. Huang et al. (2004) introduce the spatial phase gradient as a method of categorizing wave spatial structure, which circumvents these issues to some extent. The main issue is that the lowest non-DC Fourier component for a discretely sampled array has a wavelength equal to the size of the array. For example, for the 4×4 mm arrays used in this thesis, we can resolve up to 5.7 mm waves at best, if they are aligned to the diagonal of the array. However, phase gradient methods introduced by Rubino et al. (2006) can extract information about ongoing spatiotemporal wave dynamics and circumvent the aperture limit. While we have limited samples in the spatial domain, we have high-resolution information in the time domain. In cases where the relevant spatiotemporal activity is confined to a narrow-band oscillation, we extract an instantaneous Hilbert phase for each electrode. If we assume that traveling wave activity moves directly along the phase gradient, we can interpret gradients in the Hilbert phase as reflecting ongoing,

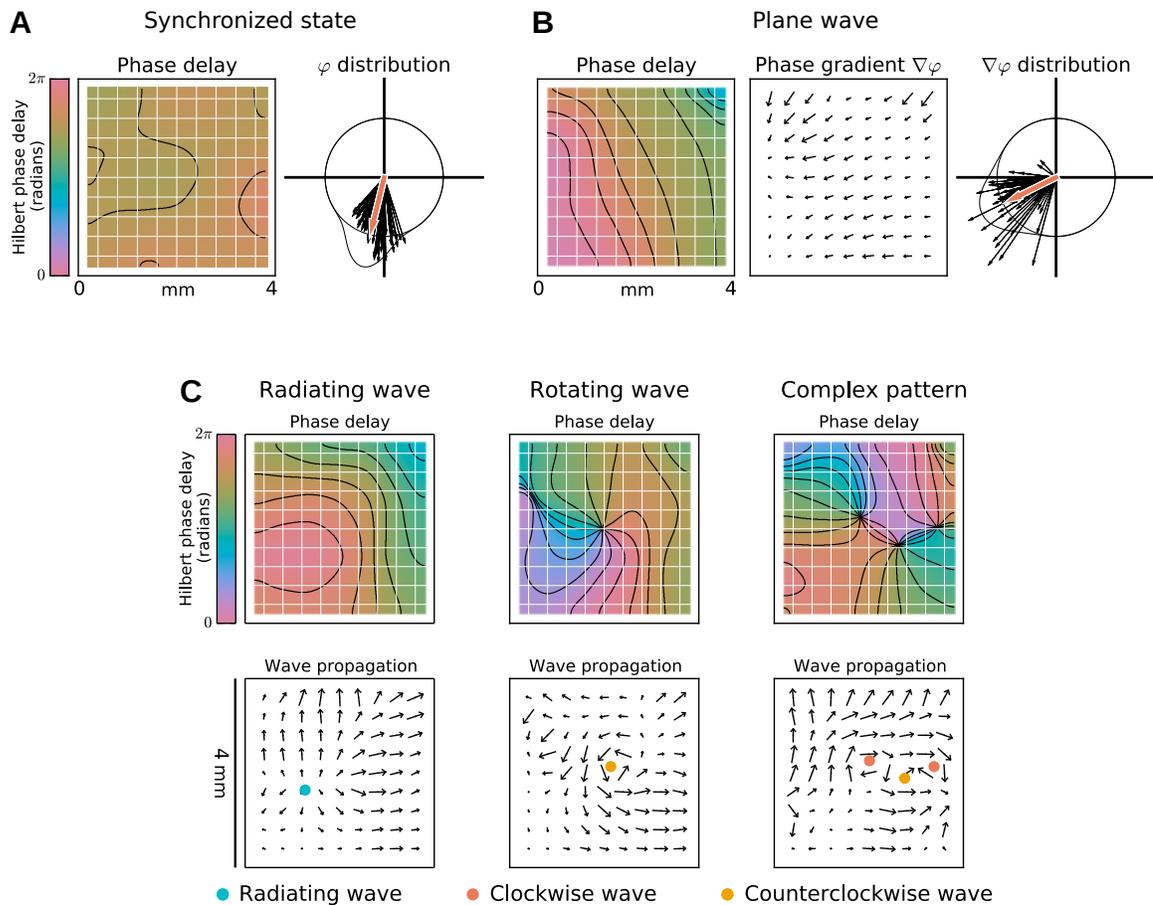


Figure 1.15: *Detection and categorization of spatiotemporal structure in LFP oscillations.* (A) States with zero lag spatial phase synchrony can be detected as times when the distribution of Hilbert phases across the MEA is unusually concentrated. The circular standard deviation of the phase distribution is one measure of phase dispersion. (B) Globally organized plane waves can be detected by taking the gradient of the Hilbert phase across the MEA and identifying times when the phase gradient vectors are aligned. Circular statistics can also be used to identify times when the distribution of phase gradients is unusually concentrated. (C) More complicated wave states, such as radiating, rotating, and complex waves, cannot be readily distinguished using the statistics of the Hilbert phases or the Hilbert phase gradients. Instead, a smoothed estimate of the Hilbert phase gradient is constructed, and complex waves are identified based on critical points in the phase gradient field. The arrows in this illustration have been reversed in these plots to demonstrate the direction of wave propagation, which is opposite that of the Hilbert phase gradient. Local maxima in the Hilbert phase map correspond to the centers of radiating waves. Saddle points and local minima can also be identified, but in our analysis do not correspond to true stable features in the organization of LFP phases and so are not considered. The centers of rotating waves can be identified as points around which the Hilbert phase cycles by $\pm 2\pi$. Such points can be found by taking a line integral of the Hilbert phase gradient around the vicinity of the point, which will be zero in the case that there is no net rotation.

traveling spatiotemporal wave phenomena (Rubino et al., 2006)

$$\nabla\varphi \cdot \mathbf{v} = \frac{\partial\varphi}{\partial x} \frac{dx}{dt} + \frac{\partial\varphi}{\partial y} \frac{dy}{dt} = -\frac{\partial\varphi}{\partial t}. \quad (1.22)$$

Identifying plane waves: Phase gradient directionality Similarly to how time periods of zero-lag spatial phase synchrony can be identified based on the dispersion of Hilbert phase vectors, traveling plane waves can be identified as times when the direction of phase gradient vectors is unusually concentrated. In analogy to assessing zero-lag spatial phase synchrony, the circular standard deviation of phase gradient directions can be used to assess phase gradient alignment and to detect plane waves. Rubino et al. (2006) define a measure of phase gradient directionality (PGD) as the magnitude of the array-averaged phase gradient normalized by the average phase gradient magnitude (Equation 1.23). This measure ranges from 0 for a disordered state to 1 in the event that all phase gradient vectors point in the same direction. See Figure 1.15B for illustration.

$$PGD(t) = \frac{|\overline{\nabla\varphi(t)}|}{|\overline{\nabla\varphi(t)}|}. \quad (1.23)$$

Identifying complex wave structures from phase gradients Circular statistics can be used to identify spatial synchrony and traveling plane waves, but spatially organized wave activity in cortex can take on many other forms. Radiating and rotating waves are possible, and the overall spatiotemporal dynamics might not exhibit global organization. Radiating and rotating waves, and combinations thereof, can be identified by treating the gradient of the Hilbert phase as if it were a vector field, and identifying critical points in this field. Radiating waves occur at local minima of the phase map, and rotating waves are centered at phase singularities, around which the oscillation phase rotates an integer multiple of 2π . Townsend et al. (2015) have recently used critical point methods to characterize the emergence of complex wave states in delta LFP. Critical point

methods are susceptible to noise, which can generate spurious local extrema and phase singularities. To address this problem, Townsend et al. (2015) applied a smoothness prior in their estimation of phase gradients. In this work, we spatially low-pass filter the array data to examine spatiotemporal activity below 2 mm in wavelength.

The Nyquist limit in spatial sampling of LFP The multielectrode arrays have a 400 μm electrode spacing, and therefore cannot resolve wavelengths shorter than 800 μm . If the LFP exhibits spatial organization at finer scales, aliasing artifacts could confound the analysis. LFPs are, to a large extent, averages over the activity of large neural populations, and also propagate (“instantaneous” volume conduction) some distance within the brain. The tendency for local collective activity to act synchronously, as well as this volume conduction, sets a minimum resolution for the LFP. However, we have anecdotally observed large LFP phase differences between adjacent electrodes, indicating that our spatial sampling is not always dense enough to fully resolve LFP spatial structure. Nevertheless, the spatial sampling density of the MEA is sufficient when the population activity itself becomes locally synchronous. The beta and gamma oscillations that we observe here tend to exhibit wavelengths larger than the electrode spacing, and so the spatial Nyquist limit is less of an issue. This is an assumption, and transient wave activity may exhibit high spatial frequency structure that could confound our analysis. Particularly dramatic artifacts occur when there is a spatially coherent oscillation higher than the Nyquist frequency, which can appear as a lower frequency wave due to aliasing. Heuristically, spatial coherence of wave phenomena in cortex scale with their wavelength. Therefore, high-frequency wave activity that is above the spatial Nyquist limit is typically localized, and the aliasing effects are more likely to result in simple per-channel ‘noise’ rather than an artifactual lower-frequency wave. Also to our benefit is the fact that, when frequencies exceed the Nyquist limit, they alias down into higher frequencies first, with only much higher frequencies aliasing into low-frequency components. If the expected wave structure is spatially low-pass it is possible to denoise the wave activity using spatial

smoothing, which removes both true noise and some of the aliasing artifacts.

Biases and confounds in spatial analysis arising from noise Presently, we have no ground truth for assessing the signal-to-noise ratio in LFP recordings. It is not so much that we do not understand the physics of the recording setup, as it is that the distinction between signal and noise is blurry. Spike trains contribute to the LFP and also tend to be highly localized, meaning that LFP effects localized to a single electrode may arise from both noise and true neural activity. For the analysis of beta waves, we can get a rough estimate of the noise level by examining the minimum beta power observed during movement-related beta suppression, in which we find it is about 10-20 μV . The presence of noise can obscure spatially organized LFP phenomena, like synchrony and plane waves. Noise can be especially problematic for estimating plane-wave wavelength using the Hilbert phase gradient. Without noise, the local Hilbert phase gradient is directly related to wavelength. Noise adds spurious high-frequency spatial structure that inflates the local estimate of the gradient. One workaround is to estimate wavelength from the global activity, by either fitting a plane-wave model to the observed data, or estimating the wavelength from the average Hilbert phase gradient, for which local noise sources average out. This approach is used by Rubino et al. (2006). However, if the wave is not perfectly planar, averaging the gradient over the array leads to an incorrectly low estimate of the average gradient, because regions of the wave propagating in different directions cancel out. This inflates the estimated wavelength.

In summary, naïvely estimating wavelength from local phase gradients can cause spuriously high estimates due to noise, and estimating wavelength from globally averaged phase gradients can cause spuriously low estimates due to deviations from the plane wave structure. If the local phase gradients are perfectly aligned, the two approaches are equivalent. Rubino et al. (2006) consider waves only above a PGD threshold of 0.5, effectively requiring that the wavelength estimated from local versus global phase gradients agree within a factor of two. In our work, we wished to examine correlations between wave-

length and amplitude, and so this confound is especially problematic. To address this, we required that phase gradient vectors be even more highly aligned, such that the potential discrepancy between the two wavelength estimation approaches is at most 30%, approximately two orders of magnitude less than the amplitude-dependent variability in wavelength observed. With this approach, biases arising from varying SNR remain small relative to the overall effect size.

Interaction between spatiotemporal filtering and apparent wave activity When using Hilbert phase gradients to estimate spatiotemporal wave properties, the time domain filtering can interact with the spatial analysis. These effects are potentially quite subtle and complex, and are not fully yet understood. The apparent spatiotemporal dynamics of the LFP can depend on frequency. In general, lower-frequency oscillations are also the most spatially coherent. Multiple narrow-band oscillations may coexist in cortex, and each may display different spatiotemporal dynamics. For example, in the examination of optogenetic stimulation in motor cortex in Chapter 4, we find four frequency bands of optogenetically induced oscillations. Two of these bands display apparently no consistent spatial organization relative to the stimulation site over time. The other two each have a different characteristic spatial organization.

Selecting appropriate filtering is essential not only for separating different spatiotemporal components from each-other, but also for separating narrow-band spatiotemporal oscillations from background noise. In all of the analyses here, we concentrate on narrow LFP peaks identified separately for each cortical area, such that the observed wave phenomena correspond to true narrow-band activity and not spatiotemporally band-pass noise. Nevertheless, in one of the studied monkeys, the beta band appeared to exhibit overtones at higher frequencies, indicating that even a single oscillation can have complex structure that might be affected by choice of filtering. In Chapter 4, we focus on spatiotemporal structure of the fundamental frequency, but how spatial structure might relate to the higher-frequency overtone remains unresolved.

Filter bandwidth can also interact with spatiotemporal analysis. The uncertainty principle creates a trade-off in our ability to resolve signals in time and frequency. Narrow-band signals have poor time resolution of amplitude and phase fluctuations. In cases where resolving amplitude fluctuations in a LFP band is important, band-pass filtering can have a dramatic effect on the analysis. More subtly, narrow band filtering can also generate a signal that is more autocorrelated, and for which the so called "instantaneous" phase estimate is not instantaneous at all, but instead results from an average over several oscillation cycles. For example, the Hilbert phase from a 2 Hz band centered at 20 Hz will be related to approximately ten periods of the underlying 20 Hz oscillation. If the "instantaneous" phase is the quantity of interest, then it is necessary to select a broader filter band in order to ensure temporal locality. Excessively narrow-band filtering can obscure the extent to which spatiotemporal wave patterns change dynamically over time. Excessively broad-band filtering increases the uncertainty in the Hilbert phase estimate, and increases the amount of noise introduced into the signal. As a heuristic, we choose bands that are no larger than a quarter of the peak frequency, meaning that the instantaneous phase is reflective of the current activity plus and minus two oscillation cycles.

Another issue to note is that filter selection interacts with spatial phase gradient analysis by virtue of the fact that each frequency band has a potentially independent contribution to the spatial distribution of the local field potential. Consider the case where the apparent spatiotemporal LFP activity represents broadband noise in time, and independent per-frequency narrow-band noise in space, resembling waves. A narrow band-pass filter will select a subset of spatially similar components, leading to apparent wave activity. On the other hand, a broad-band filter will sum over many independent spatial components and lead to less spatially organized activity. In the case of true narrow-band activity, broadening the filter band, or moving the band such that the signal-to-noise ratio becomes poorer, can introduce more noise components and lead to apparently less synchronous or more complex wave activity. This is another source of coupling between

spatial analyses and temporal filtering that must be addressed. In the case of the studies here, we focus on stable narrow-band peaks in the LFP oscillations and use fixed frequency bandwidths when comparing across conditions. We also verify that our analysis is robust to the selection of filter bandwidth.

Chapter 2

Contribution of LFP dynamics to spiking variability in motor cortex during movement execution

MICHAEL RULE, CARLOS VARGAS-IRWIN, JOHN DONOGHUE, WILSON TRUCCOLO

This chapter has been published: Rule, M. E., Vargas-Irwin, C., Donoghue, J. P., and Truccolo, W. (2015). Contribution of LFP dynamics to single neuron spiking variability in motor cortex during movement execution. *Frontiers in Systems Neuroscience*, 9:89

Abstract

Understanding the sources of variability in single-neuron spiking responses is an important open problem for the theory of neural coding. This variability is thought to result primarily from spontaneous collective dynamics in neuronal networks. Here, we investigate how well collective dynamics reflected in motor cortex local field potentials (LFPs) can account for spiking variability during motor behavior. Neural activity was recorded via microelectrode arrays implanted in ventral and dorsal premotor and primary motor cortices of non-human primates performing naturalistic 3-D reaching and grasping actions. Point process models were used to quantify how well LFP features accounted for spiking variability not explained by the measured 3-D reach and grasp kinematics. LFP features included the instantaneous magnitude, phase and analytic-signal components of narrow band-pass filtered (δ , θ , α , β) LFPs, and analytic signal and amplitude envelope features in higher-frequency bands. Multiband LFP features predicted single-neuron spiking (1 ms resolution) with substantial accuracy as assessed via receiver operating characteristic (ROC) curve analysis. Notably, however, models including both LFP and kinematics features displayed marginal improvement over kinematics-only models. Furthermore, the small predictive information added by LFP features to kinematic models was redundant to information available in fast-timescale (<100 ms) spiking history. Overall, information in multiband LFP features, although predictive of single-neuron spiking during movement execution, was redundant to information available in movement parameters and spiking history. Our findings suggest that, during movement execution, collective dynamics reflected in motor cortex LFPs primarily relate to sensorimotor processes directly controlling movement output, adding little explanatory power to variability not accounted by movement parameters.

2.1 Introduction

The variability of neuronal responses at the level of single-neuron spiking is a fundamental problem in neuroscience (Shadlen and Newsome, 1998; Churchland et al., 2010; Churchland and Abbott, 2012). Neuronal responses in neocortex to repeated stimuli presentation or behavioral tasks show substantial variability. Determining the sources of this variability is particularly important for understanding encoding of stimuli and behavioral parameters in neuronal ensembles. The issue is also critical for the development of reliable brain machine interfaces for the restoration of movement, communication, and sensory function in people with sensorimotor impairments (e.g. Hochberg et al., 2012). Beyond intrinsic stochasticity due to, for example, thermal noise and synaptic release failure (Faisal et al., 2008), variability in cortical neural responses has been proposed to arise from fluctuations in spontaneous, ongoing neural dynamics (Arieli et al., 1996; Worgotter et al., 1998; Truccolo et al., 2002; Carandini, 2004). Although often neglected, spontaneous and ongoing neural dynamics are likely to affect how neurons respond to sensory inputs or even how they modulate their activity during behavior. In this way, spontaneous neural dynamics can provide a background of contextual effects which otherwise may appear as spiking variability due to noise (Fiser et al., 2004; Goris et al., 2014; Hermes et al., 2012).

Local field potential (LFP) oscillations in different frequency bands result, to a large extent, from ongoing collective dynamics, i.e. modes of coordinated or coherent activity in neuronal populations (Buzsaki et al., 2012; Nunez and Srinivasan, 2006). Previous studies have investigated how features in multiband LFP oscillations relate to sensory stimuli and behavior and how decoding based on LFP features compares to decoding based on spiking activity (Belitski et al., 2008; Bansal et al., 2012). Additionally, some studies have examined how well LFP features predict spiking activity (e.g. Rasch et al., 2008; Montemurro et al., 2008; Kayser et al., 2009; Kelly et al., 2010; Haslinger et al., 2006). However, most of these studies have focused on early sensory cortices and most have been conducted during anesthesia, a neural state distinct from alert and active behavior. More importantly, none

of the above studies have addressed how well ongoing collective dynamics reflected in LFPs account for single-neuron spiking variability that is not explained by behavioral parameters (e.g. movement kinematics). For example, the result in Bansal et al. (2012) showing LFP and spiking activity are redundant with respect to decoding kinematics does not address the issue of excess variability in single-neuron spiking.

Here, we examined and quantified how well features in multiband LFP oscillations account for single-neuron variability not explained by behavioral parameters in a naturalistic 3-D reach and grasp task performed by rhesus macaques (*Macaca mulatta*). The behavioral task elicited diverse reaching and grasping kinematics, and included reaching to grasp different objects with different styles of grip.

We examine various LFP features, such as the amplitude envelope and phase, of several LFP bands. The frequency bands included low (<2 Hz; delta) frequency components that are common in this task, including primarily motor related responses associated with these continuously performed sequences of reach and grasp actions (Bansal et al., 2011). In addition, these low frequency signals tend to be highly correlated to the neuronal population spike count (e.g. Bansal et al., 2012, Fig. 1). Other bands included the theta (2–7 Hz), alpha (7–15 Hz), beta (15–30 Hz), gamma (30–60 Hz), high gamma (60–100 Hz), and higher frequency bands (100–200 Hz and 200–400 Hz). Although beta band oscillations dominate motor cortex LFP activity during movement preparation (Jasper and Penfield, 1949; Sanes and Donoghue, 1993; Murthy and Fetz, 1992; Broveli et al., 2004; Rubino et al., 2006), they are also characteristic during execution of isometric force tasks and other steady state conditions (Baker et al., 2001) and, less frequently, during the execution of reach and grasp actions (Reimer and Hatsopoulos, 2010). Neocortical LFP activity in higher (>100 Hz) frequency bands is generally considered to reflect fluctuations in multi-unit activity resulting from coordinated activity in neuronal populations (Buzsaki et al., 2012; Zhuang et al., 2010; Bansal et al., 2012; Scheffer-Teixeira et al., 2013), and thus is also a reflection of ongoing collective dynamics.

We quantified the amount of variability accounted for by LFP features by fitting point process models (Truccolo et al., 2005; Truccolo et al., 2010) in which the conditional intensity function (‘instantaneous’ conditional spiking rate) was modeled as a function of covariates, including the ongoing LFP features mentioned above. To assess the amount of spiking variability explained by various models, we compare the relative predictive power of LFP features, reach and grasp kinematics, intrinsic spiking history, and combinations of these covariates, using receiver operating characteristic (ROC) curve analysis.

2.2 Methods

Behavioral task. Data from three male rhesus macaques were examined in this study, C, R, and S. Data from C have been used previously in Vargas-Irwin et al. (2010). All experimental procedures were conducted as approved by the local Institutional Animal Care and Use Committee (IACUC). This study employed a task termed Free-Reach and Grasp (FRG). In FRG, an experimenter pseudo-randomly swings within the monkey’s reach one of various small (3–5 cm) objects of differing shapes attached to a string, and the monkey is rewarded for grasping the object. The FRG task was designed to elicit naturalistic and continuous three-dimensional reach and grasp behaviors that require on-line motor control (Vargas-Irwin et al., 2010). Reach and grasp actions were organized in blocks within each experimental session. We analyze data from the entire FRG block, and so our data contains a diversity of behavioral conditions, including visually guided reaching (including online corrections as the objects’ movement is unpredictable), the grasping and holding of the object, and the period of juice reward (see Vargas Irwin et al. 2010).

Kinematic feature extraction. Arm and hand kinematics were recorded at 240 frames per second using a Vicon motion capture system (Vicon Motion Systems; Oxford Metrics Group) as detailed in Vargas-Irwin et al. (2010). This system employs infrared-reflective markers to track arm and hand positions in real time, and is capable of inferring

missing data from briefly occluded markers. For our analysis, we focus on the 3-D kinematics measured at the wrist, as well as the distance between the thumb and forefinger (grasp aperture), as indicators of the kinematics related to reaching (i.e. hand position in space) and grasping. We analyzed kinematics features similar to those used in Hatsopoulos et al. (2007). These are normalized velocity trajectories of both the wrist endpoint and grip aperture in time, combined with zero lag position and speed. For comparison, we also analyze position trajectories of these markers over time. The velocities of motion capture markers were estimated using a Savitzky-Golay filter generated by fitting a 5th order polynomial to a discrete differentiation operator at the sampling rate of the kinematics. The polynomial extended 25 samples (10 ms) to either side of the current time-point. The polynomial order was selected such that frequencies higher than 20Hz were attenuated, so that the resultant velocity trajectories were sufficiently smooth to down-sample at 40 samples-per-second. Velocity trajectories were sampled from the smoothed velocity every 25 ms, starting 100 ms before the current time-point and extending 300 ms into the future. Similarly, a smoothed position estimate was extracted using Savitzky-Golay filters based on a 5th order polynomial fit to a discrete impulse. For the normalized velocity feature set, we followed the steps in Hatsopoulos et al. (2007), i.e. velocities were normalized by the L2 norm of the velocity trajectory. 3-D wrist position and grip aperture were normalized separately. The average speed and position over the trajectories were added as additional features. A separate feature set based on position trajectories was used for comparison. Position trajectories were taken from the position estimates, sampled at 40 samples-per-second, covering the same period (100 ms before to 300 ms after) as the normalized velocity trajectories.

Neural recordings. Recordings were made using microelectrode arrays (Blackrock Microsystems), as previously described in Vargas-Irwin et al. (2010). Electrodes were 1.5 mm long, likely targeting layer V of motor cortex. Monkey G was implanted with 4x4 mm 96-microelectrode arrays in both M1 and PMv. Two monkeys (R,S) were implanted

with one 96-microelectrode array in PMv, and a 3.2x2.4 mm 48-microelectrode array in each of M1 and PMd areas. Electric potential signals were recorded broadband (analog band-pass filtered to 0.3 Hz – 7.5 kHz; digitized at 16bit and sampled at 30 kilosamples per second) using a Cerebus Data Acquisition System (Blackrock Microsystems). For spike detection, recorded signals were digitally filtered with a 250 Hz fourth-order high-pass Butterworth filter. For each electrode, candidate spikes (action potentials) were identified online via threshold crossing detection in the amplitude of the high-pass filtered signal (Cerebus Data Acquisition System). Preliminary spike sorting was performed by a custom automated spike sorter (Vargas-Irwin and Donoghue, 2007), and verified using the commercial Plexon Offline Sorter (Plexon Inc.). Candidate units to be included in the analysis had a minimum signal-to-noise ratio (SNR) of 3.0 (defined as in Vargas-Irwin and Donoghue, 2007). In addition, we required that (a) the inter-spike-interval (ISI) histogram display a clear refractory period to exclude potential multi-unit clusters; (b) that there be at least 500 inter-spike-interval events smaller than 100ms within the task data, to provide accurate estimates of spike-history filters; and (c) that units be clearly separated into different clusters in the PCA feature space. Electrodes exhibiting cross-talk or excessive noise were excluded from analysis. For monkey C, LFPs were extracted during recording sessions from the broadband signal using a causal 500 Hz fourth-order low-pass Butterworth filter, and stored at two kilosamples per second. LFP data from monkeys R and S were filtered off-line to match this processing.

LFP feature extraction. For analysis of band-limited LFP oscillations, LFP signals were filtered using causal (forward only) fourth-order Butterworth low-pass and band-pass filters, with cutoff frequencies 0.3–2 Hz (δ), 2–7 Hz (θ), 7–15 Hz (α), 15–30 Hz (β), 30–60 Hz (γ), 60–100 Hz (high γ), 100–200 Hz (MUA1), 200–400 Hz (MUA2). The 0.3–2 Hz low-pass signal captures slow motor-evoked potentials (MEPs). The two highest-frequency bands are likely to reflect a substantial contribution from multi-unit activity (MUA) to neocortical LFPs, as well as other possible high-frequency source signals. For

the narrow delta, theta, alpha, and beta frequency bands, we considered four features: the instantaneous phase and amplitude of the analytic signal, and the real and imaginary component of the analytic signal. The LFP analytic signal was computed from the band-pass filtered LFP using the Hilbert transform. LFP instantaneous phase and amplitude were computed as the complex argument and modulus of the analytic signal, respectively. (The real component of the analytical signal corresponds to the band-pass LFP itself.) For the broader, higher frequency gamma and multi-unit bands, we use only the analytic signal and the amplitude envelope. Feature extraction was performed on the LFP sampled at 2 kilosamples per second, and decimated to 1 kilosample per second for neural point process modeling.

Spike contamination. In this analysis, we predict single-unit spiking from features of the LFP recorded on the same electrode as the isolated unit. Because of this, when predicting neuronal spiking from LFP features, it is important to prevent action potentials from contaminating the filtered LFP. We elected not to use existing spike removal procedures (e.g. Zanos et al. 2012) because the broadband LFP data were unavailable for monkey C, and because spike-removal methods make implicit assumptions about which features of the LFP relate to the spike waveform as opposed to collective dynamics locked to spiking. Instead, we employed causal filtering to extract LFP features. Although spike contamination can occur as low as 10 Hz (Waldert et al. 2013), causal filtering restricts this contamination to times following a spike, avoiding the confound of predicting spikes from themselves (i.e. via their contamination of the LFP features). Although the discrete Hilbert transform used to compute phase and amplitude features is non-causal, the effective filters created by the composition of the Hilbert transform with the causal band-pass filters remain predominantly causal. As a further precaution, and to guard against imprecision in localizing spike times, we added 1 ms delay to LFP features. Under this approach, the noncausal contribution to the imaginary component of the analytic signal was negligible: less than 0.14% of the impulse response (measured as the percentage of the area

under the absolute impulse response) was non-causal. Since causal filters can add amplitude and phase distortions, we addressed this concern by comparing the predictive power of causally filtered LFP and that of zero-phase (non-causal) filtered LFP, which contains no delay. We determined that the choice of causal versus zero-phase filtering did not alter the conclusions of this paper for frequencies below 30 Hz. Zero-phase filtering for higher frequencies resulted in higher predictive power in some cases, which was likely the result of action potential contamination as supported by the recent studies mentioned above.

Intrinsic spiking history. To assess the extent to which a neuron’s own spiking history explains spiking variability, and to compare its predictive power to that of kinematic and LFP features, we included features of spiking history in our modeling (Truccolo et al., 2005). In addition to intrinsic biophysical processes (refractory/recovery period, bursting, etc.), spiking history can potentially also reflect indirect neuronal network dynamics effects. For example, spiking history models are capable of capturing spiking rhythmicity that may arise as a result of oscillatory input. We used raised cosine bases in logarithmically scaled time, covering the past one-hundred milliseconds of spiking activity, to estimate temporal filters (see Truccolo et al., 2010, and Pillow et al., 2008, for more details). The resulting temporal filters were convolved with the past spiking activity to capture history effects on the spiking probability at a given time. Ten basis functions were used. More recent spiking history, typically related to after-spike refractory and recovery periods, and bursting, was modeled with more localized (finer temporal resolution) basis functions. Longer-term history effects can capture intrinsic rhythmicity and also, implicitly, network dynamics.

Stochastic neural point process models. We used a generalized linear point process model (Truccolo et al., 2005) to explore the extent to which different covariates explain spiking variability. The probability of a neuron spiking in a sufficiently small time

interval, indexed by t , of duration Δ , can be written as

$$\Pr(Y_t = 1|\lambda_t) = \lambda_t\Delta + o(\Delta), \quad (2.1)$$

where Y_t corresponds to the spiking activity at time t , $Y_t=1$ for a spike, 0 otherwise, and λ_t is the conditional intensity function (‘instantaneous conditional spiking rate’, in spikes per second) of the modeled neuron. The bin size Δ must be chosen small enough such that the probability of two spikes occurring within the same time bin is negligible. Here $\Delta=1$ ms. We used a regularized maximum likelihood approach to model the logarithm of the conditional intensity function as a linear combination of model features:

$$\ln(\lambda_t) = \mu + A \cdot X_t, \quad (2.2)$$

where X_t is the covariate vector at time t , A is a vector of model parameters, and μ is a parameter related to background activity level. X_t can refer to LFP features at time t , past and future kinematics, convolutions of intrinsic spiking history up to but not including time t with temporal filters, or combinations of these covariates. For example, for a given Hilbert-transform of an LFP band, $z(t)$, the feature vector

$$X_t = (|z_t|, \text{Re}(z_t), \text{Im}(z_t), \cos(\text{Arg}(z)), \sin(\text{Arg}(z))) \quad (2.3)$$

corresponds to a model with cosine tuning to a preferred Hilbert phase θ_0 , as well as amplitude envelope and analytic signal features, i.e.

$$\ln(\lambda_t) = \mu + a_1|z_t| + a_2 \text{Re}(z_t) + a_3 \text{Im}(z_t) + a_4 \cos(\theta_0 - \text{Arg}(z)) \quad (2.4)$$

Model fitting. Model estimation is solved by finding parameters A and μ that maximize the L2-regularized log-likelihood of the observed spiking activity (Truccolo et al.,

2005):

$$\operatorname{argmax}_{A,\mu}[\ln(\Pr(Y|X,A,\mu))] = \frac{1}{T} \sum_{t=1}^T [Y_t \ln(\lambda_t \Delta) - \lambda_t \Delta] - \alpha \|A\|^2, \quad (2.5)$$

where α is a penalty or regularization parameter. The log-likelihood is normalized by the number of samples T so that the strength of regularization does not depend on the amount of data. The parameter μ is not penalized. All features are z-scored prior to model fitting to ensure that all features are zero mean and of comparable scale, which ensures that the L2 penalty is applied equally to all features and improves numerical accuracy.

We used a gradient descent approach for the minimization of the negative log-likelihood under L2 regularization. Models were fit under a two-tier cross-validation scheme. An outer level of 10-fold cross-validation ensures that results are not over-fit. An inner level of cross-validation selects the regularization parameter α . Ten values of the regularization parameter α , base-10 logarithmically spaced between $1e-9$ and $1e2$ inclusive, as well as $\alpha=0$, were tested. On each of the 10 outer-level cross-validations, 90% of the data were taken as training data, and 10% were reserved for testing. The training data were split randomly into two equal groups. For each group, models were generated for each value of the regularization parameter α . The value of the regularization parameter that led to the best generalization (in terms of predictive power, see below) in this internal cross-validation step was selected for fitting a model on all of the training data. This model was then validated on the remaining 10% of the data that had been withheld for testing. This two-tier cross-validation procedure was repeated 10 times, such that all of the available data was used for model validation and assessing predictive power. To confirm that L2 regularization sufficiently prevented over-fitting when adding LFP features to the kinematics model, we shuffled LFP features in 100ms blocks with respect to the spiking activity. We found that adding these non-informative features to the kinematics and kinematics-history combined models reduces the predictive power very little, by at most 0.03, and with the population mean decrease ranging from 0.001 and 0.006 across all sessions. This difference is too small to alter the conclusions of our study. In preliminary

analysis, we also explored L1 regularization and also a mixed L1-L2 regularization (fitted via elastic net, Friedman et al. 2010). We found that L2 regularization outperformed these alternatives, both in terms of computational time and predictive power under generalization. Additionally, we found that L2 regularized generalized linear models (GLMs) outperformed simpler approaches, such as naive Bayes and spike-triggered average and covariance analysis.

Assessment of predictive power. Model performance was evaluated using the area under the ROC convex hull (AUC) measured on testing data, using the model (Equation 2) to compute the conditional spiking probability, $\Pr(Y_t=1|X_t, A, \mu) \approx \lambda_t(X_t, A, \mu)\Delta$, from observed covariates (Truccolo et al., 2010). ROC analysis assesses predictive power in the context of binary (in this case spike train) sequences (Fawcett, 2006). We report predictive power (PP) as $2 \times AUC - 1$, which ranges from 0 (no predictive power) to 1 (complete prediction of single-neuron spikes in 1 ms time bins). Note that this predictive power measure is based on both true and false positive rates, since is derived from the ROC analysis.

2.3 Results

We are interested in how well collective neural dynamics, as reflected by features in ongoing and evoked multi-band LFP signals, can explain motor cortex single-neuron spiking variability not accounted for by motor behavioral covariates such as reach and grasp kinematics. We first demonstrate that the examined LFP features can predict single-neuron spiking in motor cortex, then we assess the extent to which this predictive power compares and is redundant to information available in 3-D kinematics. We also assess the extent to which intrinsic spiking history, i.e. temporal dynamics or correlation in the modeled spiking activity itself, adds predictive power to kinematics, and evaluate whether LFP features remain predictive conditioned on kinematics and intrinsic spiking history.

Datasets from seven experimental sessions were used in these analyses: two each from monkey C and R, and three from monkey S. Sessions from monkeys R and S were collected within a week of each-other. The two sessions from monkey C were collected three months apart. Between three and nine reach and grasp blocks, averaging 140 seconds long, were collected on each session. Each session included 15–42 successful free-reach-to-grasp trials or reaches in each block. This yielded 7–17 minutes of FRG task data, averaging ten minutes of data per session. A detailed statistical description of the kinematics and examples of kinematic trajectories in these experimental blocks can be found in Vargas-Irwin et al. (2010) and Bansal et al. (2012), who reported some of the data from monkey C in this task. An example of kinematics trajectories with the corresponding neuronal ensemble spike raster is shown in Figure 2.1. For each array in each session, between 19 and 83 well-isolated units were identified for analysis (mean=52, $\sigma=16$). For a given monkey and area, some of the neurons are thought to be the same across sessions, for this reason we do not combine sessions when we perform statistical significance tests, as this could falsely inflate p-values due to repeated sampling of the same neurons.

2.3.1 Features of LFP oscillations predict single-neuron spiking with substantial power

First, we evaluated the ability for multiband LFP features to predict single-neuron spiking. We fit a regularized generalized linear model to predict single-unit spiking (1 ms time resolution) from multiband LFP features. Spiking probability at any given 1 ms time interval was modeled as a function (Equations 1-2, Methods) of features of ongoing LFP activity. LFP features included instantaneous phase and amplitude envelope as well as the analytic signal, extracted via a Hilbert transform, from four narrow LFP bands (Methods), δ (0.3-2 Hz, motor related potentials), θ (2–7 Hz), α (7–15 Hz), β (15–30 Hz), as well as the amplitude envelop and analytic signal for four broader, higher frequency bands: γ_1 (30–60 Hz), γ_2 (60–100 Hz), and two multi-unit activity (MUA) bands MUA₁ (100–200 Hz),

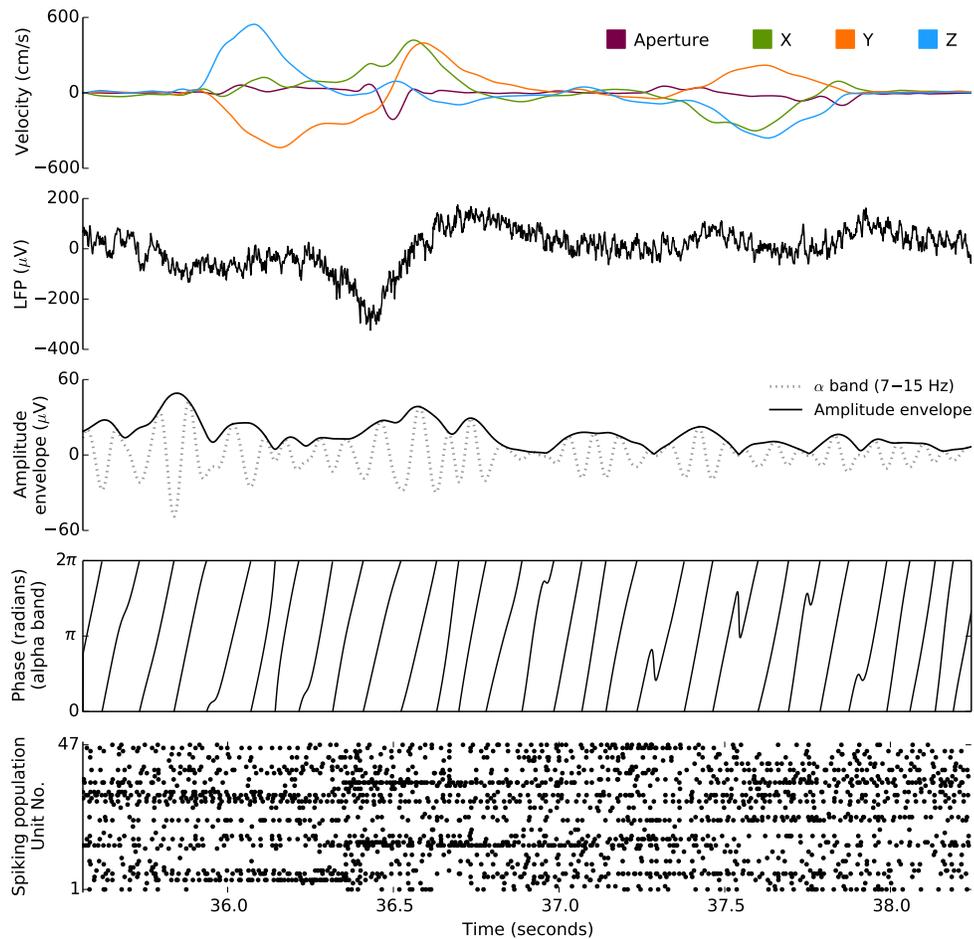


Figure 2.1: Behavioral and neural signals during the free-reach and grasp task. Kinematic, spiking, and LFP data from a single trial of the free reach to grasp task (monkey S, area M1). (a) Velocity of 3D wrist endpoint ('X', 'Y', 'Z') and distance between the thumb and forefinger ('Aperture'). (b) Broadband LFP, low-pass filtered at 500 Hz during two reach and grasp movements. (c) Instantaneous amplitude envelope (bold trace), via Hilbert transform, from the bandpass filtered LFP (grey trace). This particular example is filtered in the 7–15 Hz range. (d) Corresponding instantaneous phase extracted from the same bandpass filtered signal as in (c). (e) Spiking population raster during this trial: spikes from 47 units are plotted along the y axis. Neurons show different task-related modulations.

and MUA₂ (200–400 Hz). Figure 2.1 illustrates the signal processing steps involved in the computation of the instantaneous phase and magnitude for a single frequency band.

We report the extent to which a model explains spiking variability in terms of "predictive power" (PP). Predictive power is the normalized area under the receiver operator characteristic (ROC) curve such that 0 corresponds to chance level and 1 to perfect prediction of spike times at 1ms resolution (see Methods: Model fitting; Methods: Assessment of predictive power). Figure 2.2 shows PPs obtained from three example neurons from different monkeys and areas. For illustration and comparison, we also show the corresponding PPs based on a model including only kinematics features related to lagged 3-D velocity and position (similar to 'pathlets' in Hatsopoulos et al. 2007) and grasp aperture (Methods). The examples show a case (Figure 2.2, left) in which LFP features and kinematics both explained a substantial fraction of spiking variability (PP = 0.73, 0.75, respectively) and two other examples in which LFP features did better and worse than kinematics, respectively. Overall, we found that LFP features were typically predictive, and that for some neurons LFP accounted for a substantial (i.e. PP>0.5) fraction of variability. We performed a permutation test to assess chance level LFP predictive power by shuffling LFP features in 100ms blocks relative to spiking, and found that the 95% chance predictive power ranged from 0.03 to 0.06 across sessions. Across sessions and areas, between 85% and 100% of units showed LFP predictive power higher than this chance level. As shown in Figure 2.3, high predictive power from LFP was consistent across all monkeys, motor cortical areas (PMv, PMd and M1) and sessions. This finding demonstrates that collective dynamics reflected in ongoing and evoked LFP oscillations can account for a substantial fraction of single-neuron spiking variability.

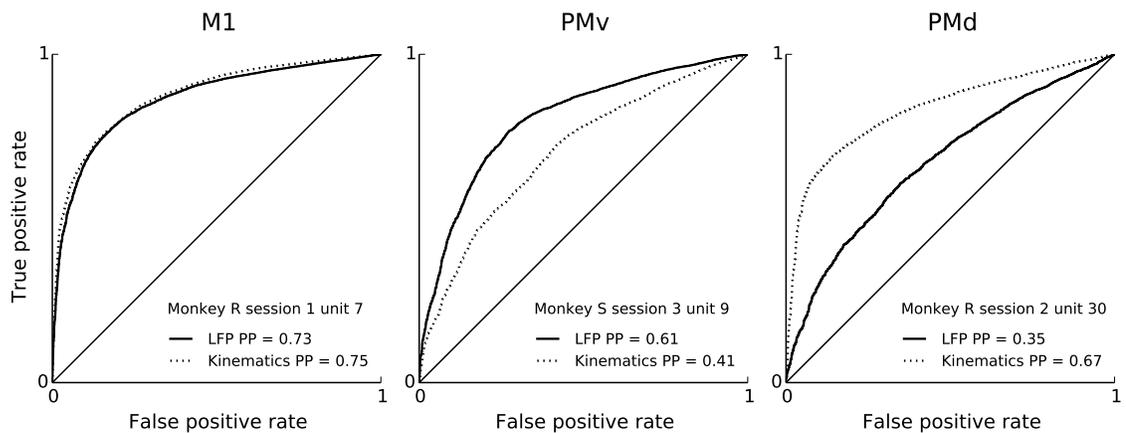


Figure 2.2: Features of ongoing and evoked multiband LFP oscillations predict single-neuron spiking: examples. ROC curves (solid) for neural point process models based on multiband LFP features (Methods: LFP feature extraction). Examples correspond to different neurons, areas, monkeys and sessions. From left to right: M1, PMv, and PMd. For comparison, ROC curves (dashed) corresponding to models based on kinematics features are shown. Both LFP features and kinematic features achieve substantial single-unit spiking predictive power. LFP features include the instantaneous phase, amplitude envelope, and analytic signal, in the $\delta=0.3-2$ Hz, $\theta=2-7$ Hz, $\alpha=7-15$ Hz, and $\beta=15-30$ Hz LFP bands, as well as amplitude envelope and analytic signal in the $\gamma_1=30-60$ Hz, $\gamma_2=60-100$ Hz, $MUA_1=100-200$ Hz, and $MUA_2=200-400$ Hz bands (see Methods: LFP feature extraction). Predictive power is the area under the ROC curve normalized, i.e. $2 \times AUC - 1$, such that 0 is chance level and 1 is perfect prediction. Predictive power was evaluated under 10-fold cross-validation.

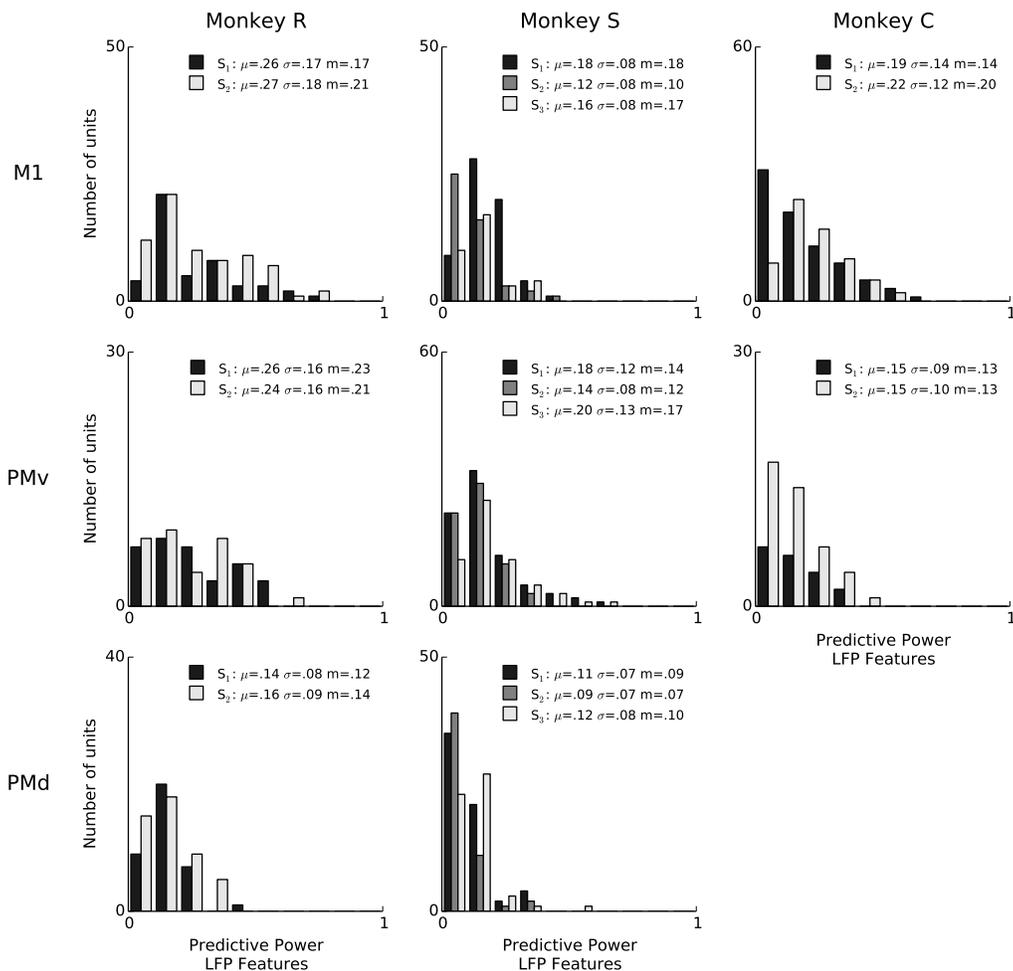


Figure 2.3: Features of ongoing and evoked multiband LFP oscillations predict single-neuron spiking: summary across animals and areas. Histogram counts of the predictive power of point-process models based on LFP features, for all isolated units. LFP features include the instantaneous phase, amplitude envelope, and analytic signal, in the $\delta=0.3-2$ Hz, $\theta=2-7$ Hz, $\alpha=7-15$ Hz, and $\beta=15-30$ Hz LFP bands, as well as amplitude envelope and analytic signal in the $\gamma_1=30-60$ Hz, $\gamma_2=60-100$ Hz, $MUA_1=100-200$ Hz, and $MUA_2=200-400$ Hz bands (see methods: LFP feature extraction). LFP was consistently predictive of spiking variability for a subset of neurons in all sessions. In the plot legends, “S” indicates the session, μ is the mean of each distribution, m is the median, and σ is standard deviation. Sessions have different numbers of units, and the differences in bar height also reflect differences in sample size (e.g. monkey C area PMv).

