

Causal and distributed mechanisms of cortical computation

Anna J. Li

A dissertation

submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Nicholas A. Steinmetz, Chair

Eric T. Shea-Brown

Fred Rieke

Program Authorized to Offer Degree:

Neuroscience

© Copyright 2026

Anna J. Li

University of Washington

Abstract

Causal and distributed mechanisms of cortical computation

Anna J. Li

Chair of the Supervisory Committee:

Nicholas A. Steinmetz

Department of Neurobiology and Biophysics

Perception and behavior emerge from the coordinated activity of neurons within local circuits and across distributed brain regions. This view is supported by extensive recurrent and long-range connectivity, as well as large-scale recordings of brain-wide activity. Open questions include which cortical interactions arise from direct causal relationships, and the relative magnitude of shared versus private activity in sensory areas. In this dissertation, we address both questions using simultaneous mesoscale calcium imaging and cortex-wide optogenetic manipulation, high-density electrophysiology, and statistical modeling of population activity and spike counts. In Chapter 1, we review the evidence that cortical activity is distributed across areas and the supporting anatomical connectivity, and then discuss methodological advances for circuit manipulation. In sensory areas, movement, arousal, and other cognitive signals have been proposed as candidate sources of neural variability in sensory responses. Next, we introduce a system for simultaneous imaging and optogenetic manipulation across dorsal cortex (Chapter 2). We show that direct input did not account for homotopic correlations, both at the level of single-neurons and population coupling. In Chapter 3, we developed a Poisson reduced-rank regression framework to partition trial-to-trial variability in primary visual cortex (V1) spike counts, and found that a low-dimensional pattern of activity shared among V1 neurons accounted for most of their variability beyond Poisson. Cortex-wide activity, measured with simultaneous electrophysiology or mesoscale

imaging, predicted more than half of this shared variability and subsumed the variability predicted by behavior. Variability in V1 therefore substantially reflects activity shared with the rest of dorsal cortex rather than fluctuations private to V1. Finally, in Chapter 4, we discuss the implications of these findings for distributed neural computation, the limitations of measuring global signals in dorsal cortex alone, and future directions for dissecting causal pathways along which cortical activity is shared.

Table of Contents

Acknowledgments.....	4
Author contributions.....	6
Chapter 1: Introduction.....	7
1.1 Distributed cortical activity.....	7
1.1.1 From localization to distributed function.....	7
1.1.2 Evidence from large-scale recordings.....	8
1.1.3 Anatomical substrates of distributed activity.....	9
1.2 Correlational vs. causal approaches to cortical interaction.....	11
1.3 Trial-to-trial variability in sensory cortex.....	12
1.3.1 From noise to structured fluctuations.....	12
1.3.2 Approaches to partitioning variability.....	13
1.4 Main objectives.....	14
1.5 General methods.....	15
1.5.1 Subjects.....	15
1.5.2 Surgeries.....	15
1.5.3 Electrophysiology and spike sorting.....	16
1.5.4 Widefield calcium imaging and optogenetics.....	17
Chapter 2: Cortex-wide measurement and manipulation reveals causal organization of cortical dynamics..	18
2.1 Introduction.....	18
2.2 Methods.....	21
2.2.1 Microscope design.....	21
2.2.2 Genetic strategy.....	22

2.2.3 Optogenetic experiments.....	22
2.2.4 Analysis.....	23
2.3 Results: System characterization.....	25
2.3.1 Validation with simultaneous imaging, electrophysiology, and optogenetic stimulation.....	26
2.3.2 Long-term, high-resolution optogenetic manipulation with flexible spatial control.....	29
2.3.3 Compatibility with subcortical electrophysiology.....	31
2.4 Results: Causal