2.3.2 LFP features contributing to prediction of single-neuron spiking

We examined whether some of the multiband LFP features contribute more to prediction of single neuron spiking than others. Analysis based on estimated model coefficients is complicated because of the nonlinear (multiplicative) interactions between different features (amplitude, phase, or analytical signal) in different frequency bands. Instead, we performed an analysis based on fitting a single model for each feature separately and assessing how well each separate model and feature predicted spiking. This allowed easy visualization of the predictive power of each individual LFP feature.

This analysis revealed some trends common to all animals and motor cortex areas, but also some variations (Figure 2.4). Consistently across animals and areas, low-frequency local field potentials (δ , 0.3–2 Hz) showed predictive power in the time domain signals and phases, but not amplitude envelopes. Additionally, the amplitude envelope in the multi-unit (100–200 Hz and 200–400 Hz) bands was predictive, more-so than the signal. The low-frequency <2 Hz analytic signal was the most predictive for 49% of units (486 units out of 991), and the instantaneous amplitude envelope and phase the most predictive for 13% and 8% of units (133 and 82 units out of 991), respectively. The amplitude envelope in the 200–400 Hz band was the most predictive for 14% of units (142 units out of 991). Features from intermediate 2–100 Hz bands generally performed poorly, with the exception of beta (15–30 Hz) amplitude, which although less predictive than the aforementioned features, was still amongst the top 4 most predictive features for 23% of units (227 units out of 991). The finding that LFP amplitude was predictive for the beta-frequency LFP was strongest in monkey R for areas M1 and PMv. To understand in more detail how the predictive power in beta amplitude varies across monkeys and areas, we examined the distribution of the model parameter weights for beta amplitude (in the case of the amplitude-only model, the parameter matrix A in Equation 2 is simply a single scalar parameter). Model weights

for beta amplitude in monkey R areas M1 and PMv were more negative (mean \pm 2SD for M1 and PMv were -0.18 ± 0.4 and -0.23 ± 0.4) than those from other monkeys and areas (-0.04 ± 0.16), indicating that a reduction in beta amplitude is typically associated with an increase in firing rate.

2.3.3 Predictive power of kinematics during naturalistic reach and grasp movements

We next quantified the predictive power of motor behavior, specifically kinematic features of the 3-D reach and grasp movements. We found that kinematics trajectories can also predict single-neuron spiking with substantial accuracy, at times achieving predictive power levels around 0.8 (e.g. Figure 2.5, monkey R area M1). However, similarly to LFP features, there was considerable diversity in the extent to which kinematics predicted spiking, with some units being predicted poorly. Similar results were obtained by using position trajectories (i.e. position at multiple time lags with respect to spiking; Methods). The 95% chance level predictive power for kinematics ranged from 0.03 to 0.08, as assessed by shuffling kinematics features in 100ms blocks relative to spiking. Across sessions and areas, between 49% and 100% of units showed LFP predictive power higher than this chance level. These effects were consistent across all animals, sessions, and motor areas (Figure 2.5), with mean predictive power ranging between 0.16 to 0.36 across sessions and areas. The fact that the task kinematics predict single-unit spiking variability confirms that we are recording from motor cortex populations that exhibit task-modulation and tuning to motor output.

Figure 2.6 directly compares the predictive power of LFP and kinematics features. Overall, the predictive power of LFP features was typically less than that of kinematics during this free reach and grasp task: the difference between the predictive power of models based on kinematics and LFP features ranged from -0.20 to 0.45. With exception of monkey S area PMv, units for which LFP features explain more variability than kinematics

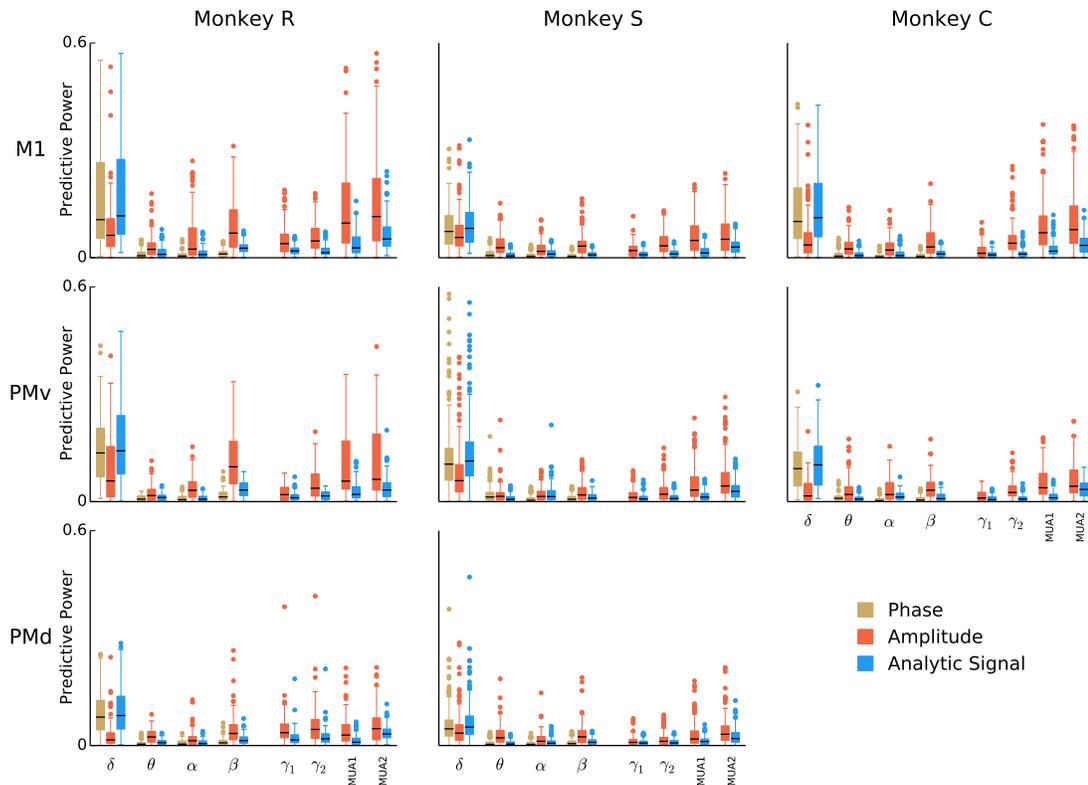


Figure 2.4: Breakdown of LFP predictive power by frequency band and LFP feature. Box-plots over the population of isolated units (all sessions combined) showing the predictive power of models based on phase, amplitude, or analytic signal features in isolation from each of eight bands. To better assess the individual predictive power of each LFP feature, models were fitted for each feature separately. Certain features, such as the instantaneous phase and analytic signal for the 0.3 – 2 Hz band, as well as the analytic signal amplitude modulation above 100 Hz, consistently predict spiking across all animals and areas.

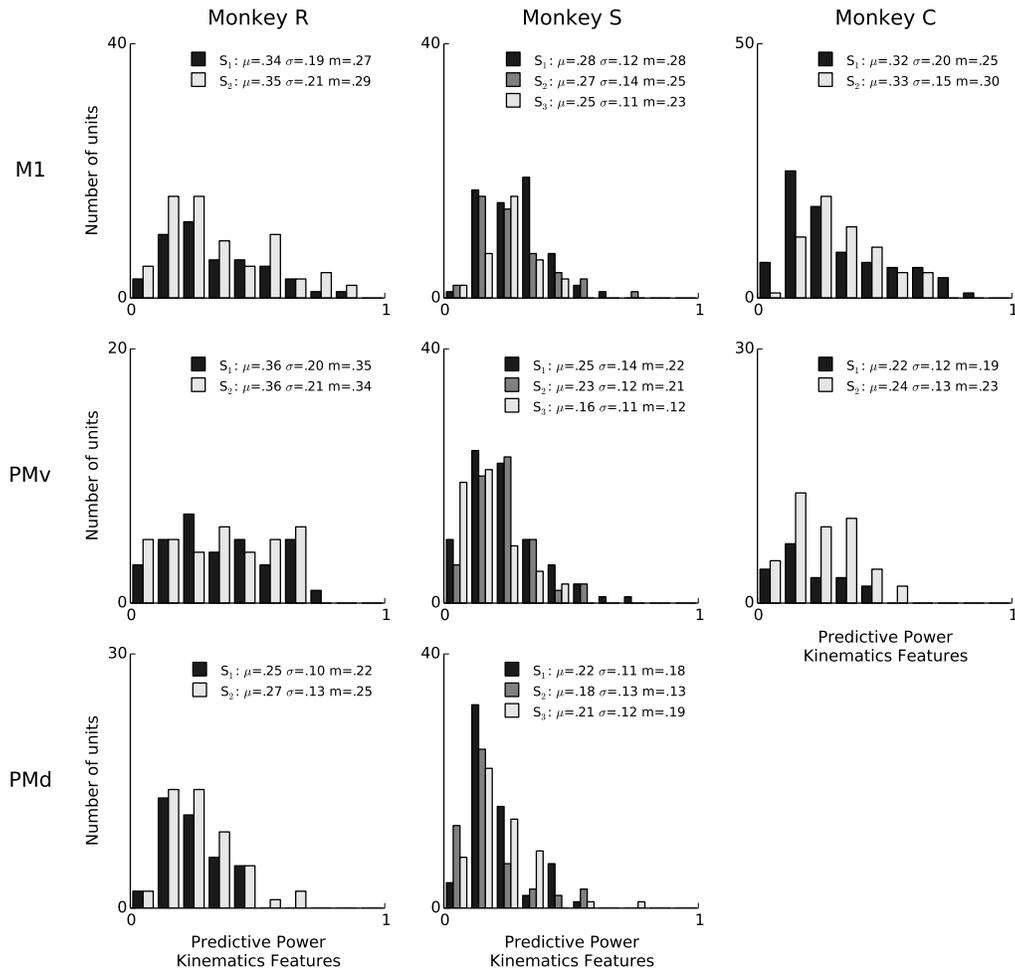


Figure 2.5: Kinematic features predictive power for single-neuron spiking: summary across animals and areas. Histogram counts of the predictive power of point-process models based on kinematics features, for all isolated units. Kinematics features are normalized velocity trajectories of wrist endpoint and grip aperture, extending from 100 ms in the past to 300 ms in the future, sampled every 25 ms, as well as the average speed and zero-lag position for wrist endpoint and grip aperture (Methods: Kinematic feature extraction). In the plot legends, “S” indicates the session, μ is the mean of each distribution, m is the median, and σ is standard deviation. Sessions have different numbers of units, and the differences in bar height also reflect differences in sample size (e.g. monkey C area PMv).

are rare. The mean difference in predictive power within each session ranged from -0.04 to 0.14, and the median difference from -0.02 to 0.12, with all (session, area) pairs except monkey S area PMv session 3 displaying significantly better median predictive power for kinematics. (Wilcoxon signed-rank test with $p < 0.05$, corrected for multiple comparisons using Bonferroni correction for 19 (session, area) pairs.) Furthermore, across all monkeys and areas, predictive power of LFP was highly correlated with that of kinematics: the Pearson correlation coefficient between the predictive power of kinematics and LFP features ranged from 0.52 to 0.96, with a mean of 0.86 and a median of 0.88. This raises the possibility that LFP and kinematics shared some common effect, which we address below.

2.3.4 LFP features are mostly redundant to kinematics when explaining single-neuron spiking variability

In order to determine whether LFP features can account for single-neuron spiking variability not explained by kinematics, we asked whether the predictive information carried in the examined LFP features about single-neuron spiking variability was redundant to the predictive information carried in kinematic features. To assess redundancy we compare the predictive power of point process models based only on kinematics features to models that included both kinematics and LFP features. We used L2 regularization to control for overfitting to the training data due to the larger number of parameters in the models that combined both kinematics and LFP features (Methods). We found that forgoing L2 regularization led to overfitting, in which the larger number of parameters in the combined LFP-kinematics model generalized less well to the evaluation data. Tests using shuffled LFP features confirmed that the L2 regularization approach adequately prevented overfitting (Methods).

Figure 2.7 compares, on a unit-by-unit basis, the relative predictive powers of kinematics and LFP features. The analysis reveals that, with a few exceptions (e.g. some neurons in PMv in monkey S), LFP features added little predictive power to kinematics. This find-

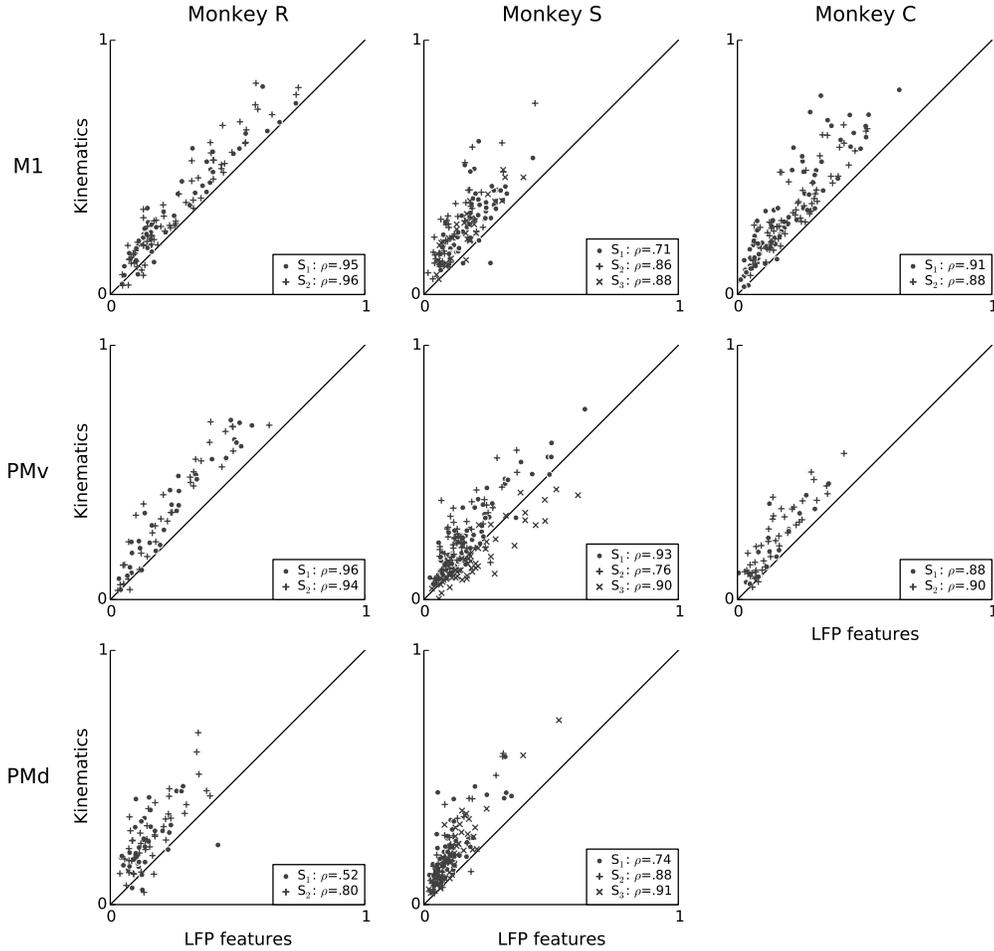


Figure 2.6: Kinematics and features of ongoing and evoked multiband LFP oscillations achieve similar predictive power on a neuron by neuron basis. Scatter plots compare the predictive power of LFP features (x-axis) to that of kinematic features (y-axis). LFP features include the instantaneous phase, amplitude envelope, and analytic signal, in the $\delta=0.3\text{--}2$ Hz, $\theta=2\text{--}7$ Hz, $\alpha=7\text{--}15$ Hz, and $\beta=15\text{--}30$ Hz LFP bands, as well as amplitude envelope and analytic signal in the $\gamma_1=30\text{--}60$ Hz, $\gamma_2=60\text{--}100$ Hz, $\text{MUA}_1=100\text{--}200$ Hz, and $\text{MUA}_2=200\text{--}400$ Hz bands (see methods: LFP feature extraction). Each data-point is a single unit from one session. The diagonal line indicates equality. Most units lie above the diagonal line, indicating that kinematic features better predict single-unit spiking variability. In the plot legends, “S” indicates the session, ρ is the Pearson correlation coefficient between the predictive power for kinematics and LFP features. Predictive power from both LFP and kinematics are highly correlated, i.e. neurons that are well predicted by LFP features tend also to be well predicted by kinematics.

ing suggests that although LFP features were able to account for a substantial fraction of spiking variability, this information was highly redundant to information available in the examined kinematics features. This finding confirms the conjecture raised earlier, based on the high correlation between LFP and kinematics predictive power (Figure 2.6), that the information available in these two signals was redundant in terms of prediction of single-neuron spiking activity. Nevertheless, for each session, after adding LFP features, the mean change in predictive power was positive, ranging from 0.008 to 0.07, and the median change in predictive power ranged from 0.006 to 0.06. This median increase was statistically significant for all sessions. (Wilcoxon signed-rank test, $p < 0.05$ corrected for multiple comparisons using Bonferroni correction for 19 (session, area) pairs).

2.3.5 Intrinsic spiking history adds substantial power to kinematics in the prediction of single-neuron spiking variability

The above findings suggest that the examined features of LFP collective dynamics during movement reflect primarily sensorimotor processes related to motor representations and computations associated with the measured reach and grasp kinematics. Nevertheless, these LFP features accounted for a small, but statistically significant, fraction of neural variability not accounted by the examined kinematics features. To investigate the potential sources of this additional predictive power, we take a detour in this section and consider first the predictive power of a neuron's own spiking history. Here we focused on the preceding 100 ms spiking history, which can capture fast intrinsic biophysical processes such as refractory and recovery periods after an action potential, and also bursting dynamics, which are common in certain types of motor cortex neurons (Chen and Fetz, 2005). In addition, temporal autocorrelations within single-neuron's spiking activity can be induced, for example, by both intrinsic rhythmicity and rhythmicity due to ongoing neuronal network dynamics affecting spiking. We used temporal filters to capture the effects of intrinsic spiking history. Temporal filters were estimated with semi-parametric

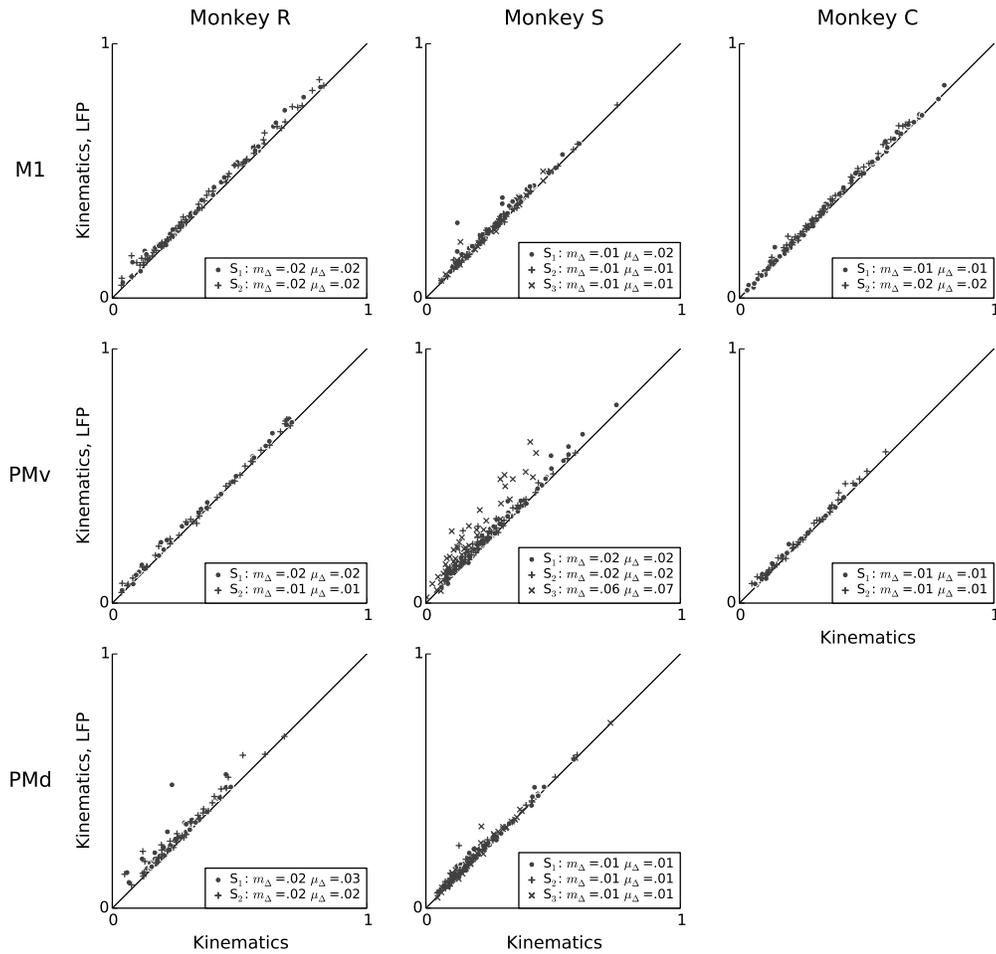


Figure 2.7: LFP features add little predictive power to a kinematics model. Scatter plots comparing, on a unit-by-unit basis, the predictive power of the kinematics model (x-axis), to the predictive power of a model that uses both LFP and kinematics features (y-axis). Although LFP features by themselves can achieve high predictive power for single-neuron spiking (Figure 2.3), their combination with kinematics typically results in only a small increase in predictive power, suggesting that most predictive information in the examined LFP features is redundant to predictive information in kinematics features. In the plot legends, “S” indicates the session, μ_{Δ} is the mean change in predictive power, and m_{Δ} is the median change in predictive power.

models using raised cosine functions (Truccolo et al., 2010; Pillow et al., 2008; Methods). Ten logarithmically-spaced raised cosine functions on the past 100 ms were used.

When information about a single neuron’s own spiking history is added to the model, there is a significant improvement in predictive power compared to a model including only kinematics features (Figure 2.8). Within each session and area, the mean increase in predictive power ranged from 0.04 to 0.17. The median increase in predictive power ranged from 0.03 to 0.16, and was statistically significant for all sessions and motor areas (Wilcoxon signed-rank test, $p < 0.05$ with Bonferroni correction for 19 (session, array) multiple tests). This result demonstrates that fast-timescale spiking history can explain variability in single-neuron spiking that is not redundant to variability examined by the kinematic features in this motor task.

2.3.6 Conditioned on spiking history, contribution of LFP features to kinematic models is further reduced

Having demonstrated that intrinsic spiking history adds predictive power to kinematic models, we finally assessed whether LFP features can account for variability in single-neuron spiking not accounted for by kinematics and intrinsic history features. Figure 2.9 shows that, across monkeys and motor areas, adding LFP features to models based on kinematics and intrinsic spiking history did not lead to a large improvement in predictive power. Across sessions and areas, the mean change in predictive power when adding LFP features to a model containing both kinematics features and intrinsic spiking history features ranged from 0.002 to 0.02, and the median change ranged from 0.001 to 0.02. Nevertheless, this median improvement was statistically significant for all but one session (monkey C area PMv session 1) (Wilcoxon signed-rank test, $p < 0.05$ with Bonferroni correction for 19 (session, array) multiple tests).

This result demonstrates that LFP predictive power that was non-redundant to kinematics was primarily redundant to information available in the recent 100 ms spiking

history in this motor task. In other words, additional single-neuron variability not explained by kinematics seems to be better explained by fast-timescale features in intrinsic spiking history than by the examined motor cortex LFP features in this reach and grasp task.

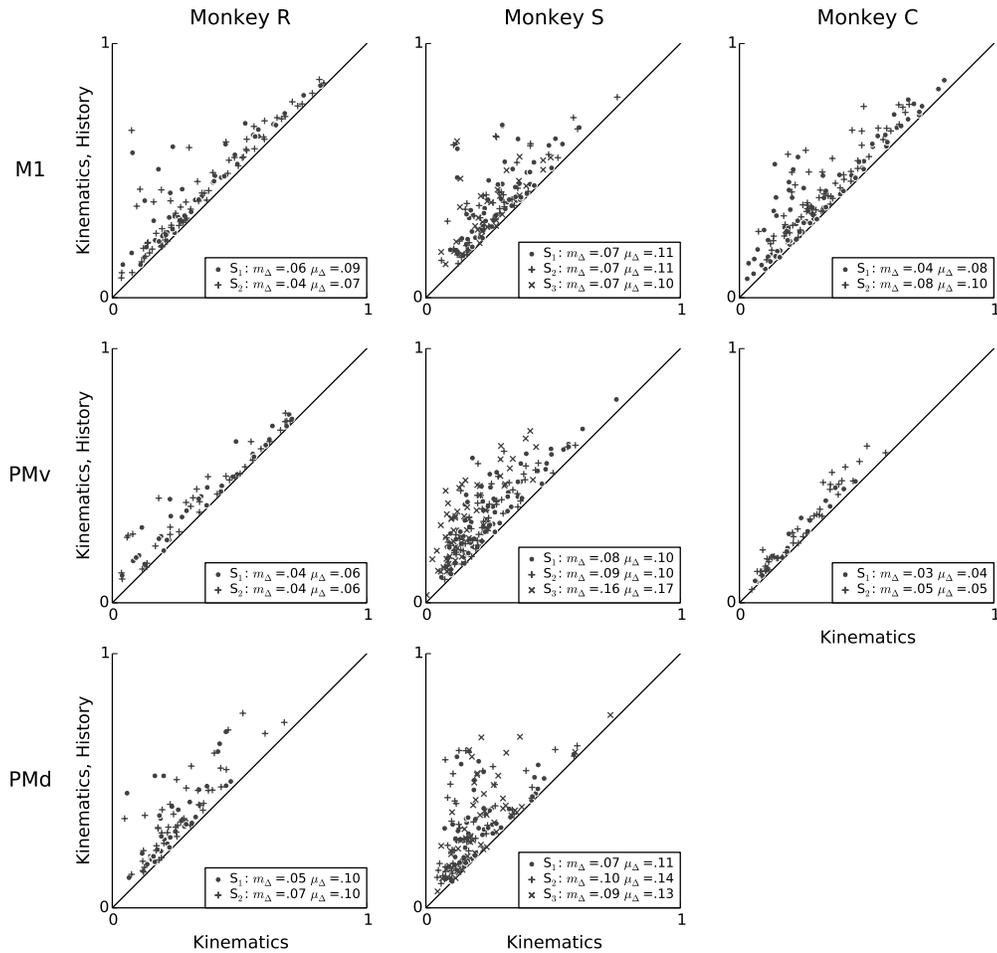


Figure 2.8: Intrinsic spiking history carries complementary information to kinematics features. Scatter-plots showing a unit-by-unit comparison of the predictive power of the kinematics model (x-axis) to that of a model that uses both kinematics features and intrinsic spiking history features (y-axis). Adding 100 ms of intrinsic spiking history information improves prediction substantially for almost all units, consistently across animals and areas. In the plot legends, “S” indicates the session, μ_{Δ} is the mean change in predictive power, and m_{Δ} is the median change in predictive power.

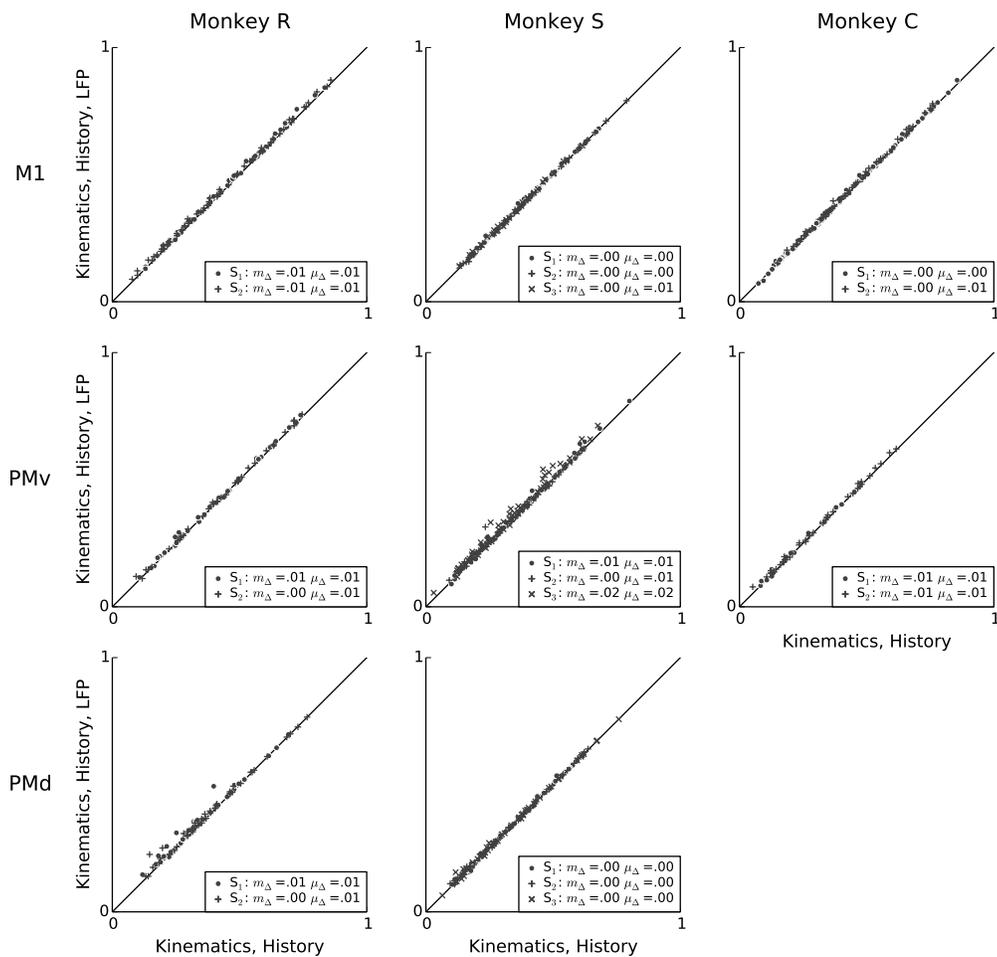


Figure 2.9: Conditioned on intrinsic spiking history, the contribution of LFP features to kinematic models is redundant. Scatter plots comparing the predictive power of a model based on kinematics and intrinsic spiking history (x-axis) to that of a model based on kinematics, intrinsic history, and LFP features (y-axis). LFP features add negligible predictive power after accounting for behavioral and intrinsic spiking history effects. In the plot legends, “S” indicates the session, μ_{Δ} is the mean change in predictive power, and m_{Δ} is the median change in predictive power.

2.4 Discussion

Neocortical neurons are embedded in large networks possessing highly recurrent connectivity. Recurrent connectivity typically leads to rich spontaneous collective dynamics. The extent to which these spontaneous dynamics contribute to single neuron variability in awake behaving primates, and how these dynamics interact with sensory inputs and behavioral outputs, is an important open question in neuroscience. Here we examined this problem in the context of collective dynamics reflected in LFP oscillations at multiple frequencies in three different areas of motor cortex in monkeys performing naturalistic 3-D reach and grasp actions. These LFPs are thought to result, to a large extent, from collective modes of activity driving spatially coherent postsynaptic potentials at multiple spatiotemporal scales (Nunez and Srivivasan, 2006; Buzsaki et al., 2012). LFP features (e.g. amplitude envelope, phase, and analytic signal) in eight different frequency bands predicted single neuron spiking (1 ms time resolution) with significant predictive power for many neurons in all of the three examined motor cortex areas (PMv, PMd and M1). Neurons for which LFP predictive power was high tended also to show high kinematics predictive power. In fact, this relationship was close to linear (Pearson correlation coefficient ranging from 0.52 to 0.96 across all the studied areas, monkeys and sessions). More importantly, predictive information in the examined LFP features was mostly redundant to the predictive information available in kinematics. In other words, models combining both LFP features and kinematics typically improved only marginally over models using only kinematics in the studied 3-D reach and grasp task. These results should not be dismissed as overfitting artifacts since they were obtained under well controlled L2 regularization aiming to preserve generalization of models with larger number of parameters. Furthermore, in the few cases for which LFP features seemed to add predictive information with respect to kinematics, this information turned out to be redundant to the information available in short term correlations in the intrinsic spiking history. Overall, our findings suggest that multiband LFP oscillations in motor cortex of alert behaving

primates, although predictive of single-neuron spiking during movement execution, are primarily related to collective dynamics controlling aspects of motor output (e.g. kinematics) rather than other potential ongoing dynamics not directly related to the task (e.g. arousal levels).

Several previous studies have looked at the relationship between single-neuron spiking and features of LFP oscillations, mostly in sensory cortices and during anesthesia (e.g. Kelly et al., 2010; Haslinger et al., 2006). Recent work by Ecker et al. (2014) has shown that previously reported high correlations between neuronal pairs and strong phase locking to ongoing LFPs in primary visual cortex during stimulation were highly dependent on the anesthesia state, with neuronal ensemble spiking becoming much more asynchronous during awake stimulation tasks. Our analysis goes beyond previous studies by examining motor cortex LFP and spiking in awake behaving non-human primates. Furthermore, to our knowledge, this is the first time that the redundancy between the information available in multiband LFP features and the information available in behavioral output (kinematics) has been systematically assessed in motor cortex. It remains to be seen how much of the residual variability is inherent to stochastic aspects of the biophysics (e.g. noise due to synaptic failure and amplification effects during spike generations; Carandini, 2004), to other motor-related covariates (e.g. torques and muscle activations) not examined in this paper, or to network dynamics not faithfully reflected in LFP features. In the latter, it is possible that the cortical layer from which the electrode tips recorded (likely layer V in our data) may impact LFP predictive power. For example, LFPs recorded from layers 2/3 of motor cortex may potentially exhibit different spike prediction performance and different levels of redundancy with respect to kinematics. In addition, we note that typically recorded LFPs might not be as ‘localized’ as previously thought (Kajikawa and Schroeder, 2011). In particular, rhythmic oscillations in electric potentials recorded intracellularly and on broad extracellular fields may share similar frequencies, and yet show very different phase-locking dynamics with respect to neuronal spiking (e.g. Harvey et al., 2009).

Thus, the broader LFP spatial average might result in signals that are less predictive of single-neuron spiking and more related to population activity.

The relationship between single-neuron spiking and ongoing LFP oscillations, in particular the locking of neuronal spiking to the phase of oscillations in specific frequency bands, might be highly dependent on the neuron types (e.g. pyramidal vs. fast spiking interneurons; Buzsaki et al. 2012). Recent work by Vigneswaran et al. (2011) has demonstrated that certain types of pyramidal neurons in primary motor and premotor cortices can show features of action potential waveforms and spiking statistics that are indistinguishable from features in inhibitory interneurons. Therefore an analysis based on such putative classification would remain highly questionable in our motor cortex data.

In our analysis, low frequency (0.3–2 Hz) and higher (>100 Hz) frequency LFP bands tended to contribute the most to prediction of neuronal spiking. The former relate to motor evoked potentials, which are known to be highly correlated with the population spiking (Bansal et al., 2012), and the latter to multi-unit activity, whose movement-related modulation also reflects correlated spiking in neuronal populations. Intermediate frequency bands tended to contribute little during movement execution in this type of task. One could raise the possibility that the relationship between LFP features and single-neuron spiking in these intermediate frequency bands could be much more transient than the relationship between spiking and kinematics during movement execution. For example, beta oscillations, even during movement preparation, typically occur in transient, not sustained, events lasting a few or several cycles. Thus, one would like to build models in which spiking phase-locking should be obviously conditioned on the amplitude of the beta filtered LFPs, so that these transients can be properly captured. In this regard, we note that the neural point process models used here should capture such dependence on beta amplitude, since the log-additive form of the models allows for (nonlinear) multiplicative effects and interactions among different terms (e.g. beta amplitude and phase) in the models. We also note that, although more complex LFP features and models could po-

tentially improve spike prediction, the same could be said about improving the predictive power of motor behavioral covariates by using more complex or a larger set of kinematic features, including for example kinetics (torques) and muscle activation covariates. We hope to be able to examine more complex LFP and motor behavior-related features in the future.

The results reported here on the redundancy between motor-cortex multiband-LFP features and motor behavior are specific to execution of motor tasks in non-human primates who were highly engaged during movement execution. Multiband LFP features also provide reliable biomarkers for broader brain states and their changes. For example, the relationship between single-neuron spiking activity and ongoing LFPs is likely to change substantially depending on anesthesia, drowsiness, resting vs. awake states, attentional and volitional states, as well as stages during motor tasks (e.g. preparation vs. execution). In this broader context, including a larger variety of neural states than examined in this study, we expect multiband LFP features will be an important independent signal to account for neuronal spiking variability not explained by stimuli or behavioral covariates.

Variability in single-neuron spiking activity has often been characterized as of two types: private and shared (e.g. DeWeese and Zador, 2004; Litwin-Kumar and Doiron, 2012; Churchland and Abbott, 2012; Goris et al., 2014). Private variability is likely to reflect chaotic nonlinear dynamics in highly recurrent neuronal networks (Litwin-Kumar and Doiron, 2012). Amplification of membrane potential fluctuations by the spiking generation process (Carandini, 2004) in addition to local stochastic factors such as thermal fluctuations and synaptic failure (Faisal et al., 2008) are also important contributors. On the other hand, shared variability in neuronal ensembles is thought to evolve on slower time scales and reflect representational and computational states in neuronal networks (Litwin-Kumar and Doiron, 2012; Churchland and Abbott, 2012). The examined fluctuations in multiband LFP oscillations seem primarily to be related to this shared variability.

Multiband oscillatory LFP activity results in large part from coherent or shared dynamics in neuronal networks. In addition, features in these oscillations that are predictive of single-neuron spiking seemed mostly redundant to parameters in motor behavior. Overall, our finding was that information in the examined multiband LFP features directly relates to these shared representational and computational dynamics across neural populations in motor cortex. Single-neuron activity in motor cortex populations has been shown to be dominated by latent low-dimensional collective dynamics (Churchland et al., 2012; Truccolo et al., 2010). We hope in the future to investigate the relationship between multiband oscillatory LFP activity, in particular slow fluctuations, and latent low-dimensional rhythmic dynamics (Churchland et al., 2012) in motor cortex.

Acknowledgements. We thank Corey Triebwasser, Michelle Nevor, John Murphy, and Allan Rydberg for their assistance with animal care and instrumentation design. This research is supported by: the National Institute of Neurological Disorders and Stroke (NINDS), R01 NS25074 (to JPD; co-inv.: WT), K01 Career Award NS057389 (to WT); Defense Advanced Research Projects Agency (DARPA REPAIR N66001-10-C-2010, Co-PI: JPD, WT); National Science Foundation predoctoral fellowship (MER); and the Pablo J. Salame '88 Goldman Sachs endowed Assistant Professorship in Computational Neuroscience (WT).

Chapter 3

Dissociation between single-neuron spiking β -rhythmicity and transient β -LFP oscillations during movement preparation in primate motor cortex

MICHAEL RULE

The experimental data examined here were collected in the Donoghue Lab. We thank Carlos Vargas-Irwin, Lachlan Franquemont and John Donoghue for sharing the data.

Abstract

The nature and strength of the coupling between single-neuron spiking activity and collective dynamics reflected in ongoing local field potential (LFP) oscillations is an important problem in neuroscience. In primate motor cortex, beta (~ 20 Hz) frequency oscillations dominate LFP activity during the planning and preparation of reach and grasp actions, while being largely suppressed during movement execution, except during the execution of isometric-force tasks. Previous studies have focused on single units that show some level of phase coupling with the ongoing β -LFP signal. Here, we report a different and striking phenomenon wherein many (43%, 125/292) well-isolated single units showed beta spiking rhythmicity but no apparent or weak phase-coupling with ongoing β -LFP oscillations. The relative phase between single-neuron spiking and β -LFP oscillations seemed to drift randomly. This phenomenon was common in units recorded via multiple microelectrode arrays from primary motor cortex (M1), as well as ventral (PMv) and dorsal (PMd) premotor cortices in monkeys during steady-state movement preparation periods prior to reach and grasp movements. Importantly, while beta rhythmic spiking and firing rates were commonly sustained during steady-state movement preparation periods, β -LFP oscillations emerged as transient events with considerable trial-to-trial variability. Single neuron firing rates within and outside these β -LFP transient events showed no differences during steady-states, and no consistent correlation was found between the beta oscillations' amplitude and firing rates. These findings indicate that the modulations in β -LFP power were not coupled to changes in the firing rates of the underlying population, as was the case for movement and visual cue related beta suppression in this task. Additionally, the population spiking activity was only weakly coupled to the ongoing β -LFP. Our findings suggest that the relationship between single-neuron spiking and collective dynamics as reflected in ongoing LFP oscillations can be more complex than previously thought. Although beta spiking rhythmicity indicates strong beta oscillatory synaptic inputs, it does not imply significant phase-coupling with ongoing beta LFP oscillations, or even β -LFP recorded in the

same microelectrode site. This dissociation points to two different sources of variation in the level of beta activity in motor cortex, one that impacts the single-neuron spiking dynamics, and another related to the generation of the beta LFP signal reflecting the collective dynamics in the population.

3.1 Introduction

Sensorimotor cortex beta (β -) LFP oscillations reflect the collective dynamics in neuronal populations embedded in local and large-scale brain networks. In the specific case of motor cortex, β -LFP oscillations are especially evident during movement preparation, planning, and also during the execution of isometric-force grip tasks (Baker et al. 1997, 2001, 2003; Jackson et al. 2003; Murthy and Fetz 1992, 1996a,b; Sanes and Donoghue 1993). The relationship between single-neuron spiking and β -LFP oscillations is an important issue towards revealing the nature and role of these oscillations in the primate motor cortex. In addition, addressing this issue may be critical for the development of new therapies for movement disorders, such as Parkinson's disease (Beuter et al. 2014; Gale et al. 2008; Yang et al. 2014), and for the development of brain machine interfaces for people with paralysis.

Most initial studies focused on spike-triggered averages of LFPs to examine the relationship between neuronal spiking and ongoing β -LFPs. Based on this approach, several studies have shown some level of phase coupling between spikes and LFP (e.g. Murthy and Fetz (1996b)). However, an assessment of the coupling strength or magnitude based on spike-triggered averages is difficult. To address this issue, Baker et al. (2003) used spike-field coherence measures and showed that, although statistically significant, the coupling magnitude between pyramidal tract neurons and β -LFP oscillations is typically weak during isometric force tasks (average coherence values around ~ 0.05). However, this study did not examine the coupling strength between neuronal spiking and β -LFP oscillations during movement preparation in reach and grasp actions. In addition, the transient nature of β -LFP oscillations was not taken into account. More recently, preliminary analyses in Denker et al. (2007) has shown that phase coupling may occur primarily during beta transients during movement preparation periods, but did not explicitly investigate how this relates to the level of rhythmicity in single neuron spiking.

The relationship between single-neuron beta (β -) rhythmicity and β -LFPs is also an important issue. Previous studies have shown that β -LFP power is highest in layer V of

motor cortex (Murthy and Fetz 1996a; Witham and Baker 2007) and that pyramidal tract layer V neurons tend to fire rhythmically in the beta frequency (Wetmore and Baker, 2004). However, Witham and Baker (2007) have found that different motor areas show different levels of single-neuron spiking β -rhythmicity and β -LFP power and, more importantly, that the level of spiking β -rhythmicity in a given motor area does not predict the level of β -LFP power in the same area. Furthermore, it remains unclear the extent to which beta rhythmic spiking coincides with β -LFP transients, and the extent to which beta rhythmic spiking occurs outside periods of elevated beta power. Clarifying these issues is an important step for understanding the function and mechanisms of beta oscillations in motor cortex. We note that previous work has addressed the coexistence of coherent population gamma oscillations despite underlying single-neuron asynchronous and irregular spiking activity (e.g. Ardid et al. 2010; Brunel and Wang 2003; Geisler et al. 2005; Hoseini and Wessel 2015). The problem examined here, instead, is that of weak coupling between transient coherent population oscillations and sustained regular or rhythmic spiking.

This chapter addresses these issues in the context of a visually cued reaching and grasping task with instructed delays. Single units were simultaneously recorded via multiple microelectrode arrays from areas M1, PMd and PMv, while the monkey performed 3-D reach and grasp actions. We focused on examining the relationship between well-isolated single units and β -LFP activity during steady-state movement preparation stages of this task, which may potentially differ from synchronization dynamics previously studied in association with isometric force during precision grip. Overall, we found a striking phenomenon that has been overlooked in previous studies. While β -LFP oscillations tended to appear as short transients, even during steady-state movement preparation, beta spiking rhythmicity, evident in the inter-spike time interval (ISI) distributions and autocorrelation functions, was sustained. Furthermore, different spike-LFP phase coupling measures revealed that single-neuron beta-rhythmic spiking was at most only weakly coupled to the β -LFP oscillations, even when the analysis was restricted to transient periods of high

β -LFP power. We observed this phenomenon in many single units from three recorded cortical areas. In addition, single units clustered into two groups (narrow and wide extracellular action potentials), with groups showing differences in firing statistics, but no consistent differences in the strength of their phase coupling to β -LFP oscillations.

3.2 Methods

The CGID task The Cued Grasp with Instructed Delay (CGID) task investigates neural activity in motor cortex associated with sensory integration, working memory across instructed delays, and planning of upcoming reach and grasp (Figure 3.1; see Vargas-Irwin et al. (2015) for additional details). The task requires a monkey to reach out and grasp one of two objects, using one of two possible grips. A sequence of visual cues instructs the monkey which object to grasp, and which grip to use. When the task begins, the lights in the room are turned off, and one of the two objects is rotated into place. One second later, that object is illuminated. The monkey now knows which object to grasp, but not which grip to perform. One second after that, a cue light (red or yellow, left or right position) is illuminated, specifying the grip. If the light is red, the monkey is to perform a power grip. If the light is yellow, the monkey is to perform a precision grip or a key grip, depending on the object. Two seconds after the ‘Grip’ cue, a ‘Go’ cue (green light, middle position) is given. The monkey is now allowed to reach out and grasp the object. If the monkey moves before the ‘Go’ cue or uses the incorrect grip on the object, the trial is voided. If the monkey uses the correct grip, he receives a juice reward.

In this Chapter, we refer to the task epochs preceding the ‘Go’ cue as the planning and preparatory period. Movement periods were defined as the time from when the monkey lifts his hand from the holding position to the time when the monkey contacts the object, as detected by capacitive touch sensors. We concentrate our analysis on two steady-state periods. The first period is the one second between the start of the trial and when the first visual cue is presented. In this epoch, the monkey is waiting attentively and has not

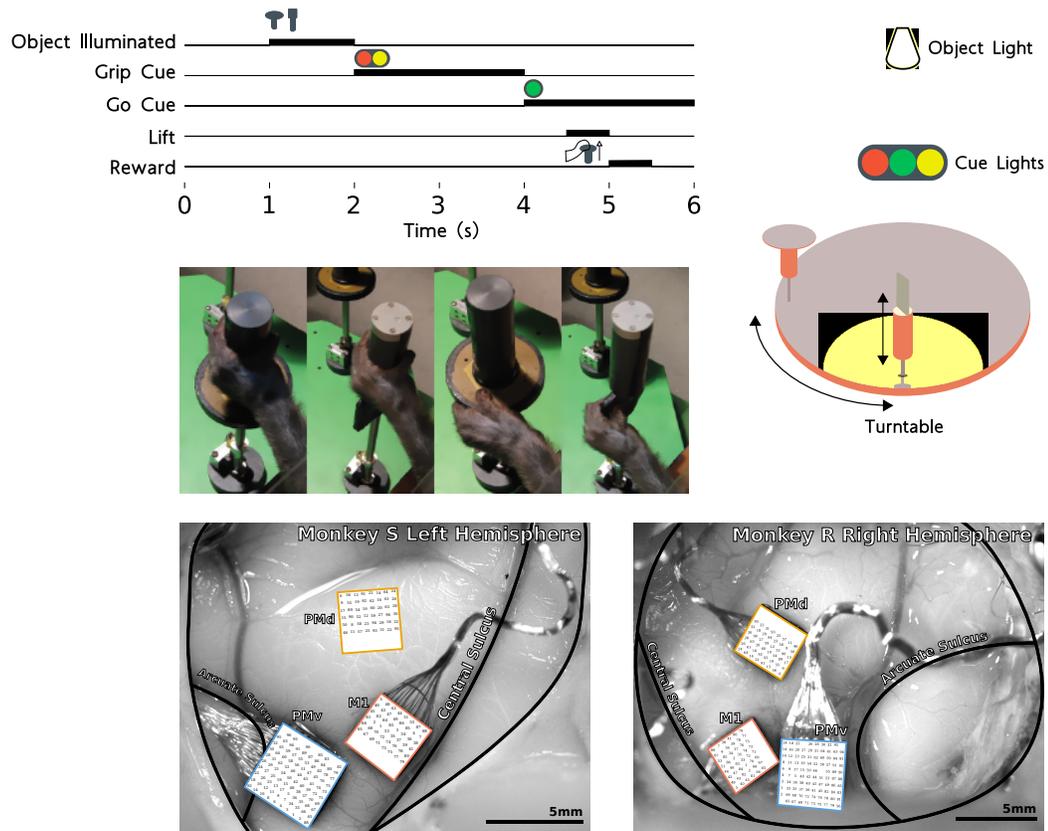


Figure 3.1: *The Cued Grasp with Instructed Delay (CGID) task.* The CGID task investigates integration of visual cues and planning of upcoming reach and grasp movements across instructed delays. When the task begins, one of two objects is presented. One second later, a cue light instructs the monkey to use either a power grip (left panels in the photo inset) or precision grip (right panels) to lift the object. A two second planning period follows, after which a 'Go' cue signals that the monkey may reach and grasp the object. The planning and delay periods reliably elicit beta oscillations. Touch sensors detect when the monkey begins to move. LFP data were recorded on 10x10 (ventral premotor cortex PMv) and 6x8 (dorsal premotor cortex PMd and primary motor cortex M1) Blackrock arrays with 0.4 mm electrode spacing. Data from two monkeys (R and S) were analyzed. Broadband LFPs recorded at 30 kilosamples/s (0.3 Hz - 7.5 kHz) were downsampled (zero-phase 4th order Butterworth, ≤ 250 Hz) to 1 kilosample/s for analysis. In this study, we focused on two steady-state motor periods: the 1 second before object presentation and the 1 second preceding the go cue. Data were collected by Carlos Vargas-Irwin and Lachlan Franquemont in the Donoghue lab.

yet received the information needed to plan or prepare for movement. The second period is the one second preceding the ‘Go’ cue. In this period, the monkey has been cued with the information needed to plan the reaching and grasping action, and the transient neural activity associated with the visual cues has passed. It is important to note that the visual cue lights were present until the ‘Go’ cue, so this second steady-state epoch represents a motor preparatory state and not a state that explicitly required working memory.

Neural recordings Data were recorded from triple microelectrode arrays (MEAs; Blackrock Microsystems, Utah), with an electrode depth of 1.5 mm targeting layers II/III-V of motor cortex. Neuronal spiking and LFP data were recorded on 10×10 (ventral premotor cortex PMv) and two 6×8 (dorsal premotor cortex PMd and primary motor cortex M1) arrays with 0.4 mm electrode spacing. Data from two monkeys (R and S) were analyzed (see Vargas-Irwin et al. (2015) for additional details). Broadband LFPs recorded at 30 kilosamples/s (0.3 Hz - 7.5 kHz) were down-sampled (zero-phase 4th order Butterworth, \leq 250 Hz MATLAB filtfilt) to 1 kilosample/s for analysis.

Spike sorting For each electrode, candidate spikes (extracellular action potentials) were identified online via threshold crossing in the amplitude of the high-pass filtered signal (250 Hz 4th order high-pass Butterworth filter, Cerebus Data Acquisition System, Blackrock). Preliminary spike sorting was performed by a custom automated spike sorter (Vargas-Irwin and Donoghue, 2007), and verified using the commercial Plexon Offline Sorter (Plexon Inc.). Candidate units included in the analysis had a minimum signal-to-noise ratio (SNR) of 3.0, defined as in Vargas-Irwin and Donoghue (2007). Additionally, we required that: (1) the inter-spike-interval (ISI) histogram display a clear refractory period to exclude potential multi-unit clusters; (2) that the units exhibit at least 100 inter-spike interval events during each of the defined one-second steady-state periods of the CGID task within a session, to provide for adequate estimation of ISI distributions; and (3) that units be clearly separated into different clusters in the waveform PCA feature

space. Electrodes exhibiting cross-talk or excessive noise were excluded from analysis.

ISI histogram statistics Isolated single units showed diverse firing characteristics as assessed by the inter-spike interval (ISI) distribution and related statistics, both across time and across units. For a given unit, these statistics were computed from the ISI distribution from all inter-spike intervals pooled over all trials for a given one-second epoch of the CGID task. We computed mean firing rates, the ISI mode, and the coefficient of variation (CV; i.e. the standard deviation of the ISI distribution divided by the corresponding mean). We quantified the tendency of units to fire bursts as the percentage of ISIs shorter than 10 ms.

We estimated a single unit's preferred firing frequency (in Hz), by computing the inverse of the ISI mode, henceforth referred to mode frequency. The mode firing frequency was identified for unimodal and bimodal ISI histograms using kernel density estimation (Python `scipy.stats.gaussian_kde`). Because some units exhibited an ISI distribution with an additional mode corresponding to bursts, and since we were interested in slower 'rhythmicities', we considered only ISI events longer than 10 ms when estimating the mode firing frequency. Because ISI distributions were right-skewed, we applied kernel density estimation to the transformed variable $\log(5 \text{ ms} + \text{ISI})$. The shift of 5 ms was added to improve numerical stability when estimating the log ISI distribution close to zero, which was an issue in the subset of units that fired bursts of spikes.

Unit categorization Units were manually categorized based on features of their ISI distributions during the movement preparation steady-state periods of the CGID task. Units exhibiting a clear mode in the ISI distribution between 10 and 100 ms were classified as unimodal. Units that showed an additional peak below 10 ms in the ISI histogram were further classified as bimodal (bursting/rhythmic) cells. Units exhibiting exponential ISI distributions (allowing for refractoriness) were classified as Poisson-like. Units displaying a mixture of these features, e.g. some amount of bursting, with an exponen-

tial ISI distribution exhibiting a long recovery period, were classified as “intermediate”. We restricted spike-field phase coupling analysis to well-isolated single-units classified as unimodal or bimodal (bursting/rhythmic) that also displayed an ISI mode frequency between 10 and 45 Hz in at least one of the steady-state movement preparation epochs, and a mean rate at least one fifth the mode firing frequency. Allowing these low firing rates permitted analysis of single units whose spiking was coupled to the beta phase, but did not fire in every beta cycle. The distribution of mean firing rates across units during these epochs concentrated below 30 Hz.

Units were also classified as thin- and thick-spike based on their mean extracellular action potential waveform. In order to precisely align spikes, we upsampled waveforms using sinc interpolation. To minimize edge effects during up sampling, the linear trend in the waveform was removed, the de-trended waveforms were upsampled with reflected boundary conditions, and the linear trend restored. We extracted mean waveforms by averaging peak-aligned upsampled waveforms. Waveforms were clustered based on the voltage of the mean waveform 300 μ s after the spike peak. We found this feature led to better cluster separation than using the waveform width, since at this time post-spike, thin spike cells have recovered (and may exhibit afterhyperpolarization), while thick-spike cells remain depolarized. Average waveforms from all areas, sessions, and monkeys, were combined for clustering. Clustering was performed using a 1D Gaussian mixture model, and units were assigned as either thin- or thick-spike based on likelihood ratio.

Beta phase extraction and transient identification For analysis, raw LFP traces (30 kilosamples per second) were low-pass filtered at 250 Hz using a zero-phase 4th order Butterworth, ≤ 250 Hz MATLAB `filtfilt`, and down-sampled to 1 kHz (Matlab `decimate`). In the Generalized Linear Model (GLM; Truccolo et al. (2005)) assessment of spike-LFP phase coupling, the beta band was identified separately for each session and channel, and also separately for the two different steady-state movement preparation periods. Beta was selected as the 5 Hz band surrounding the highest peak between 15

and 30 Hz in the multitaper LFP spectrogram (5 Hz bandwidth) during the one second preceding the ‘Go’ cue. Beta was extracted using a 4th-order Butterworth band-pass filter applied forwards and backwards. Beta phase was extracted using the Hilbert transform (SciPy `hilbert`, Oliphant (2007)), which generates a beta analytic signal $z(t)$ consisting of a real component $x(t)$ (the filtered beta signal) and an imaginary component $i \cdot y(t)$ which is a $\pi/2$ phase shifted copy of the filtered beta signal. The instantaneous phase $\varphi(t)$ and amplitude $|z(t)|$ can be extracted from the analytic signal $z(t) = |z(t)| \cdot \exp(i\varphi(t))$.

We extracted transient periods of elevated beta power by examining the amplitude envelope of the beta analytic signal. First, the amplitude envelope $|z|$ was smoothed with a 50 ms boxcar filter. Events for which this smoothed amplitude signal exceeded 1.5 times the standard deviation of the filtered beta signal for at least 40 ms (approximately one beta wavelength) were designated as high-beta events.

Spectral estimation We summarize the power spectra for a given epoch using multitaper spectral estimation (Bokil et al. 2010; Percival and Walden 1993). Multitaper (5 Hz bandwidth) spectral estimates were computed separately for each trial for a given epoch, then averaged over all trials. We visualize single-trial LFP activity using wavelet spectrograms, which enabled good time-resolution for higher frequencies while also maintaining good frequency resolution at lower frequencies (Torrence and Compo, 1998). We use a Morlet continuous wavelet transform with a time-bandwidth parameter of 5.

Spike triggered LFP averages We estimate the spike triggered averages between spikes and the 250 Hz low-pass filtered (Butterworth, 4th-order, forward-backwards, Matlab `filtfilt`) LFP sampled at 1 kilosample on the same electrode. This approach preserves the spiking artifact in the LFP. (The section “Spike contamination” below address this concern.) Stable phase coupling of neuronal spiking to ongoing LFP oscillations will appear as oscillatory components in the spike-triggered averages (STAs). However, estimation of spike-LFP phase coupling is susceptible to several biases. These biases are

exacerbated if both the spike trains and the LFP signals exhibit autocorrelations at similar time scales. For example, if a rhythmic spike train (~ 20 Hz) co-occurs with a burst of 20 Hz LFP oscillations, it might appear that the single unit is phase coupled to the LFP even if there is no relationship. Additional biases may emerge if changes in firing rates are correlated with changes in LFP power. In addition, the STA does not offer a direct assessment of the phase coupling magnitude, as it reflects both phase-coupling and amplitude effects, and poorly visualizes variability around the mean trend. Because of the biases inherent in the STA, we also use two complementary approaches to get unbiased estimates of spike-LFP phase coupling: pairwise phase consistency (PPC), and GLM point process models for assessment of phase coupling.

Pairwise phase consistency Pairwise Phase Consistency (PPC; Vinck et al. 2010) is an estimate of spike-LFP phase coupling that is not biased by the firing rate or correlated modulations in LFP power and firing rate. Vinck et al. (2010) define PPC as the average dot product between all pairs of spike-triggered phase measurements. We computed PPC using the equivalent expression (Aydore et al. 2013; Equation 11)

$$\text{PPC} = \frac{N}{N-1} \left(|\bar{z}|^2 - \frac{1}{N} \right), \quad (3.1)$$

where $|\bar{z}|$ is the magnitude of the average over spike-triggered β -LFP phase vectors $\bar{z} = \frac{1}{N} \sum_{k=1}^N \exp(i\varphi_k)$, where φ_k represents the phase measurement at a given spike time and k indexes over spikes. To compute PPC, we extracted instantaneous LFP phase estimates for a range of frequencies by taking the Fourier transform of the LFP in a ± 100 ms window surrounding each spike. Each LFP segment was mean-subtracted and multiplied by a Hanning window to reduce boundary effects. To attenuate temporal dependencies among samples, spikes that occurred fewer than 200 ms after a previous spike were excluded. We report the PPC value at the peak beta frequency, identified separately for each monkey, session, channel, and task epoch. The PPC bias correction relies on the assumption that

successive samples are independent. Although we reduced temporal correlations between successive samples by removing events for which the spike-triggered LFP segments would overlap, residual correlations may remain in both the spike trains and LFP. Therefore, we estimated the chance level empirically by phase randomizing LFP segments (Mammen et al., 2009), which preserves the autocorrelation structure of the LFP.

Point-process GLMs for assessing phase coupling between single neuron spiking and ongoing β -LFP oscillations

We use a discrete-time point-process generalized linear model (GLM) framework (Truccolo et al., 2005) to detect spike-LFP phase coupling. We used 1 ms time bins. This is similar to the approach used in Lepage et al. (2013), Zhou et al. (2015), and Rule et al. (2015). GLM point-process based estimators explicitly model the conditional intensity function (CIF) $\lambda(t)$ and include an offset parameter μ as a separate regression term, therefore providing an estimate of spike-LFP phase coupling that is less susceptible to variations in firing rate. We consider GLM point process models of the form

$$\ln(\lambda(t)) = \mu + \alpha \cos(\varphi_{\text{LFP}}(t) - \varphi_0) = \mu + \beta_1 \cos \varphi_{\text{LFP}}(t) + \beta_2 \sin \varphi_{\text{LFP}}(t), \quad (3.2)$$

where $\lambda(t)$ is the conditional intensity function of the neuron’s measured point process, μ is a mean-rate offset parameter, φ_0 is the preferred phase of firing relative to the LFP, $\varphi_{\text{LFP}}(t)$ is the time-varying instantaneous Hilbert phase of the LFP signal, and α is a parameter measuring the strength of phase coupling. In this study we assess the predictive power of the model using receiver operating characteristic (ROC) curve analysis (Fawcett, 2006; Rule et al., 2015; Truccolo et al., 2010). The area under the ROC curve (AUC) summarizes the accuracy of spike times predicted based on the model, and typically ranges from 0.5 (chance level) to 1.0 (perfect prediction). We report predictive power (PP) as normalized AUC values such that 0 is chance level and 1 is perfect prediction. Chance level was estimated using phase randomized LFP (Mammen et al., 2009) and by shuffling the

LFP trial blocks relative to the spike trains.

Point-process GLMs for relating single neuron spiking to population spiking activity

We also used CIF models to relate single units to the population spiking activity. This population spiking activity was defined as the total number of spikes across all of the recorded single units (except the predicted neuron) in a given motor area within 1 ms time bins, followed by a 25 ms boxcar filter. The CIF model consisted of:

$$\ln(\lambda(t)) = \mu + \beta \cdot A(t), \quad (3.3)$$

where $A(t)$ corresponds to the population spiking activity. As a second measure of population activity, we also considered multi-unit activity (MUA) recorded in the same electrode as the single unit. MUA was defined as the amplitude envelope of > 250 Hz LFP, further bandpass filtered in the 5 Hz band surrounding the peak beta frequency in order to focus on the relationship to beta LFP oscillations.

Assessing coupling between population spiking activity and ongoing β -LFP

We assessed the relationship between the population spiking activity and the ongoing β -LFP oscillations by computing their cross-correlation functions. Population spiking activity was defined as above, except that in this case, all well-isolated single units were included (for the spiking population history model, the unit being predicted was excluded from the population rate). Statistical tests were applied to the peak of the cross-correlation functions computed for time lags ranging over one beta cycle (± 25 ms).

Spike contamination In this study, we examined statistical relationships between neuronal spiking activity and local field potentials recorded on the same electrode. In this case, the spikes themselves contribute to LFP power, even at frequencies as low as the ~ 20 Hz beta band investigated here (Waldert et al., 2013). Waldert et al. (2013) find that the spiking contribution to low-frequency LFPs can arise from both low-frequency compo-

nents of the spike waveform, including slow afterhyperpolarization potentials (AHPs), as well as spike-train rhythmicity at low-frequencies. We elected not to use spike removal procedures like those of Zanos et al. (2011) for several reasons. First, we are primarily interested in the phenomenon of apparently weak spike-field coherence. In this case, a bias toward detecting coherence resulting from spiking contamination is conservative. Additionally, it is possible that there is ambiguity between spike-locked local network oscillations and low-frequency components of the extracellular spike waveform (e.g. slow afterhyperpolarization potentials). Since it is possible that low-frequency components of the spike waveform relate to the origins of LFP, we wish to avoid erroneously removing a true contributor to β -LFP. Nevertheless, we can distinguish between action potential contamination and other spike-LFP phase interactions by inspecting the PPC spectrum. True β spike-LFP phase coupling leads to a PPC peak at the beta band, whereas spiking contamination leads to a broad-band monotonically increasing PPC spectrum.

3.3 Results

We analyzed three CGID task sessions each from two monkeys (R, S) with simultaneous MEA implants in three motor areas (M1, PMv, PMd) (Methods ‘The CGID task’, Figure 3.1). Each session yielded between 46 and 114 correctly executed seven-second CGID trials, collected over twenty minutes to one hour. For each session, each MEA yielded between 7 and 48 well-isolated and high signal to noise ratio single units, for a total of 699 units. Of these, 292 exhibited sufficient firing rates during the steady-state delay periods of the task to permit further analysis.

We first review single unit firing statistics during the steady-state periods of the CGID task, and show that many units displayed rhythmic spiking during steady-state periods close to the beta (~ 20 Hz) frequency range. Steady-state periods corresponded to an attentive waiting period in the first second before object presentation, and a movement preparation period one-second before the go-cue (Methods).

We then show that, concurrent with the emergence of β -rhythmic spiking, β -LFP oscillations are enhanced during steady state periods in the CGID task. We examine the phase coupling between single neuron spiking and the β -LFP oscillations. Examples where the preferred phase of firing of these β -rhythmic units appears to drift relative to the β -LFP are presented, showing no stable phase coupling (Figures 3.8, 3.9). Summary statistics across neurons demonstrate that β spike-LFP phase coupling was typically negligible, even for single units that showed strong β -rhythmic spiking. We demonstrate that β -LFP power was substantially more variable than β -rhythmic spiking, and that the weak spike-LFP coupling persisted even when analysis was restricted to periods when beta power was high. Finally, we show that spike coupling to basic measures of spiking population activity was also weaker during the steady-state periods of the CGID task than during movement execution.

3.3.1 Single units show sustained β -rhythmic spiking during steady-state movement preparation periods

We categorized units based on features of their ISI distributions (Methods ‘ISI histogram statistics’) during the steady-state movement preparation periods (Figure 3.2a,b). 699 units exhibited well-isolated spiking. Of those, 71% (499/699) met the minimum SNR cutoff of 3.0 for inclusion in the analysis, 54% (377/699) exhibited at least 100 ISI events during the task steady-state epochs, and 42% (292/699) met both conditions and were suitable for analysis. Out of these 292 well-isolated single units that satisfied the inclusion criteria, 66% (192/292) showed a unimodal peak in ISI events longer than 10 ms during the two steady-state movement preparation periods. A subset of units (25%, 72/292) exhibited bursting as evidenced by bimodal ISI distributions with a second peak in short latency (<10 ms) ISI events, while also exhibiting an overall slower rhythmicity. A minority of units (7%, 21/292) showed low firing rates and irregular Poisson-like spiking, or had an ISI distribution that could not be clearly categorized (2%, 7/292).

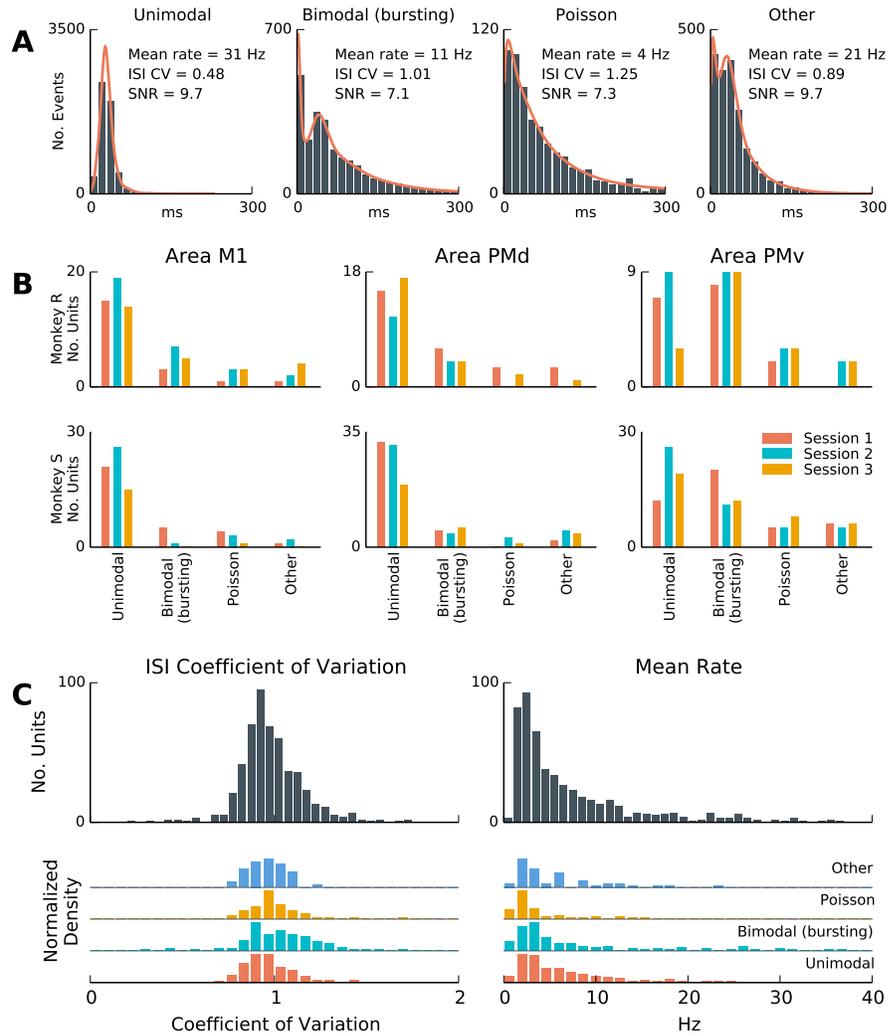


Figure 3.2: A subset of units fire rhythmically during steady-state movement preparation periods of the CGID task. **(A)** Inter-spike interval (ISI) distributions from selected well-isolated units during the steady-state periods of the CGID task. In each plot, from left to right, we see rhythmically firing units, units that exhibit both bursting and rhythmicity, units that exhibit Poisson-like firing, and units that exhibit intermediate ISI distributions. The ISI coefficient of variation (CV) reflects the dispersion of the ISI distribution, with low CV correlating with rhythmicity; SNR = signal to noise ratio for unit waveform. **(B)** Single units were manually categorized based on ISI features (Methods) as unimodal (rhythmic), bimodal (bursting and rhythmic), Poisson process-like (i.e. exponential with refractory period), or intermediate ISI distributions. In both monkeys and all areas, single units with unimodal and bimodal ISIs were most prevalent. **(C)** A summary of ISI mean and CV statistics for the same units. Statistics of ISI distributions varied continuously and did not form discrete clusters. Mean rate was variable, with 25% of units exhibiting mean rates higher than 10 Hz. ISI CV reflects the dispersion of the ISI distribution, with low CV correlating with rhythmicity. Because some rhythmic units start and stop firing during the steady-state epochs, and because the rhythmic frequency may change over time and across trials, the effective CVs can be larger than might be expected for sustained rhythmic firing at a single frequency.

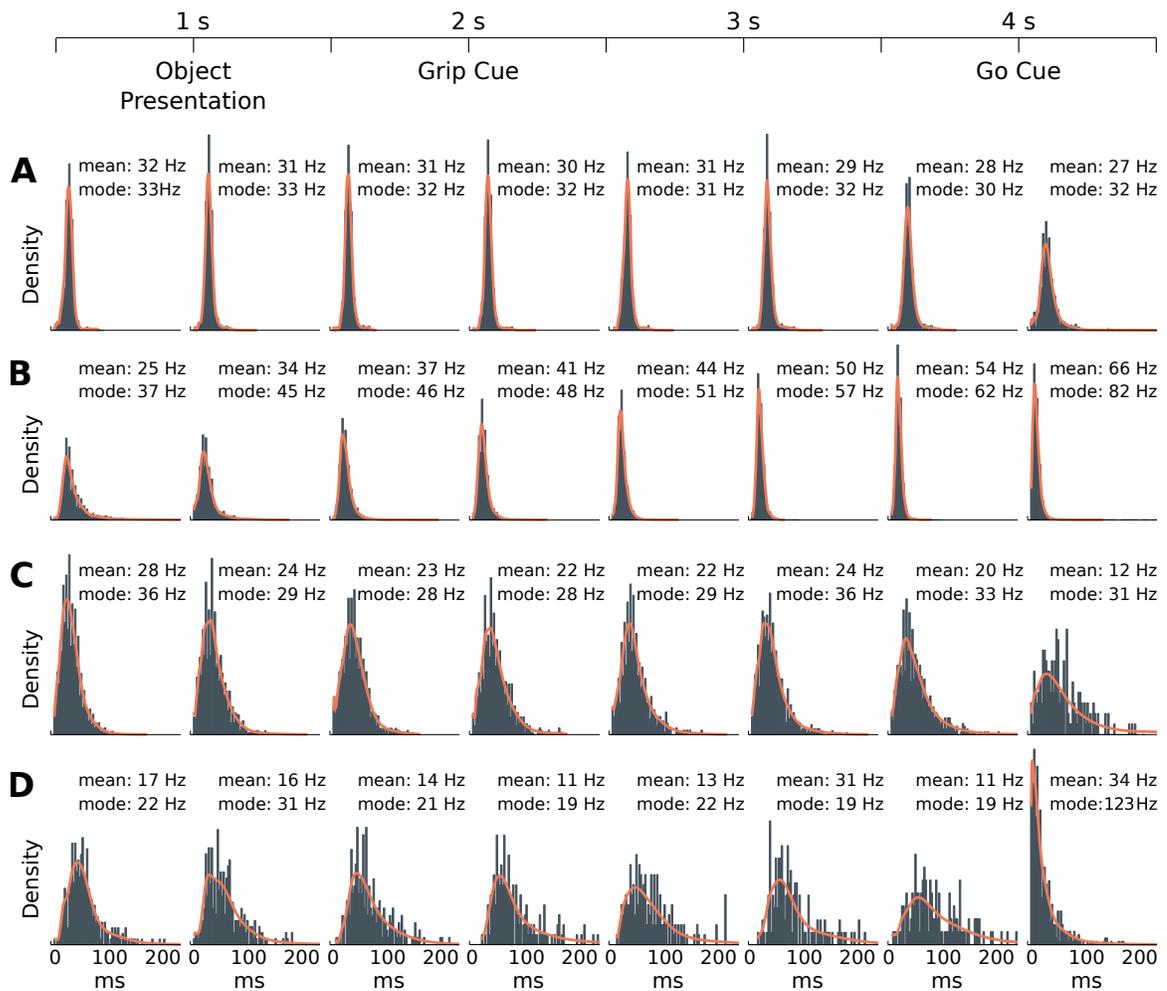


Figure 3.3: Single-unit ISI statistics change across different stages of the task. Shown here are four examples of how the ISI distributions change for well-isolated units over the course of the CGID task. The insets specify the mean rate μ and the mode of the ISI distribution. Each ISI histogram was computed based on non-overlapping one-half second time windows of the CGID task. All trials within a session were combined. Examples, from top to bottom, illustrate: **(A)** a highly rhythmic unit (monkey R session 2 unit 101) that decreased its mean firing rate during the movement epoch ((1/2) second after ‘Go’ cue), without changing its mode; **(B)** A highly rhythmic unit (monkey S session 1 unit 74) that steadily increased both its ISI mode frequency and mean firing rate, transitioning gradually from $\mu=25$ Hz at the trial outset to $\mu=66$ Hz during the movement epoch; **(C)** A unit (monkey R session 3 unit 92) whose firing became more variable, with a slight decrease in mean rate, only during the movement epoch; **(D)** A unit (monkey R session 1 unit 88) that switched from rhythmic firing at beta frequency ~ 11 -17 Hz, to Poisson-like firing at a much higher rate (123 Hz). These examples emphasize that the rhythmicity observed in a subset of units during the steady-state movement preparation periods of the CGID task was unlikely to arise entirely from intrinsic neuronal properties (e.g. subthreshold resonance). Instead, this rhythmicity likely reflected and was modulated by the collective network state. The colored traces represent the transformed KDE estimate of the distributions used to determine the ISI mode, and is shown to confirm that the mode estimation procedure approximates well the location of the ISI mode firing frequency (Methods ‘ISI histogram statistics’).

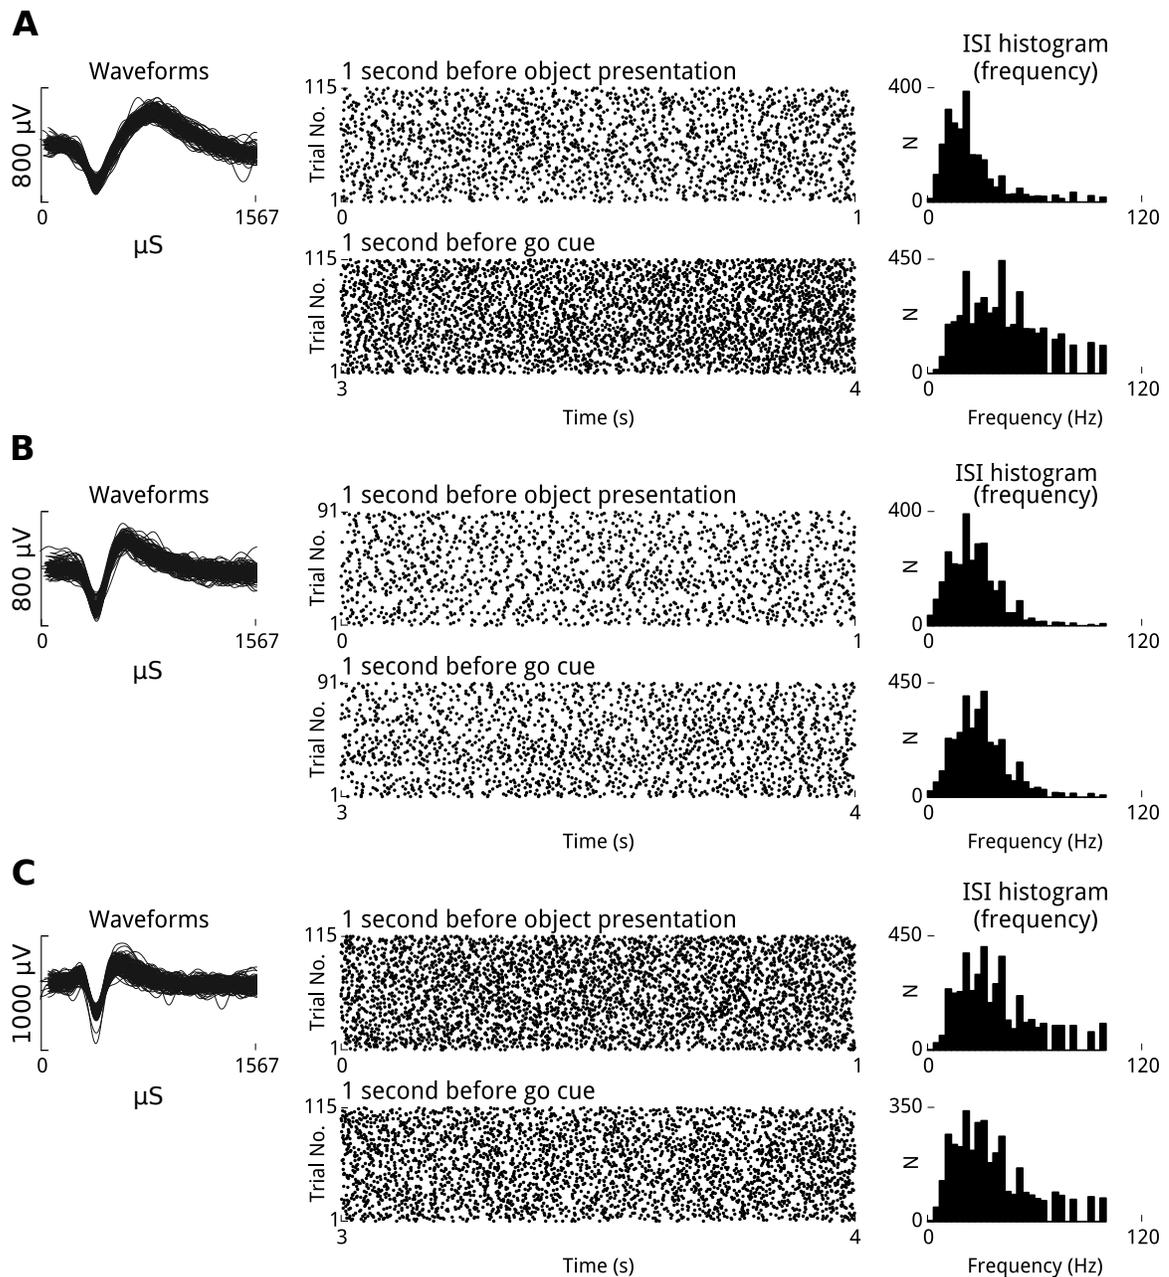


Figure 3.4: Well-isolated single units fire rhythmically at beta frequency, and firing rates are stable across trials. Shown here are three example well-isolated units that exhibited β -rhythmicity during the steady-state movement preparation periods of the CGID task. Spike rasters, which show trial number on the y-axis and task time on the x-axis for the two steady-state epochs, reveal that these units fired in a rhythmic manner that was reliable over trials and sustained across the steady-state periods. We also note that in several cases the mode frequency differed between the steady state period at the beginning of the trial, before visual cues have been provided, and the one second period preceding the 'Go' cue. The modes of the ISI distributions for these units, expressed in terms of frequency, show that these units were firing with a preferred frequency in the beta range of 15 to 30 Hz. (A) unit 43 from area PMd, monkey S session 3. (B) unit 49 from area PMd, monkey S session 2. (C) unit 20 from area M1, monkey S session 3.

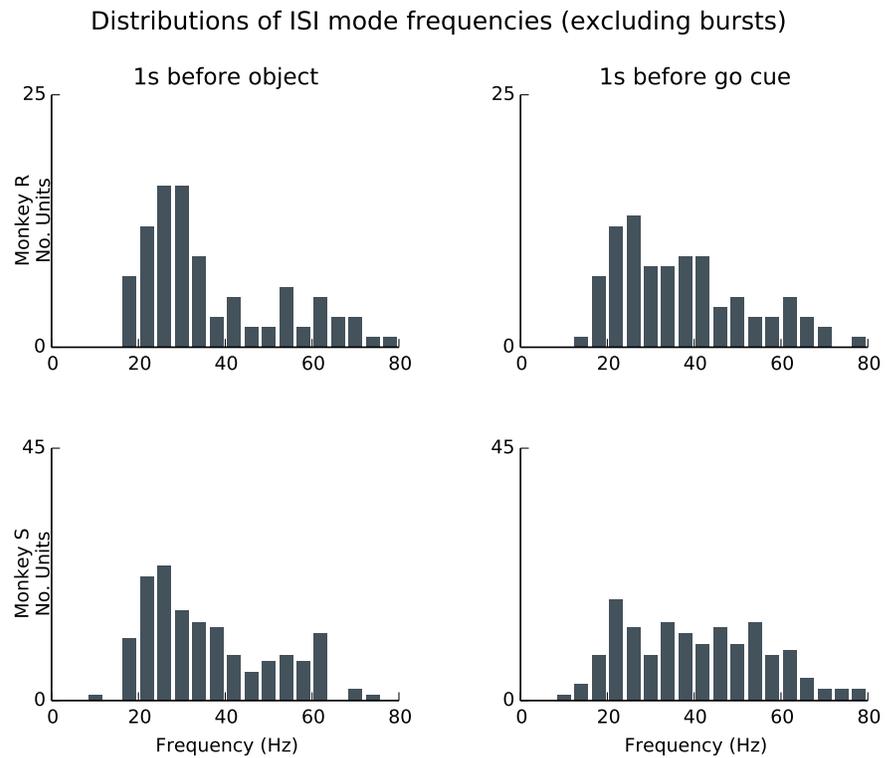


Figure 3.5: *The preferred firing frequency of rhythmic units varies, but typically falls within the beta band.* Shown here are summary distributions, pooled over all sessions and areas, for both monkeys during the two steady-state movement preparation epochs for units that showed unimodal and bimodal ISI distributions. Mode firing frequency for isolated single units ranged between 10 and 80 Hz, but for each monkey and epoch between 60

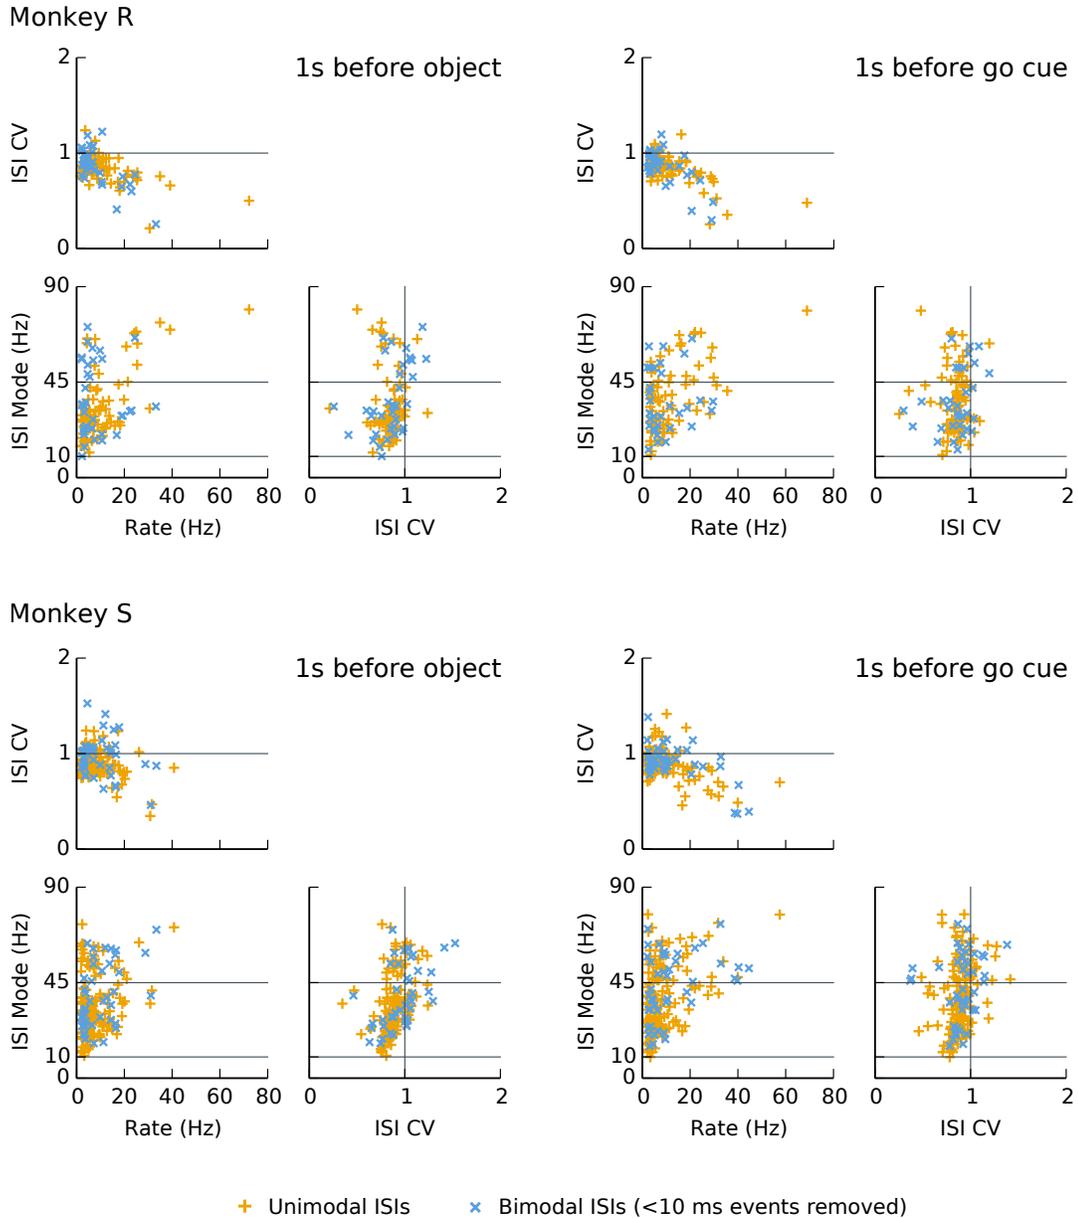


Figure 3.6: *Unimodal and bimodal units display similar statistics in the slow rhythmic component of the inter-spike interval distribution.* Shown here is a summary of the statistics of the subpopulation of units that exhibited signatures of a slower, rhythmic preferred firing frequency, i.e. units with unimodal ($n = 192$; yellow) or bimodal ISI ($n = 72$; rhythmic/bursting; blue) histograms. To focus on the lower mode frequency, the statistics for the bursting units were computed with spikes closer than 10 ms removed. The statistics then reflect the mean and dispersion of the slower (non-bursting) ISI peak. Unit properties are broadly similar between monkeys and task epochs. Units tend to exhibit a $CV < 1$ indicating some level of regularity in their firing patterns. Because some rhythmic units start and stop firing during the steady-state epochs, and because the rhythmic frequency may change over time and across trials, the effective CVs can be larger than might be expected for sustained rhythmic firing at a single frequency. Mode frequency and mean rates are highly variable, with mean rates uniformly lower than mode firing frequency. In the spike-LFP β phase coupling analyses, we focused on those rhythmic units displaying a mode frequency between 10 and 45 Hz, and a firing rate of at least 20% of their mode frequency to allow units that may skip oscillation cycles or cease firing during parts of the trial.

We considered identifying the above three classes, (refractory) Poisson-like spiking, bimodal (bursting/rhythmic), and unimodal (rhythmic) units with the three neuron types I, II and III described in Chen and Fetz (2005), which each exhibit different characteristic spike waveforms. However, Baranyi et al. 1993a,b describe a larger number of neuronal subtypes in motor cortex, with potentially overlapping firing statistics and spike waveform shapes, and we found that 63% (185/292) of units exhibited ISIs that could not be clearly identified with any of the categories in Chen and Fetz (2005). We tentatively identified 38% (21/55) irregular Poisson-spiking units with type I, 31% (22/72) bursting units with type II, and 33% (64/192) units exhibiting fast regular spiking with type III.

We further separated extracellular waveforms into ‘thin’ (42%, 124/292) and ‘thick’ (58%, 168/292) classes (Methods, ‘Unit categorization’, Figure 3.15a,b). We observed a partial agreement between ISI features and the extracellular waveform categorizations consistent with Chen and Fetz (2005): 62% (13/21) of putative type I neurons exhibited broad spikes, and 86% (19/22) of type II (bursting) neurons, and 72% (46/64) of putative type III (fast rhythmic) neurons exhibited thin spikes. We note that Chen and Fetz (2005) suggest that the rhythmic firing observed in the bursting neurons in their study was likely to arise from network interactions and not intrinsic neuronal properties, as is the case for the type III neurons.

The presence of intermediate examples suggested that the overlap between the distributions of firing statistics for each neuronal subtype in our data was too large to allow classification. Previous work has highlighted that intrinsic neural properties can also be heterogeneous (Battaglia et al., 2013). Because of these ambiguities in identifying neuronal subtypes based on spike train statistics, we focused our analysis on units that exhibited a clear mode in the ISI between 20 and 100 ms, which may potentially exhibit rhythmicity at the same frequencies as β -LFP. Two summary statistics: the ISI coefficient of variation (CV) and mean firing rate are shown in Figure 3.2c. Further exploration of the relationship between mean rates, ISI mode frequencies (Methods), and ISI coefficients of variability, is

illustrated for the unimodal and bimodal units in Figure 3.6.

Of units that exhibited rhythmic firing during the movement preparation periods, several dramatically changed their firing statistics during movement execution (Figure 3.3). Following the ‘Go’ cue, many units increased or decreased their firing rates (examples 1,3,4; Figure 3.3). Some units did not show abrupt changes following the ‘Go’ cue, but rather a gradual shift over the course of the preparatory period (example 2, Figure 3.3). Some units that exhibited unimodal/bimodal ISI distributions (a potential signature of rhythmic firing) during the preparatory period shifted to more Poisson-like spiking following the ‘Go’ cue (example 4, Figure 3.3). This suggests that rhythmic spiking need not be a fixed subthreshold resonance property of these neurons, and instead reflects the network state during the preparatory and delay periods.

We observed that rhythmically firing units tended to fire in a sustained manner during the examined steady-state periods, with high reproducibility across trials in terms of mean firing rates and ISIs (e.g. Figure 3.4). Inspection of the firing mode frequency for rhythmic units (Figures 3.5 and 3.6) revealed that the preferred firing frequencies were concentrated between 10 and 45 Hz, overlapping the β range. In the first steady-state epoch preceding the visual cues, 76% (78/103) of units showed an ISI mode frequency between 10 and 45 Hz for monkey R, and 74% (119/161) for monkey S. In the second steady-state epoch following the visual cues and preceding the ‘Go’ cue, 73% (75/103) of units in monkey R and 60% (96/161) of units in monkey S fell between 10-45 Hz. Some units (e.g. Figure 3.4a, Figure 3.3b) increased their mode frequencies to frequencies higher than 45 Hz during the pre-go steady-state period. Mode frequencies increased somewhat between the pre-cued and post-cued movement preparation periods, with the median mode frequency shifting from 30 to 34 Hz for monkey R, and from 32 to 39 Hz for monkey S. (this increase was statistically significant $p < 0.5$ in 5/6 sessions after a Benjamini-Hochberg correction for a false discovery rate of $\alpha = 0.05$, Benjamini and Hochberg (1995)).

3.3.2 β -LFP during steady-state movement preparation periods can be dissociated from β -rhythmic spiking

Given that a majority of isolated single units exhibited sustained rhythmicity at beta frequencies during the steady-state movement preparation periods of the CGID task, we next investigated the extent to which this β -rhythmicity was evident in local field potential (LFP) oscillations. In both monkeys, the LFP showed task-related changes in its power spectrum (Figure 3.7), especially in the beta band. (Chapter 4 addresses task-related changes in the spatial organization of these β -LFP oscillations.) Consistent with previous studies, the movement period was associated with suppression of β -LFP power. Importantly, beta was also transiently suppressed following the visual cues. In contrast, beta was elevated during steady-state movement preparation periods of the CGID task, including the first second of the task before object presentation, and the one second leading up to the 'Go' cue. For monkey S, the beta peak was identified between 22 and 25 Hz for all areas and sessions. Monkey R exhibited two beta frequency peaks, ~ 18 Hz and ~ 32 Hz in areas M1 and PMd. These two different beta frequencies may potentially correspond to the beta 1 and beta 2 oscillations previously examined in experimental and computational studies (Kopell et al. 2011; Roopun et al. 2008, 2006). The intended distinction in these studies between beta 1 and 2 is different from the common distinction of motor mu and beta rhythms in the EEG literature. Roopun et al. (2008) suggest that low frequency beta emerges as a result of a concatenation of one period of a high frequency beta with one period of a gamma oscillation. The dual β bands observed here are unlikely to be an example of this phenomenon, as the high and low bands exhibited phase coupling and appeared to be harmonically related. Because this phenomenon was not reproduced across subjects or motor areas, we focused the analyses on the low beta band for monkey R.

Given that beta power was elevated during the steady-state movement preparation periods of the CGID task, during which we observed sustained rhythmic spiking at beta

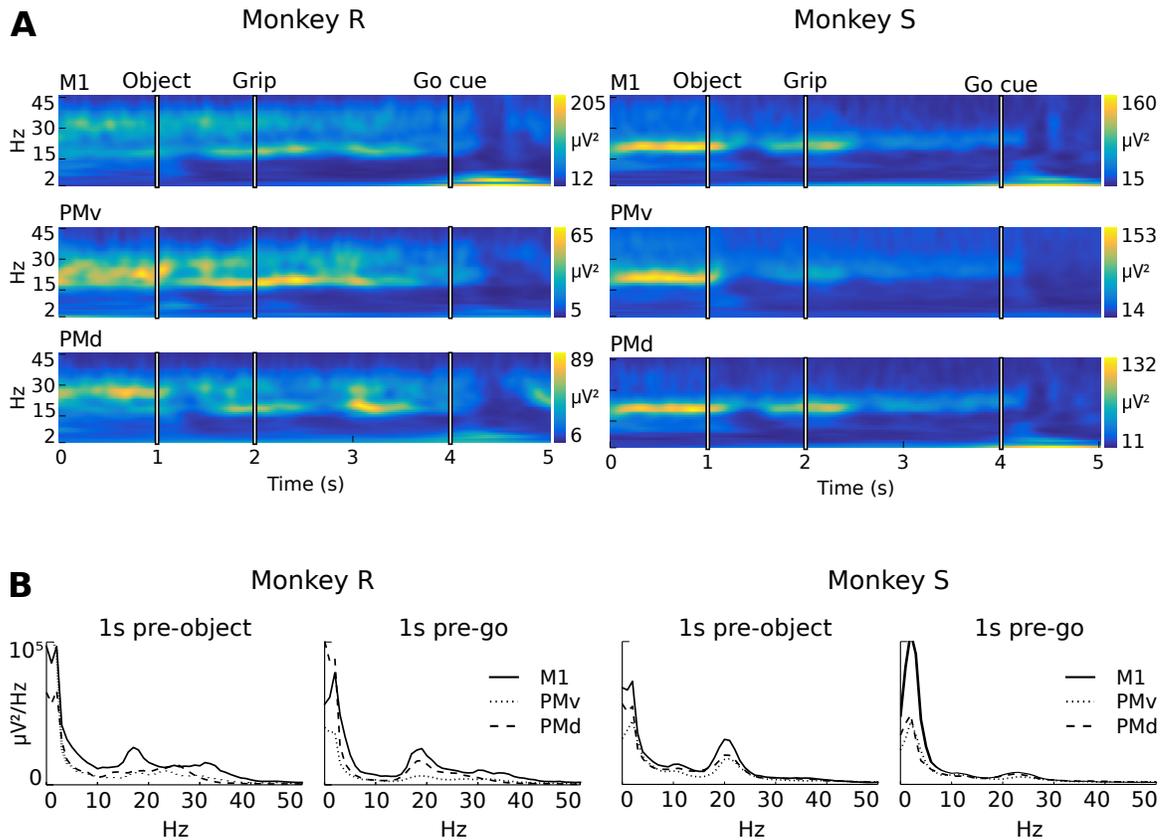


Figure 3.7: ~ 20 Hz beta oscillations are enhanced during the steady-state movement preparation periods of the CGID task. **(A)** Trial and channel averaged spectrograms show sustained beta power during the first second of the task, as well as the one second leading up to the ‘Go’ cue. Although on average beta power is enhanced during steady-state periods, LFP activity on individual trials is more variable. Additionally, monkey R showed more variability in beta power overall than monkey S. **(B)** Multitaper power spectral densities, averaged over channels and trials, show clear peaks in the beta band for both epochs. The beta oscillations in monkey R area PMd presented a complex spectrum with multiple peaks. We focused the analyses on the low beta peak at ~ 19 Hz. Monkey S displays both alpha (~ 10 Hz) and beta oscillations. For both subjects, we defined the beta band as a 5 Hz window around the peak beta frequency during the one second before the ‘Go’ cue. -

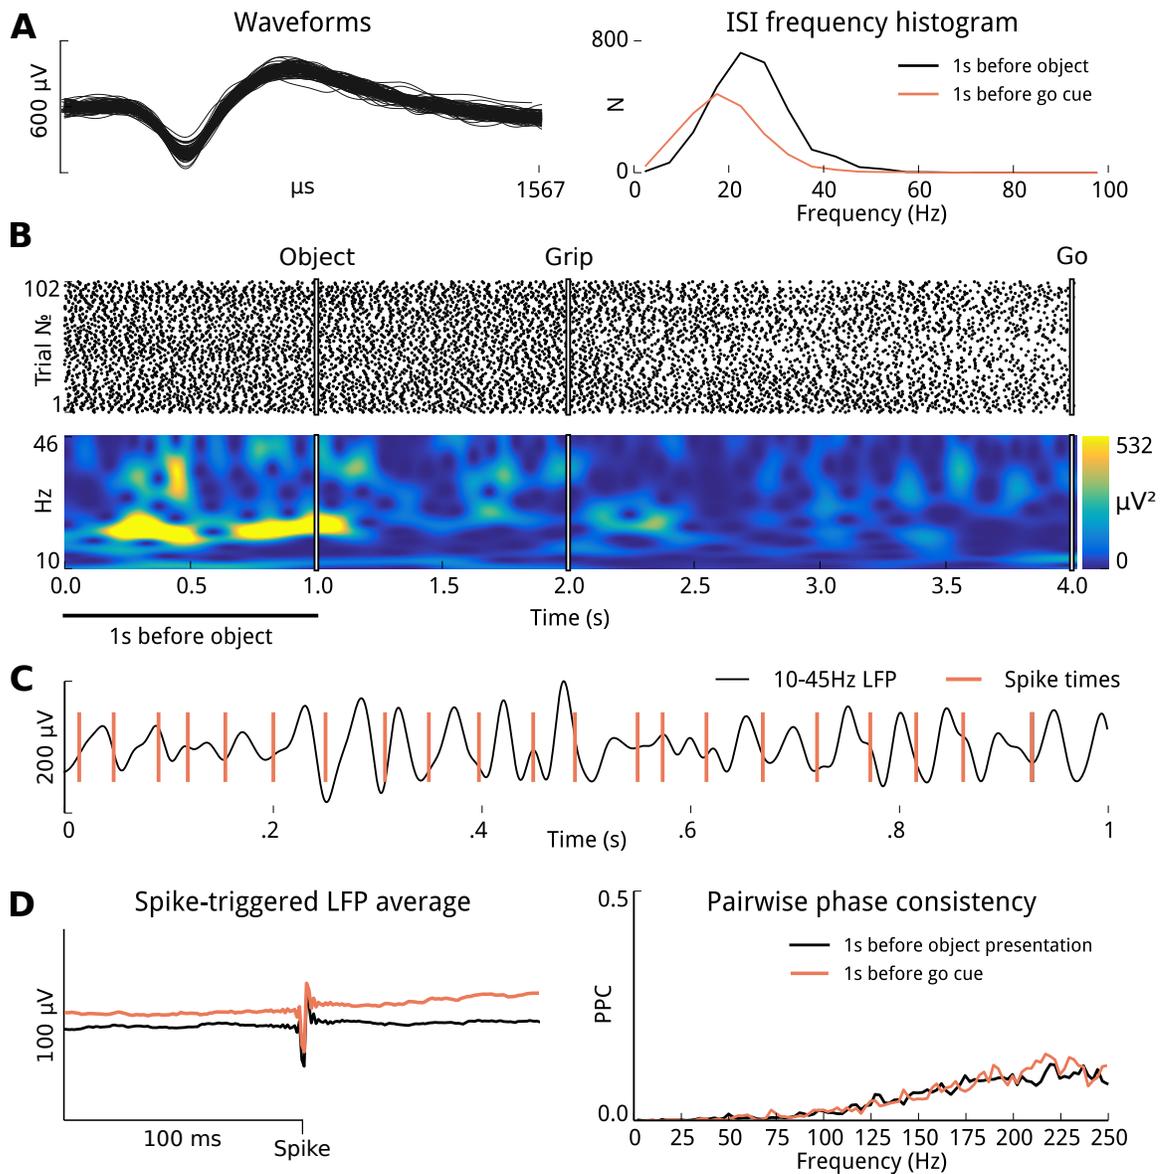


Figure 3.8: Single units exhibit β -rhythmicity that appears dissociated from the phase of the β -LFP oscillations. **(A)** Shown here is an example well-isolated unit recorded from primary motor cortex that displayed rhythmic firing at beta frequency (~ 20 Hz) during the steady-state movement preparation periods of the CGID task. **(B)** The spike raster plot shows reliable and steady single unit firing during the steady-state movement preparation periods. In contrast, the example single trial β -LFP spectrogram plot shows transient β -LFP events. **(C)** An inspection of neuronal spiking and β -LFP oscillations during the first second of this trial reveals that the phase at which single units fired relative to the β -LFP oscillations drifted, and that β -rhythmic spiking remained steady while β -LFP power fluctuated. **(D)** Spike-triggered LFP averages show primarily an artifact from spike contamination (Methods), and reveal no beta phase coupling either during the first second of the task or the one second before 'Go' cue. Pairwise phase consistency plots corroborated this finding, showing only a broad-band increase in high frequency phase coupling associated with contamination of the LFPs by extracellular action potentials (Methods).

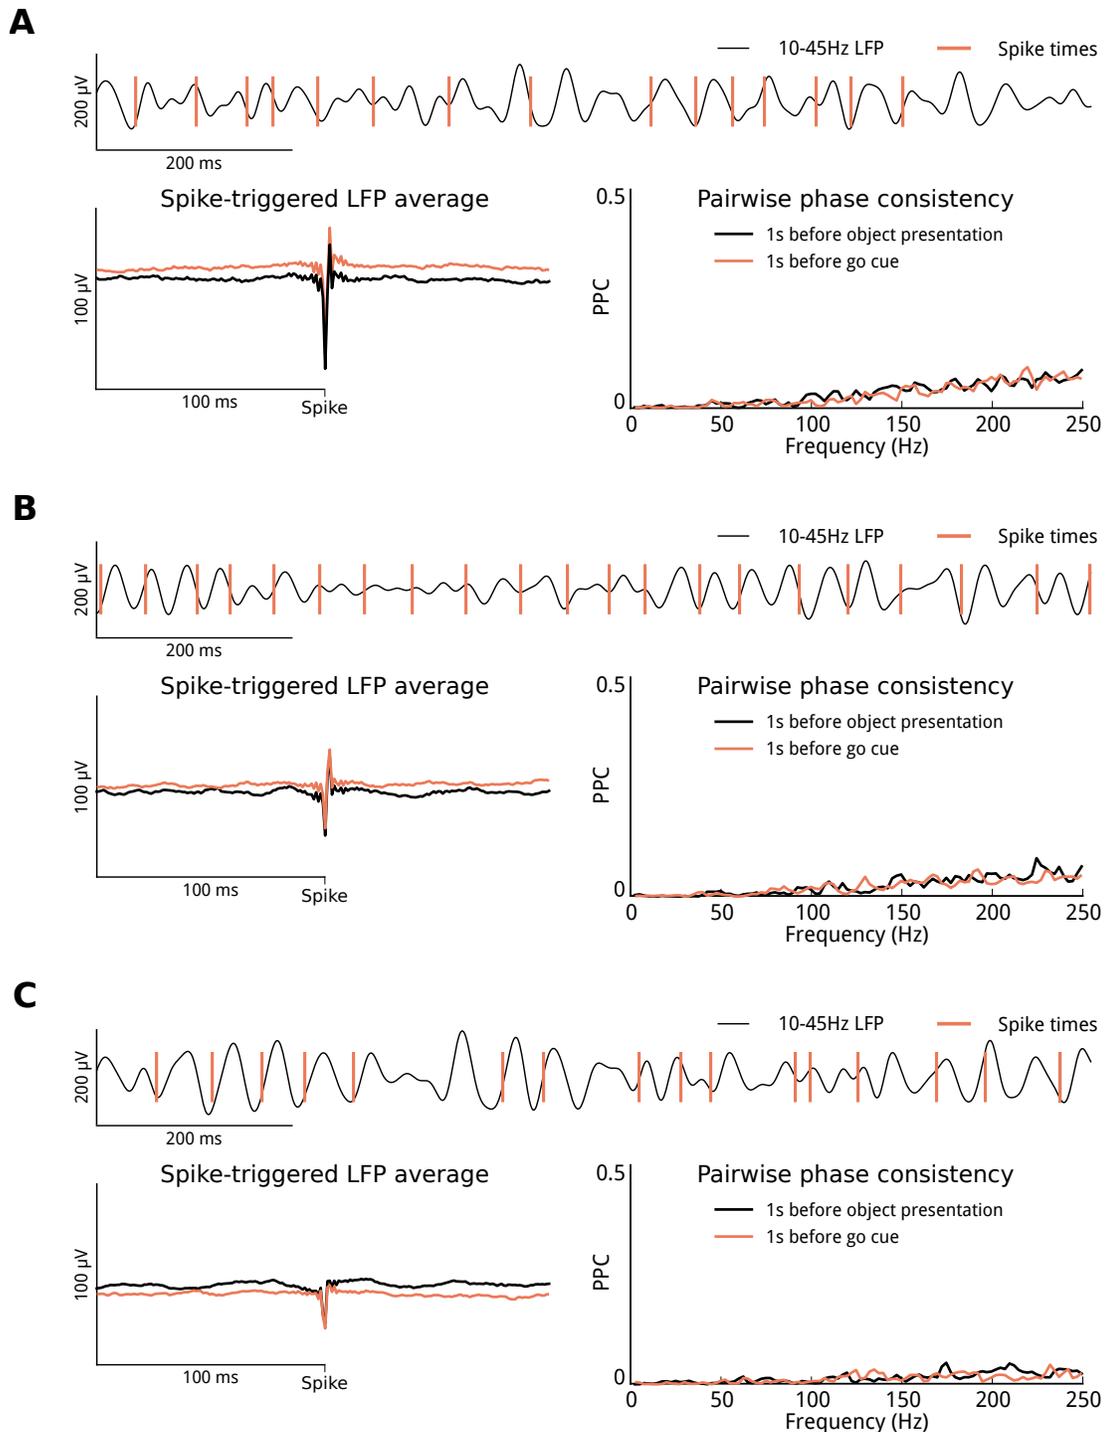


Figure 3.9: Single units exhibit β -rhythmicity that is dissociated from the phase of the β -LFP oscillations: Additional examples.. Shown here are three additional examples (A, B, C) of well isolated single units in motor cortex. Inspection of neuronal spikes and β -LFP oscillations during the first second of the task reveals that the phase at which single units fire relative to ongoing β -LFP drifts. Spike-triggered LFP averages during steady-state periods show only artifacts from spike contamination (Methods), not phase coupling to ongoing beta. Pairwise phase consistency plots confirm this finding, and show no spike-LFP phase coupling at beta frequencies. All units' spikes consisted of well isolated extracellular potentials with high signal to noise ratio and no contamination from other units. Inter-spike interval histograms, both during the first second of the task and during the one second before the 'Go' cue, showed that units fired with a preferred (mode) frequency at ~ 20 Hz.

frequency, we next assessed the extent to which single-units phase locked to the ongoing β -LFP oscillations. Inspection of example single-units showed that the β -LFP phase at which single units spiked appeared to drift randomly (Figures 3.8, 3.9). Under visual inspection it appeared that single units phase locked to the β -LFP oscillations during short transients, and that the phase precessed during these transient events. However, this could simply reflect chance events resulting from the fact that both signals exhibited strong autocorrelations at similar frequencies. Indeed, inspection of spike-triggered averages revealed little reliable phase relationship, and examples, shown in Figures 3.8 and 3.9, instead showed only a spiking artifact.

To comprehensively quantify the relationship between single-unit firing and the phase of ongoing β -LFP oscillations, we used two measures of spike-field coupling that are designed to avoid the biases inherent in STA and spike-field coherence approaches: the pairwise phase consistency (Vinck et al., 2010) (Methods: Pairwise phase consistency), and generalized linear (GLM) point-process models that expressed the conditional intensity (instantaneous spiking rate) as a function of the phase of the ongoing β -LFP oscillations (Methods: Point-process GLMs for assessing phase coupling between single neuron spiking and ongoing β -LFP oscillations). Pairwise phase consistency assesses the tendency of a neuron to fire at the same phase of the ongoing β -LFP oscillation. It ranges from 0 for no phase coupling to 1 for perfect phase coupling.

For assessing spike-LFP phase coupling, we focused on single units that showed unimodal or bimodal ISI distributions, and exhibited a preferred firing frequency (ISI mode frequency) between 10 and 45 Hz. We observed that mean firing rates were typically lower than 10 Hz, and on inspection found that rhythmic single units could skip some beta cycles (e.g. Fig. 3.9a,c). For this reason, we also required that units exhibit mean rates of at least 20% their mode frequency. Overall, 47% (125/264) of units were selected as exhibiting beta rhythmicity under these criteria. Of the selected, 23% (29/125) exhibited bimodal (bursting/rhythmic) ISIs and 77% (96/125) had unimodal ISIs. Of the units

with unimodal ISIs, 40% (38/96) exhibited oscillations in their autocorrelation functions, 40% (38/96) exhibited a non-oscillatory post-recovery rebound, and 21% (20/96) exhibited irregular Poisson-spiking with a long recovery period that placed their mode frequency in the beta range.

We found that PPC values were typically marginal (Figure 3.10a), with the median PPC for each session, area, and task epoch ranging from 0 to 0.12. Overall, 95% (118/125) of units had a PPC value smaller than 0.03 during the pre-object period and less than 0.01 during the pre-go period. No unit had a PPC value that exceeded the 95% confidence interval for the null-hypothesis PPC distribution, assessed by computing PPC between spikes and trial-shuffled LFPs. PPC values were surprisingly weak, given that one might expect the β -LFP and the β -rhythmic spiking to relate to the same ongoing network phenomenon.

As a complementary approach, we summarized phase coupling between single neuron spiking and β -LFP oscillations by assessing the conditional intensity function (CIF) phase models' ability to predict the timing of spikes (Methods). We report a measure of model performance 'predictive power' (PP), which ranges from 0 for no prediction and 1 for perfect prediction (Methods 'Point-process GLMs for assessing phase coupling between single neuron spiking and ongoing β -LFP oscillations'). In terms of phase coupling, a predictive power of zero implies no coupling, and a predictive power of 1 implies perfect phase coupling. We observed predictive power values ($PP = 2 \times AUC - 1$) as high as 0.24. During the steady state epoch preceding object presentation, 39% (49/125) units exceeded the 95% confidence interval for the null PP distribution, and during the steady state epoch preceding the 'Go' cue 19% (24/125) of units exceeded their 95% chance level. This suggests that at least some true phase coupling is present, and that the point process GLM approach exhibited greater statistical power than the PPC analysis. Although the predictive power was sometimes statistically significant, it remained extremely low for the vast majority of units, with 95% (118/128) of units exhibiting a GLM phase model predictive power less than 0.01. Thus, consistent with the PPC results, the CIF phase model found relatively

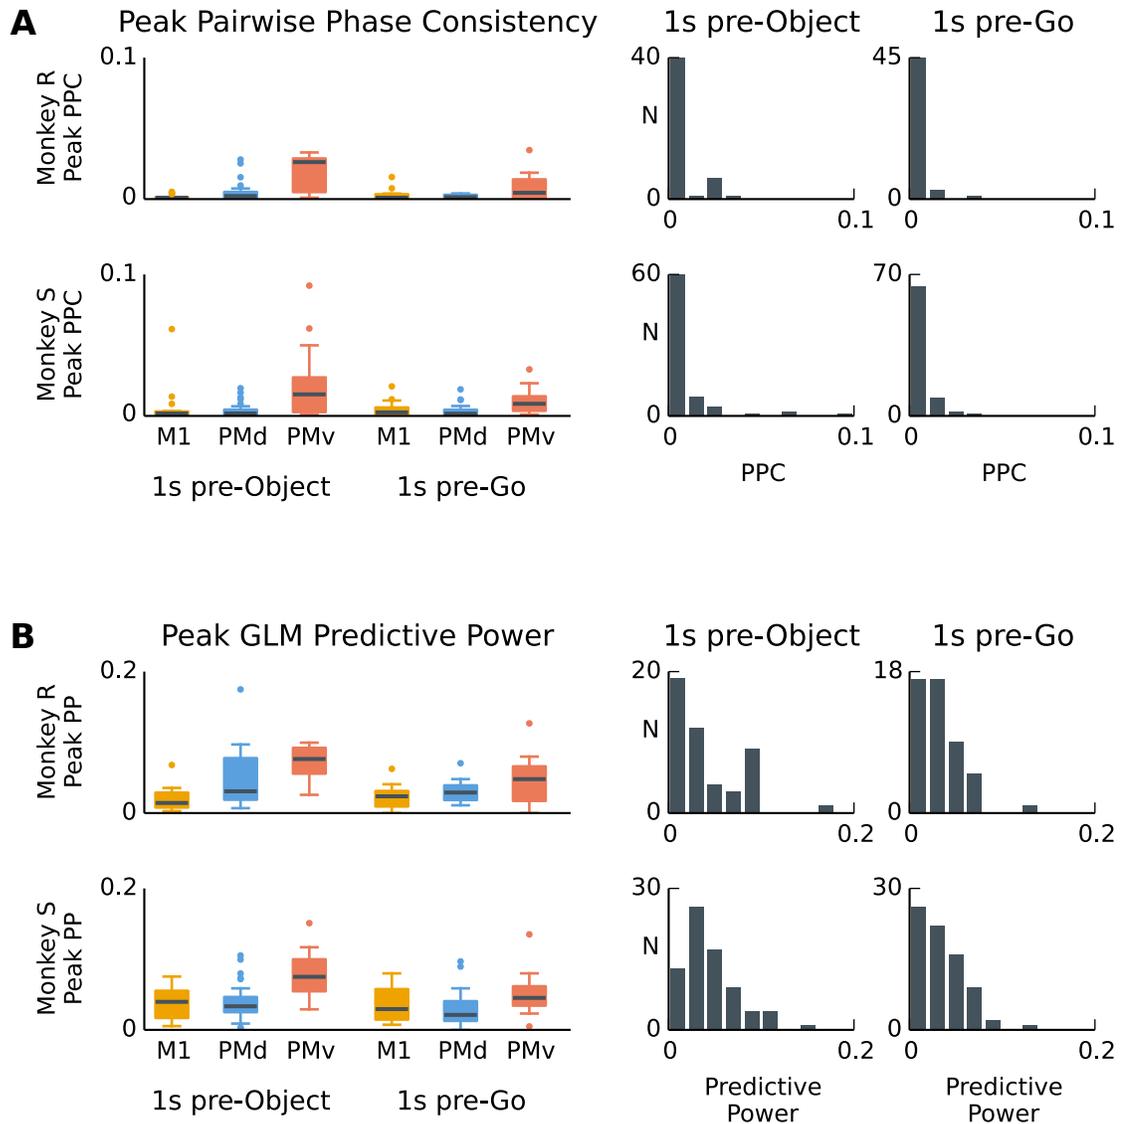


Figure 3.10: Spike-LFP phase coupling at the peak beta frequency is typically small during the steady-state movement preparation periods. **(A)** Left: box plots summarizing the magnitude of the pairwise phase consistency (PPC) value at the beta peak (Methods: ‘Pairwise phase consistency’). Each area is summarized separately for each monkey, and for two steady-state periods: the first second of the task, and the one second before ‘Go’ cue. Right: histograms representing the distribution of PPC values for each monkey in the two task epochs. All sessions and areas are combined here. Despite the spiking rhythmicity at beta and elevated β -LFP power, PPC values between spikes and LFP were typically negligible, with 95% of units showing PPC values below 0.04 for any given session or area. No units showed PPC above the 95% chance level as assessed by phase randomizing the LFP signals. **(B)** Spike-LFP phase coupling assessed by the predictive power of point process GLM phase models (see Methods: ‘Point-process GLMs for assessing phase coupling between single neuron spiking and ongoing β -LFP oscillations’) was also marginal. Although select units displayed predictive power as high as 0.24, predictive power was less than 0.1 for 95% (118/125) of units during both epochs. During the first steady-state period (one second before object presentation), the predictive power exceeded the 95% chance level confidence interval for 39% (49/125) of the units. During the second steady-state period (one second before go cue), the predictive power exceeded the 95% chance level confidence interval for 19% (24/total) of the units. We report these numbers without correcting for multiple comparisons, so 5% of units are expected to be above the 95% chance level. Point process GLM phase models were able to detect weak phase coupling that the PPC did not.

little stable phase coupling of spikes to β -LFP oscillations (Figure 3.10).

3.3.3 β -LFP appears as transient bursts

Inspection of single-trial LFP traces revealed that the dynamics of the β -LFP oscillations differed from that of the β -rhythmic spiking observed in many single units. Notably, β -LFP oscillations were not sustained across the steady state movement preparation periods of the CGID task. Instead, β -LFP oscillations occurred in transient bursts, which exhibited considerable trial-to-trial variability (Figure 3.11). We identified periods of elevated power (beta transients) as periods when the amplitude envelope of the band-pass filtered β -LFP exceeded 1.5 standard deviations of the filtered signal amplitude for at least 40 ms, i.e. approximately more than one beta cycle (Methods: Beta phase extraction and transient identification). The mean threshold for distinguishing high versus low beta power was 39 μ V with a standard deviation of 7.5 μ V. We found that there was considerable variability in the durations of these beta transients, with a mean duration of approximately 80 ms in monkey R and 90 ms in monkey S. Nonetheless, not all beta events were short, and several trials showed elevated beta power lasting for several hundred milliseconds. Given the transient nature of the β -LFP oscillations, we investigated the possibility that spike-LFP phase coupling was confined to these beta transient events.

We were unable to detect statistically significant differences between PPC values restricted to periods of high versus low beta power during the two movement preparation epochs (Figure 3.12a). Two sessions appeared to show increase significant at the $p < 0.05$ level (Wilcoxon signed-rank test). However, after correcting for multiple comparisons (Benjamini-Hochberg procedure for 6 comparisons and a false discovery rate of $\alpha = 0.05$, Benjamini and Hochberg (1995)), these increases were significant only in monkey S, session 3, for the pre-object period. Although some individual units did show larger PPC values (Figure 3.12b), only 4 out of the 125 units showed PPC increases larger than 0.1.

Additionally, we found that there was relatively little difference in firing rate statis-

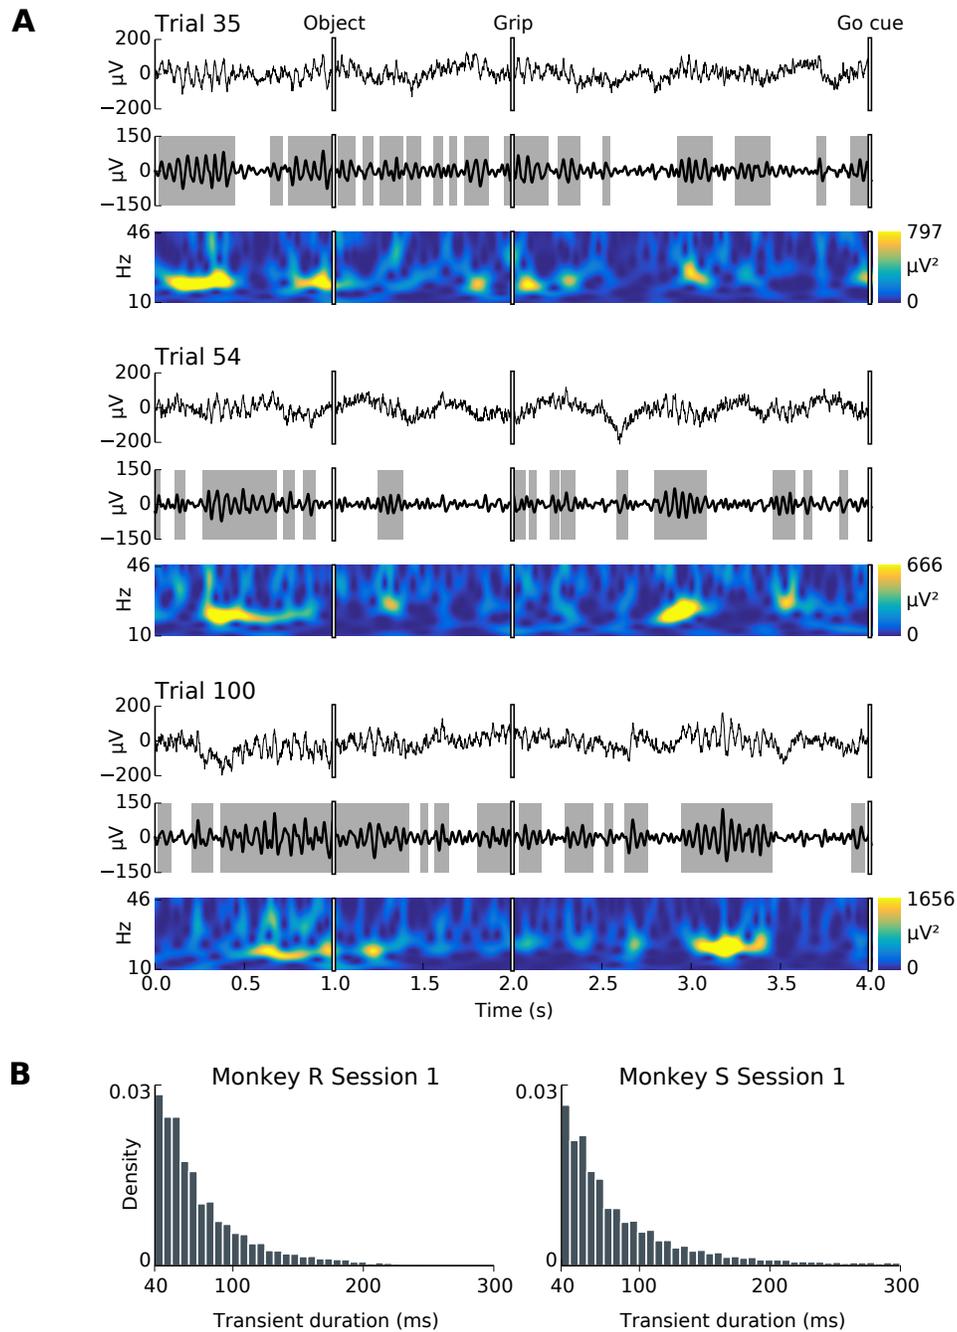


Figure 3.11: β -LFP oscillations typically occur in transients and exhibit high trial-to-trial variability. **(A)** Shown here are three representative example trials from a single session, monkey S, areas M1, session 1. In each example, the top plot shows the ‘raw’ LFP, the middle plot the bandpass filtered β -LFP, and the bottom plot shows the spectrogram. Transient beta events were defined as periods for which β -LFP amplitude was elevated ($>1.5\sigma$, shaded in gray). Inspection of β -LFP activity in single trials revealed that beta oscillations were rarely sustained, occurring as transients lasting commonly a few oscillation cycles. **(B)** However, as evidenced in histograms of the durations of high (≥ 1.5 standard deviations) beta transients, there was no characteristic duration for these transients, and periods of sustained beta oscillations lasting up to 8 or more beta cycles (e.g. > 200 milliseconds) were also observed in many trials.

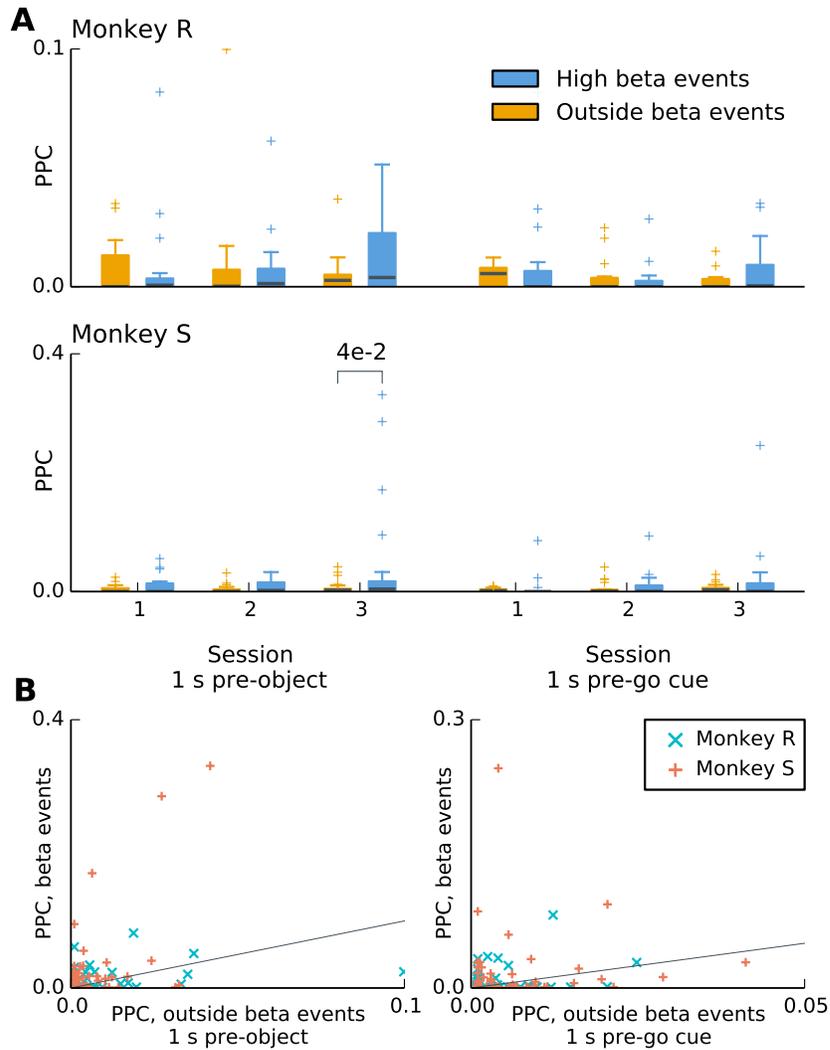


Figure 3.12: Pairwise phase consistency increases only marginally for some units when the analysis is restricted to high beta transient events. The plots show PPC values at 25 Hz computed based only during (transient) beta oscillations with high power. One session showed a significant increase in PPC during the pre-object epoch when restricting to beta events. P-values were computed using the Wilcoxon signed-rank test for difference of medians, and corrected with the Benjamini-Hochberg procedure for 12 comparisons with a false discovery rate of $\alpha=0.05$.

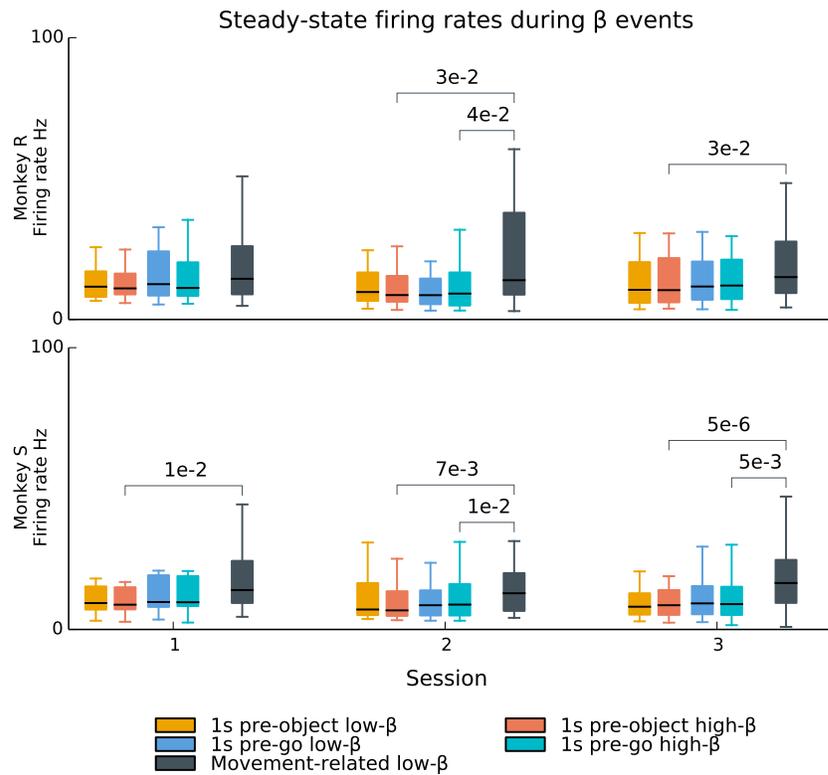


Figure 3.13: Single unit firing rates during steady-state movement preparation periods are not affected by beta transients. During the steady state movement preparation periods of the CGID task, β -LFP oscillations occurred as transient events, with the beta power exhibiting transient bursts. In contrast, rhythmic single-unit spiking at beta frequencies was sustained. Single-unit firing rates did not change between high-beta (Hilbert amplitude $> 1.5\sigma$) and low-beta time periods (colored bars) during these steady-state periods. On the other hand, beta suppression associated with movement execution (low beta periods in the 1 s post ‘Go’ cue, black bars) was typically associated with increased firing. P-values were computed using the Wilcoxon signed-rank test for difference in medians, and corrected for 24 multiple comparisons (4 comparisons per session: high-low beta within each steady state epoch, and high beta in each epoch to movement-related low-beta) using the Benjamini-Hochberg procedure for a false discovery rate of $\alpha=0.05$. This result suggests that the transient beta power fluctuations during steady-state movement preparation periods may arise from a different mechanism than the power fluctuations (beta suppression) associated with visual cue presentation and movement execution.

tics during beta transients compared to periods outside beta transients (Figure 3.13). In contrast, firing rates were significantly higher during movement-related beta suppression (low-beta periods in the 1 s post ‘Go’ cue), showing statistically significant increases between the pre-object and movement period in 5 out of 6 sessions, and between the pre-go period and movement in 3 out of 6 sessions. (Wilcoxon signed-rank tests for difference in the median, corrected for 24 multiple comparisons using the Benjamini-Hochberg procedure for a false discovery rate of $\alpha=0.05$.) This finding indicates that the modulations in β -LFP power during steady-state movement preparation periods were not coupled to changes in the firing rates of the underlying neuronal population, as was the case during the movement execution and visual cue related beta suppression.

3.3.4 Population spiking activity also shows only weak coupling to ongoing β -LFP oscillations

We examined the possibility that the phase coupling between spiking and the ongoing β -LFP oscillations could be too weak to be detected, but much stronger if assessed at the level of the population spiking activity. Population spiking activity was defined here as the total number of spikes (1 ms time bins) summed across the well-isolated single units within a given motor area, smoothed by a 25 ms boxcar filter (Methods). For each motor area we computed the cross-correlation function between the population spiking activity and the β -LFP averaged across the channels in the area. Cross-correlation functions were computed for time lags ranging over one beta cycle (± 25 ms). A cross-correlation function was computed for each motor area, epoch, session and subject.

The extrema of the cross-correlation functions between population spiking activity and the mean β -LFP were small, ranging from 0.0039 to 0.042. After correcting for 36 (monkey, session, area, epoch) comparisons using the Benjamini-Hochberg procedure with a false discovery rate of $\alpha=0.05$ (Benjamini and Hochberg, 1995), three correlations were statistically significant, all in monkey S. Monkey S area PMv showed significant

correlations of 0.036 and 0.033 for sessions 1 and 3, and monkey S area M1 showed a significant correlation of 0.042 for session 3. P-values were obtained from a chance level distribution: crosscorrelation function peaks were computed from resampled data generated by shuffling the LFP trials (2000 resamples).

3.3.5 Single units show weak coupling to measures of population activity

Previous studies in sensorimotor cortex have demonstrated strong coupling of single neuron spiking to both the population spiking activity (Aghagolzadeh and Truccolo, 2014, 2015; Okun et al., 2015) and ensemble spiking histories (Truccolo et al., 2010) during sensory stimulation and execution of motor actions. In particular, Aghagolzadeh and Truccolo 2015 showed that, in the same datasets examined here, single neuron spiking was strongly coupled to low-dimensional representations of the neuronal ensemble activity during movement execution of the CGID task. For completeness, we thus also considered the possibility that spiking could be only weakly coupled to the ongoing β -LFP, but at the same time show strong coupling to other measures of the population activity during the movement preparation epochs.

Using point process GLM analysis (Methods ‘Point-process GLMs for relating single neuron spiking to population spiking activity’), we found that single neuron spiking was only weakly related to the population spiking activity during the steady state movement preparation periods (Figure 3.14A). In contrast, and consistent with our previous work ((Rule et al., 2015); Chapter 2), predictive power was higher during the one second movement phase following the ‘Go’ cue.

Qualitatively similar results were obtained when using a different measure of population activity consisting of multi-unit activity (MUA), defined as >250 Hz LFP amplitude fluctuations, bandpass filtered in the 5 Hz band surrounding the peak beta frequency (Methods). Specifically, median predictive power (PP) values during the two movement

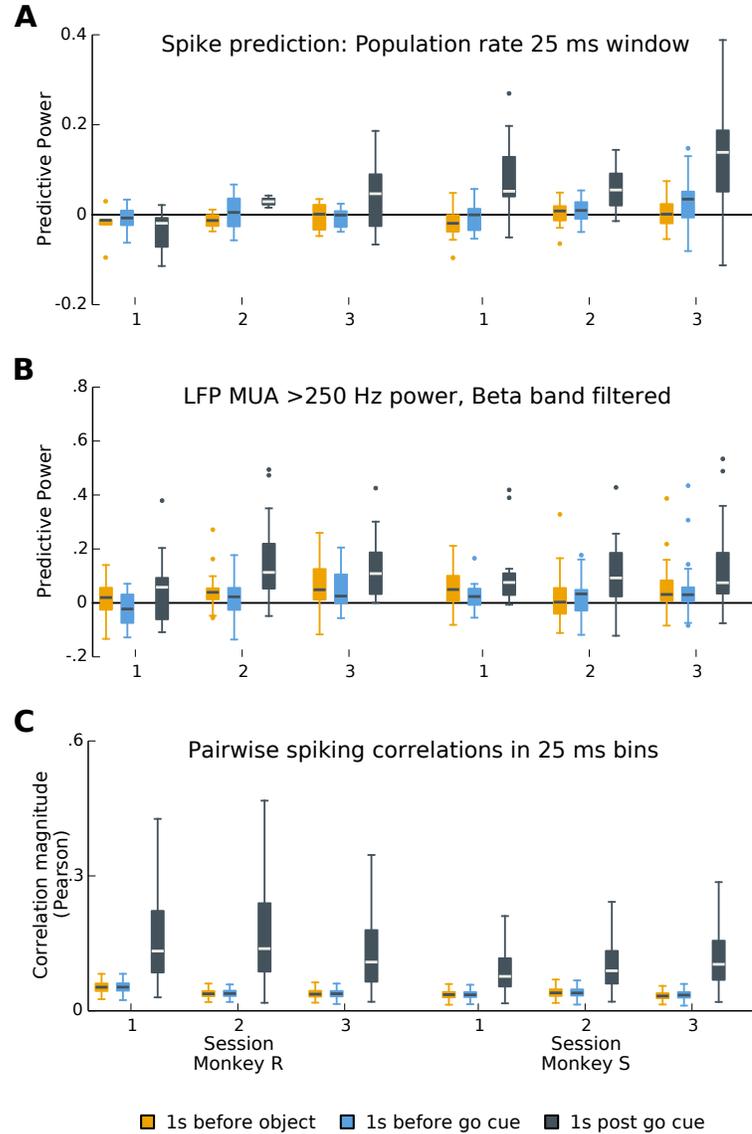


Figure 3.14: *Contrasting collective neural dynamics between steady-state movement preparation and movement execution periods.* (A,B) Predictive power of point process GLM-CIF models for β -rhythmic single neuron spiking based on the population spiking activity (Methods; under cross-validation). (A) Spiking prediction based on the population spiking activity measured on the same MEA (excluding the unit being predicted). Each box plot summarizes the distribution of predictive power values for one session and epoch. (Values have been pooled across motor areas.) Predictive power during steady-state movement preparation periods (colored bars) was typically distributed around chance. In contrast, population spiking activity predicted single unit spiking above chance levels during the 1 s movement execution period following the ‘Go’ cue. (black bars). (B) Single neuron spiking prediction based on MUA using the same point process GLM approach as above. MUA was defined as the > 250 Hz LFP amplitude envelope, further bandpass filtered in the 5 Hz band surrounding the peak beta frequency. Predictive power trends are similar to those obtained in (A) for the population spiking activity. (C) Box-plots summarizing the distribution of pairwise (Pearson) correlation coefficients. Spiking activity for each single unit was represented in 25 ms binned spike counts. Pairwise cross-correlation functions were then computed and the extremum of each function was chosen as the pairwise correlation coefficient. (Outliers are not shown). Pairwise correlations were weaker during steady-state movement preparation periods compared to the movement execution period.

preparation epochs (pooled across motor areas) ranged from 0 to 0.05 (pre-object period) and pre-go period median PP ranged and from -0.02 to 0.03 (pre go cure period). During the movement period, median predictive power values ranged from 0.06 to 0.11. Predictive power values during movement were statistically significantly higher than those in the pre-object period in two sessions for monkey S; and higher than those in the pre-go period in one session in monkey R and all sessions in monkey S. (Wilcoxon signed-rank test with Benjamini-Hochberg correction for a FDR of 0.05 for 12 comparisons.) This analysis confirms that single-unit spiking can be predicted from measures of population activity in the CGID task, but that this predictive information was relatively weaker during the steady-state movement preparation periods.

Finally, we investigated whether multi-unit activity might show more substantial phase coupling to β -LFP. We examined two measures of multi-unit activity: (1) all threshold crossings (unsorted spikes) occurring on the same channel and the four nearest neighbor channels (spiking-MUA), summed in 1ms bins and (2) The amplitude envelope in >250 Hz filtered LFP as described previously (LFP-MUA). Beta coherence between these measures of multi-unit activity and the β -LFP on the same channel were weak: We found a statistically significant coherence peak between β -LFP and LFP-MUA in 4

Overall, the above results show a stark contrast between collective dynamics during steady-state movement preparation periods, where spiking activity appears to be much more asynchronous, and collective dynamics during movement execution, where both the ability of population activity to predict single neuron spiking and pairwise correlations are much higher.

3.3.6 Thin- and thick-spike cells show similar weak phase-coupling to β -LFP oscillations during preparatory steady-states

We found that isolated single units in our recordings could be clearly separated (as assessed by visual inspection) into a population of thin- and thick-spike cells (Figure 3.15a,b;

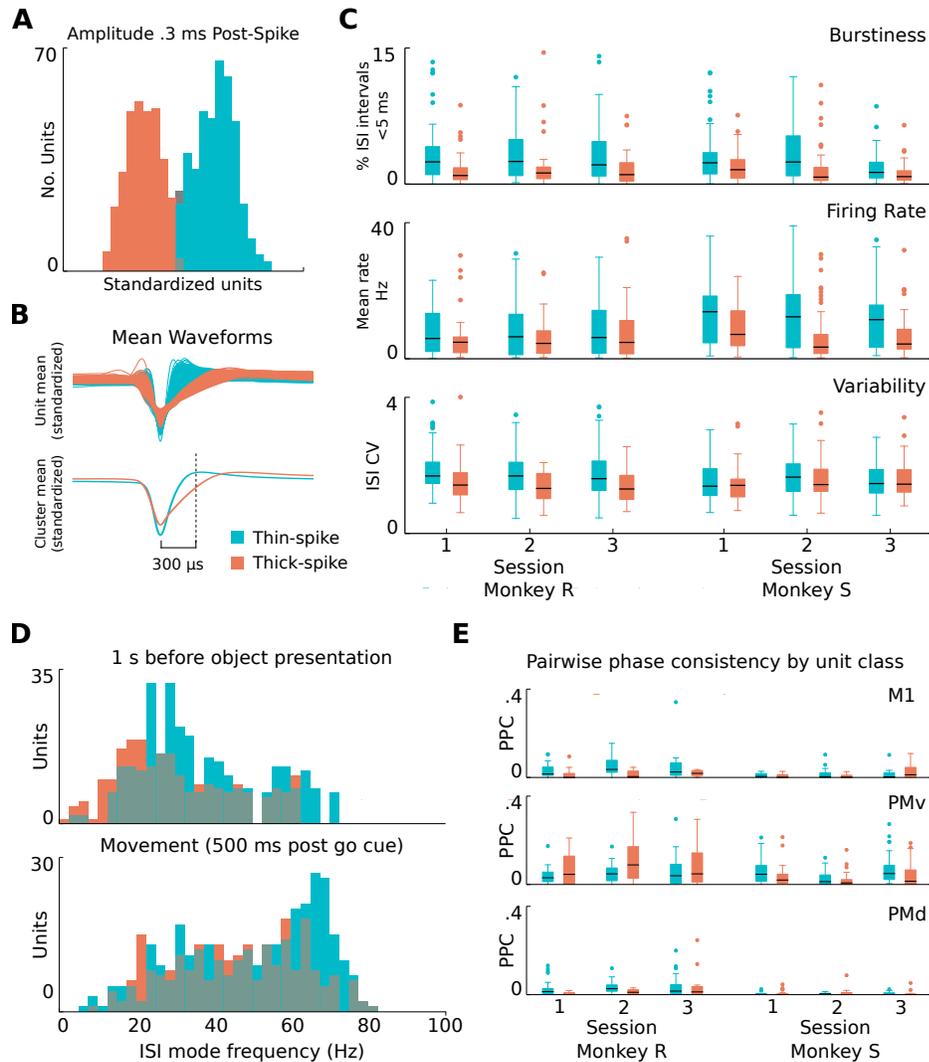


Figure 3.15: Single units cluster into thin- and thick-spike groups, but these groups show no consistent differences in PPC values. Spike waveforms recorded in motor cortex exhibit a diversity of spike widths that cluster into two main groups. **(A)** Histograms show the clustering of well isolated units according to spike widths, including all sessions, areas, and monkeys. Clustering based on the amplitude of the waveform 300 μ s after the spike peak provided better separation than the traditional approach of estimating the spike width at half maximum. Thin-spike units are denoted in blue, thick-spike in red. **(B)** Traces of the mean waveform for thin-spike and thick spike units illustrate the differences between the unit classes. **(C)** Thin- and thick-spike units appeared to show consistent differences in firing statistics. On average, thin-spike units fired more bursts (top), fired at a higher mean rates (middle), and exhibited higher coefficients of variation (bottom). However, none of the apparent differences were statistically significant. (Mann-Whitney U test with Benjamini-Hochberg correction for a FDR of 0.05, for dependent samples and 18 comparisons.) **(D)** Preferred spiking frequency (ISI mode) of thin-spike and thick-spike units changed with different CGID task stages. Units for all sessions, monkeys, and areas were combined in these summary histograms. During the first second of the task, both thin- and thick-spike units fired rhythmically around beta frequency. During movement execution, firing rates increased on average, but the increase was most notable for thin-spike units. **(E)** Beta-peak PPC values did not show any consistent trend in the differences for thin versus thick spike units. Furthermore, none of the differences were statistically significant. (Mann-Whitney U test with Benjamini-Hochberg correction for a FDR of 0.05, for dependent samples and 18 comparisons.)

Methods: Unit categorization). These two classes appeared to exhibit differences in firing statistics. Overall, thin spike neurons exhibited more short-ISI events (<10 ms) indicative of bursting, fired at higher rates, and greater coefficients of variation (Figure 3.15c). The ISI mode frequency of thin-spike units appeared typically slightly higher during the steady-state movement preparation periods. In addition, thin-spike units appeared to show a higher increase in their mode firing frequency during movement as compared to thick-spike neurons (Figure 3.15d). While these apparent differences between the two classes, even though consistent across monkeys and sessions, were not found to be statistically significant. (Mann-Whitney U test with Benjamini-Hochberg correction for multiple comparisons for positively dependent samples and a FDR of 0.05). Additionally, no consistent differences were found between thin and thick spike units with respect to spike and β -LFP phase coupling. Trends were broadly stable across sessions for a given area in a given monkey, but were not consistent between monkeys and areas: for example, some areas (e.g. monkey R area PMv) showed greater spike-LFP phase coupling for thick spike neurons, while other areas (e.g. monkey S area PMv) showed the opposite (Figure 3.15e).

3.4 Discussion

In this study, we found a remarkable dissociation between steady-state β -rhythmic spiking and transient β -LFP oscillations in motor cortex of rhesus macaques during the movement preparation periods of a visually cued reaching and grasping task. We observed that single-unit firing during steady states was rhythmic and sustained. In contrast, β -LFP oscillations appeared as transients that exhibited high trial-to-trial variability. The fact that single neuron firing rates were not affected by, or correlated with, the occurrence of transient β -LFP events or β -LFP amplitude suggests that the modulations in β -LFP power during these steady-states did not result from changes in the level of beta rhythmicity in the underlying neuronal population, as is the case for movement and visual cue related beta suppression. Furthermore, two complementary measures of spike-LFP phase cou-

pling (pairwise phase consistency and predictive power of point process GLMs) showed that the coupling was at chance level for the majority of the neurons. This dissociation between steady rhythmic spiking and β -LFP oscillations has implications for understanding the multi-scale (single neuron and ensembles) dynamics underlying the generation of β -LFP activity, and for understanding the functional role of beta oscillations in motor cortex, including putative roles for beta in modulating communication among cortical areas and phase coding. Our findings also contribute to the characterization of the statistical properties of neocortical electrical signals recorded via microelectrode arrays.

Precedence for dissociation between single neuron spiking activity and narrowband LFP oscillations Previous experimental and theoretical work (Ardid et al., 2010; Brunel and Wang, 2003; Geisler et al., 2005; Hoseini and Wessel, 2015; Truccolo et al., 2014, 2011) has found that sustained gamma LFP oscillations can coexist with irregular and asynchronous spiking activity, in which the oscillation emerges as a collective mean field. We emphasize that the phenomenon reported here differs, in that both spiking and LFP exhibit strong rhythmicity in a narrow band, and yet show relatively weak phase-coupling. We also emphasize that the phase coupling measures adopted here would detect strong phase coupling even if single neuron spiking was strongly phase locked to a LFP oscillation, but skipped most cycles of the oscillation.

A precedent for the dissociation between single-unit rhythmicity and ongoing LFP oscillations in primate motor cortex reported here has been found before in the hippocampus and has since puzzled the field. During spatial navigation, the place cells of the rat hippocampus spiking drifts with respect to the phase of the ongoing theta LFP oscillations. This phenomenon has been studied as the “theta phase precession” (e.g. Harvey et al. (2009)). The spiking phase relative to theta LFP depends on animals past, present, and planned location. Averaged over time in this case, units show relatively weak phase coupling. In particular, Harvey et al. (2009) has found that, despite the spiking phase drift with respect to the ongoing narrow-band theta LFP oscillations, a place cell’s spiking is

strongly phase coupled to the theta oscillations of its intracellular membrane potential. Their finding indicates that the “local” oscillation that impinges upon single unit spiking can be dissociated from the population oscillatory signal reflected in the LFP.

In addition, Witham and Baker (2007) found that the level of β -LFP power in a given area need not correlate with the corresponding single unit rhythmicity in the same area, and Baker et al. (2003) observed relatively weak spike-field coherence in beta during a precision grip task. We emphasize that our work goes beyond these studies in examining the movement preparation and instructed delay periods preceding naturalistic reaching and grasping actions, as opposed to an isometric precision grip task. We also focused our analysis on units that exhibited sustained spiking rhythmicity, and find that phase coupling remains weak even when the analysis is restricted to transient periods of elevated beta activity. The coexistence of sustained β -rhythmic spiking with β -LFP transients, as well as the relatively weak phase coupling of single units to the β -LFP, and to the mean population activity, are important features that should be recapitulated in computational models of motor cortex.

Statistical caveats When both LFP and spikes exhibit strong autocorrelations in the form of narrow-band oscillations, there is risk of detecting apparent phase coupling by chance. This is true even for estimators such as the pairwise phase consistency, which correct for spike-rate biases. We addressed this problem by obtaining empirical chance level distributions through phase-randomization and shuffling of trials. Nevertheless, spurious contributions of these temporal correlations to inflated phase coupling assessments would only reinforce the points made here. As stated above, both the PPC and point process GLM phase coupling assessments are capable of detecting a preferred phase of firing relative to the β -LFP even when cells do not spike on every cycle. This is because the PPC relies primarily on the distribution of spike-triggered LFP phases (this also true for point process GLMs, but for slightly different reasons), and a phase locked unit that fires only occasionally will still exhibit a concentrated spike-triggered distribution of LFP phase.

Localization of β -LFP A natural question is whether the apparent dissociation between spiking and β -LFP oscillations could result from LFPs being nonlocal (Kajikawa and Schroeder, 2011). "However, we observe large phase differences between adjacent electrodes during complex beta wave events, as it will be demonstrated in Chapter 4.

The lateral localization of the β -LFP in our MEA recordings is likely quite good during such events, on the order of the resolution of the MEA (400 μm or smaller). As a cautionary note, however, this does not exclude the possibility that local beta oscillations mix with remote sources during more globally synchronous states. Another potential issue is that the single units and the sources of the LFP signal were localized in different cortical layers. Identification of the laminar origin of β -LFP is not possible with the MEA recording setup employed here due to the possibility of LFP conducting between layers. Laminar probe recordings and modeling work may be necessary to precisely localize the origins of β -LFP, although previous studies have suggested that current sinks and sources in layer V pyramidal neurons are the origin of motor cortex beta (Murthy and Fetz 1996a; Witham and Baker 2007). Given the uncertainty about the depth of the MEA implant, it is possible that the single units we recorded were from layer II-III, and that single-unit spiking activity could then be dissociated from β -LFP arising in layer V. This alternative scenario of dissociation between beta rhythmic spiking in motor cortex superficial layers and β -LFP in deep layers would nevertheless remain an important finding. Despite these caveats, our findings raise important questions about the interpretation of single neuron spiking activity and LFPs recorded from MEAs, with relevant implications for understanding the origin of β -LFP oscillations.

Origins of β -LFP transients Our data point to an important feature of motor cortex beta oscillations: β -LFP power fluctuates during steady-state movement preparation periods in our task, even while the firing rates of beta-rhythmic single neurons remains constant. We conjecture that the observed fluctuations in β -LFP power during movement preparation could arise from changes in the synchronization among more local sources

of β -rhythmic network activity. It is possible that β -LFP power fluctuations represent transient synchronization of a large population of weakly coupled single units, such that, although the macroscopic LFP power exhibits a transient amplitude increase, individual spike-LFP phase coupling remains weak. Alternatively, beta transients may reflect additional synaptic input from other areas that weakly entrain ongoing beta rhythmic spiking. These scenarios are to be contrasted with the attenuation of β -LFP power observed when neuronal activity undergoes substantial changes in firing rates during movement execution. During movement execution, the majority of units exhibit large excursions in firing rate and many rhythmic single units shift their firing frequencies up and out of the beta frequency band, and other units may switch from rhythmic to Poisson firing. Therefore, it is likely that movement-related beta suppression relates to a reduction of total β -rhythmic network activity. This points to the possibility that there are two processes governing variability of β -LFP power in motor cortex: an overall modulation of the level of β -rhythmicity that is evidenced by changes in single-unit firing properties during movement execution, and an additional source of variability that gives rise to the transient fluctuations in β -LFP power despite sustained β -rhythmicity at the level of single neuron spiking during movement preparation steady-states.

Neuronal types and the origins of the SUA β -rhythmicity The origin and function of sustained β -rhythmic spiking across movement preparation remains puzzling. The microelectrode arrays used in this study record preferentially, but not exclusively, from pyramidal neurons. One possibility is that inhibitory interneurons might show a stronger coupling with the ongoing β -LFP oscillations. The clustering of single units into the two classes of thin- and thick-spikes shown here is commonly associated with the classification into putative inhibitory interneurons and principal cells, respectively (Barthó et al., 2004; McCormick et al., 1985). However, the unique features of pyramidal tract neurons (PTNs) makes identifying putative inhibitory interneurons vs. excitatory pyramidal cells from extracellular spike width and firing properties challenging.

Some PTNs can show higher firing rates and thin spikes and can be mistaken for fast spiking interneurons (Vigneswaran et al., 2011). More advanced approaches that identify or manipulate specific neuronal subtypes will be needed to clarify the relation between single-unit beta rhythmic spiking and β -LFP.

Previous computational and experimental studies on the origin of beta oscillations have emphasized a variety of mechanisms ranging from the role of thalamic inputs (e.g. Jones et al. (2009) in somatosensory cortex) to more local or intrinsic features of cortical dynamics (e.g. Kopell et al. (2011), Roopun et al. (2006)). Regarding the latter, Kopell et al. (2011) proposes that the beta1 rhythm (~ 15 Hz) in rat association cortex arises as a consequence of rebound from inhibition, and can therefore be maintained without strong collective activity, in contrast to the pyramidal-interneuron gamma (PING) mechanism of gamma oscillations which requires feedback from inhibitory interneurons. Roopun et al. (2006) (in vitro neocortical slices from rats) also find a 20-30 Hz rhythm in layer V pyramidal tract neurons that depends on intrinsic currents, and is synchronized by gap junctions. Thus, it is plausible that β -rhythmicity may be supported by the subthreshold dynamics of single-units, potentially related to the slow afterhyperpolarizations identified by Chen and Fetz (2005) in type III rhythmic neurons. Additionally, it is possible that beta oscillations are mediated by collective network dynamics and reverberative E-I feedback, but that these dynamics are highly local and only weakly coupled to the population activity observed in the LFP. The possibility that β -rhythmic spiking arises from very local or even intrinsic unit properties opens another potential explanation of the transient nature of β -LFP oscillations. In the more general scenario of weakly coupled oscillators, Popovych and Tass (2011) found that, when oscillators with heterogeneous frequencies are driven by a common oscillatory input, transient power fluctuations are expected to result from momentary synchronization between oscillators, in a mechanism akin to the beats heard from two slightly out of tune notes.

Future work

Identifying functional or computational roles of β -rhythmicity in motor cortex remains a major open problem. Vargas-Irwin et al. (2015) examined the same datasets considered here from the perspective of neural decoding, and demonstrated that information in neural firing rates about planned upcoming movements emerges following visual cue presentation, and is sustained across the steady-state epochs. We have seen here that steady-state movement preparation epochs are associated with transient β -LFP oscillations, which are apparently dissociated from single-unit spiking β -rhythmicity. The decoding results demonstrate that differential firing rates encode task-relevant information during steady-state periods where we also see widespread β -rhythmicity and β -LFP transients. If information is encoded in the distribution of rates across a population of neurons, those neurons must necessarily exhibit diverse firing rates. If neural firing is rhythmic, as we observed here, then the relative phases of neurons at different rates will necessarily drift continuously in and out of phase. It is an intriguing problem whether the widespread β -rhythmic firing somehow has a facilitatory role in maintaining the preparatory distribution of firing rates while preventing the transition into dynamical trajectories of movement execution. Future studies should examine the relationship between transient β -LFP events and the temporal evolution of the encoding of motor plans during movement preparation.

Further, it will be important to examine the coupling between single neuron activity and β -LFP oscillations in instructed delay tasks that enforce the need for working memory, something not required in the task examined here, since the visual cues remained available until the go cue. This additional instructed delay condition might elucidate which features of beta activity relate to the active maintenance of the preparatory state versus simply the hold condition prior to movement execution.

The extent to which motor cortex beta rhythmic spiking arises from nonlocal oscil-

latory network inputs, local recurrent dynamics, or the intrinsic electrical properties of single neurons, remains unclear. Combined extracellular and intracellular in-vivo recordings akin to those performed by Harvey et al. (2009) may be illuminating.

In summary, the dissociation of single-unit β -rhythmicity and β -LFP reported here, both in terms of power modulation and phase coupling, is a surprising finding which has not been thoroughly investigated. It is possible that the nature of the beta states revealed here allows multiple cell assemblies, each resonant at slightly different frequencies, to coexist with relatively little interference or competition. Future work is needed to evaluate the functional importance of beta phase and frequency diversity during preparatory steady-states in motor cortex, especially with respect to evaluating potential roles for this diversity in encoding, and in gating or binding together functional assemblies of neurons (e.g. Maris et al. 2016).

Chapter 4

Phase diversity and spatiotemporal wave dynamics in primate motor cortex local field potentials

MICHAEL RULE

The experimental data examined here were collected in the Donoghue Lab. We thank Carlos Vargas-Irwin, Lachlan Franquemont, Jonas Zimmermann, and John Donoghue for sharing the data.

Abstract

Transient spatiotemporal patterns in beta (~ 20 Hz) local field potential (LFP) oscillations are a prominent feature of primate motor cortex. β -LFP power is typically elevated during movement preparation, suppressed around movement onset, and enhanced during isometric force tasks. Characterizing the spatiotemporal dynamics of these beta oscillations may shed light on their role on both normal and pathological brain activity. Previous studies have focused on traveling beta plane waves. However, little is known about the diversity in phase relations and variety of beta wave spatiotemporal patterns. Here, we systematically characterized motor cortex β -LFP spatiotemporal activity using microelectrode arrays (MEAs) in *m. mulatta* during a cued reaching and grasping task with instructed delays. Data from two monkeys were analyzed, each with a 96-MEA in ventral premotor cortex (PMv), and two 48-MEAs in the primary motor cortex (M1) and dorsal premotor cortex (PMd), respectively. Our main findings are fourfold: (1) In contrast to previous studies, we find in M1, PMv and PMd a variety of transient beta spatiotemporal patterns ranging from states with no clear wave dynamics, to states with complex wave dynamics, to globally organized states consisting of traveling waves or homogeneous synchrony. (2) Wave patterns appeared to transition continuously from desynchronized states, to more complex wave activity, to globally synchronized states. (3) Beta waves appeared to lack a characteristic wavelength. Instead, the spatial scale of beta waves depended on the beta state. There was a continuum between desynchronized activity at low amplitudes and spatially synchronous states at higher beta amplitudes. (4) We compared the above spontaneous beta wave dynamics to activity induced by optogenetic 1-second constant stimulation pulses in a third subject. Induced spatiotemporal patterns were globally organized, consisting of simultaneous radiating waves and localized plane waves at ~ 50 Hz and ~ 100 Hz, respectively. In contrast to beta waves, induced waves exhibited a characteristic wavelength of 2-4 mm. This finding is consistent with the hypothesis that spatiotemporal beta oscillations result from the synchronization dynamics of ongoing, local beta oscillations, while the gamma

waves are true traveling waves induced by a local optogenetic perturbation. Overall, our findings demonstrate that beta spatial dynamics are much more diverse than previously characterized. Spontaneous transitions among patterns of beta activity may result from (a) fast modulation of lateral connectivity in a system of coupled oscillators, (b) changes in spatiotemporal inputs to motor cortex, and (c) stochastic transitions between spatiotemporal attractors in a multi-stable system. The diversity in beta spatiotemporal dynamics may have implications for theories of spiking information transfer during beta wave, and for constraining spatially extended models of motor cortex dynamics.

4.1 Introduction

Beta-frequency oscillations at ~ 20 Hz are a prominent feature of the motor cortex local field potential, appearing during postural maintenance, isometric force, motor readiness, and during instructed delays. Rubino et al. (2006) found that beta oscillations organize as traveling plane waves, and hypothesize that such waves correlate with directed communication between motor areas. Subsequently, Takahashi et al. (2015) demonstrated that traveling waves correlate with directed functional connectivity, supporting the hypothesis that spatiotemporal organization of beta waves may play a role in communication in motor cortex. Although previous studies of motor cortex beta oscillations have focused on traveling plane waves, spatially organized oscillatory activity in cortex can take on many forms, including zero-lag spatial phase synchrony, standing waves, traveling waves, rotating spiral waves, radiating waves, and complex patterns comprised of mixtures of the above. For example, Prechtl et al. (1997) describe stimulus evoked waves at 10-20 Hz in turtle visual cortex, and found plane waves, spiral waves, and complex waves. Huang et al. (2004) observe a similar phenomenon in tangential slides of rat visual cortex spanning layers III-V, and report plane, radiating, and complex waves. Benucci et al. (2007) observe a wave response with aspects of both a standing and traveling wave. In visual cortex, ongoing inputs may trigger directed spiking information transfer correlated with a phase shift in gamma oscillations consistent with a traveling wave phenomenon (Besserve et al., 2015). (See Chapter 1 for an in depth introduction to spatiotemporal wave patterns in cortex, their statistical analysis and modeling.)

Furthermore, some studies have observed that spatiotemporal dynamics change rapidly over time. Schiff et al. (2007) report a characteristic sequence of wave patterns in slice that begin as chaotic, irregular, and complex waves, and then transition into transient plane or rotating waves, before returning to complex activity. They note that these changes might reflect an underlying biological order parameter. Similarly, Townsend et al. (2015) characterize delta waves in marmoset during anesthesia, and find that complex waves are associ-

ated with higher single-unit firing rates, and that spatially synchronized states transition into complex waves by passing through intermediate wave patterns like plane, rotating, or radiating waves. Ermentrout and Kleinfeld (2001) explore wave dynamics in sensory cortices from the perspective of coupled oscillators, and note that traveling waves, spiral waves, and spatially synchronized solutions are all possible. They note that spatiotemporal dynamics correlate with the system state, with synchronized activity being correlated with the presentation of a strong stimulus. The interest in the variety of transient spatiotemporal beta wave patterns is also motivated by recent work hypothesizing the role of phase relation diversity in spatiotemporal neural dynamics on cortico-cortical communication. This phase diversity would enhance selectivity and support the segregation of parallel information streams in the cortex (e.g. Maris et al. (2016), "Diverse phase relations among neuronal rhythms and their potential function", Trends in neuroscience, in press.)

A diverse number of processes may yield superficially similar LFP spatiotemporal dynamics. For example, traveling waves may arise from a gradient of conduction delays from a common input, from propagating waves in an excitable system, or from phase gradients in a population of coupled oscillators (Ermentrout and Kleinfeld, 2001). Traveling plane waves in beta may be a correlate of local information transfer, of coordination of relative phases between cortical and subcortical motor areas for communication or encoding, or of external inputs arriving at cortex with varying latency. In addition, specially constructed stimuli can lead to spontaneous emergence of wave patterns (Rule et al., 2011). The waves could be an epiphenomenal correlate of underlying neural order parameters that depend on the dynamic lateral interactions between local circuits.

The diversity of beta wave spatial patterns in primate motor cortex has not been systematically characterized. This characterization may provide important constraints for information flow in motor cortex and the development of spatially extended biophysical models (e.g. Heitmann and Ermentrout 2015; Heitmann et al. 2012) of its dynamics. Here, we examined in detail the spatiotemporal structure of beta oscillations in three mo-

tor cortex areas (M1, PMv and PMd) of *m. mulatta* during movement preparation before and after visual cues in a cued reaching and grasping task. We extended the approach of classifying critical points in the phase gradient field to categorize wave states. Beta spatiotemporal patterns were diverse, and exhibited differences both among motor areas, and between different phases of the motor task. In contrast to M1 and PMd (the focus of previous studies), plane waves were relatively rare in PMv, even though beta transient events were frequent in this area. Furthermore, motor cortex neural dynamics spontaneously transitioned among different transient spatiotemporal beta wave patterns. We provide statistical summaries of the spatiotemporal wave dynamics, and reveal correlations that are important for constraining future models of motor cortex wave activity. Finally, we compare the statistics of spontaneous beta (β -) LFP waves to the statistics of optogenetically induced traveling waves reported in Lu et al. (2015), and reveal several differences between these phenomena that point to important differences in the underlying neural dynamics.

4.2 Methods

4.2.1 The CGID task

The Cued Grasp with Instructed Delay (CGID) task investigates neural activity in motor cortex associated with sequential visual cues, instructed delays, and the planning of upcoming reaching and grasping movements. The CGID task has been described in detail in Vargas-Irwin et al. (2015). Briefly, it requires a monkey to reach out and grasp one of two objects, using one of two possible grips. A sequence of visual cues instruct the monkey in which object to grasp and which grip to use. When a trial begins, the lights in the room are turned off, and the subject is at rest with hands resting on capacitative touch sensors. One of the two objects is rotated into the reach target position. One second later, that object is illuminated. The monkey now knows which object to grasp, but not which grasp

to perform. One second after that, a cue light (red or yellow) is illuminated, specifying the grip. If the light is red, the monkey is to perform a power grip. If the light is yellow, the monkey is to perform a precision grip or a key grip, depending on the object. Two seconds after the ‘Grip’ cue, a ‘Go’ cue (green light) is given. The monkey is now allowed to reach out and grasp the object. If the monkey moves before the ‘Go’ cue or uses the incorrect grip on the object, the trial is voided. If the monkey executes the reach and grasp action after the ‘Go’ cue using the correct grip, he receives a juice reward. Movement periods were identified as the times between when the monkey lifts his paw from the holding position to the time when the monkey contacts the object, as detected by capacitative touch sensors. In this study, we examine beta oscillations during the 1-second preceding object presentation and during the 1-second preceding the grip cue.

4.2.2 Neural recordings

Data were recorded from triple microelectrode arrays (MEAs) (Blackrock Microsystems, Utah), with an electrode depth of 1.5 mm targeting layers 3-5 of motor cortex. For the CGID datasets (Figure 3.1), LFP data were recorded on 10×10 (ventral premotor cortex PMv) and 6×8 Blackrock (dorsal premotor cortex PMd and primary motor cortex M1) arrays with 0.4 mm electrode spacing. Data from two monkeys (R and S) were analyzed. See Vargas-Irwin et al. (2015) for additional details. For all datasets, broadband LFPs recorded at 30 kilosamples/s (0.3 Hz - 7.5 kHz) were down-sampled (zero-phase 4th order Butterworth, ≤ 250 Hz MATLAB `filtfilt`) to 1 kilosample/s for analysis.

4.2.3 Optogenetically induced spatiotemporal waves in motor cortex

We compared spontaneous beta spatiotemporal wave patterns to waves induced by optogenetic stimulation performed in a third monkey (subject T in Lu et al. (2015)) during a rest state (Figure 4.10). Data were recorded from a single 10×10 array implanted in primary

motor cortex, modified to contain an optical fiber at the center for delivery of laser light. Neurons were transfected with the chimeric opsin C1V2_{T/T} under the CaMKII α promoter, preferentially expressing in excitatory cells with some expression likely also in inhibitory cells (Lu et al., 2015). In the data presented here, constant pulse optical stimulation at 561 nm and 6 mW was provided for one second during rest. See Lu et al. (2015) for details of the optogenetic recordings.

4.2.4 Signal processing and LFP phase extraction

The beta band was identified separately for each monkey and area. Multitaper spectra (5 tapers) were computed from the first second of the CGID task for all channels and trials, and the beta band was defined as a 5 Hz band centered at the largest spectral peak between 15 and 30 Hz during the one second preceding the ‘Go’ cue. Monkey R area PMd exhibited two prominent beta peaks. We restricted analysis to the lower band, which was closest in frequency to the beta oscillations in other areas. Narrow-band gamma oscillation peaks were identified from wavelet spectrograms averaged over all stimulation trials. Instantaneous phase and amplitude envelope for narrow-band LFP oscillations were extracted similarly for both the beta and the induced gamma analysis. Raw LFP data at 1 kilosample/s were band-pass filtered using a 4th order Butterworth filter, applied forwards and backwards (`filtfilt()` in Matlab). The resultant narrow-band signal $a(t)$ was passed through the Hilbert transform to generate an analytic signal

$$z(t) = a(t) + ib(t) = |z(t)| \exp(i \cdot \varphi(t)), \quad (4.1)$$

which has both an instantaneous magnitude $|z(t)|$ and phase

$$\varphi(t) = \text{Arg}(z(t)). \quad (4.2)$$

4.2.5 Assessing spatial synchrony and coherence

Spatial coherence (Figure 4.13) was computed for each time-frequency bin as the largest eigenvalue of the matrix of pairwise coherences between channels on the array (Mitra and Bokil, 2007). Pairwise coherence was calculated using `scipy.signal.coherence`, with a sliding window 512 ms long with 256 ms segments. To assess spatial synchrony over the recorded neocortical patch, we used two related statistics: Kuramoto's order parameter and the circular standard deviation. Kuramoto's order parameter

$$\text{Kuramoto's order parameter} = \frac{1}{N} \sum_{i=1}^N e^{i\varphi(t)} \quad (4.3)$$

reflects the degree of phase synchrony in a population of oscillators, ranging from 0 for a completely asynchronous population to 1 for perfectly phase synchronized populations (Kuramoto, 1984). We found that Kuramoto's order parameter was unsuitable for visualizing the relationship between phase synchrony and other wave statistics (Figures 4.6, 4.16, 4.17), as it displayed highly nonlinear relationships and saturated close to 1, concealing much of the variability. For these comparisons we use the circular standard deviation of the instantaneous LFP Hilbert phases

$$\text{Standard deviation of Hilbert phase} = S_\varphi(t) = \sqrt{-2 \ln(R(t))}. \quad (4.4)$$

The above corresponds to the standard deviation of a circularly wrapped normal distribution, and can be obtained by a transformation of the first moment for circularly distributed data. Due to noise, samples with smaller LFP amplitude yield less reliable phase estimates, so we weight each vector's contribution to this statistic by its magnitude

$$R(t) = \frac{|\sum_k z_k(t)|}{\sum_k |z_k(t)|}, \text{ where } k \text{ indexes channels.} \quad (4.5)$$

We define spatially synchronized states as those times for which $S_\varphi(t) < \pi/4$, such that approximately 95% of phase vectors are concentrated within one-quarter cycle.

4.2.6 Spatial gradient of the Hilbert phase

The local phase gradient provides an estimate of the wave propagation velocity. To estimate this, we compute the gradient of the phase $\nabla\varphi(t)$ as the (circularly-wrapped) discrete difference of the Hilbert phase between adjacent electrode sites. The propagation velocity \mathbf{v} and wavelength λ can also be extracted from the phase gradient using the relationship in Equation ?? (see Rubino et al. (2006) for details)

$$\frac{\partial\varphi}{\partial t} = -\nabla\varphi \cdot \mathbf{v}. \quad (4.6)$$

Propagation speed was estimated as the average magnitude of the local phase gradient (in radians/mm) divided by the median angular frequency (in radians/second) to yield an instantaneous estimate of the local propagation speed in mm/second.

Plane waves were detected as moments when the distribution of phase gradient vector directions was concentrated. Rubino et al. (2006) defined a phase gradient directionality (PGD) measure

$$PGD(t) = \frac{|\sum_{i=1}^N \nabla\varphi_i(t)|}{\sum_{i=1}^N |\nabla\varphi_i(t)|}, \quad (4.7)$$

which is similar to Kuramoto's order parameter and measures the concentration of direction of the phase gradient (Equation 4.7). We report a related measure, the circular standard deviation of the Hilbert phase gradient. A version of this that is weighted by the amplitudes of the phase gradient vectors can be computed as the following transformation of the PGD

$$S_{\nabla\varphi}(t) = \sqrt{-2 \ln(PGD(t))}. \quad (4.8)$$

Plane waves were classified as those times for which $S_{\nabla\varphi}(t) < \pi/4$, i.e. 95% of phase gradi-

ent vectors point within $\pi/2$ radians of each-other. This corresponds to a *PGD* threshold of approximately 0.73. This cutoff is more stringent than that used in Rubino et al. (2006) and was necessary to address biases in estimating wavelength arising from noise and non-planar wave activity.

For the analysis of plane waves in β -LFP, the wavelength was estimated from the magnitude of the average phase gradient. Taking the magnitude after averaging the phase gradient vectors reduces the influence of noise on the estimation of plane-wave wavelength, as opposed to averaging the local phase gradient magnitudes. Asymmetric radiating waves with a high signal-to-noise ratio that could not be described as a simple planar or radiating wave were common in the optogenetically induced activity, and in this case wavelength was estimated from the local gradient of the Hilbert phase and then averaged over the array.

4.2.7 Critical point analysis

We classified complex waves structures based on critical points in a smoothed estimate of the phase gradient field. Data were denoised and smoothed by convolving the analytic signals in each frame with a sinc function to attenuate structure finer than 2 mm. Signals were interpolated via sinc upsampling on the analytic signal. Convolution and upsampling are performed in the frequency domain using reflected boundary conditions via a discrete cosine transform type I. The sinc kernels are constructed in the frequency domain with anti-aliasing to reduce numerical artifacts related to the small spatial domain. The phase gradient is extracted by convolving the Hilbert phase with 2×2 discrete difference kernels in the x and y directions.

Minima, maxima, and saddle points, were identified by locating zeros in the phase gradient. These points can be distinguished by the direction of sign change in the gradient. Equivalently, one may examine the determinant and trace of the Jacobian at each critical point (Townsend et al., 2015). The centers of rotating waves are identified by taking a line

integral of the Hilbert phase gradient surrounding each point. Points around which this line integral equals nonzero multiples of 2π are singularities in the Hilbert phase gradient and are the centers of rotating waves. Numerically, this integral is computed for every point by representing the Hilbert phase gradient as a complex number, and representing the line integral around each point as a convolution. We carefully addressed some caveats to this analysis. When beta oscillations are asynchronous, the spatial structure of waves may be finer than the electrode spacing, and interpolation may create spurious critical points. Conversely, smoothing may eliminate some structure. Critical point analysis is most accurate when signal-to-noise ratio is high and high spatial frequency structure is attenuated.

4.2.8 Averaged phase maps for optogenetically induced spatiotemporal waves

The wave response to optogenetic stimulation was reliable across trials, which allowed us to compute summaries of wave activity averaged over trials and time. For computing the average phase gradient, local phase gradient fields were smoothed by convolution with a sinc kernel to remove structure finer than 1.8 mm. These phase gradients were averaged over time (Figures 4.14a,b rightmost panels). To compute averaged phase delay maps, phases were unwrapped relative to a reference electrode (channel 25, Figure 4.12) near the center of the radiating wave. Average phase delay was then computed as the average over time of the phase delay relative to this channel.

4.3 Results

We aim to characterize the spatiotemporal dynamics of waves in motor cortex beta (~ 20 Hz) LFP oscillations in order to better understand the collective neural dynamics underlying these oscillations. We first demonstrate that movement preparation stages in the CGID

task reliably evoke beta oscillations, which occur as transients during steady-state periods with considerable trial-to-trial variability in terms of transient onset times and durations. We show that these beta oscillations correspond to diverse types of wave dynamics, and summarize the prevalence of various wave patterns. Next, we examine summary statistics of beta spatiotemporal oscillations and relationships among these statistics. Finally, we conclude by contrasting these statistics to those of traveling waves induced by 1-second constant pulse optogenetic stimulation in motor cortex.

We analyzed three CGID datasets each from areas M1, PMv, and PMd in monkeys R and S. For both monkeys, all examined datasets were recorded within eight days of each-other. Each session yielded between 46 and 114 successful seven-second CGID trial repetitions, collected over twenty minutes to one hour. For the comparison to optogenetically induced traveling waves, we analyze the spatiotemporal waves in primary motor cortex observed in monkey T, previously reported in Lu et al. (2015).

4.3.1 Motor evoked potentials and beta oscillations in the CGID task

The CGID task reliably elicited task-related activity in all three motor areas (M1, PMd, PMv) from both monkeys studied. The movement period of the CGID task was marked by slow motor evoked potentials (Figure 4.1a), increased single-unit firing rates (Figure 4.1b), and beta suppression (Figures 3.7, 4.1c). Motor evoked potentials did not consist entirely of a slow <2 Hz component, but also contained peaks in the theta (2-7 Hz) band, in all areas except monkey S area PMv (Figure 3.7). The CGID task reliably elicited beta oscillations during the planning and preparatory phases. In monkey S, the beta band peaked at 21 Hz. In Monkey R, the beta band peaked at 17 Hz (Figure 3.7). Beta oscillations were especially stable during the first second of the task, and during the one second before ‘Go’ cue. In monkey S, beta oscillations were strongest before object presentation. In monkey R, beta oscillations showed greater power during the delay between the grip cue and the ‘Go’ cue

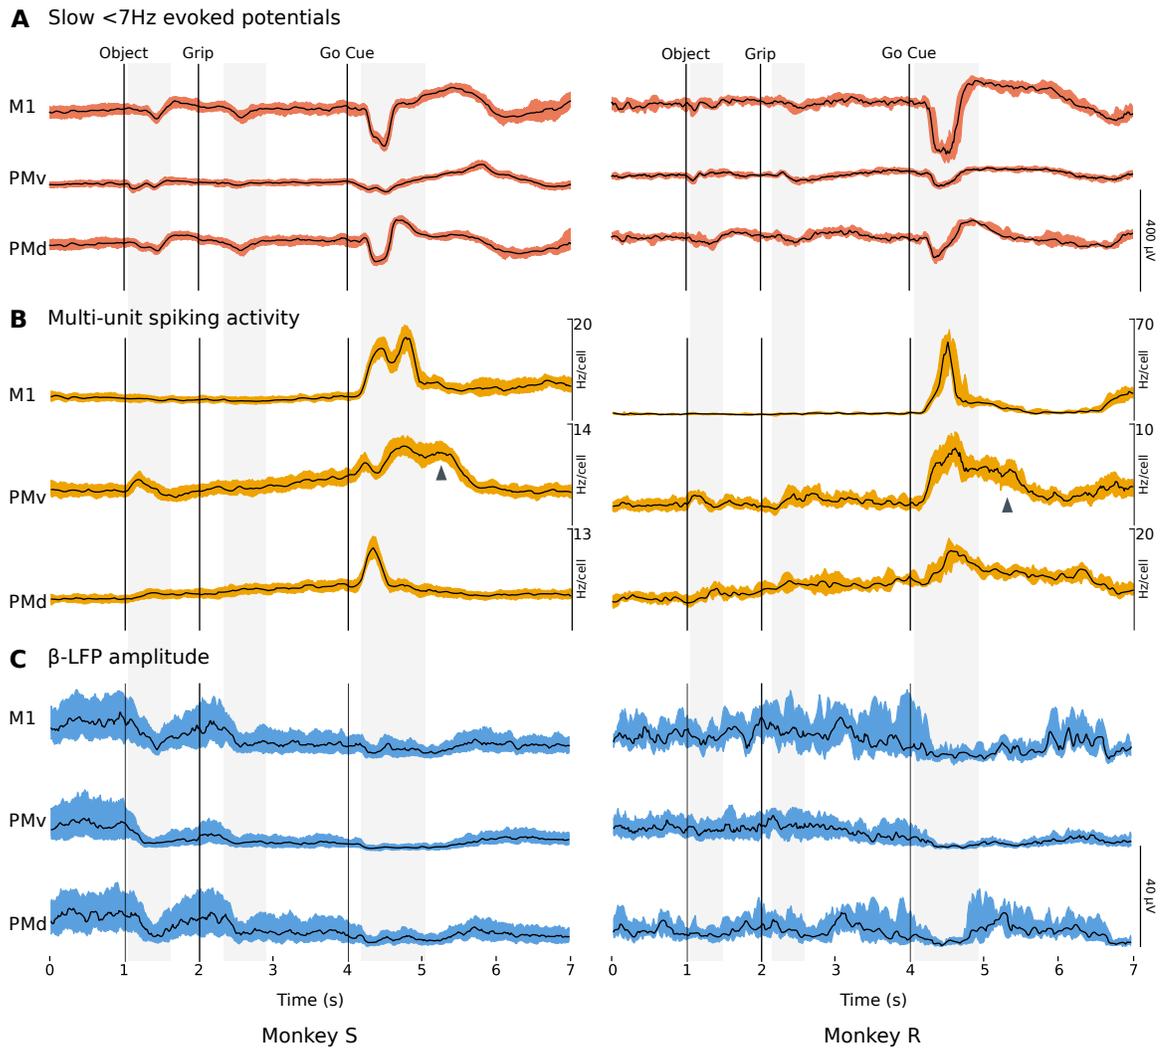


Figure 4.1: Evoked potentials and modulations in multi-unit spiking activity and β -LFP in the CGID task. Statistics are summarized over all sessions and trials for each monkey. Colored regions represent the interquartile range, black lines the median. Time periods displaying evoked potentials are shaded in gray. **(A)** Sensory and motor evoked potentials in the CGID task. Both monkeys showed task-related evoked potentials in the LFP (shown here averaged over good channels within a single area). In addition to motor evoked potentials in all areas, we also see visual-cue evoked potentials. **(B)** Task-related modulations in multi-unit spiking activity varied between motor areas. In both monkeys R and S and in all motor areas, multi-unit spiking activity modulates with different stages of the CGID task. Area PMv displayed multi-unit modulation after object presentation. All motor areas displayed movement-evoked increases in firing rates, with area PMv showing elevated spiking extending for up to two seconds after 'Go' cue (arrows). **(C)** Task-related modulations in β -LFP magnitude. Periods of beta suppression occur during evoked potentials and increases in multi-unit firing rates. Beta oscillations were strongest during the first second of the task for monkey S, and strongest during the one second before the 'Go' cue for monkey R. Beta amplitude was also more variable in monkey R. For all plots, data shown are from the first CGID dataset for each monkey, averaged over all good electrodes within a MEA.

(Figure 3.7).

There were several interesting features of the task-related LFP activity. Visual cues evoked beta suppression and low frequency evoked potentials, much like movement execution. There was also a short spike in alpha (7-15 Hz) LFP activity immediately following object presentation in all monkeys and areas with exception of monkey R area PMd (Figure 3.7). Beta was transiently suppressed in all areas between approximately 250 and 750 ms following the object presentation, and 400 and 900 ms following the 'Grip' cue. For monkey S only, the beta suppression evoked by object presentation was stronger in area PMv than in the other motor areas. The duration of movement-related beta suppression was longer in areas M1 and PMv than in area PMd, consistent with the conjecture that PMd is concerned with the visually guided reaching, which is completed earlier than the grip and object manipulation movements.

Although beta power was enhanced during the steady-state periods of the CGID task, we found that it occurred in transient events with variable duration and exhibited considerable trial-to-trial variability (Figure 3.11; see also Chapter 3). The duration of beta transients was broadly distributed (Figure 3.11), with some trials showing sustained elevation of β -LFP for as long as one second.

During visual inspection of the data, we observed diverse wave patterns. For illustration, we have collected examples of wave activity in Figure 4.2. We observed states ranging from spatially synchronous states, to plane waves, to more complex activity like radiating or rotating waves, to complex states that cannot be easily categorized, showing that beta events can be associated with a diversity of wave patterns. Importantly, wave patterns appeared to transition continuously from desynchronized states, to more locally synchronized wave activity, to globally synchronized states. In the following sections, we quantify and elaborate upon these observations.

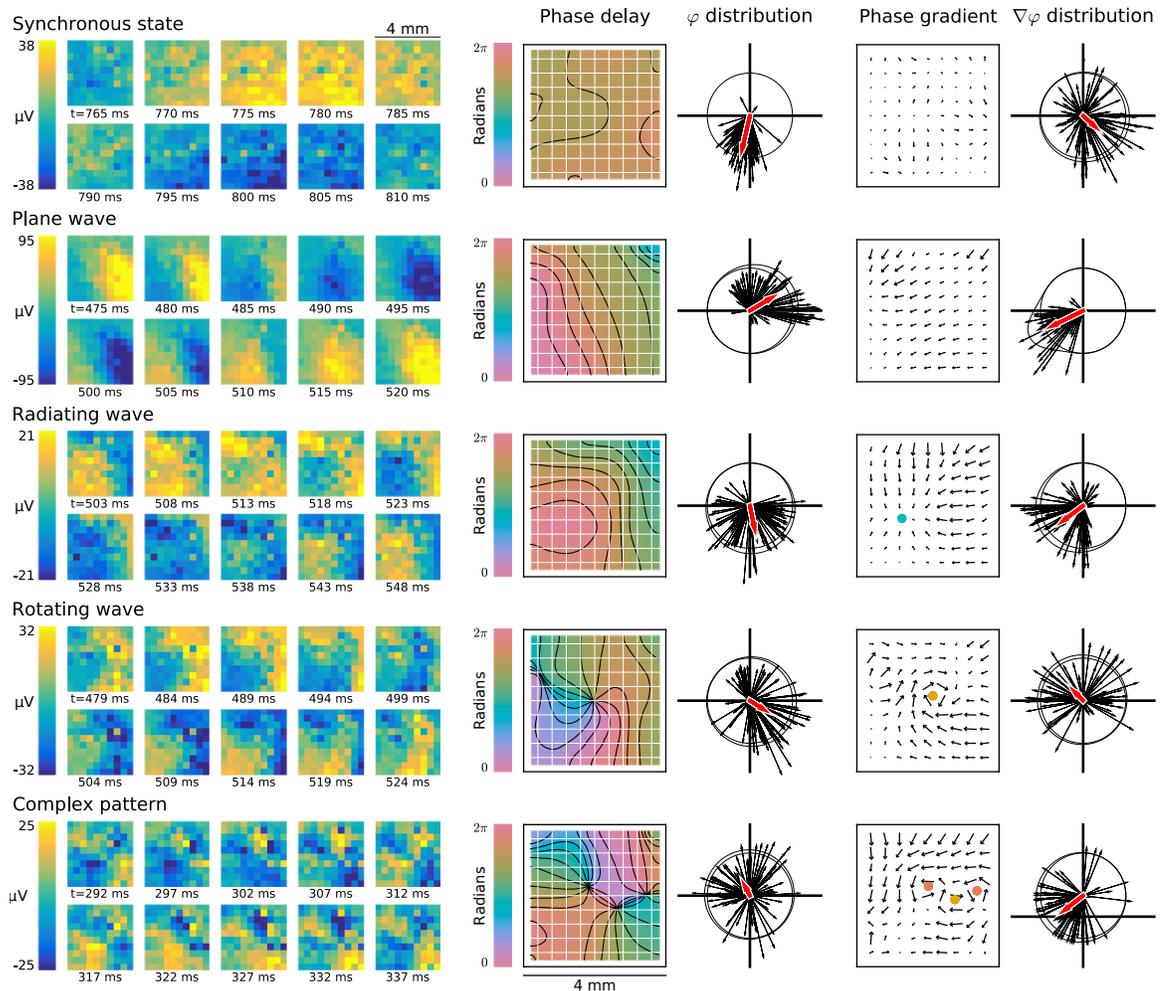


Figure 4.2: *Transient beta spatiotemporal wave patterns are diverse.* In addition to traveling plane waves, we observe synchronous states, radiating and rotating waves, and other more complex wave patterns. Each of these examples was taken from the 4×4 mm area sampled by the 10×10 multielectrode array in area PMv of monkey S. Missing electrodes were interpolated from nearest neighbors. Average phase delay maps were computed by unwrapping Hilbert phases at the median frequency of the wave event before computing the average analytic signal. The mean analytic signal was smoothed at a scale of 2 mm and upsampled to generate the phase delay maps pictured here. The smoothed and upsampled Hilbert phase was differentiated to extract critical points from the wave dynamics (Methods: ‘Critical point analysis’) Spatially synchronized states were detected as spatial patterns where the angular distribution of analytic signals was especially concentrated (first example, Methods: ‘Assessing spatial synchrony and coherence’). Plane wave states were detected as spatial patterns where the angular distribution of the phase gradient direction was especially concentrated (second example, Methods: ‘Spatial gradient of the Hilbert phase’).

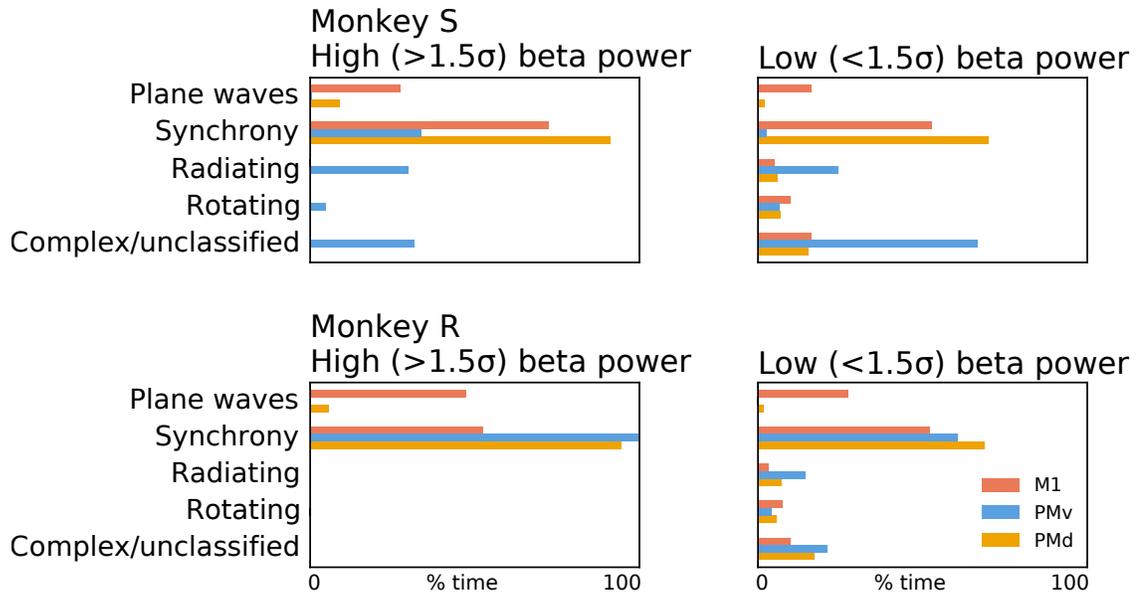


Figure 4.3: Summary of the prevalence of different beta wave patterns the steady movement preparation period. Categorization of spatiotemporal beta waves during the 1 second before the ‘Go’ cue shows a variety of patterns. Time-points are separated into high and low beta power by thresholding the instantaneous amplitude, with amplitudes larger than $1.5\times$ the standard deviation of the beta signal being classified as high beta. High beta power tends to be associated with spatial synchrony or plane waves. More complex wave activity tends to occur at lower power. Some motor areas showed little traveling plane wave activity, even at high beta power (Monkey S area PMd, Money R and S area PMv).

4.3.1.1 Beta wave categorization

Figure 4.3 shows a breakdown of the prevalence of different classes of wave patterns for all areas in both monkeys R and S. We identified synchronous states as spatial patterns where the circular standard deviation of β -LFP phase was less than $\pi/4$, such that approximately 95% of electrodes on the array fell within the same quarter-cycle ($\pi/2$ radians) of the oscillation (Methods: ‘Assessing spatial synchrony and coherence’). Similarly, we classified plane waves as those spatial patterns where the circular standard deviation of the phase gradient direction was less than $\pi/4$ (Methods ‘Spatial gradient of the Hilbert phase’). For the purpose of assigning each beta cycle a unique classification, waves were classified as synchronous only if they were not also classified as plane waves. More complex wave patterns like rotating and radiating waves were classified by identifying critical points in the spatially smoothed phase gradient (Methods: ‘Critical point analysis’).

The spatiotemporal structure varied considerably across motor areas. Area PMd in

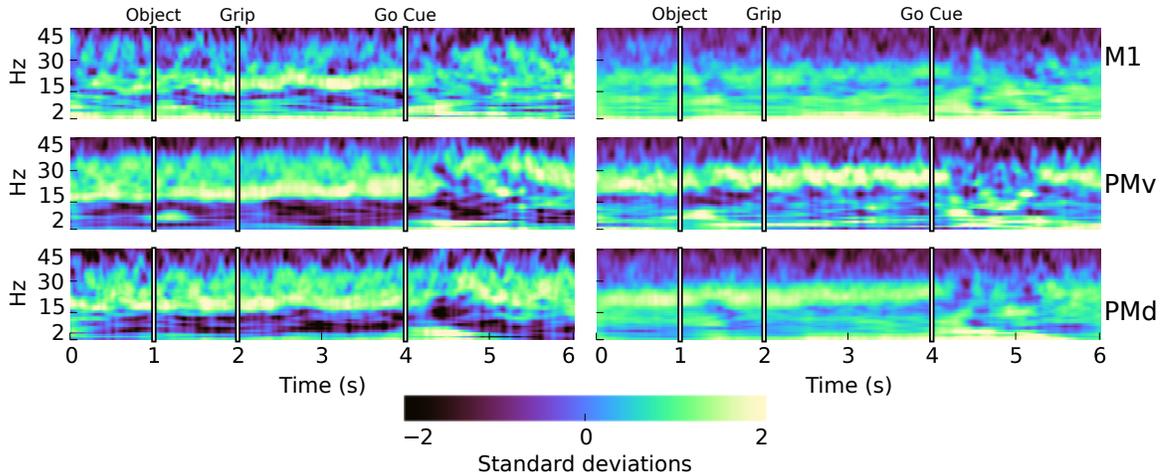
both monkeys was typically highly synchronous, displaying few traveling waves. Area PMv in monkey S showed more complex wave activity, even during periods of high beta power. In both monkeys, organized plane wave activity was apparently rarer in area PMv, an area that had not been examined in previous studies. However, area PMv in monkey S displayed other organized traveling wave states such as radiating and rotating waves (Figure 4.3). These differences between areas were not due simply to the different sizes of the arrays, as the classification trends did not change when the 10×10 PMv arrays were cropped to match the 6×8 size of the M1 and PMd arrays. High beta power was strongly correlated with spatially synchronized states, while more complex activity were seen at lower beta power. Because the signal-to-noise ratio is lower when beta power is low, there is a possibility that some of the wave activity was classified as complex due to noise (Methods: ‘Critical point analysis’). However, complex wave patterns were still identified in some areas even when the signal-to-noise ratio was high, indicating that these patterns are a true mode of the spatiotemporal activity in motor cortex.

While we adopted a classification of beta wave activity into the above discrete types, we also emphasize an important aspect of the spatiotemporal dynamics that should not be overlooked: beta spatiotemporal patterns varied continuously, and there were no clear boundaries or abrupt transitions between different patterns. Furthermore, although the thresholds for classification of synchrony and plane waves reflect intuition about synchronized and plane-wave activity, the data are not clustered into "synchronous" and "asynchronous" or "planar" and "nonplanar" groups. For example, long wavelength plane waves were also spatially synchronous (over the limited aperture of the MEA).

4.3.1.2 Task-related modulations in beta spatiotemporal activity

To more accurately quantify wave activity, we analyzed summary statistics of the wave dynamics. The spatiotemporal properties of beta oscillations varied throughout the stages of the CGID task. Beta synchrony was elevated during the planning periods, but was espe-

A LFP synchrony



B LFP trial phase locking

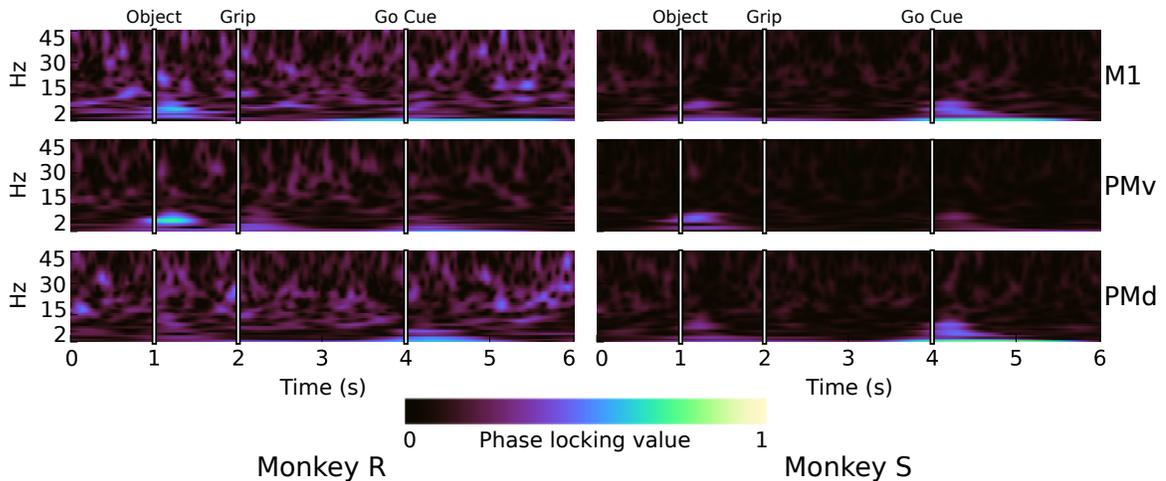


Figure 4.4: Changes in beta spatial synchrony correlate with the CGID task stages. **(A)** In the CGID task, LFP spatial synchrony occurs in the alpha, beta, and low-frequency bands. Beta oscillations are typically more synchronous over each area when beta power is high. Periods of beta suppression correlate with reduced spatial synchrony. In both monkeys a brief period of alpha (5-10 Hz) synchrony appears at the start of beta suppression, most prominently in area PMv after object presentation. Spatial synchrony in low-frequency (<7 Hz) LFP is present throughout the entire task, during both the planning and movement periods. Spatial synchrony has been z-scored to visualize the variations in synchrony over time, frequency, and area. **(B)** Beta spatial phase synchrony is not locked to the CGID task. Shown here are phase locking values (0 = no phase locking, 1 = complete phase locking) over all sessions and channels. We see phase-locking of synchronous alpha oscillations following object presentation and 'Go' cue, as well as reproducible phase in motor evoked potentials, but no phase locking in beta. For all plots, data shown are from the first CGID dataset for each monkey, averaged over all good electrodes within a MEA.

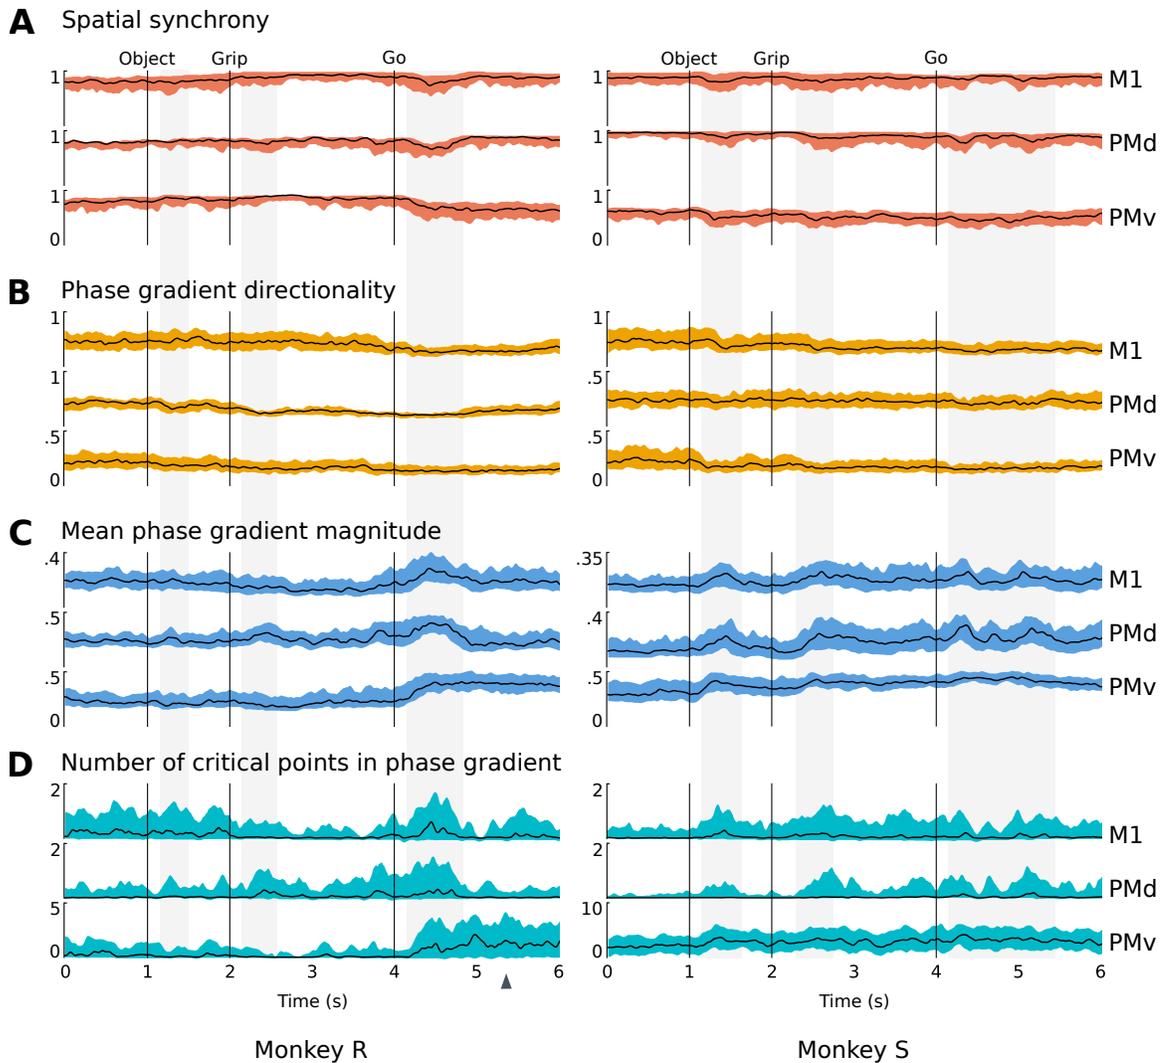


Figure 4.5: *Modulations in beta spatiotemporal activity are task-related.* Shown here are four statistics reflecting spatiotemporal dynamics in beta oscillations for both monkeys and all three motor areas. Data from all trials in three sessions are combined in these summary plots. Statistics are computed from the instantaneous phase of the beta analytic signal, and smoothed with a 50 ms boxcar filter. Colored regions represent the inter-quartile range, and black lines the median, over all trials and sessions. The timing of visual cues (object presentation, grip cue, go cue) are marked with vertical black lines. Times displaying putative cue-evoked changes in beta spatiotemporal properties are highlighted in grey. **(A)** Spatial synchrony, as assessed by Kuramoto's order parameter, ranges from 0 for asynchronous to 1 for perfectly synchronized activity. Reductions in spatial synchrony can be observed in all areas following visual cues in monkey S, and during the movement period in both monkeys. **(B)** Phase gradient directionality assesses the extent to which beta activity resembles a traveling plane wave. Both monkeys show a reduction in phase gradient directionality over the course of the CGID task, and in monkey S the shifts appear to relate to visual cues. **(C)** The average phase gradient magnitude is a proxy for spatial wavelength. Higher phase gradients correspond to higher spatial frequency waves. The task-related modulations in beta phase gradient magnitude differ between the monkeys, but both exhibit an increase in average gradient magnitude shortly after the go cue, indicating a reduction in large spatial scale phase organization. **(D)** The complexity of spatiotemporal wave activity can be assessed by the number of critical points (rotating, radiating centers) in the (smoothed) phase gradient map. The complexity of wave activity in area PMv differed between the monkeys, with monkey R exhibiting highly synchronous activity during the preparatory phases, and sustained complex wave activity during the movement period that continues longer than the desynchronization seen in areas M1 or PMd (arrow). In contrast, monkey S exhibited complex wave activity in area PMv throughout the task.

cially strong surrounding the visual cues, and absent during periods of beta suppression (Figure 4.4a, 4.5a). Although beta oscillations in monkey R were composed of two frequency bands, it was the lower 15-25 Hz band frequency that displayed the greater spatial synchrony.

Figure 4.5 details the task-related changes in the spatiotemporal structure of beta oscillations in two example areas. Statistics of beta spatiotemporal activity exhibited trial-to-trial variability, but tended to track task-related changes in the magnitude of beta oscillations. During task phases in which beta power was elevated, spatial synchrony was also higher (Figure 4.5a), and beta phase gradient magnitude, which is inversely related to wavelength, was lower (Figure 4.5c). Phase gradient directionality, a measure of how planar wave activity is, (Methods ‘Spatial gradient of the Hilbert phase’) was also higher during times when beta power was high (Figure 4.5b). The number of critical points in the phase gradient map, a measure of wave complexity, increases during task phases with weaker beta power (Figure 4.5d). The above results show that the task-related changes in the properties of beta spatiotemporal waves correlate with variations in beta amplitude, indicating that changes in beta spatiotemporal organization may be closely related to modulations in beta power. However, β -LFP activity was highly variable. In our inspection of the data, the precise nature and timing of beta spatiotemporal patterns was seldom identical across trials. To better understand the spatiotemporal structure of β -LFP oscillations, we assessed the instantaneous relationship between various statistics of spatiotemporal activity on a trial-by-trial basis.

4.3.1.3 Beta wave statistics correlate with amplitude

Having surveyed the spatiotemporal structure of beta oscillations, we next asked how these spatiotemporal statistics correlate with each-other and with β -LFP amplitude. We compared the average instantaneous Hilbert amplitude of beta across the array to the circular standard deviations of both the Hilbert phase, and the Hilbert phase gradient. Figure

4.6 explores the correlated variation in wave statistics and reveals how the spatiotemporal dynamics in beta vary with beta amplitude. We found that beta amplitude and beta synchrony were positively correlated for all areas in both monkeys. The correlation between beta amplitude, and the standard deviation of the phase gradient direction, varied across areas. Overall, traveling wave states existed on a continuum between asynchronous and spatially synchronized states. We note that, because of the inherent challenges in differentiating signal and noise in these LFP recordings, there is the possibility for changes in signal-to-noise ratio to create artifactual correlations between apparent spatial synchrony or phase gradient and amplitude. To address this confound, we further analyzed wavelength from planar wave events that showed an extremely high degree of phase gradient order unlikely to occur by chance.

4.3.1.4 Beta wavelength correlates with amplitude

We analyze the correlation between planar traveling wavelength and beta amplitude in the M1 arrays of both monkeys, which displayed the largest number of planar traveling waves. For both monkeys, organized plane wave activity was rare in area PMd, which tended to be highly synchronous, and area PMv, which tended to show more complex wave activity in monkey S and synchrony in monkey R (see Figure 4.3). We note that our results are not inconsistent with Rubino et al. (2006), which reported traveling waves in area PMd, because we are using a more stringent inclusion criterion to avoid bias in our estimation of the correlation of amplitude and wavelength (Methods: ‘Spatial gradient of the Hilbert phase’).

With this approach (Figure 4.7) we found that beta wavelength and beta amplitude were interrelated: larger amplitudes corresponded to longer wavelengths (Pearson $\rho = 0.64$ for monkey R, $\rho = 0.53$ for monkey S). In monkey R, wavelengths ranged from 3.2 mm to 28 mm, with a median of 7.2 mm. In monkey S, wavelengths ranged from 3.5 mm to 31 mm, with a median of 10 mm. High beta power (amplitude envelope >1.5 standard de-

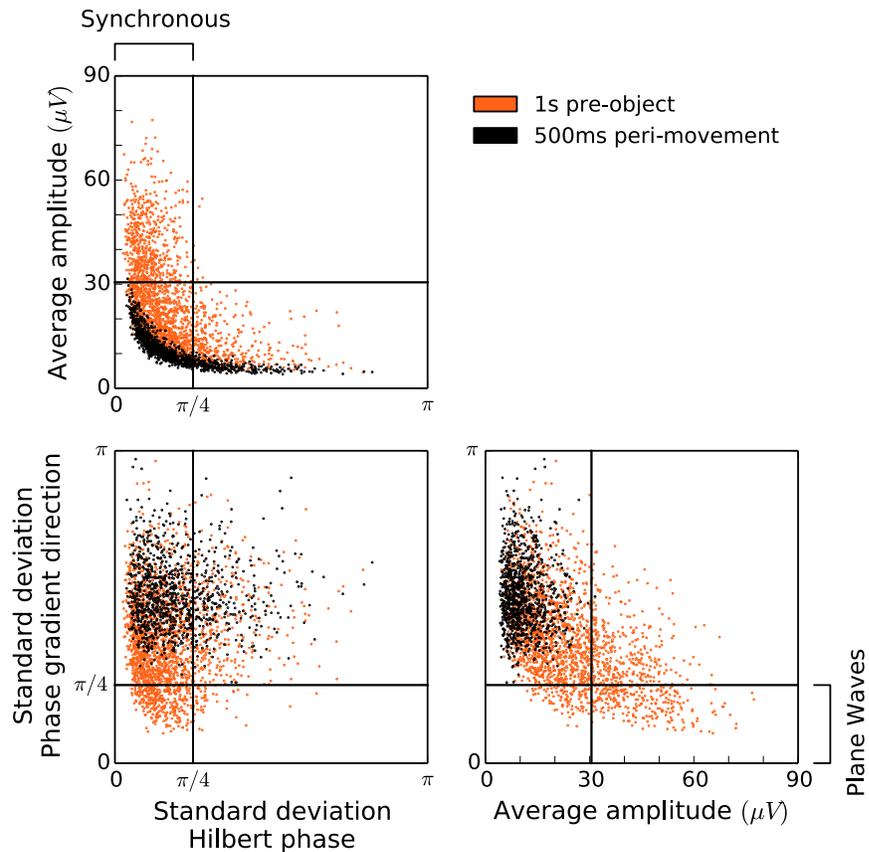


Figure 4.6: Amplitude, phase gradient alignment, and spatial synchrony of spontaneous β -LFP oscillations are related. Wave statistics were taken every 50 ms from the first second (movement preparation) of the CGID task. Higher beta power was associated with increased synchrony and also increased tendency toward plane waves. Beta spatiotemporal statistics varied continuously and no clear boundaries separating plane wave states and synchronized states from more complex wave activity were observed. This emphasizes that spatiotemporal activity in motor cortex is a heterogeneous phenomenon, and raises the question of whether variations in beta spatiotemporal state correlate with changes in the functional organization of motor cortex neural ensembles. Amplitude was taken as the average Hilbert amplitude over the array. Phase gradient alignment was assessed by the circular standard deviation of the gradient of the Hilbert phase. Spatial synchrony was assessed by the circular standard deviation of zero-lag oscillation phase, for which smaller values correspond to more synchronous states.

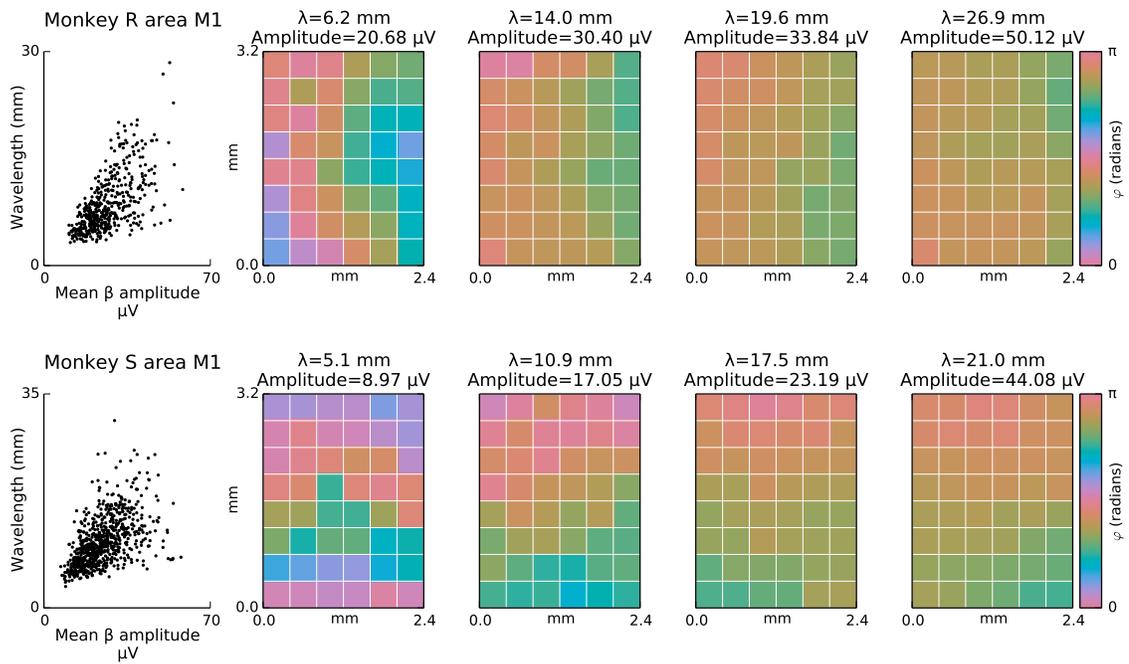


Figure 4.7: *The amplitude and wavelength of traveling β -LFP plane waves are correlated.* Neocortical patches showing planar traveling waves exhibited a correlation between beta amplitude and wavelength. Scatter plots comparing the wavelength of plane-wave beta events with the average beta amplitude over the array reveal a strong correlation (Pearson $\rho = 0.64$ for monkey R, $\rho = 0.53$ for monkey S) between beta amplitude and wavelength. Samples of the array statistics were taken every beta cycle (50 ms), and only those time-points displaying a high degree of alignment in the phase gradient were considered for the analysis (see Methods: ‘Spatial gradient of the Hilbert phase’). Three sessions for each monkey were combined in the scatter plots. Specific examples are also shown, illustrating the progression, left to right, from shorter wavelengths at lower amplitudes, to longer wavelengths at higher amplitudes. The hue scale is doubled to resolve the longer-wavelength patterns more clearly. Very long wavelengths cannot be distinguished from spatial synchrony due to the size of the microelectrode array. These plots indicate that traveling waves tended to transition smoothly into synchronized states, with the LFP phase tending toward greater spatial homogeneity at higher amplitudes.

viations (σ) of the beta-band signal) was associated with a median wavelength of 11 mm in monkey R and 13 mm in monkey S, and low beta power ($<1.5\sigma$) was associated with a median wavelength of 6.8 in monkey R and 8.8 in monkey S. In summary, the spatial scale of beta waves is not fixed but depends on the beta state. There was a continuum between desynchronized activity at low amplitudes and spatially synchronous states at higher amplitudes. We were only able to analyze planar traveling waves with this approach, as correlations between amplitude and the spatial scale of complex patterns could be confounded with variations in signal-to-noise ratio (see Discussion).

4.3.1.5 Spatial variations in both the phase and amplitude of β -LFP may be important for understanding beta wave events

So far, studies of spatiotemporal LFP phenomena have focused on relationships between LFP phase or amplitude separately. However, the filtered analytic LFP signal exhibits variations in both phase and amplitude, and understanding how these are related is important for understanding the origin of spatiotemporally organized LFP signals. All of the statistics focused on here so far discard the amplitude except as a weighting factor to reflect that low-amplitude signals have a less certain phase estimate.

Figure 4.8 illustrates the evolution of analytic signals across all three MEAs during a plane wave event. It is evident that beta amplitude varies across electrodes, and evolves in time jointly with the distribution of beta phases. The wave event begins with an increase in dispersion of the amplitude, with this dispersion in amplitude transforming into a phase gradient as the wave event initiates. It is not the case that plane wave events emerge as a result of increased phase dispersion at a fixed amplitude, as would be predicted in models based on Kuramoto oscillators (Breakspear et al. 2010; Heitmann and Ermentrout 2015; Heitmann et al. 2012). We emphasize, however, that the analytic LFP signal reflects a summed contribution of a larger population of neural oscillators, and amplitude fluctuations can reflect both modulations in the amplitude of local oscillations

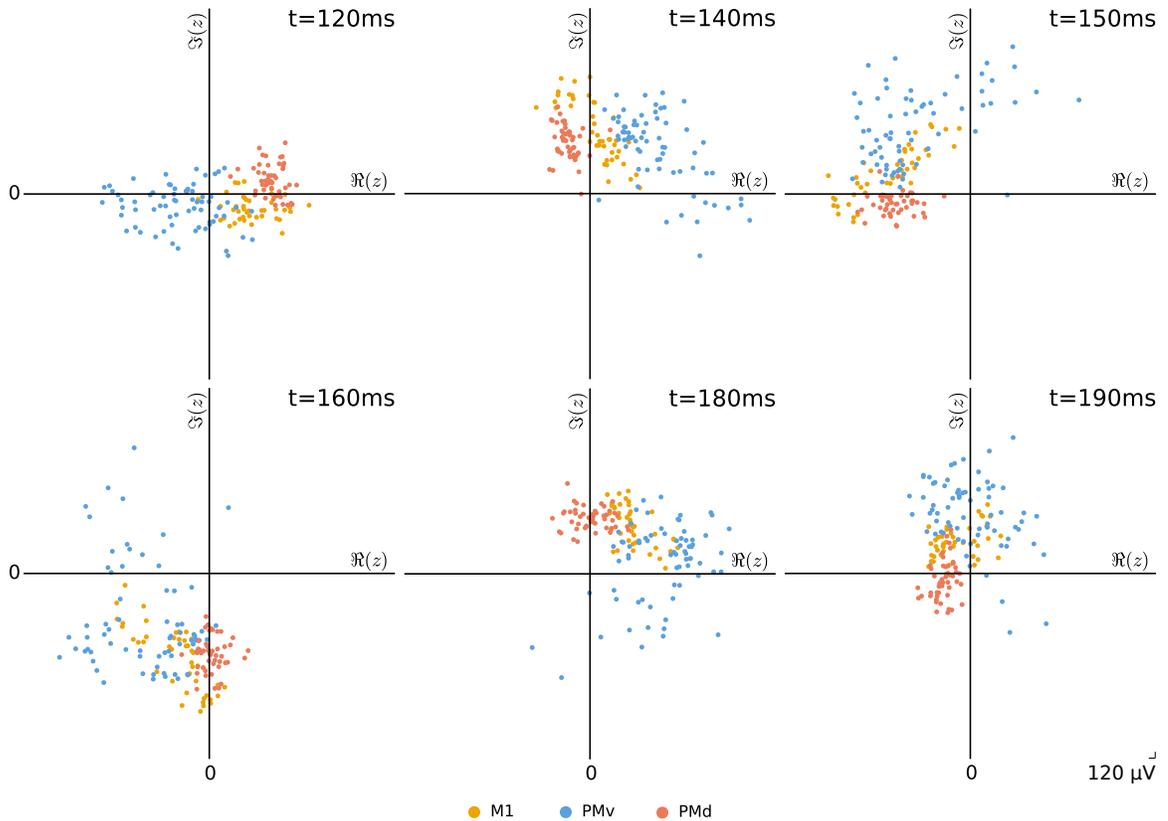


Figure 4.8: *Evolution of a wave event in the analytic signal phase space.* Beta wave events tended to begin by excitation of oscillations in each channel. This excitation of amplitudes shifted, creating a plane wave where, in this case, the phase of area PMv lags behind that in areas PMd and M1. This event lasted only for a few cycles of the beta oscillation, before collapsing back to an asynchronous state. Note that although this is a traveling wave event, it is not a plane wave event. In this instance, the PMv activity was classified as a radiating wave, and the M1 and PMd activities were classified as synchronous. It is possible that we are observing in each cortical patch part of a larger wave phenomenon. This particular wave example was taken from trial 2 of monkey S, session 1. $\Re(z)$ and $\Im(z)$ denote the real and imaginary parts of the complex-valued analytic LFP signal, respectively.

(e.g. dendritic or subthreshold currents), or changes in synchronization between local oscillators. Furthermore, modulations in amplitude and phase are affected by filtering and the Hilbert transform procedure. It is important to emphasize that even the early pre-processing steps of generating the beta analytic signal are subject to ambiguity, a fact which should be more clearly addressed with further methodological research.

Figure 4.9 compares two attempts to jointly summarize changes in amplitude and phase. A model that treats phase and amplitude separately fails to describe the distribution of beta analytic signals in a number of cases. In particular, spatial patterns that exhibit a correlation between amplitude and phase occur near the initiation or termina-

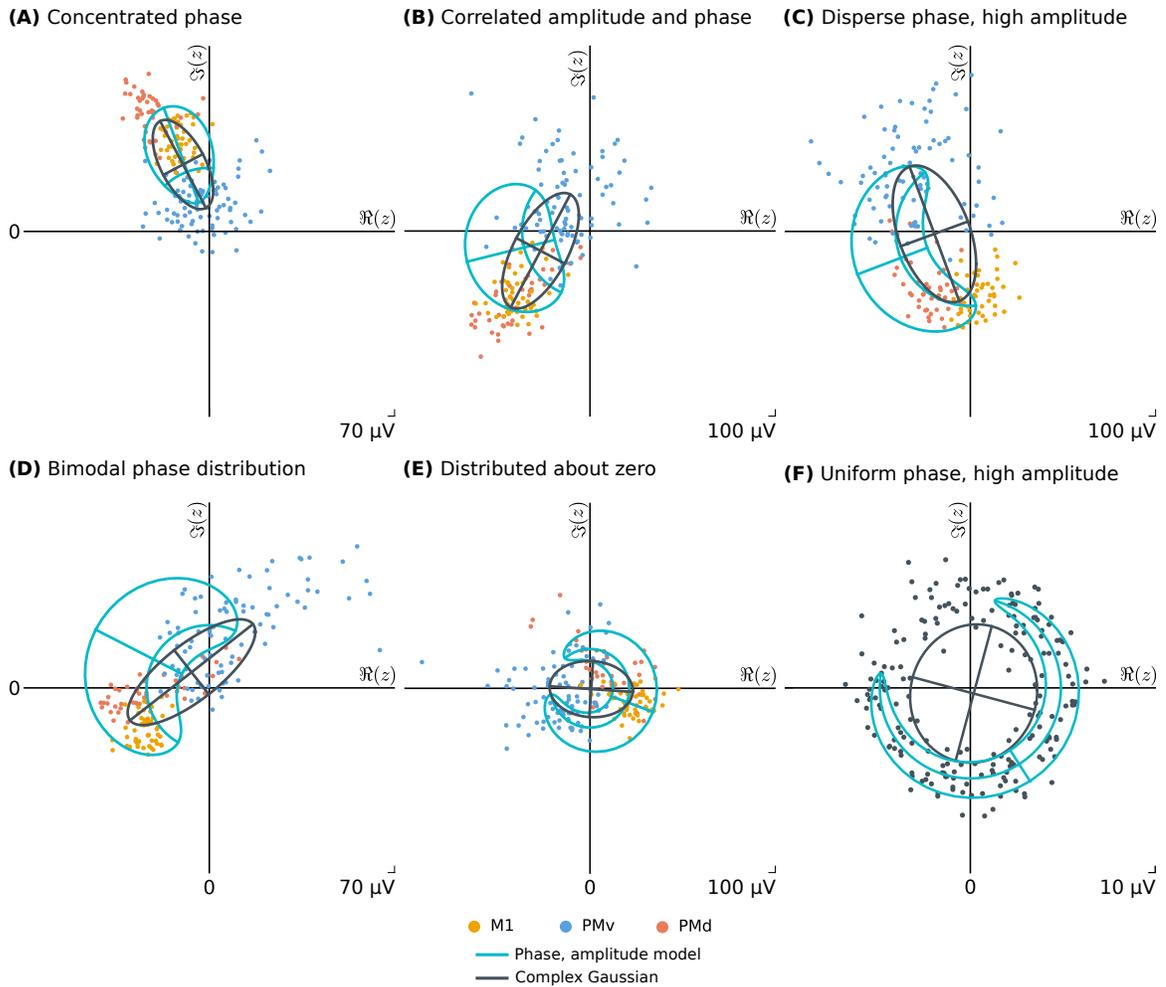


Figure 4.9: *Limitations of phase statistics in summarizing LFP population dynamics.* In these plots, the black ellipse represents a complex Gaussian model of the population of complex-valued β -LFP analytic signals, with the ellipse boundary at one standard deviation and ellipse axes representing the eigenvectors of the covariance matrix. The cyan contours represent a log-polar description of the data, which uses a circular Gaussian model for the phase and a log-normal distribution for the amplitude, with the ellipse boundary representing one standard deviation, transformed from log-polar coordinates using conformal mapping. $\Re(z)$ and $\Im(z)$ denote the real and imaginary parts of the complex-valued analytic LFP signal, respectively. **(A)** During zero-lag spatially synchronous wave events, in which phase is concentrated and not correlated with amplitude, both the log-polar statistics and the complex Gaussian distribution describe the distribution of analytic signals well. **(B)** When phase and amplitude are correlated, as we observe occurs at the initiation and termination of traveling wave events, the log-polar model cannot capture the phase-amplitude interactions. **(C)** During traveling wave events, signal amplitude is high, and there is dispersion in phase. In practice we do not observe highly disperse phase distributions at high amplitudes, and both the log-polar and complex Gaussian distributions describe the analytic signal distribution well. **(D)** Traveling wave events often evolve from states that show a mixture of synchrony and standing wave dynamics. The log-polar statistics break down when the phase distribution is bimodal, but the complex Gaussian can describe these states gracefully. **(E)** At low signal amplitudes, the system is often asynchronous, and the mean phase in the log-polar model is difficult to interpret. The complex Gaussian model gracefully describes the asynchronous state. **(F)** Although absent in our data, a hypothetical distribution with uniform phase and concentrated amplitude could occur, say, during traveling wave events with short wavelength. In this case, the complex Gaussian model is unsuitable, highlighting that while the complex Gaussian appears most suitable to describing our data in practice, neither model is capable of gracefully describing all states that might occur in spatiotemporal wave dynamics.

tion of plane wave events, and are poorly captured by a model in which phase and amplitude are treated separately. Additionally, asynchronous states in which the amplitude is distributed around zero and the phase is uniformly distributed, are poorly captured. In contrast, modeling the distribution of analytic signals as a bivariate Gaussian distribution appears promising. It can handle cases in which modulations of phase and amplitude are correlated, and gracefully describes the asynchronous case as a symmetric Gaussian with zero mean. However, the complex Gaussian description fails to capture the distribution of signals when amplitude is elevated and phases are dispersed, as can be the case for a high amplitude short wavelength traveling wave. Nevertheless, this particular case did not occur in our data and we had to illustrate it using simulated data.

4.3.2 Optogenetically evoked traveling waves

To contrast the statistics of spontaneous beta spatiotemporal oscillations with those of a clear example of traveling waves, we examined propagating waves at gamma frequency induced by optogenetic stimulation as first reported in primate motor cortex by Lu et al. (2015). Lu et al. (2015) showed that 1-second constant pulse of optogenetic stimulation elicits ~ 50 Hz gamma oscillations throughout the cortical patch recorded by the MEA. These gamma oscillations organized primarily as traveling radial waves.

Here, we analyzed the same data further and report that the spatial and frequency structure of the LFP response to optogenetic stimulation is in fact significantly more complex (Figure 4.11). In addition to the induced ~ 50 Hz oscillation, there is a distinct ~ 100 Hz high-gamma oscillation local to the ~ 1 mm of tissue receiving direct optical stimulation. We emphasize that the power around the higher 100 Hz frequency is not a Fourier harmonic. Instead, ~ 100 Hz oscillations can be easily detected via visual inspection in the time domain. Additionally, there was a broad ~ 70 Hz band evident on several channels (e.g. channel 88), as well as a low-amplitude band between 140 and 150 Hz. These induced oscillations appear to coexist with ongoing background β -LFP oscillations at ~ 20

Hz, which was addressed in Lu et al. (2015) and is also evident as a ~ 20 Hz beta peak that does not modulate with stimulation in Figure 4.11.

To further resolve the spatial distribution of the LFP response, we plot a time-frequency spectrogram for each channel in (Figure 4.12). Each of these plots show the frequency response over the 1-second optogenetic stimulation pulse. Unlike the 70 Hz and 100 Hz oscillations, the 50 Hz activity extends away from the site of optical stimulation and therefore is mediated by network interactions. We found that the induced LFP response was highly reproducible. Figure 4.14 explores two statistics of the average wave activity: the average phase delay, and the average zero-lag spatial phase gradient (Methods: ‘Averaged phase maps for optogenetically induced spatiotemporal waves’). The ~ 50 Hz and ~ 100 Hz bands showed reproducible spatial structure (Figure 4.14), (~ 70 Hz and ~ 140 Hz oscillations did not). To visualize the time-domain evolution of the 50 and 100 Hz wave components, we computed averages of LFP activity triggered on the phase of ongoing narrow band oscillations. Figure 4.15 displays phase triggered averages for both the ~ 50 and ~ 100 Hz gamma bands, as well as a single average of broad-band gamma activity triggered on the phase of the ~ 50 Hz oscillations, which combines the structure of the high and low gamma waves. Although the lower ~ 50 Hz oscillation appears to dominate the broad-band activity, the ~ 100 Hz alters the overall spatiotemporal structure during some phases of the oscillation

4.3.2.1 Correlations in the summary statistics of optogenetically induce 50 Hz gamma waves

Similarly to the beta analysis, we examined the correlations between the amplitude of the induced gamma oscillations and the degree of spatial synchrony, as assessed by the standard deviation of the Hilbert phase angles across the array (Figure 4.16, top left). Because this is not a plane-wave phenomenon, we did not examine correlations between the amplitude and the degree of alignment of phase gradient vectors. However, there

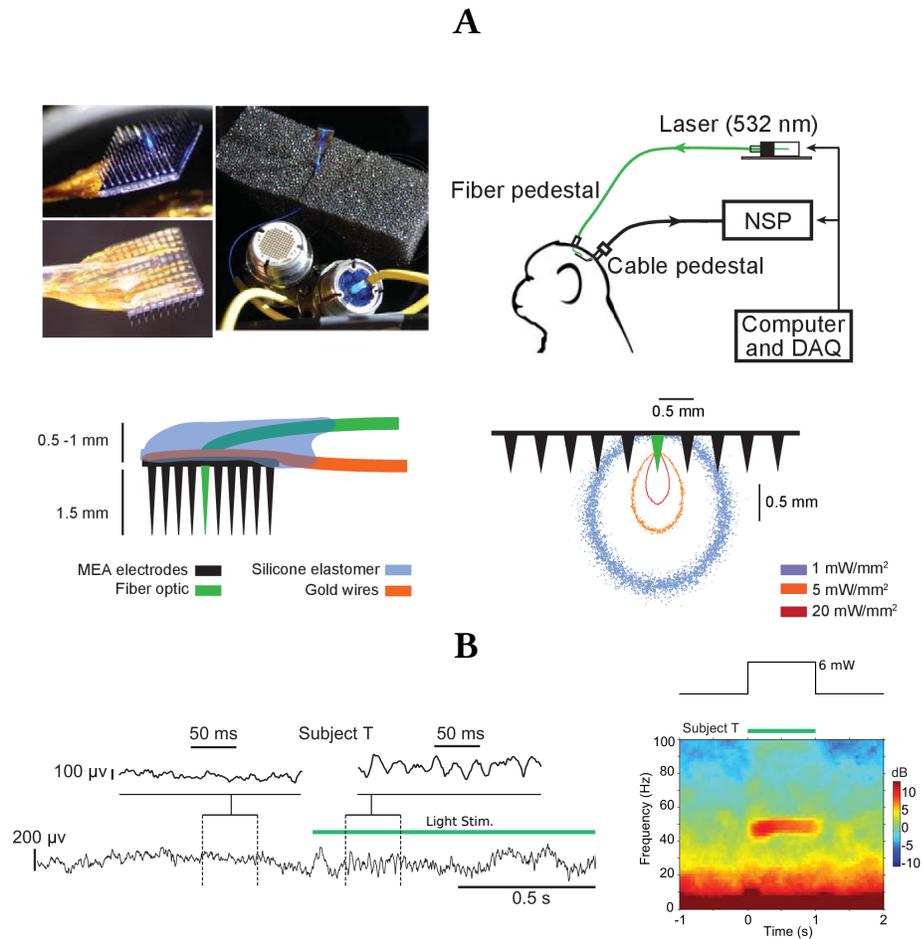


Figure 4.10: *Optogenetically induced spatiotemporal gamma waves in motor cortex.* (A) For optogenetic stimulation, a 10×10 Blackrock multi-electrode array was modified to contain a tapered optical fiber. The multi-electrode array output and fiber optic cable were routed through two separate pedestals. The region of direct optical stimulation was approximately 1 mm across. (B) 50 Hz gamma LFP oscillations are a prominent feature of the response to constant 1-second pulse optogenetic stimulation. When a constant pulse of 6 mW light stimulation was delivered, ~ 50 Hz oscillations emerge abruptly, as visible in both the example LFP trace (left), and the LFP spectrogram (right). *Reproduced from Lu et al. (2015) with permission.*

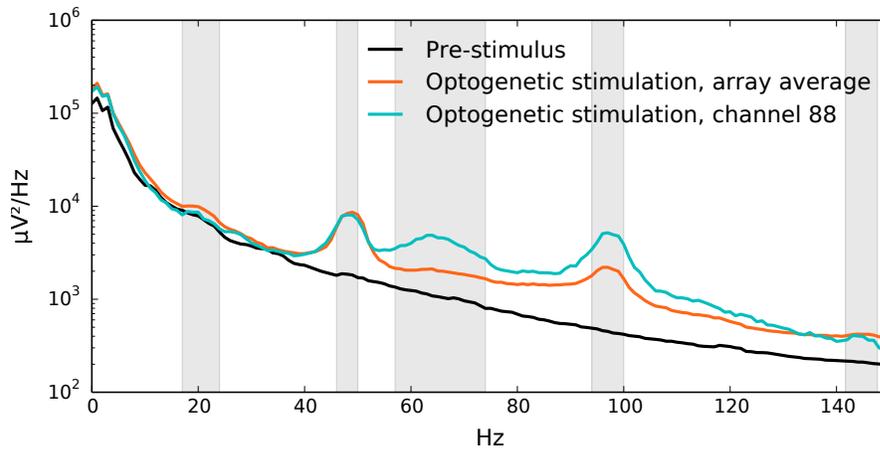


Figure 4.11: *Constant 1-second optogenetic stimulation of motor cortex induces complex oscillatory states.* Compared to background activity, power in ~ 50 Hz and ~ 100 Hz oscillations were elevated. Lu et al. (2015) examined only the ~ 50 Hz band. On electrodes close to the light stimulation (e.g. channel 88), an additional ~ 70 Hz band was also induced. Induced gamma (~ 50 , ~ 70 , ~ 100 Hz) oscillations coexisted with spontaneous beta (~ 20 Hz) oscillations. The ~ 50 Hz and ~ 100 Hz bands showed reproducible spatial structure (Figure 4.14), while the ~ 70 Hz oscillations showed no consistent phase delays relative to the stimulation site. Lu et al. (2015) conjectured that transition into spatiotemporal gamma oscillations occurred via a Hopf bifurcation. (Dataset from subject T in Lu et al. (2015).)

was very good signal-to-noise ratio during the induced gamma oscillations, such that we were able to estimate the local phase gradient accurately. This allowed us to extract a local estimate of wavelength (Methods: ‘Spatial gradient of the Hilbert phase’), and the spatial scale of the induced wave could be estimated by averaging this wavelength over the array.

Figure 4.16 summarizes the relationship between the amplitude, wavelength, and the standard deviation of phase for the optogenetically induced ~ 50 Hz gamma waves. Notably, there is clear clustering distinguishing the spatiotemporal statistics in the stimulated and the non-stimulated conditions. The induced gamma waves show a characteristic 2-4 mm wavelength, elevated amplitude, and reduced spatial synchrony. Compared to the background gamma activity, the induced waves show reduced variability in the estimated wavelength. Additionally, the optogenetic stimulation disrupted zero lag spatial phase synchrony over the cortical patch, which need not have been the case. These statistics were consistent with abrupt appearance of a true traveling wave with a characteristic wavelength, and are to be contrasted with the spatiotemporal statistics of the spontaneous ongoing beta oscillations detected during the movement preparation stages of the

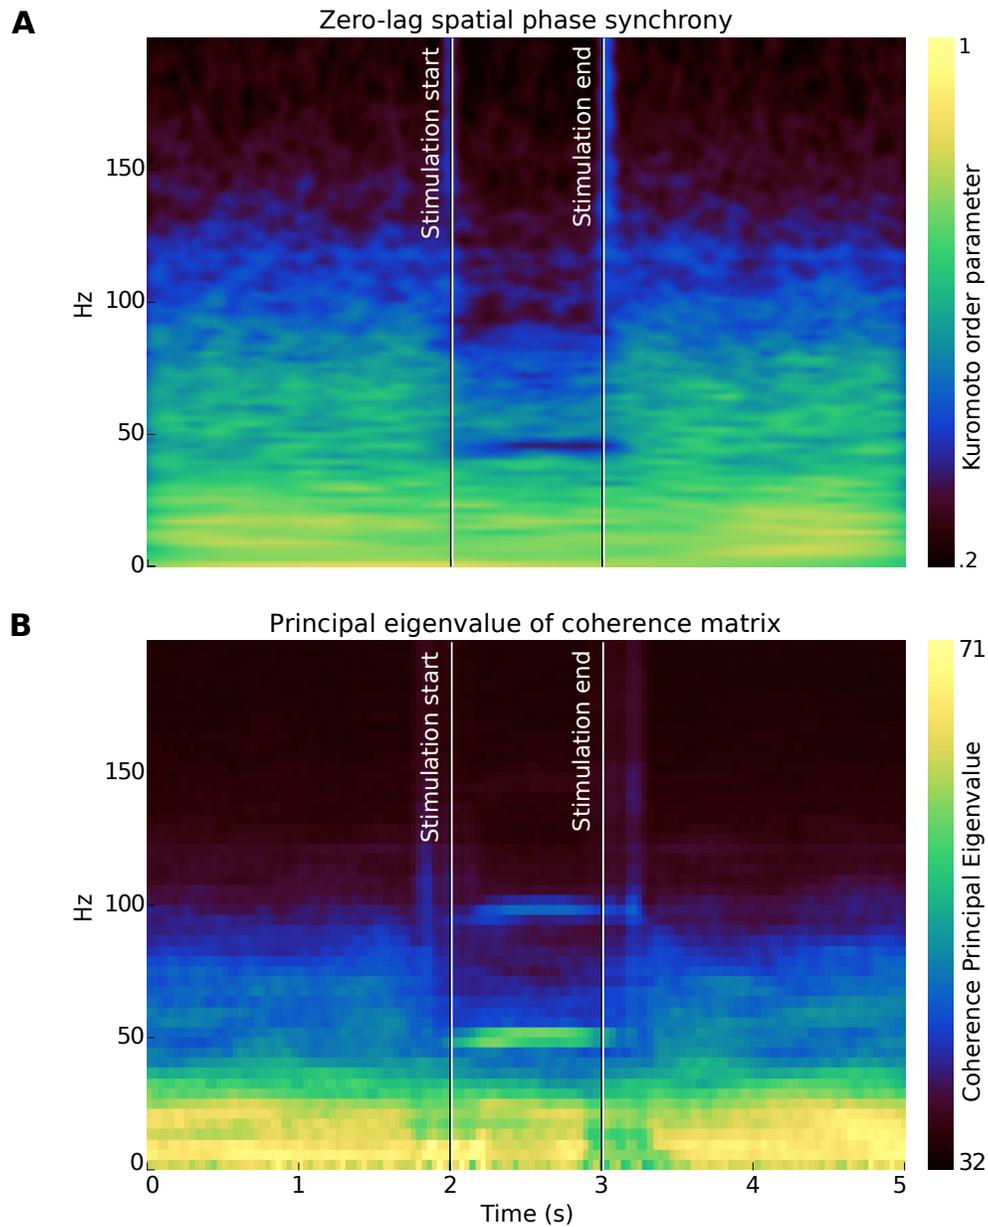


Figure 4.13: *Spatial synchrony decreases and coherence increases during optogenetic stimulation.* **(A)** Shown here is a plot of Kuramoto’s order parameter over the array channels, as it varies across frequencies during stimulation. Optogenetic stimulation occurred between 2 and 3 seconds. Frequency is on the y-axis. The Kuramoto order parameter was computed as $k(\omega, t) = |E_{channels}[\exp(i \cdot \phi(\omega, t))]|$. It consists of the magnitude of an average over channels of unit-length phase vectors representing the instantaneous phase for a given frequency ω , time t , and channel. Instantaneous phase for all times and frequencies was extracted using a Morlet wavelet transform. **(B)** Coherence, calculated using Welch’s method over a sliding 512 ms window. Pairwise coherence values are averaged over trials, and the largest eigenvalues of this trial-averaged pairwise coherence matrix are taken as a summary of the spatial coherence across the array. The spatial coherence increases during optogenetic stimulation, indicating collective participation in the induced ~50 Hz and ~100 Hz oscillations.

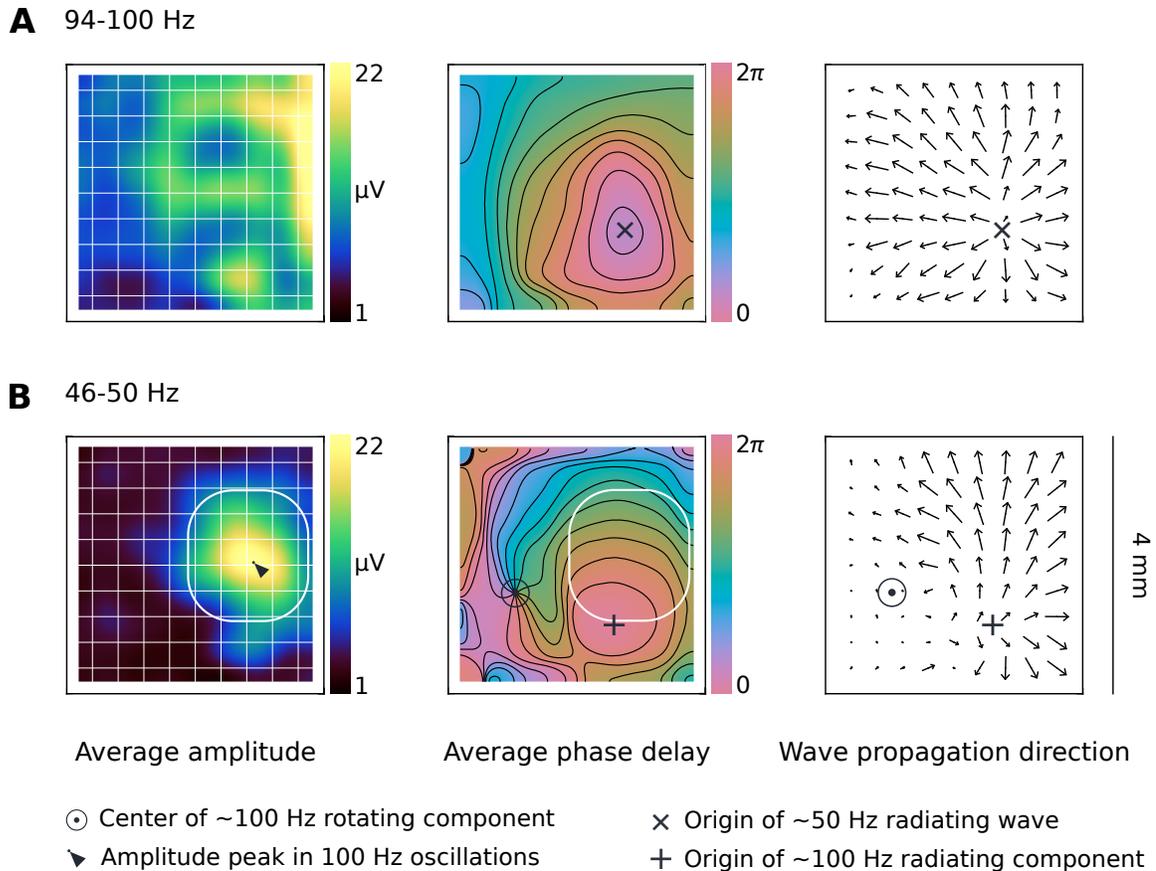


Figure 4.14: *Induced gamma oscillations exhibit spatiotemporal wave structure.* **(A)** ~50 Hz LFP amplitude is increased across the multielectrode array (left panel). The optical stimulation area was on the order of 1 mm across Lu et al. (2015), implying that these spatially extended ~50 Hz oscillations involved network interactions. Unwrapping phases relative to a reference electrode (channel 25) near the center of optical stimulation permits creating an average phase delay map (center panel). The induced ~50 Hz gamma oscillations organized as radiating waves, with an origin that was off-center from the source of the optical stimulation. The average phase gradient direction (right panel) was consistent with the average phase delay map, and confirms that the primary structure in these ~50 Hz oscillations was a radiating wave. **(B)** Induced ~100 Hz oscillations also showed spatiotemporal structure. ~100 Hz oscillations were localized to the stimulation area (left panel). Measuring the average phase delay relative to an electrode near the optical fiber reveals that these ~100 Hz oscillations exhibited spatiotemporal structure. Within the region where ~100 Hz oscillations have the highest amplitude, traveling plane waves propagated across the stimulated region (white box), starting at a point on the boundary of the stimulated region (black '+', panels 2 and 3). There was an associated spiral wave center (white dot). The average phase gradient map (right panel) confirms that there was a traveling wave structure in the 100 Hz band.

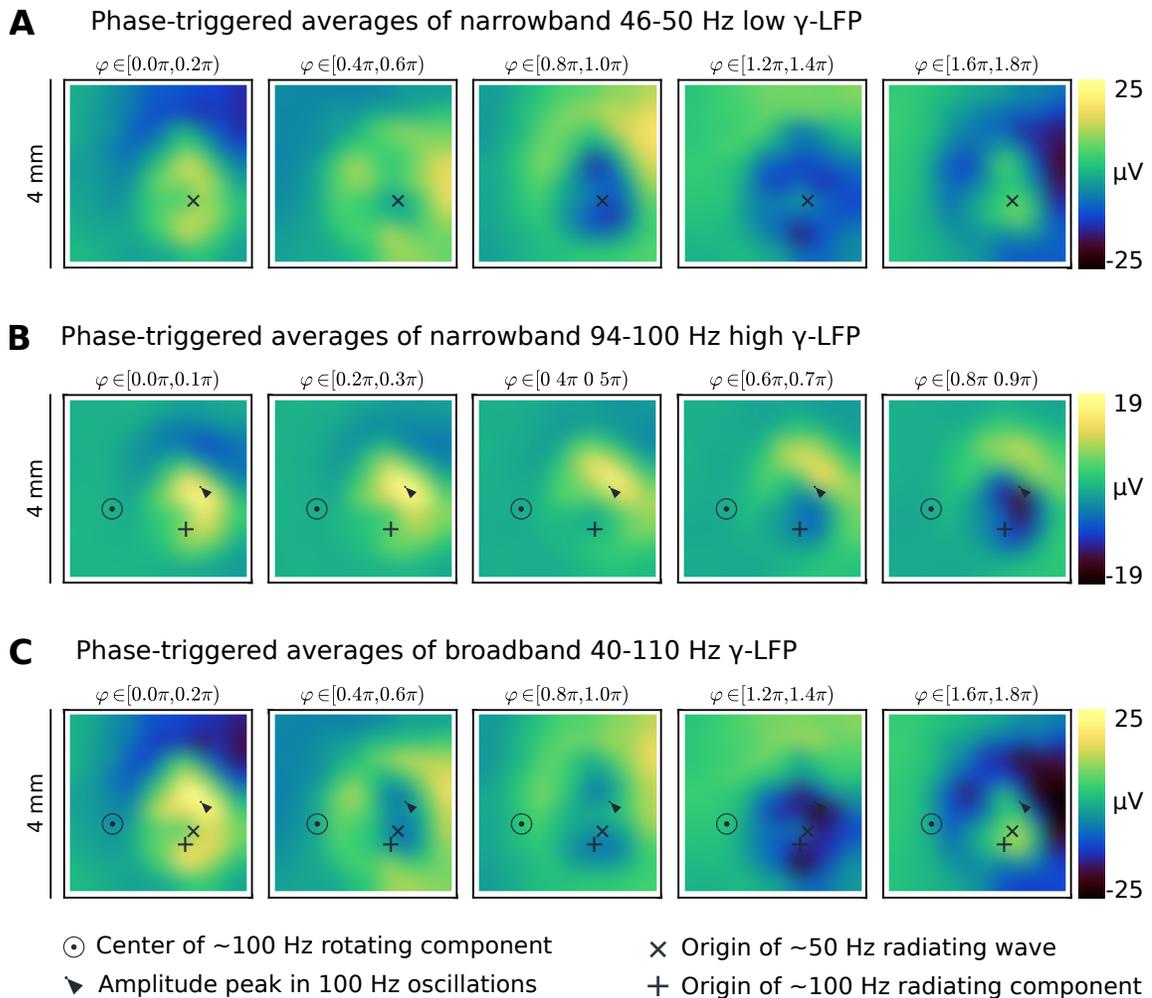


Figure 4.15: Phase-triggered averages of multi-band γ -LFP waves. Averages over all optical stimulation trials, triggered on the LFP phase (φ) at channel 25 (see Figure 4.12). **(A)** The average ~ 50 Hz wave component radiates from a point near the lower boundary of the optically stimulated region of cortex. **(B)** The average ~ 100 Hz oscillation (one-half oscillation cycle shown) displays traveling waves locally, seemingly emerging from a pair of counterclockwise and clockwise rotating centers close to the origin of the ~ 50 Hz radiating wave. **(C)** The sum of the ~ 50 and ~ 100 Hz waves, triggered on the phase of the ~ 50 Hz wave, displays a complex structure, superficially most similar to the ~ 50 Hz wave, but showing a rebound in LFP voltage in the center of the stimulated region midway through the cycle, related to the ~ 100 Hz component. Both the lower ~ 50 Hz and upper ~ 100 Hz γ -LFP contribute to the overall spatiotemporal response, highlighting that analysis of isolated LFP bands may overlook aspects of spatiotemporal dynamics.

CGID task.

4.3.2.2 Comparison of spontaneous beta waves with optogenetically induced gamma

We highlight the differences between the statistics of the spontaneous beta and the induced gamma oscillations in Figure 4.17, which shows that the induced gamma oscillations were associated with an abrupt increase in amplitude and decrease in spatial synchrony, compared to the background activity. In contrast, the spontaneous beta showed a continuous variation in amplitude and a tendency for *increased* spatial synchrony at higher amplitudes.

Figure 4.17 demonstrates that the induced traveling gamma waves have a characteristic fixed wavelength that varies little with amplitude relative to the background activity. In contrast, the wavelength of beta oscillations increased linearly with amplitude, such that the largest amplitude beta events were effectively highly synchronous over the recorded area of cortex. This highlights that simple summary statistics like amplitude and spatial synchrony can distinguish different types of traveling wave phenomena, and reflect differences in their underlying dynamics. Our findings are consistent with the hypothesis that spatiotemporal beta oscillations during movement preparation result from the synchronization dynamics of ongoing, local beta oscillations, while the gamma waves are true traveling waves induced by a local optogenetic perturbation.

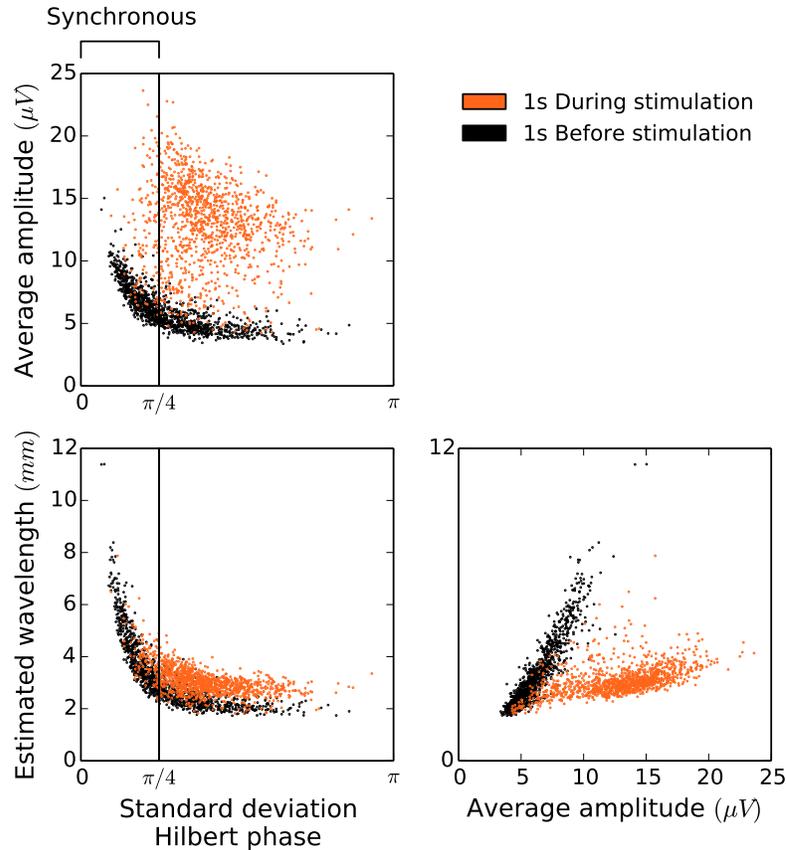


Figure 4.16: Amplitude, wavelength, and synchrony of optogenetically induced γ -LFP spatiotemporal oscillations. Scatter plots compare three summary statistics of optogenetically induced spatiotemporal gamma oscillations: amplitude, wavelength, and spatial synchrony. Amplitude was taken as the average Hilbert amplitude over the array. Wavelength was computed from the average local phase gradient magnitude. Spatial synchrony was assessed by the circular standard deviation of zero-lag oscillation phase, for which smaller values correspond to more synchronous states. Each dot in the scatter plot represents a snapshot of the spatiotemporal dynamics during one gamma cycle (20 ms). All stimulation trials during the session are combined. The statistics of optogenetically induced gamma waves (orange) were distinctly different from the background ~ 50 Hz activity (black). Note that the wavelength statistic does not have a physical interpretation outside of the stimulation window, since organized wave activity was absent. Instead, this can be interpreted as a statistic proportional to the inverse of the average local phase gradient magnitude – a statistic that exhibits broad variation and linear dependence on amplitude outside of the stimulation epoch. During optogenetic stimulation, ~ 50 Hz gamma activity exhibited a 2-4 mm wavelength and increased in amplitude as compared to the background activity. Importantly, the induced oscillations lacked spatial synchrony.

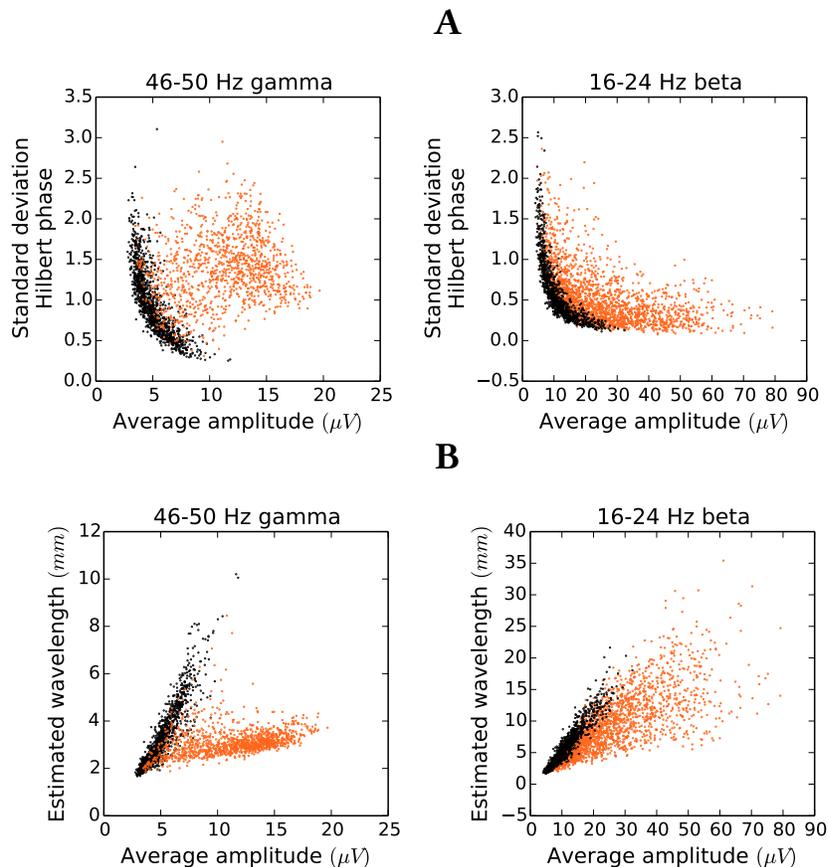


Figure 4.17: Comparison of statistics between spontaneous β -LFP and induced γ -LFP spatiotemporal waves. **(A)** Synchrony was assessed by the circular standard deviation of zero-lag oscillation phase, for which smaller values correspond to more synchronous states. Optogenetically induced ~ 50 Hz gamma oscillations in subject T (left panel) Lu et al. (2015) were associated with abrupt emergence of spatial structure in the form of radiating waves, creating a state that was *less* spatially synchronized than the background activity. In contrast, spontaneous beta ~ 20 Hz oscillations in M1 of subject S (right panel) showed a tendency toward *increased* spatial phase synchrony at higher amplitudes. This contrast suggests that the statistics of beta oscillations differ from those of an induced traveling wave in an excitable medium, indicating that beta traveling waves are only one of many possible collective modes of beta oscillations in motor cortex. **(B)** A comparison of the estimated wavelength of traveling waves between induced gamma in subject T (left) and spontaneous beta in subject S area M1 (right): The analysis of traveling plane waves at beta frequency was restricted to unambiguous plane wave events (circular standard deviation of the phase gradient direction less than $\pi/4$, see Methods: ‘Spatial gradient of the Hilbert phase’). Induced gamma oscillations emerged abruptly with a characteristic wavelength of 2-4 mm. In contrast, spontaneous beta plane waves showed varied wavelengths that correlate strongly with amplitude. Importantly, longer wavelength plane waves displayed a high degree of spatial phase synchrony, suggesting that beta plane waves transitioned smoothly from more complex wave patterns into globally synchronized activity.

4.4 Discussion

In this Chapter , we characterized two spatiotemporal phenomena in motor cortex LFP: spontaneous beta waves appearing during the steady-state periods of an instructed delay task, and traveling gamma waves induced by optogenetic stimulation. To our knowledge, our study is the first to report such systematic analysis of the range of spatiotemporal waves in motor cortex. We found that beta wave activity in motor cortex was significantly more diverse than previously reported. In particular, we observed that global synchronization of beta phase occurred more frequently than traveling plane waves, and that complex wave states like radiating and rotating waves were also present. Spatiotemporal beta waves appeared as transient events with varied durations (sometimes up to several hundred milliseconds). During a single transient, beta oscillations could transition spontaneously among different spatial wave patterns, suggesting stochastic transitions between spatiotemporal attractors in a multi-stable system. The analysis of datasets from optogenetically induced spatiotemporal gamma waves in primate motor cortex, first reported in Lu et al. (2015), provided a clear reference for truly traveling waves in excitable media to which to compare the spontaneous beta waves. These induced spatiotemporal patterns were globally organized, consisting of radial waves or synchronous gamma oscillation patterns. While spatiotemporal beta waves showed no characteristic spatial wavelength, induced traveling waves exhibited a characteristic wavelength of 2-4 mm. This finding is consistent with the hypothesis that spatiotemporal beta oscillations during movement preparation result from the synchronization dynamics of ongoing, local beta oscillations, while the gamma waves are true traveling waves induced by a local optogenetic perturbation.

We examined beta spatiotemporal dynamics in area PMv, a premotor area thought to be involved in grasping but not previously studied in terms of beta spatiotemporal waves, and found that it differed from that of areas M1 and PMd. In monkey S area PMv exhibited complex wave activity, even during transient events in which beta power was

elevated. In contrast, monkey R exhibited synchronous beta oscillations in PMv, showing few plane waves. These results are intriguing as previous research has suggested that beta phase relationships between motor cortical areas may relate to neural communication, but the potential role of beta oscillations in mediating intra- and inter-areal communication between PMv and M1 during grasping has not been studied. Area PMv does not appear to exhibit the plane wave phenomena previously reported in M1 and PMd (Rubino et al., 2006), but does exhibit organized radiating and rotating waves.

Variations in the statistics of spatiotemporal beta oscillations appeared to be closely related to variation in beta amplitude. High beta states showed a tendency to be more synchronous and exhibit more traveling wave states, while radiating, rotating, and complex wave states occurred more frequently at lower beta power. Combined with the observation that the wavelength of plane waves correlates with beta amplitude in area M1, this suggests that modulations in beta power are closely related to changes in lateral synchronization, with higher beta power associated with more synchronous states. As the origins of β -LFP in motor cortex remain unclear, we speculate that these observations may be consistent with a number of scenarios which should be disambiguated in further studies.

It is possible that the recorded neocortical β -LFP oscillations have multiple origins, reflecting both local sources of beta rhythmicity (e.g. Kopell et al. (2011), Roopun et al. (2006)), as well as cortico-thalamic interactions and beta rhythmic inputs from other brain areas (e.g. Jones et al. (2009) for the case of beta in somatosensory areas). Changes in the apparent spatial organization of motor cortex β -LFP may reflect changes in the relative contributions of these sources to the β -LFP signal. For example, a remote synchronous source may mix with local, asynchronous beta oscillations. If fluctuations in beta power reflect mainly fluctuations in this synchronous source, it could explain the correlation between beta synchrony and amplitude. However, the fact that we also observe that the wavelength of plane waves in M1 correlates with beta amplitude would seem at odds with this. These globally organized states cannot be explained as chance occurrences in

an otherwise asynchronous population, and it is therefore likely that changes in beta spatiotemporal patterns reflect changes in local network properties that influence the degree of synchrony and the wavelength of traveling waves. It may be that beta-rhythmic external inputs, perhaps from thalamus or sensory cortex, act to synchronize ongoing beta activity in motor cortex. These inputs would appear as transient increases in β -LFP power, reflecting currents associated with synaptic input, which may explain the correlation between the level of β -LFP power and the degree of spatial synchronization.

Alternatively, the variations in β -LFP power and beta spatiotemporal patterns may arise intrinsically in motor cortex. For example, coupling between local ongoing sources of beta may be higher when beta power is increased, leading to greater overall synchronization. Conversely, when local beta oscillations become more synchronous, they would be expected to interfere constructively and give rise to a larger LFP signal. Understanding the mechanisms that give rise to variability in beta power and in beta spatiotemporal pattern is important for understanding the role of beta oscillations in motor control, and for better understanding how beta oscillations are induced or suppressed within the motor system. If changes in motor cortex beta spatiotemporal LFP reflect inputs from other brain areas, they could be an important signature of coordinated activity of the motor system. If these changes arise from local population dynamics, it is an interesting challenge to understand how local dynamics could shift between different beta states, and how these local dynamics are modified by inputs to motor cortex.

We examined optogenetically induced gamma waves in greater detail, and found several phenomena not previously reported. We emphasize that these spatiotemporal gamma patterns were not induced by periodic optogenetic drive of neuronal populations, but by 1-second *constant* light stimulation pulses. This fact suggest that gamma oscillations are also a “natural mode” of the dynamics in motor cortex and that sustained gamma activity may arise via a Hopf bifurcation (Lu et al., 2015). Furthermore, the optogenetic response was extremely complex, consisting of up to four frequency bands at approx-

imately 50, 70, 100, and 140 Hz. The frequency multiples 50 and 100 were not simply Fourier harmonics, but consisted of the actual narrow band oscillations obvious to visual inspection. Only the ~50 Hz oscillation propagated away from the stimulation site as a traveling wave, although spatial structure local to the stimulation site also existed in the ~100 Hz band. These induced oscillations do not appear to interfere with ongoing beta oscillations. Previous studies have suggested that beta oscillations arise from layer V and gamma oscillations from layer II/III (Chen and Fetz, 2005; Murthy and Fetz, 1996a; Roopun et al., 2006; Wetmore and Baker, 2004; Witham and Baker, 2007). One possibility is that the optogenetic response involved mainly superficial layers and did not interfere strongly with ongoing beta activity in layer V. It is an open question whether existing models of gamma oscillations can capture these spatial and temporal properties of such complex multi-band optogenetic response. In addition, modeling the impact of optogenetic stimulation on motor cortex populations will be an important next step in understanding and controlling stimulation effects in motor cortex.

The contrast between the statistics of spontaneous beta and induced gamma waves highlights differences in the neural mechanisms underlying these phenomena and may provide constraints for spatially extended models of neural dynamics in motor cortex. As mentioned above, ~50 Hz gamma waves emerged with a 2-4 mm wavelength that varied little with amplitude, and the emergence of gamma waves disrupted gamma-frequency phase synchrony across the MEA. Spontaneous beta waves showed the opposite trend of increased synchrony at higher amplitudes, and lacked a characteristic wavelength, instead showing a gradual lengthening of plane waves at progressively larger amplitudes. This observation is important because it is seemingly inconsistent with neural field models of waves that rely on Mexican-hat style lateral interactions and predict wave activity emerging at a fixed spatial frequency (e.g. Ermentrout and Cowan 1979; Heitmann and Ermentrout 2015; Rule et al. 2011). The wide variability in the spatial scale of beta spatiotemporal activity is, however, consistent with a model in which beta waves arise from

transient organization of the relative phases of local, ongoing beta oscillations. In contrast, the traveling gamma waves are consistent with a model of wave propagation in an excitable medium, one where the wavelength is set by the properties of lateral interactions.

The difference between spontaneous beta and induced gamma experiments should, nevertheless, be interpreted cautiously. Some of the differences in the statistics of these phenomena arise from the discrete and localized nature of optogenetic stimulation. The optical stimulation abruptly creates a distinct network state, while neural activity during spontaneous beta oscillations varies continuously. The optogenetic stimulation is localized, while changes in network state giving rise to changes in beta dynamics may be global. The reduced synchrony during optogenetic stimulation may be due to the fact that the inputs to the system are spatially inhomogeneous. The lack of variability in the scale of gamma waves may be particular to the fixed amplitude of stimulation, and further studies should be conducted to thoroughly explore the relationship between stimulation amplitude and the spatiotemporal structure of the optogenetic response.

Multielectrode array recording approaches have limitations that should be considered when interpreting analyses of spatiotemporal wave dynamics. Because neocortex is curved, the depth of the electrodes could change significantly across the array. The phase of motor cortex beta varies across layers (Murthy and Fetz, 1996a), and so changes in the depth of the electrodes could translate to a phase gradient in the MEA recordings that mimics a traveling wave. This scenario cannot yield rotating or complex waves, and so cannot explain away several of the important beta and gamma wave patterns we reported here. It may, however, bias estimates of spatial structure in LFP oscillations. Additionally, there remains some ambiguity as to the thickness of *m. mulatta* motor cortex, as well as the extent to which these MEA implants may compress or cause thinning of cortex. Future histology of the implanted neocortical patches may resolve this issue.

β -LFP oscillations exhibited considerable variability in amplitude, and LFP phase is

difficult to interpret at lower amplitudes, both due to reduced signal-to-noise and because changes in amplitude and phase of the analytic signal are ambiguous when the amplitude is changing rapidly. Neural signals from MEAs are complex, and exhibit structure in time, frequency, space, and spatial frequency, which are interrelated. Further work is necessary to clarify the methods surrounding LFP phase extraction, and it remains unclear whether descriptions of spatiotemporal dynamics based on phase alone are sufficient. In light of the complex multi-band spatiotemporal dynamics observed in optogenetically induced gamma waves, the phase of a single narrow-band oscillation may not always be sufficient to capture spatiotemporal wave dynamics. In contrast to previous studies, here we have explicitly, if briefly, taken into account potential relationships between amplitude and phase in spatiotemporal beta waves. Further research addressing this issue, and improvements in the joint analysis of spatiotemporal changes in phase and amplitude, would be beneficial. Our preliminary analysis suggests that the complex Gaussian distribution over channels may be a promising alternative to analyzing population activity in neural oscillations with fluctuating amplitudes and phase organizations. Another possibility would be to incorporate a phase-amplitude coupling term. The approach used in Kempster et al. (2012) to describe joint relationships between linear position and the phase of hippocampal theta oscillations may be useful in jointly describing amplitude and phase. A more thorough study of how such approach may complement phase gradient methods would be welcome.

It has been suggested that spatiotemporal waves in cortex may mediate and coordinate information flow (e.g. Lubenov and Siapas (2009); Rubino et al. (2006); Takahashi et al. (2011); Wu et al. (2008a)), or at least be an epiphenomenal correlate of neural communication (e.g. Besserve et al. (2015)). Spatiotemporal waves occur in motor cortex in both spontaneous beta oscillations and the induced gamma oscillations studied here. Previous studies have linked spontaneous beta waves to spiking communication (Rubino et al., 2006; Takahashi et al., 2015, 2011). The diversity of beta spatiotemporal states observed

here, as well as the paucity of traveling waves in area PMv in both subjects, calls into question the simple explanation that planar beta traveling waves mediate inter-area communication during the preparation of reaching and grasping actions. Further work is needed to investigate the implications of this diversity, especially understanding the functional roles of traveling waves vs globally synchronized states, which both occur during high beta power. Additionally, the relationship between spatiotemporal organization of beta and modulation in beta amplitude may hold important clues as to how beta power is modulated in motor cortex, and how changes in beta power impact neural computation and communication. The observation that motor cortex beta phase can exhibit complex spatial dynamics has implications for our understanding of collective dynamics and single-unit activity in motor cortex. In particular, the relationship between complex phase dynamics and single-unit spiking still needs to be clarified (Okun et al., 2015). Additionally, traveling waves may contribute to single-unit spiking variability. The findings reported here are consistent with the hypothesis that beta traveling waves arise from phase gradients between local oscillators. Further research should clarify the extent to which beta spatiotemporal activity represents a coupled oscillator phenomenon as opposed to a propagating wave phenomenon. In addition, understanding the neuronal network response to optogenetic perturbations will be an important challenge for incorporating optogenetic tools into primate motor cortex research.

Chapter 5

Conclusion

MICHAEL RULE

This thesis investigated collective neural dynamics in primate motor cortex during two different motor control states: movement execution and preparatory steady-states. Chapters 2 and 3 addressed the relationship between single-unit activity and collective dynamics reflected in local field potentials during movement and steady-states, respectively. Chapters 3 and 4 explored collective dynamics related to beta oscillations in preparatory steady-states, with Chapter 4 focusing on the organization of spatiotemporal wave patterns in beta phase across motor cortex.

To review the main results, Chapter 2 found that, although information available in multi-band LFPs can predict single neuron spiking with substantial accuracy during movement, this information is largely redundant to the information available in the movement kinematics. In particular, during movement, collective neural dynamics reflected in local field potentials predominantly relate to two components of neuronal variability: fast-timescale spiking history effects, and slow dynamics related to movement generation. Chapter 3 found that rhythmic spiking in individual neurons can be dissociated from population beta oscillations, and raised several interesting questions about steady-state neural dynamics. Chapter 4 found that spatiotemporal dynamics in motor cortex local

field potentials are more complex than previously appreciated, and reports several observations that should be explored in future studies as well as additional methodological and theoretical research.

To restate the context of this work, it is hypothesized that the motor system acts collectively as a dynamical system to generate the slow, low-dimensional trajectories needed for movement execution. During such kinetic states, narrow-band beta oscillations are suppressed, and large deviations in multi-unit activity and local field potentials can be observed. In contrast, steady-states in motor cortex reflect attention, posture, stabilized force, and instructed delays, and exhibit transient bursts of ~20 Hz beta oscillations. Previous studies have found that the output neurons in motor cortex can support intrinsic beta rhythmicity, but the significance of this mechanism remains unresolved as several other mechanisms may contribute to beta oscillations.

Both motor kinetic and steady-states are of general interest. The mechanisms by which stochastic fast-timescale spiking acts collectively to generate robust slow-timescale dynamics are a subject of ongoing research. Slow dynamics of movement execution may be representative of spontaneous and stimulus-evoked processing in frontal cortex. Movement preparation steady-states constitute stabilized neural activity over long timescales, but whether such states are related to the beta oscillations that have been proposed to support working memory is unknown. Although beta oscillations are a correlate of the 'status quo', how beta oscillations facilitate both stabilization of motor steady-states and maintenance of rate-coded information about planned upcoming movements remains to be clarified. Steady-states employ cortical "reflex arcs" for stabilization and therefore must be compatible both with active sensorimotor processing and the stabilization of local firing rate information.

Chapter 2 found that the predictive information about spiking variability contained in motor cortex local field potentials was redundant to information available in kinematics trajectories ('pathlets') and intrinsic spiking history. The existence of kinematics-related

variability in both LFPs and single neuron spiking does not imply that kinematics is the only source of shared variability between the two. Shared variability between LFP and spiking may arise from intrinsic, and spontaneous neural dynamics. Understanding the relationship between LFP and spiking is of particular interest as local field potentials are thought to contain signatures of synaptic input averaged over a population. Because motor cortex presumably performs nontrivial transformation on its inputs, it is surprising that LFPs did not contain complementary information explaining spiking variability. Some have theorized that motor cortex forms a recurrent dynamical loop with deep brain structures that operates collectively for pattern generation. If this is the case, then inputs to motor cortex reflected in LFP would relate to this dynamical state. The LFP could (at least partially) reflect cortical inputs, but if the sources of these inputs are part of a recurrent dynamical system for movement generation, such signals would relate to past and upcoming kinematic output reflected in ‘pathlets’, and therefore be redundant to kinematics features for the purpose of explaining neural variability. Chapter 2’s result that stochastic fast-timescale effects can be accounted for by both local field potentials and spiking history is also relevant to brain machine interfaces, which at present typically smooth spiking signals to extract the slow-timescale signals related to motor output. Lower latency decoders may benefit from actively modeling fast-timescale effects to dissociate movement-related and stochastic-history related variability.

An open question raised in Chapters 2 and 3 is the significance of neurons in motor cortex that exhibit thin extracellular spikes. Narrow spikes are assumed to relate to fast repolarizing currents that allow cells to sustain high firing rates, a feature associated with inhibitory interneurons. Future studies are needed to determine whether motor cortex thin-spike cells are in fact PTNs related to driving fast-twitch muscle activity. Methods for separating fast-spiking pyramidal tract neurons from putative inhibitory interneurons are needed. While thin- and thick-spike cells showed little difference in their relationship to collective dynamics investigated here, they may show very different relationships to

kinematic output, which may be important for decoding and brain-machine interfaces.

We find in Chapter 3 that preparatory steady states (1) exhibit transient periods of beta-LFP power of variable duration, (2) are associated with sustained rhythmic spiking distributed between 10 and 45 Hz, which (3) does not show a stable phase relationship to the concurrent beta-LFP. We have also shown in Chapter 4 that movement preparation steady-states (4) exhibit changes in beta spatiotemporal dynamics, including increased spatial synchrony, smaller phase differences between adjacent electrodes, and less complex wave activity, and (5) exhibit a correlation between beta amplitude and the spatial scale of planar wave activity. All of these observations will need to be accounted for in future work toward models of how the motor system solves a variety of motor control problems, including (1) goal-dependent slow firing rate trajectories for movement execution, (2) steady-state stabilization of information in local firing during preparatory steady-states, and (3) facilitation of reflexive, stabilizing sensorimotor computations for motor steady-states. The mechanisms underlying the stabilization of firing rates during preparatory steady-states in particular remain elusive. It is unclear whether the observed LFP dynamics emerge as spatial averages over the beta rhythmic single-neurons recorded on the same electrode, and further experiments with laminar probes and current source density analysis will be important for addressing this. An intriguing finding in Chapter 3 was that transient fluctuations in beta-LFP oscillations do not appear to affect firing rates of single units exhibiting beta-rhythmic spiking, indicating that separate neural mechanisms modulate the degree of beta oscillations at the level of single neurons versus the local field potential.

The results in Chapters 3 and 4 lead to a number of conjectures about the nature of collective dynamics at beta frequency during preparatory steady-states. It is possible that beta-LFP oscillations arise as averages over the activity of weakly coupled beta-rhythmic single neurons, each firing at a slightly different rate, and that the high amplitude beta events arise when these phases transiently align. The primary reason that we see more

synchrony at higher amplitudes is that synchronized local oscillations sum constructively to yield a large amplitude signal. However, we did not find increases in single-unit phase coupling during high-beta states, either pairwise between neurons or between neurons and the LFP. It is possible that we did not have enough statistical power to detect weak effects in single-neurons that create a large effect at the population level. Conversely, beta amplitude fluctuations may have a causal role in modulating spatial phase diversity in LFP across motor cortex. Changes in the magnitude of beta oscillatory activity may alter the strength of lateral phase-resetting interactions, promoting or inhibiting spatially synchronized states. This hypothesis is similar to the models proposed in Heitmann et al., and the similarities and differences between the beta dynamics we observe, and the phase model outlined in the Heitmann papers, warrant further discussion. There remains the possibility that an extrinsic, synchronized source of beta oscillations mixes with local asynchronous beta activity. This would be the case if beta transients reflect shared sub-cortical or cortico-cortical input. However, the beta wave activity in monkey S area PMv displays a diversity of globally organized states, including complex states at high beta amplitudes. We would not expect such a diversity in beta spatiotemporal patterns at high amplitudes if the variation in beta amplitude were due entirely to changes in the magnitude of a synchronized source. Additionally, plane waves involve a high degree of order, and are incompatible with random asynchronous activity. These plane waves exhibit longer wavelengths at higher amplitudes, indicating that the correlation between beta synchrony and amplitude may influence the lateral interactions that give rise to plane waves. The different spatiotemporal states, (complex waves, plane waves, synchrony) may represent various possible modes of the beta dynamics, and the synchronized states simply exhibit higher amplitudes. This might be expected in the presence of self-excitatory connections, where phase-locking of lateral excitation increases its efficacy. Further experiments and modeling are needed to evaluate these hypotheses. In particular, the relative contribution of intrinsic spiking rhythmicity versus local and lateral network interactions in sustaining

beta oscillations needs to be clarified.

The characterization of beta spatiotemporal wave activity in Chapter 4 also constrains candidate mathematical and biophysical models of motor cortex, particularly the relationships between various statistics of beta spatiotemporal wave structure. We observed variability in the spatial scale of beta spatiotemporal activity, seemingly inconsistent with any models that require fixed center-surround inhibitory interactions for wave emergence. We suggest that the diverse spatiotemporal patterns seen in beta waves represent transient spatial organization of the relative phases of local ongoing beta oscillations, and hypothesize that the statistics of the observed beta traveling waves are inconsistent with true traveling waves propagating through an excitable medium. Models for transitions between wave states that involve not only phase but also shifts in the joint amplitude-phase distribution should be investigated. The characterization of optogenetically induced gamma oscillations presented here is of interest both as an example of complex, multi-band spatiotemporal neural activity, and as a contrasting example to suggest that various wave mechanisms (traveling, oscillatory phase gradients, propagation delays) may be distinguished by observing their spatiotemporal statistics. Clarifying the extent to which the statistics of spatiotemporal wave activity can constrain underlying network models is a promising avenue of future theoretical research.

In summary, this thesis addressed questions related to collective neural activity in primate motor cortex during movement execution and preparatory steady-states, examining the relationship between spiking and LFP, as well as spatiotemporal structure in LFP across motor cortex. Open questions remain in understanding the relationship between spiking and LFP, especially during preparatory steady-states. Of personal interest to me is understanding how rhythmic neural spiking at heterogeneous firing frequencies can give rise to the collective beta oscillations in local field potentials, and whether the lack of stable phase relationships and diverse firing frequencies has functional importance, perhaps in phase coding or in binding together neural assemblies for maintenance of preparatory

states. Additionally, I would like to understand whether the complex spatiotemporal wave dynamics that we observe in beta LFP are functionally significant. Neuroprosthetics based on cortical stimulation would benefit from understanding how to control the variations in beta phase across motor cortex, especially if the spatial organization causally impacts the neural dynamics. Likewise, I am excited by the possibilities of optogenetic stimulation to investigate population responses in motor cortex. This thesis demonstrates that the response to optogenetic stimulation in neocortical patches is nontrivial, and further experiments and modeling work will be needed before we can begin to use closed-loop optogenetic stimulation in a controlled way in motor cortex.

Overall, I view this thesis as work toward better treatments for movement disorders that use closed-loop spatiotemporally patterned electromagnetic and optogenetic stimulation. By understanding collective dynamics of the motor system, I hope that we can develop interventions that restore normal motor function with fewer side-effects. Of course, understanding the mechanisms of primate motor control also has broader philosophical implications. Understanding the neural basis of spontaneous movement touches on the nature of volition, and can inform our development of artificial intelligences that mimic the complex behavior of animals. To me, personally, these broader impacts are secondary to the potential for neuroscience research to improve the quality of life for patients with neurological impairments that are too complex to treat with current methods. The research presented here is several steps removed from translational applications, but I believe that we must invest in research at all levels of abstraction if we are to make progress. The results presented in this thesis contribute to our basic understanding of the motor system, and are a step forward toward future clinical investigations.

References

- Abbott, L. and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural computation*, 11(1):91–101.
- Adrian, E. D. (1936). The spread of activity in the cerebral cortex. *The Journal of physiology*, 88(2):127–161.
- Adrian, E. D. and Matthews, B. (1934). The interpretation of potential waves in the cortex. *The Journal of physiology*, 81(4):440.
- Aghagolzadeh, M. and Truccolo, W. (2014). Latent state-space models for neural decoding. In *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE*, pages 3033–3036. IEEE.
- Aghagolzadeh, M. and Truccolo, W. (2015). Inference and decoding of motor cortex low-dimensional dynamics via latent state-space models.
- Akam, T. and Kullmann, D. M. (2014). Oscillatory multiplexing of population codes for selective communication in the mammalian brain. *Nature Reviews Neuroscience*, 15(2):111–122.
- Alegre, M., Gurtubay, I., Labarga, A., Iriarte, J., Valencia, M., and Artieda, J. (2004). Frontal and central oscillatory changes related to different aspects of the motor process: a study in go/no-go paradigms. *Experimental Brain Research*, 159(1):14–22.
- Alegre, M., Imirizaldu, L., Valencia, M., Iriarte, J., Arcocha, J., and Artieda, J. (2006). Alpha and beta changes in cortical oscillatory activity in a go/no go randomly-delayed-response choice reaction time paradigm. *Clinical neurophysiology*, 117(1):16–25.
- Androulidakis, A. G., Doyle, L. M., Gilbertson, T. P., and Brown, P. (2006). Corrective movements in response to displacements in visual feedback are more effective during periods of 13–35 hz oscillatory synchrony in the human corticospinal system. *European Journal of Neuroscience*, 24(11):3299–3304.
- Androulidakis, A. G., Doyle, L. M., Yarrow, K., Litvak, V., Gilbertson, T. P., and Brown, P. (2007). Anticipatory changes in beta synchrony in the human corticospinal system and associated improvements in task performance. *European Journal of Neuroscience*, 25(12):3758–3765.

- Ardid, S., Wang, X.-J., Gomez-Cabrero, D., and Compte, A. (2010). Reconciling coherent oscillation with modulation of irregular spiking activity in selective attention: Gamma-range synchronization between sensory and executive cortical areas. *The Journal of Neuroscience*, 30(8):2856–2870.
- Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science*, 273(5283):1868–1871.
- Aroniadou, V. A. and Keller, A. (1993). The patterns and synaptic properties of horizontal intracortical connections in the rat motor cortex. *Journal of Neurophysiology*, 70(4):1553–1569.
- Asanuma, H. and Rosén, I. (1973). Spread of mono- and polysynaptic connections within cat's motor cortex. *Experimental Brain Research*, 16(5):507–520.
- Aumann, T. D. and Prut, Y. (2015). Do sensorimotor β -oscillations maintain muscle synergy representations in primary motor cortex? *Trends in neurosciences*, 38(2):77–85.
- Averbeck, B. B., Latham, P. E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5):358–366.
- Averbeck, B. B. and Lee, D. (2006). Effects of noise correlations on information encoding and decoding. *Journal of Neurophysiology*, 95(6):3633–3644.
- Aydore, S., Pantazis, D., and Leahy, R. M. (2013). A note on the phase locking value and its properties. *Neuroimage*, 74:231–244.
- Azouz, R. and Gray, C. M. (1999). Cellular mechanisms contributing to response variability of cortical neurons in vivo. *The Journal of neuroscience*, 19(6):2209–2223.
- Baker, S., Kilner, J., Pinches, E., and Lemon, R. (1999). The role of synchrony and oscillations in the motor output. *Experimental Brain Research*, 128(1-2):109–117.
- Baker, S., Olivier, E., and Lemon, R. (1997). Coherent oscillations in monkey motor cortex and hand muscle EMG show task-dependent modulation. *The Journal of Physiology*, 501(1):225–241.
- Baker, S., Spinks, R., Jackson, A., and Lemon, R. (2001). Synchronization in monkey motor cortex during a precision grip task. I. task-dependent modulation in single-unit synchrony. *Journal of Neurophysiology*, 85(2):869–885.
- Baker, S. N. (2007). Oscillatory interactions between sensorimotor cortex and the periphery. *Current Opinion in Neurobiology*, 17(6):649–655.
- Baker, S. N., Pinches, E. M., and Lemon, R. N. (2003). Synchronization in monkey motor cortex during a precision grip task. II. effect of oscillatory activity on corticospinal output. *Journal of Neurophysiology*, 89(4):1941–1953.

- Banerjee, A., Dean, H. L., and Pesaran, B. (2012). Parametric models to relate spike train and LFP dynamics with neural information processing. *Frontiers in Computational Neuroscience*, 6.
- Bansal, A. K., Truccolo, W., Vargas-Irwin, C. E., and Donoghue, J. P. (2012). Decoding 3-D reach and grasp from hybrid signals in motor and premotor cortices: spikes, multiunit activity, and local field potentials. *Journal of Neurophysiology*, (5):1337–1355.
- Bansal, A. K., Vargas-Irwin, C. E., Truccolo, W., and Donoghue, J. P. (2011). Relationships among low-frequency local field potentials, spiking activity, and three-dimensional reach and grasp kinematics in primary motor and ventral premotor cortices. *Journal of Neurophysiology*, 105(4):1603–1619.
- Baranyi, A., Szente, M. B., and Woody, C. D. (1993a). Electrophysiological characterization of different types of neurons recorded in vivo in the motor cortex of the cat. I. patterns of firing activity and synaptic responses. *Journal of Neurophysiology*, 69(6):1850–1864.
- Baranyi, A., Szente, M. B., and Woody, C. D. (1993b). Electrophysiological characterization of different types of neurons recorded in vivo in the motor cortex of the cat. II. membrane parameters, action potentials, current-induced voltage responses and electrotonic structures. *Journal of Neurophysiology*, 69(6):1865–1879.
- Barthó, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K. D., and Buzsáki, G. (2004). Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *Journal of neurophysiology*, 92(1):600–608.
- Battaglia, D., Karagiannis, A., Gallopin, T., Gutch, H. W., and Cauli, B. (2013). Beyond the frontiers of neuronal types. *Frontiers in neural circuits*, 7.
- Bedard, C. and Destexhe, A. (2013). Reply to gratiy et al. *J Neurophysiol*, 109(1683).
- Belitski, A., Gretton, A., Magri, C., Murayama, Y., Montemurro, M. A., Logothetis, N. K., and Panzeri, S. (2008). Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *The Journal of Neuroscience*, 28(22):5696–5709.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 289–300.
- Benucci, A., Frazor, R. A., and Carandini, M. (2007). Standing waves and traveling waves distinguish two circuits in visual cortex. *Neuron*, 55(1):103–117.
- Berdyeva, T. K. and Reynolds, J. H. (2009). The dawning of primate optogenetics. *Neuron*, 62(2):159–160.
- Berger, H. (1929). Über das elektrenkephalogramm des menschen. *European Archives of Psychiatry and Clinical Neuroscience*, 87(1):527–570.

- Besserve, M., Lowe, S. C., Logothetis, N. K., Schölkopf, B., and Panzeri, S. (2015). Shifts of gamma phase across primary visual cortical sites reflect dynamic stimulus-modulated information transfer. *PLoS Biol*, 13(9):e1002257.
- Best, M. D., Takahashi, K., Suminski, A. J., Ethier, C., Miller, L. E., and Hatsopoulos, N. G. (2016). Comparing offline decoding performance in physiologically defined neuronal classes. *Journal of Neural Engineering*, 13(2):026004.
- Beuter, A., Lefaucheur, J.-P., and Modolo, J. (2014). Closed-loop cortical neuromodulation in Parkinson’s disease: An alternative to deep brain stimulation? *Clinical Neurophysiology*, 125(5):874–885.
- Bishop, W., Chestek, C. C., Gilja, V., Nuyujukian, P., Foster, J. D., Ryu, S. I., Shenoy, K. V., and Byron, M. Y. (2014). Self-recalibrating classifiers for intracortical brain–computer interfaces. *Journal of neural engineering*, 11(2):026001.
- Bokil, H., Andrews, P., Kulkarni, J. E., Mehta, S., and Mitra, P. P. (2010). Chronux: a platform for analyzing neural signals. *Journal of Neuroscience Methods*, 192(1):146–151.
- Bosch-Bouju, C., Hyland, B. I., and Parr-Brownlie, L. C. (2015). Motor thalamus integration of cortical, cerebellar and basal ganglia information: implications for normal and parkinsonian conditions. *Basal ganglia: physiological, behavioral, and computational studies*.
- Boussaoud, D., Ungerleider, L. G., and Desimone, R. (1990). Pathways for motion analysis: cortical connections of the medial superior temporal and fundus of the superior temporal visual areas in the macaque. *Journal of Comparative Neurology*, 296(3):462–495.
- Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nature neuroscience*, 8(9):1263–1268.
- Braendgaard, H., Evans, S., Howard, C., and Gundersen, H. (1990). The total number of neurons in the human neocortex unbiasedly estimated using optical disectors. *Journal of Microscopy*, 157(3):285–304.
- Breakspear, M., Heitmann, S., and Daffertshofer, A. (2010). Generative models of cortical oscillations: neurobiological implications of the kuramoto model. *Frontiers in Human Neuroscience*, 4.
- Brittain, J.-S. and Brown, P. (2014). Oscillations and the basal ganglia: motor control and beyond. *Neuroimage*, 85:637–647.
- Brittain, J.-S., Sharott, A., and Brown, P. (2014). The highs and lows of beta activity in cortico-basal ganglia loops. *European Journal of Neuroscience*, 39(11):1951–1959.
- Brovelli, A., Ding, M., Ledberg, A., Chen, Y., Nakamura, R., and Bressler, S. L. (2004). Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by granger causality. *Proceedings of the National Academy of Sciences of the United States of America*, 101(26):9849–9854.

- Brown, E. N., Frank, L. M., Tang, D., Quirk, M. C., and Wilson, M. A. (1998). A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *The Journal of Neuroscience*, 18(18):7411–7425.
- Brown, P. (2007). Abnormal oscillatory synchronisation in the motor system leads to impaired movement. *Current opinion in neurobiology*, 17(6):656–664.
- Brunel, N. (2003). Dynamics and plasticity of stimulus-selective persistent activity in cortical network models. *Cerebral Cortex*, 13(11):1151–1161.
- Brunel, N. and Wang, X.-J. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? I. synaptic dynamics and excitation-inhibition balance. *Journal of Neurophysiology*, 90(1):415–430.
- Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012). The origin of extracellular fields and currents— EEG, ECoG, LFP and spikes. *Nature reviews neuroscience*, 13(6):407–420.
- Buzsáki, G. and Draguhn, A. (2004). Neuronal oscillations in cortical networks. *science*, 304(5679):1926–1929.
- Byron, M. Y., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. In *Advances in neural information processing systems*, pages 1881–1888.
- Calvin, W. H. and Sybert, G. W. (1976). Fast and slow pyramidal tract neurons: an intracellular analysis of their contrasting repetitive firing properties in the cat. *Journal of Neurophysiology*, 39(2):420–434.
- Camperi, M. and Wang, X.-J. (1998). A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *Journal of computational neuroscience*, 5(4):383–405.
- Canolty, R. T., Ganguly, K., and Carmena, J. M. (2012). Task-dependent changes in cross-level coupling between single neurons and oscillatory activity in multiscale networks. *PLoS Comput Biol*, 8(12):e1002809.
- Capaday, C., van Vreeswijk, C., Ethier, C., Ferkinghoff-Borg, J., and Weber, D. (2011). Neural mechanism of activity spread in the cat motor cortex and its relation to the intrinsic connectivity. *The Journal of Physiology*, 589(10):2515–2528.
- Carandini, M. et al. (2004). Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS biology*, 2:1483–1493.
- Cardin, J. A., Palmer, L. A., and Contreras, D. (2007). Stimulus feature selectivity in excitatory and inhibitory neurons in primary visual cortex. *The Journal of Neuroscience*, 27(39):10333–10344.

- Cassidy, M., Mazzone, P., Oliviero, A., Insola, A., Tonali, P., Di Lazzaro, V., and Brown, P. (2002). Movement-related changes in synchronization in the human basal ganglia. *Brain*, 125(6):1235–1246.
- Cavanaugh, J., Monosov, I. E., McAlonan, K., Berman, R., Smith, M. K., Cao, V., Wang, K. H., Boyden, E. S., and Wurtz, R. H. (2012). Optogenetic inactivation modifies monkey visuomotor behavior. *Neuron*, 76(5):901–907.
- Chatrian, G. E., Petersen, M. C., and Lazarte, J. A. (1959). The blocking of the rolandic wicket rhythm and some central changes related to movement. *Electroencephalography and clinical neurophysiology*, 11(3):497–510.
- Chen, C. C., Litvak, V., Gilbertson, T., Kühn, A., Lu, C. S., Lee, S. T., Tsai, C. H., Tisch, S., Limousin, P., Hariz, M., et al. (2007). Excessive synchronization of basal ganglia neurons at 20 hz slows movement in parkinson’s disease. *Experimental neurology*, 205(1):214–221.
- Chen, D. and Fetz, E. E. (2005). Characteristic membrane potential trajectories in primate sensorimotor cortex neurons recorded in vivo. *Journal of Neurophysiology*, 94(4):2713–2725.
- Chhatbar, P. Y. and Francis, J. T. (2013). Towards a naturalistic brain-machine interface: hybrid torque and position control allows generalization to novel dynamics. *PloS one*, 8(1):e52286.
- Churchland, M. M. and Abbott, L. (2012). Two layers of neural variability. *Nature Neuroscience*, 15(11):1472–1474.
- Churchland, M. M., Byron, M. Y., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., et al. (2010a). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*, 13(3):369–378.
- Churchland, M. M. and Cunningham, J. P. (2015). A dynamical basis set for generating reaches. In *Cold Spring Harbor symposia on quantitative biology*, page 024703. Cold Spring Harbor Laboratory Press.
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., and Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, 487(7405):51–56.
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., and Shenoy, K. V. (2010b). Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron*, 68(3):387–400.
- Compte, A., Brunel, N., Goldman-Rakic, P. S., and Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, 10(9):910–923.

- Cruz, A. (2012). *An entropy-based investigation of underpinnings and impact of oscillations in a model of PD*. PhD thesis, UCL (University College London).
- Dai, J., Brooks, D. I., and Sheinberg, D. L. (2014). Optogenetic and electrical microstimulation systematically bias visuospatial choice in primates. *Current Biology*, 24(1):63–69.
- Denker, M., Roux, S., Lindén, H., Diesmann, M., Riehle, A., and Grün, S. (2011). The local field potential reflects surplus spike synchrony. *Cerebral Cortex*, 21(12):2681–2695.
- Denker, M., Roux, S., Timme, M., Riehle, A., and Grün, S. (2007). Phase synchronization between LFP and spiking activity in motor cortex during movement preparation. *Neurocomputing*, 70(10):2096–2101.
- Destexhe, A. and Bedard, C. (2012). Do neurons generate monopolar current sources? *Journal of neurophysiology*, 108(4):953–955.
- Deweese, M. R. and Zador, A. M. (2004). Shared and private variability in the auditory cortex. *Journal of Neurophysiology*, 92(3):1840–1855.
- Diester, I., Kaufman, M. T., Mogri, M., Pashaie, R., Goo, W., Yizhar, O., Ramakrishnan, C., Deisseroth, K., and Shenoy, K. V. (2011). An optogenetic toolbox designed for primates. *Nature neuroscience*, 14(3):387–397.
- Dockstader, C., Gaetz, W., Cheyne, D., Wang, F., Castellanos, F. X., and Tannock, R. (2008). Meg event-related desynchronization and synchronization deficits during basic somatosensory processing in individuals with adhd. *Behav Brain Funct*, 4(8):1–13.
- Donoghue, J. P., Sanes, J. N., Hatsopoulos, N. G., and Gaál, G. (1998). Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *Journal of Neurophysiology*, 79(1):159–173.
- Doyle, L. M., Yarrow, K., and Brown, P. (2005). Lateralization of event-related beta desynchronization in the eeg during pre-cued reaction time tasks. *Clinical neurophysiology*, 116(8):1879–1888.
- Dugué, L., Marque, P., and VanRullen, R. (2011). The phase of ongoing oscillations mediates the causal relation between brain excitation and visual perception. *The Journal of Neuroscience*, 31(33):11889–11893.
- Dum, R. P. and Strick, P. L. (2002). Motor areas in the frontal lobe of the primate. *Physiology & behavior*, 77(4):677–682.
- Ecker, A. S., Berens, P., Cotton, R. J., Subramaniyan, M., Denfield, G. H., Cadwell, C. R., Smirnakis, S. M., Bethge, M., and Tolias, A. S. (2014). State dependence of noise correlations in macaque primary visual cortex. *Neuron*, 82(1):235–248.
- Engel, A. K. and Fries, P. (2010). Beta-band oscillations—signalling the status quo? *Current Opinion in Neurobiology*, 20(2):156–165.

- Ermentrout, G. B. and Cowan, J. D. (1979). A mathematical theory of visual hallucination patterns. *Biological cybernetics*, 34(3):137–150.
- Ermentrout, G. B. and Kleinfeld, D. (2001). Traveling electrical waves in cortex: insights from phase dynamics and speculation on a computational role. *Neuron*, 29(1):33–44.
- Eusebio, A. and Brown, P. (2009). Synchronisation in the beta frequency-band—the bad boy of parkinsonism or an innocent bystander? *Experimental neurology*, 217(1):1–3.
- Faisal, A. A., Selen, L. P., and Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4):292–303.
- Fawcett, T. (2006). An introduction to ROC analysis. *Pattern recognition letters*, 27(8):861–874.
- Feingold, J., Gibson, D. J., DePasquale, B., and Graybiel, A. M. (2015). Bursts of beta oscillation differentiate postperformance activity in the striatum and motor cortex of monkeys performing movement tasks. *Proceedings of the National Academy of Sciences*, 112(44):13687–13692.
- Ferraina, S., Battaglia-Mayer, A., Genovesio, A., Marconi, B., Onorati, P., and Caminiti, R. (2001). Early coding of visuomanual coordination during reaching in parietal area pec. *Journal of Neurophysiology*, 85(1):462–467.
- Ferraina, S., Garasto, M., Battaglia-Mayer, A., Ferraresi, P., Johnson, P., Lacquaniti, F., and Carniniti, R. (1997). Visual control of hand-reaching movement: activity in parietal area 7m. *European Journal of Neuroscience*, 9(5):1090–1095.
- Fiser, J., Chiu, C., and Weliky, M. (2004). Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature*, 431(7008):573–578.
- Foehring, R. C., Lorenzon, N. M., Herron, P., and Wilson, C. J. (1991). Correlation of physiologically and morphologically identified neuronal types in human association cortex in vitro. *Journal of Neurophysiology*, 66(6):1825–1837.
- Freeman III, W. J. (1972). Waves, pulses, and the theory of neural masses. *Progress in theoretical biology*, 2(1).
- Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1.
- Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in cognitive sciences*, 9(10):474–480.
- Fries, P., Nikolić, D., and Singer, W. (2007). The gamma cycle. *Trends in neurosciences*, 30(7):309–316.
- Fujioka, T., Ross, B., and Trainor, L. J. (2015). Beta-band oscillations represent auditory beat and its metrical hierarchy in perception and imagery. *The Journal of Neuroscience*, 35(45):15187–15198.

- Fujioka, T., Trainor, L. J., Large, E. W., and Ross, B. (2012). Internalized timing of isochronous sounds is represented in neuromagnetic beta oscillations. *The Journal of Neuroscience*, 32(5):1791–1802.
- Gale, J. T., Amirnovin, R., Williams, Z. M., Flaherty, A. W., and Eskandar, E. N. (2008). From symphony to cacophony: pathophysiology of the human basal ganglia in Parkinson disease. *Neuroscience & Biobehavioral Reviews*, 32(3):378–387.
- Galvan, A., Hu, X., Smith, Y., and Wichmann, T. (2012). In vivo optogenetic control of striatal and thalamic neurons in non-human primates. *PloS one*, 7(11):e50808.
- García-Cabezas, M. Á. and Barbas, H. (2014). Area 4 has layer iv in adult primates. *European Journal of Neuroscience*, 39(11):1824–1834.
- Geisler, C., Brunel, N., and Wang, X.-J. (2005). Contributions of intrinsic membrane dynamics to fast network oscillations with irregular neuronal discharges. *Journal of Neurophysiology*, 94(6):4344–4361.
- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., and Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of Neuroscience*, 2(11):1527–1537.
- Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233(4771):1416–1419.
- Gerits, A., Farivar, R., Rosen, B. R., Wald, L. L., Boyden, E. S., and Vanduffel, W. (2012). Optogenetically induced behavioral and functional network changes in primates. *Current Biology*, 22(18):1722–1726.
- Gerits, A. and Vanduffel, W. (2013). Optogenetics in primates: a shining future? *Trends in Genetics*, 29(7):403–411.
- Gilbertson, T., Lalo, E., Doyle, L., Di Lazzaro, V., Cioni, B., and Brown, P. (2005). Existing motor state is favored at the expense of new movement during 13-35 hz oscillatory synchrony in the human corticospinal system. *The Journal of neuroscience*, 25(34):7771–7779.
- Goldman, M. S., Levine, J. H., Major, G., Tank, D. W., and Seung, H. (2003). Robust persistent neural activity in a model integrator with multiple hysteretic dendrites per neuron. *Cerebral cortex*, 13(11):1185–1195.
- González-Ramírez, L., Ahmed, O., Cash, S., Wayne, C., and Kramer, M. (2015). A biologically constrained, mathematical model of cortical wave propagation preceding seizure termination. *PLoS Computational Biology*, 11(2):e1004065–e1004065.
- Goris, R. L., Movshon, J. A., and Simoncelli, E. P. (2014). Partitioning neuronal variability. *Nature Neuroscience*, 17(6):858–865.

- Goulet, J. and Ermentrout, G. B. (2011). The mechanisms for compression and reflection of cortical waves. *Biological cybernetics*, 105(3-4):253–268.
- Grasse, D. W. and Moxon, K. A. (2010). Correcting the bias of spike field coherence estimators due to a finite number of spikes. *Journal of neurophysiology*, 104(1):548–558.
- Gredal, O., Pakkenberg, H., Karlsborg, M., and Pakkenberg, B. (2000). Unchanged total number of neurons in motor cortex and neocortex in amyotrophic lateral sclerosis: a stereological study. *Journal of neuroscience methods*, 95(2):171–176.
- Grefkes, C. and Fink, G. R. (2005). Review: The functional organization of the intraparietal sulcus in humans and monkeys. *Journal of anatomy*, 207(1):3–17.
- Guler, S. D. and Rule, M. E. (2013). Invent-abling: enabling inventiveness through craft. In *Proceedings of the 12th International Conference on Interaction Design and Children*, pages 368–371. ACM.
- Haegens, S., Nácher, V., Hernández, A., Luna, R., Jensen, O., and Romo, R. (2011). Beta oscillations in the monkey sensorimotor network reflect somatosensory decision making. *Proceedings of the National Academy of Sciences*, 108(26):10708–10713.
- Hafting, T., Fyhn, M., Bonnevie, T., Moser, M.-B., and Moser, E. I. (2008). Hippocampus-independent phase precession in entorhinal grid cells. *Nature*, 453(7199):1248–1252.
- Hammond, C., Bergman, H., and Brown, P. (2007). Pathological synchronization in parkinson’s disease: networks, models and treatments. *Trends in neurosciences*, 30(7):357–364.
- Han, X. (2012). Optogenetics in the nonhuman primate. *Progress in brain research*, 196:215.
- Han, X., Qian, X., Bernstein, J. G., Zhou, H.-h., Franzesi, G. T., Stern, P., Bronson, R. T., Graybiel, A. M., Desimone, R., and Boyden, E. S. (2009). Millisecond-timescale optical control of neural dynamics in the nonhuman primate brain. *Neuron*, 62(2):191–198.
- Harvey, C. D., Collman, F., Dombeck, D. A., and Tank, D. W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature*, 461(7266):941–946.
- Haslinger, R., Ulbert, I., Moore, C. I., Brown, E. N., and Devor, A. (2006). Analysis of LFP phase predicts sensory response of barrel cortex. *Journal of Neurophysiology*, 96(3):1658–1663.
- Hasson, U., Malach, R., and Heeger, D. J. (2010). Reliability of cortical activity during natural stimulation. *Trends in cognitive sciences*, 14(1):40–48.
- Hatsopoulos, N. G. and Amit, Y. (2012). Synthesizing complex movement fragment representations from motor cortical ensembles. *Journal of Physiology-Paris*, 106(3):112–119.
- Hatsopoulos, N. G., Xu, Q., and Amit, Y. (2007). Encoding of movement fragments in the motor cortex. *The Journal of Neuroscience*, 27(19):5105–5114.

- Havenith, M. N., Yu, S., Biederlack, J., Chen, N.-H., Singer, W., and Nikolić, D. (2011). Synchrony makes neurons fire in sequence, and stimulus properties determine who is ahead. *The Journal of Neuroscience*, 31(23):8570–8584.
- Heitmann, S., Boonstra, T., and Breakspear, M. (2013). A dendritic mechanism for decoding traveling waves: Principles and applications to motor cortex. *PLoS Computational Biology*, 9(10).
- Heitmann, S. and Ermentrout, G. B. (2015). Synchrony, waves and ripple in spatially coupled kuramoto oscillators with mexican hat connectivity. *Biological cybernetics*, pages 1–15.
- Heitmann, S., Gong, P., and Breakspear, M. (2012). A computational role for bistability and traveling waves in motor cortex. *Frontiers in Computational Neuroscience*, 6.
- Hermes, D., Siero, J. C., Aarnoutse, E. J., Leijten, F. S., Petridou, N., and Ramsey, N. F. (2012). Dissociation between neuronal activity in sensorimotor cortex and hand movement revealed as a function of movement rate. *The Journal of Neuroscience*, 32(28):9736–9744.
- Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., Haddadin, S., Liu, J., Cash, S. S., van der Smagt, P., et al. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375.
- Hoseini, M. S. and Wessel, R. (2015). Coherent and intermittent ensemble oscillations emerge from networks of irregular spiking neurons. *Journal of Neurophysiology*, pages jn-00578.
- Huang, X., Troy, W. C., Yang, Q., Ma, H., Laing, C. R., Schiff, S. J., and Wu, J.-Y. (2004). Spiral waves in disinhibited mammalian neocortex. *The Journal of Neuroscience*, 24(44):9897–9902.
- Huang, X., Xu, W., Liang, J., Takagaki, K., Gao, X., and Wu, J.-y. (2010). Spiral wave dynamics in neocortex. *Neuron*, 68(5):978–990.
- Huntley, G. W. and Jones, E. G. (1991). Relationship of intrinsic connections to forelimb movement representations in monkey motor cortex: a correlative anatomic and physiological study. *Journal of neurophysiology*, 66(2):390–413.
- Hwang, K., Ghuman, A. S., Manoach, D. S., Jones, S. R., and Luna, B. (2014). Cortical neurodynamics of inhibitory control. *The Journal of Neuroscience*, 34(29):9551–9561.
- Hyva, J. et al. (1981). Regional distribution of functions in parietal association area 7 of the monkey. *Brain research*, 206(2):287–303.
- Jackson, A., Gee, V. J., Baker, S. N., and Lemon, R. N. (2003). Synchrony between neurons with similar muscle fields in monkey motor cortex. *Neuron*, 38(1):115–125.
- Jasper, H. and Penfield, W. (1949). Electrocticograms in man: effect of voluntary movement upon the electrical activity of the precentral gyrus. *Archiv für Psychiatrie und Nervenkrankheiten*, 183(1-2):163–174.

- Jasper, H. H. and Andrews, H. L. (1938). Electro-encephalography: Iii. normal differentiation of occipital and precentral regions in man. *Archives of Neurology & Psychiatry*, 39(1):96–115.
- Jazayeri, M., Lindbloom-Brown, Z., and Horwitz, G. D. (2012). Saccadic eye movements evoked by optogenetic activation of primate v1. *Nature neuroscience*, 15(10):1368–1370.
- Jenkinson, N. and Brown, P. (2011). New insights into the relationship between dopamine, beta oscillations and motor function. *Trends in neurosciences*, 34(12):611–618.
- Jensen, O. and Bonnefond, M. (2013). Prefrontal alpha-and beta-band oscillations are involved in rule selection. *Trends in cognitive sciences*, 17(1):10–12.
- Jensen, O. and Lisman, J. E. (2000). Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding. *Journal of neurophysiology*, 83(5):2602–2609.
- Johnson, P. B., Ferraina, S., Bianchi, L., and Caminiti, R. (1996). Cortical networks for visual reaching: physiological and anatomical organization of frontal and parietal lobe arm regions. *Cerebral Cortex*, 6(2):102–119.
- Jones, E., Oliphant, T., Peterson, P., et al. (2001). SciPy: Open source scientific tools for Python. [Online; accessed 2016-02-04].
- Jones, S. R., Pritchett, D. L., Sikora, M. A., Stufflebeam, S. M., Hämäläinen, M., and Moore, C. I. (2009). Quantitative analysis and biophysically realistic neural modeling of the MEG mu rhythm: rhythmogenesis and modulation of sensory-evoked responses. *Journal of Neurophysiology*, 102(6):3554–3572.
- Kajikawa, Y. and Schroeder, C. E. (2011). *Neuron*, 72(5):847–858.
- Kaneko, T. (2013). Local connections of excitatory neurons in motor-associated cortical areas of the rat. *Frontiers in neural circuits*, 7.
- Kaufman, M. T., Churchland, M. M., Ryu, S. I., and Shenoy, K. V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nature neuroscience*, 17(3):440–448.
- Kaufman, M. T., Churchland, M. M., Santhanam, G., Byron, M. Y., Afshar, A., Ryu, S. I., and Shenoy, K. V. (2010). Roles of monkey premotor neuron classes in movement preparation and execution. *Journal of Neurophysiology*, 104(2):799–810.
- Kayser, C., Montemurro, M. A., Logothetis, N. K., and Panzeri, S. (2009). Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron*, 61(4):597–608.
- Kelly, R. C., Smith, M. A., Kass, R. E., and Lee, T. S. (2010). Local field potentials indicate network state and account for neuronal response variability. *Journal of computational neuroscience*, 29(3):567–579.

- Kempster, R., Leibold, C., Buzsáki, G., Diba, K., and Schmidt, R. (2012). Quantifying circular-linear associations: Hippocampal phase precession. *Journal of Neuroscience Methods*, 207(1):113–124.
- Khanna, P. and Carmena, J. M. (2015). Neural oscillations: beta band activity across motor networks. *Current opinion in neurobiology*, 32:60–67.
- Kilavik, B. E., Ponce-Alvarez, A., Trachel, R., Confais, J., Takerkart, S., and Riehle, A. (2012). Context-related frequency modulations of macaque motor cortical lfp beta oscillations. *Cerebral cortex*, 22(9):2148–2159.
- Kilavik, B. E., Zaepffel, M., Brovelli, A., MacKay, W. A., and Riehle, A. (2013). The ups and downs of beta oscillations in sensorimotor cortex. *Experimental neurology*, 245:15–26.
- Kilner, J. M., Baker, S. N., Salenius, S., Hari, R., and Lemon, R. N. (2000). Human cortical muscle coherence is directly related to specific motor parameters. *The Journal of Neuroscience*, 20(23):8838–8845.
- Kilpatrick, Z. P. and Ermentrout, B. (2013). Wandering bumps in stochastic neural fields. *SIAM Journal on Applied Dynamical Systems*, 12(1):61–94.
- Koike, H., Mano, N., Okada, Y., and Oshima, T. (1970). Repetitive impulses generated in fast and slow pyramidal tract cells by intracellularly applied current steps. *Experimental brain research*, 11(3):263–281.
- Koike, H., Okada, Y., and Oshima, T. (1968). Accommodative properties of fast and slow pyramidal tract cells and their modification by different levels of their membrane potential. *Experimental Brain Research*, 5(3):189–201.
- Kopell, N., Whittington, M., and Kramer, M. (2011). Neuronal assembly dynamics in the beta1 frequency range permits short-term memory. *Proceedings of the National Academy of Sciences*, 108(9):3779–3784.
- Krack, P., Dostrovsky, J., Ilinsky, I., Kultas-Ilinsky, K., Lenz, F., Lozano, A., and Vitek, J. (2002). Surgery of the motor thalamus: problems with the present nomenclatures. *Movement disorders*, 17(S3):S2–S8.
- Kristeva, R., Patino, L., and Omlor, W. (2007). Beta-range cortical motor spectral power and corticomuscular coherence as a mechanism for effective corticospinal interaction during steady-state motor output. *Neuroimage*, 36(3):785–792.
- Kühn, A. A., Doyle, L., Pogosyan, A., Yarrow, K., Kupsch, A., Schneider, G.-H., Hariz, M. I., Trottenberg, T., and Brown, P. (2006). Modulation of beta oscillations in the subthalamic area during motor imagery in parkinson’s disease. *Brain*, 129(3):695–706.
- Kühn, A. A., Tsui, A., Aziz, T., Ray, N., Brücke, C., Kupsch, A., Schneider, G.-H., and Brown, P. (2009). Pathological synchronisation in the subthalamic nucleus of patients with parkinson’s disease relates to both bradykinesia and rigidity. *Experimental neurology*, 215(2):380–387.

- Kuramoto, Y. (1984). Chemical oscillations, waves, and turbulence, ser. *Springer Series in Synergetics*. Berlin/Heidelberg, Germany: Springer-Verlag, 19.
- Kuramoto, Y. (2012). *Chemical oscillations, waves, and turbulence*, volume 19. Springer Science & Business Media.
- Laing, C. R. and Chow, C. C. (2001). Stationary bumps in networks of spiking neurons. *Neural Computation*, 13(7):1473–1494.
- Lepage, K. Q., Gregoriou, G. G., Kramer, M. A., Aoi, M., Gotts, S. J., Eden, U. T., and Desimone, R. (2013). A procedure for testing across-condition rhythmic spike-field association change. *Journal of Neuroscience Methods*, 213(1):43–62.
- Leventhal, D. K., Gage, G. J., Schmidt, R., Pettibone, J. R., Case, A. C., and Berke, J. D. (2012). Basal ganglia beta oscillations accompany cue utilization. *Neuron*, 73(3):523–536.
- Leão, A. A. (1944). Spreading depression of activity in the cerebral cortex. *Journal of neurophysiology*, 7(6):359–390.
- Li, D., Wang, Q., and Shen, Y. (2015). Intelligent greedy pursuit model for sparse reconstruction based on l_0 minimization. *Signal Processing*.
- Lim, S. and Goldman, M. S. (2013). Balanced cortical microcircuitry for maintaining information in working memory. *Nature neuroscience*, 16(9):1306–1314.
- Lim, S. and Goldman, M. S. (2014). Balanced cortical microcircuitry for spatial working memory based on corrective feedback control. *The Journal of Neuroscience*, 34(20):6790–6806.
- Litwin-Kumar, A. and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience*, 15(11):1498–1505.
- Lu, Y., Truccolo, W., Wagner, F. B., Vargas-Irwin, C. E., Ozden, I., Zimmermann, J. B., May, T., Agha, N., Wang, J., Nurmikko, A. V., et al. (2015). Optogenetically-induced spatiotemporal gamma oscillations and neuronal spiking activity in primate motor cortex. *Journal of Neurophysiology*, pages jn-00792.
- Lubenov, E. V. and Siapas, A. G. (2009). Hippocampal theta oscillations are travelling waves. *Nature*, 459(7246):534–539.
- MacLeod, K., Bäcker, A., and Laurent, G. (1998). Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature*, 395(6703):693–698.
- Mallet, N., Pogosyan, A., Sharott, A., Csicsvari, J., Bolam, J. P., Brown, P., and Magill, P. J. (2008). Disrupted dopamine transmission and the emergence of exaggerated beta oscillations in subthalamic nucleus and cerebral cortex. *The Journal of neuroscience*, 28(18):4795–4806.

- Mammen, E., Nandi, S., Maiwald, T., and Timmer, J. (2009). Effect of jump discontinuity for phase-randomized surrogate data testing. *International Journal of Bifurcation and Chaos*, 19(01):403–408.
- Maris, E., Fries, P., and van Ede, F. (2016). Diverse phase relations among neuronal rhythms and their potential function. *Trends in Neurosciences*.
- Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*, 5(10):793–807.
- Maynard, E., Hatsopoulos, N., Ojakangas, C., Acuna, B., Sanes, J., Normann, R. A., and Donoghue, J. (1999). Neuronal interactions improve cortical population coding of movement direction. *The Journal of Neuroscience*, 19(18):8083–8093.
- Mazurek, M. E. and Shadlen, M. N. (2002). Limits to the temporal fidelity of cortical spike rate signals. *Nature Neuroscience*, 5(5):463–471.
- McComas, A. J. (1991). Invited review: motor unit estimation: methods, results, and present status. *Muscle & nerve*, 14(7):585–597.
- McCormick, D. A., Connors, B. W., Lighthall, J. W., and Prince, D. A. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of neurophysiology*, 54(4):782–806.
- McIntyre, C. C. and Hahn, P. J. (2010). Network perspectives on the mechanisms of deep brain stimulation. *Neurobiology of disease*, 38(3):329–337.
- Mehanna, R. and Lai, E. C. (2013). Deep brain stimulation in parkinson’s disease. *Translational neurodegeneration*, 2(1):1.
- Menzer, D. L., Rao, N. G., Bondy, A., Truccolo, W., and Donoghue, J. P. (2014). Population interactions between parietal and primary motor cortices during reach. *Journal of neurophysiology*, 112(11):2959–2984.
- Mitra, P. and Bokil, H. (2007). *Observed brain dynamics*. Oxford University Press.
- Montemurro, M. A., Rasch, M. J., Murayama, Y., Logothetis, N. K., and Panzeri, S. (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Current biology*, 18(5):375–380.
- Moser, E. I., Kropff, E., and Moser, M.-B. (2008). Place cells, grid cells, and the brain’s spatial representation system. *Annu. Rev. Neurosci.*, 31:69–89.
- Murthy, V. N. and Fetz, E. E. (1992). Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proceedings of the National Academy of Sciences*, 89(12):5670–5674.

- Murthy, V. N. and Fetz, E. E. (1996a). Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. *Journal of Neurophysiology*, 76(6):3949–3967.
- Murthy, V. N. and Fetz, E. E. (1996b). Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *Journal of Neurophysiology*, 76(6):3968–3982.
- Nain, A. S., Chung, F., Rule, M., Jadlovec, J., Campbell, P. G., Amon, C., Sitti, M., et al. (2007). Microrobotically fabricated biological scaffolds for tissue engineering. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 1918–1923. IEEE.
- Neuper, C. and Pfurtscheller, G. (2001). Event-related dynamics of cortical rhythms: frequency-specific features and functional correlates. *International journal of psychophysiology*, 43(1):41–58.
- Niedermeyer, E. and da Silva, F. L. (2005). *Electroencephalography: basic principles, clinical applications, and related fields*. Lippincott Williams & Wilkins.
- Nunez, P. L. and Srinivasan, R. (2006). *Electric fields of the brain: the neurophysics of EEG*. Oxford university press.
- Ohara, S., Ikeda, A., Kunieda, T., Yazawa, S., Baba, K., Nagamine, T., Taki, W., Hashimoto, N., Mihara, T., and Shibasaki, H. (2000). Movement-related change of electrocorticographic activity in human supplementary motor area proper. *Brain*, 123(6):1203–1215.
- O’Keefe, J. and Recce, M. L. (1993). Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus*, 3(3):317–330.
- Okun, M., Steinmetz, N. A., Cossell, L., Iacaruso, M. F., Ko, H., Barthó, P., Moore, T., Hofer, S. B., Mrsic-Flogel, T. D., Carandini, M., et al. (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature*, 521(7553):511–515.
- Oliphant, T. E. (2007). Python for scientific computing. *Computing in Science & Engineering*, 9(3):10–20.
- Pasquereau, B. and Turner, R. S. (2013). Primary motor cortex of the parkinsonian monkey: altered neuronal responses to muscle stretch. *Frontiers in Systems Neuroscience*, 7.
- Percival, D. B. and Walden, A. T. (1993). *Spectral analysis for physical applications*. Cambridge University Press.
- Perge, J. A., Zhang, S., Malik, W. Q., Homer, M. L., Cash, S., Friehs, G., Eskandar, E. N., Donoghue, J. P., and Hochberg, L. R. (2014). Reliability of directional information in unsorted spikes and local field potentials recorded in human motor cortex. *Journal of neural engineering*, 11(4):046007.
- Pfurtscheller, G., Neuper, C., Brunner, C., and da Silva, F. L. (2005). Beta rebound after different types of motor imagery in man. *Neuroscience letters*, 378(3):156–159.

- Pfurtscheller, G., Neuper, C., Pichler-Zalaudek, K., Edlinger, G., and da Silva, F. H. L. (2000). Do brain oscillations of different frequencies indicate interaction between cortical areas in humans? *Neuroscience letters*, 286(1):66–68.
- Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E., and Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207):995–999.
- Pogosyan, A., Gaynor, L. D., Eusebio, A., and Brown, P. (2009). Boosting cortical activity at beta-band frequencies slows movement in humans. *Current Biology*, 19(19):1637–1641.
- Popovych, O. V. and Tass, P. A. (2011). Macroscopic entrainment of periodically forced oscillatory ensembles. *Progress in biophysics and molecular biology*, 105(1):98–108.
- Prechtl, J., Cohen, L., Pesaran, B., Mitra, P., and Kleinfeld, D. (1997). Visual stimuli induce waves of electrical activity in turtle cortex. *Proceedings of the National Academy of Sciences*, 94(14):7621–7626.
- Purkyně, J. E. (1819). *Beiträge zur Kenntnis des Sehens in subjectiver Hinsicht*. Vetterl.
- Quinn, C. J., Coleman, T. P., Kiyavash, N., and Hatsopoulos, N. G. (2011). Estimating the directed information to infer causal relationships in ensemble neural spike train recordings. *Journal of computational neuroscience*, 30(1):17–44.
- Quiroga, R. Q. and Panzeri, S. (2009). Extracting information from neuronal populations: information theory and decoding approaches. *Nature Reviews Neuroscience*, 10(3):173–185.
- Rasch, M., Logothetis, N. K., and Kreiman, G. (2009). From neurons to circuits: linear estimation of local field potentials. *The Journal of Neuroscience*, 29(44):13785–13796.
- Rasch, M. J., Gretton, A., Murayama, Y., Maass, W., and Logothetis, N. K. (2008). Inferring spike trains from local field potentials. *Journal of Neurophysiology*, 99(3):1461–1476.
- Ray, S. and Maunsell, J. H. (2010). Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron*, 67(5):885–896.
- Reimann, M. W., Anastassiou, C. A., Perin, R., Hill, S. L., Markram, H., and Koch, C. (2013). A biophysically detailed model of neocortical local field potentials predicts the critical role of active membrane currents. *Neuron*, 79(2):375–390.
- Reimer, J. and Hatsopoulos, N. G. (2010). Periodicity and evoked responses in motor cortex. *The Journal of Neuroscience*, 30:11506–11515.
- Rheinberger, M. B. and Jasper, H. H. (1937). Electrical activity of the cerebral cortex in the unanesthetized cat. *American Journal of Physiology–Legacy Content*, 119(1):186–196.
- Riddle, C. N. and Baker, S. N. (2006). Digit displacement, not object compliance, underlies task dependent modulations in human corticomuscular coherence. *Neuroimage*, 33(2):618–627.

- Rizzolatti, G. and Craighero, L. (2004). The mirror-neuron system. *Annu. Rev. Neurosci.*, 27:169–192.
- Rizzolatti, G., Fadiga, L., Gallese, V., and Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive brain research*, 3(2):131–141.
- Rizzolatti, G., Luppino, G., and Matelli, M. (1998). The organization of the cortical motor system: new concepts. *Electroencephalography and clinical neurophysiology*, 106(4):283–296.
- Roopun, A. K., Kramer, M. A., Carracedo, L. M., Kaiser, M., Davies, C. H., Traub, R. D., Kopell, N. J., and Whittington, M. A. (2008). Period concatenation underlies interactions between gamma and beta rhythms in neocortex. *Frontiers in cellular neuroscience*, 2.
- Roopun, A. K., Middleton, S. J., Cunningham, M. O., LeBeau, F. E., Bibbig, A., Whittington, M. A., and Traub, R. D. (2006). A beta2-frequency (20–30 Hz) oscillation in nonsynaptic networks of somatosensory cortex. *Proceedings of the National Academy of Sciences*, 103(42):15646–15650.
- Rubino, D., Robbins, K. A., and Hatsopoulos, N. G. (2006). Propagating waves mediate information transfer in the motor cortex. *Nature Neuroscience*, 9(12):1549–1557.
- Rule, M., Stoffregen, M., and Ermentrout, B. (2011). A model for the origin and properties of flicker-induced geometric phosphenes. *PLoS Comput. Biol.*, 7(9):e1002.
- Rule, M. E., Vargas-Irwin, C., Donoghue, J. P., and Truccolo, W. (2015). Contribution of LFP dynamics to single neuron spiking variability in motor cortex during movement execution. *Frontiers in Systems Neuroscience*, 9:89.
- Sacchet, M. D., LaPlante, R. A., Wan, Q., Pritchett, D. L., Lee, A. K., Hämäläinen, M., Moore, C. I., Kerr, C. E., and Jones, S. R. (2015). Attention drives synchronization of alpha and beta rhythms between right inferior frontal and primary sensory neocortex. *The Journal of Neuroscience*, 35(5):2074–2082.
- Saleh, M., Reimer, J., Penn, R., Ojakangas, C. L., and Hatsopoulos, N. G. (2010). Fast and slow oscillations in human primary motor cortex predict oncoming behaviorally relevant cues. *Neuron*, 65(4):461–471.
- Salenius, S. and Hari, R. (2003). Synchronous cortical oscillatory activity during motor action. *Current opinion in neurobiology*, 13(6):678–684.
- Sanes, J. N. and Donoghue, J. P. (1993). Oscillations in local field potentials of the primate motor cortex during voluntary movement. *Proceedings of the National Academy of Sciences*, 90(10):4470–4474.
- Schaefer, A. T., Angelo, K., Spors, H., and Margrie, T. W. (2006). Neuronal oscillations enhance stimulus discrimination by ensuring action potential precision.

- Scheffer-Teixeira, R., Belchior, H., Leão, R. N., Ribeiro, S., and Tort, A. B. (2013). On high-frequency field oscillations (> 100 Hz) and the spectral leakage of spiking activity. *The Journal of Neuroscience*, 33(4):1535–1539.
- Schiff, S. J., Huang, X., and Wu, J.-Y. (2007). Dynamical evolution of spatiotemporal patterns in mammalian middle cortex. *BMC Neuroscience*, 8(Suppl 2):P61.
- Schnitzler, A. and Gross, J. (2005). Normal and pathological oscillatory communication in the brain. *Nature reviews neuroscience*, 6(4):285–296.
- Schwartz, A. B., Taylor, D. M., and Tillery, S. I. H. (2001). Extraction algorithms for cortical control of arm prosthetics. *Current opinion in Neurobiology*, 11(6):701–708.
- Seung, H. S. (1996). How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences*, 93(23):13339–13344.
- Shadlen, M. N. and Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *The Journal of Neuroscience*, 18(10):3870–3896.
- Shenoy, K. V., Kaufman, M. T., Sahani, M., and Churchland, M. M. (2011). A dynamical systems view of motor preparation: implications for neural prosthetic system design. *Progress in brain research*, 192:33.
- Shenoy, K. V., Sahani, M., and Churchland, M. M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annual review of neuroscience*, 36:337–359.
- Shipp, S., Blanton, M., and Zeki, S. (1998). A visuo-somatomotor pathway through superior parietal cortex in the macaque monkey: cortical connections of areas v6 and v6a. *European Journal of Neuroscience*, 10(10):3171–3193.
- Shoham, S., Paninski, L. M., Fellows, M. R., Hatsopoulos, N. G., Donoghue, J. P., Normann, R., et al. (2005). Statistical encoding model for a primary motor cortical brain-machine interface. *Biomedical Engineering, IEEE Transactions on*, 52(7):1312–1322.
- Siapas, A. G., Lubenov, E. V., and Wilson, M. A. (2005). Prefrontal phase locking to hippocampal theta oscillations. *Neuron*, 46(1):141–151.
- Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annual review of physiology*, 55(1):349–374.
- Singer, W. and Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual review of neuroscience*, 18(1):555–586.
- Skaggs, W., McNaughton, B., Wilson, M., and Barnes, C. (1995). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6(2):149–172.
- Stein, E. and Bar-Gad, I. (2013). Beta oscillations in the cortico-basal ganglia loop during parkinsonism. *Experimental neurology*, 245:52–59.

- Stevens, C. F. and Zador, A. M. (1998). Input synchrony and the irregular firing of cortical neurons. *Nature neuroscience*, 1(3):210–217.
- Takahashi, K. (1965). Slow and fast groups of pyramidal tract cells and their respective membrane properties. *Journal of Neurophysiology*, 28(5):908–924.
- Takahashi, K., Kim, S., Coleman, T. P., Brown, K. A., Suminski, A. J., Best, M. D., and Hatsopoulos, N. G. (2015). Large-scale spatiotemporal spike patterning consistent with wave propagation in motor cortex. *Nature communications*, 6.
- Takahashi, K., Saleh, M., Penn, R. D., and Hatsopoulos, N. (2011). Propagating waves in human motor cortex. *Frontiers in Human Neuroscience*, 5:40.
- Tanné-Gariépy, J., Rouiller, E. M., and Boussaoud, D. (2002). Parietal inputs to dorsal versus ventral premotor areas in the macaque monkey: evidence for largely segregated visuomotor pathways. *Experimental Brain Research*, 145(1):91–103.
- Torrence, C. and Compo, G. P. (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological society*, 79(1):61–78.
- Townsend, R. G., Solomon, S. S., Chen, S. C., Pietersen, A. N., Martin, P. R., Solomon, S. G., and Gong, P. (2015). Emergence of complex wave patterns in primate cerebral cortex. *The Journal of Neuroscience*, 35(11):4657–4662.
- Truccolo, W., Ahmed, O. J., Harrison, M. T., Eskandar, E. N., Cosgrove, G. R., Madsen, J. R., Blum, A. S., Potter, N. S., Hochberg, L. R., and Cash, S. S. (2014). Neuronal ensemble synchrony during human focal seizures. *The Journal of Neuroscience*, 34(30):9927–9944.
- Truccolo, W., Donoghue, J. A., Hochberg, L. R., Eskandar, E. N., Madsen, J. R., Anderson, W. S., Brown, E. N., Halgren, E., and Cash, S. S. (2011). Single-neuron dynamics in human focal epilepsy. *Nature Neuroscience*, 14(5):635–641.
- Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., and Brown, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*, 93(2):1074–1089.
- Truccolo, W., Hochberg, L. R., and Donoghue, J. P. (2010). Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nature Neuroscience*, 13(1):105–111.
- Truccolo, W. A., Ding, M., Knuth, K. H., Nakamura, R., and Bressler, S. L. (2002). Trial-to-trial variability of cortical evoked responses: implications for the analysis of functional connectivity. *Clinical Neurophysiology*, 113(2):206–226.
- Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 237(641):37–72.

- Van Aerde, K. I., Mann, E. O., Canto, C. B., Heistek, T. S., Linkenkaer-Hansen, K., Mulder, A. B., Van Der Roest, M., Paulsen, O., Brussaard, A. B., and Mansvelder, H. D. (2009). Flexible spike timing of layer 5 neurons during dynamic beta oscillation shifts in rat prefrontal cortex. *The Journal of physiology*, 587(21):5177–5196.
- van Ede, F., Jensen, O., and Maris, E. (2010). Tactile expectation modulates pre-stimulus β -band oscillations in human sensorimotor cortex. *Neuroimage*, 51(2):867–876.
- van Steveninck, R. R. d. R., Lewen, G. D., Strong, S. P., Koberle, R., and Bialek, W. (1997). Reproducibility and variability in neural spike trains. *Science*, 275(5307):1805–1808.
- van Wijk, B. C., Beek, P. J., and Daffertshofer, A. (2012). Neural synchrony within the motor system: what have we learned so far? *Frontiers in human neuroscience*, 6.
- van Wijk, B. C. M., Daffertshofer, A., Roach, N., and Praamstra, P. (2009). A role of beta oscillatory synchrony in biasing response competition? *Cerebral Cortex*, 19:1294–1302.
- Varela, F., Lachaux, J.-P., Rodriguez, E., and Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nature reviews neuroscience*, 2(4):229–239.
- Vargas-Irwin, C. and Donoghue, J. P. (2007). Automated spike sorting using density grid contour clustering and subtractive waveform decomposition. *Journal of Neuroscience Methods*, 164(1):1–18.
- Vargas-Irwin, C. E., Franquemont, L., Black, M. J., and Donoghue, J. P. (2015). Linking objects to actions: Encoding of target object and grasping strategy in primate ventral premotor cortex. *The Journal of Neuroscience*, 35(30):10888–10897.
- Vargas-Irwin, C. E., Shakhnarovich, G., Yadollahpour, P., Mislow, J. M., Black, M. J., and Donoghue, J. P. (2010). Decoding complete reach and grasp actions from local primary motor cortex populations. *The Journal of Neuroscience*, 30(29):9659–9669.
- Vigneswaran, G., Kraskov, A., and Lemon, R. N. (2011). Large identified pyramidal cells in macaque motor and premotor cortex exhibit “thin spikes”: implications for cell type classification. *The Journal of Neuroscience*, 31(40):14235–14242.
- Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P., and Pennartz, C. M. (2010). The pairwise phase consistency: a bias-free measure of rhythmic neuronal synchronization. *Neuroimage*, 51(1):112–122.
- Vogels, T. P., Rajan, K., and Abbott, L. (2005). Neural network dynamics. *Annu. Rev. Neurosci.*, 28:357–376.
- Waldert, S., Lemon, R. N., and Kraskov, A. (2013). Influence of spiking activity on cortical local field potentials. *The Journal of Physiology*, 591(21):5291–5303.
- Wang, X.-J. (2003). Persistent neural activity: experiments and theory. *Cerebral Cortex*, 13(11):1123–1123.

- Weiler, N., Wood, L., Yu, J., Solla, S. A., and Shepherd, G. M. (2008). Top-down laminar organization of the excitatory network in motor cortex. *Nature neuroscience*, 11(3):360–366.
- Weinberger, M., Hutchison, W. D., and Dostrovsky, J. O. (2009). Pathological subthalamic nucleus oscillations in pd: can they be the cause of bradykinesia and akinesia? *Experimental neurology*, 219(1):58–61.
- Wetmore, D. Z. and Baker, S. N. (2004). Post-spike distance-to-threshold trajectories of neurones in monkey motor cortex. *The Journal of Physiology*, 555(3):831–850.
- Wimmer, K., Nykamp, D. Q., Constantinidis, C., and Compte, A. (2014). Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nature neuroscience*, 17(3):431–439.
- Witham, C. L. and Baker, S. N. (2007). Network oscillations and intrinsic spiking rhythmicity do not covary in monkey sensorimotor areas. *The Journal of Physiology*, 580(3):801–814.
- Witham, C. L., Riddle, C. N., Baker, M. R., and Baker, S. N. (2011). Contributions of descending and ascending pathways to corticomuscular coherence in humans. *The Journal of Physiology*, 589(15):3789–3800.
- Witte, M., Patino, L., Andrykiewicz, A., Hepp-Reymond, M.-C., and Kristeva, R. (2007). Modulation of human corticomuscular beta-range coherence with low-level static forces. *European Journal of Neuroscience*, 26(12):3564–3570.
- Woloszyn, L. and Sheinberg, D. L. (2012). Effects of long-term visual experience on responses of distinct classes of single units in inferior temporal cortex. *Neuron*, 74(1):193–205.
- Wróbel, A. (2000). Beta activity: a carrier for visual attention. *Acta neurobiologiae experimentalis*, 60(2):247–260.
- Wu, J.-Y., Huang, X., and Zhang, C. (2008a). Propagating waves of activity in the neocortex: what they are, what they do. *The Neuroscientist*, 14(5):487–502.
- Wu, S., Hamaguchi, K., and Amari, S.-i. (2008b). Dynamics and computation of continuous attractors. *Neural computation*, 20(4):994–1025.
- Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., and Black, M. J. (2006). Bayesian population decoding of motor cortical activity using a kalman filter. *Neural computation*, 18(1):80–118.
- Wörgötter, F., Suder, K., Zhao, Y., Kerscher, N., Eysel, U. T., and Funke, K. (1998). State-dependent receptive-field restructuring in the visual cortex. *Nature*, 396(6707):165–168.
- Xu, W., Huang, X., Takagaki, K., and Wu, J.-y. (2007). Compression and reflection of visually evoked cortical waves. *Neuron*, 55(1):119–129.

- Yamamoto, Y. and Hughson, R. L. (1993). Extracting fractal components from time series. *Physica D: Nonlinear Phenomena*, 68(2):250–264.
- Yamawaki, N., Borges, K., Suter, B. A., Harris, K. D., and Shepherd, G. M. (2015). A genuine layer 4 in motor cortex with prototypical synaptic circuit connectivity. *Elife*, 3:e05422.
- Yang, A. I., Vanegas, N., Lungu, C., and Zaghoul, K. A. (2014). Beta-coupled high-frequency activity and beta-locked neuronal spiking in the subthalamic nucleus of Parkinson’s disease. *The Journal of Neuroscience*, 34(38):12816–12827.
- Zanos, T. P., Mineault, P. J., and Pack, C. C. (2011). Removal of spurious correlations between spikes and local field potentials. *Journal of Neurophysiology*, 105(1):474–486.
- Zhou, P., Burton, S. D., Snyder, A. C., Smith, M. A., Urban, N. N., and Kass, R. E. (2015). Establishing a statistical link between network oscillations and neural synchrony. *PLoS Computational Biology*, 11(10):e1004549.
- Zhuang, J., Truccolo, W., Vargas-Irwin, C., and Donoghue, J. P. (2010). Decoding 3-D reach and grasp kinematics from high-frequency local field potentials in primate primary motor cortex. *Biomedical Engineering, IEEE Transactions on*, 57(7):1774–1784.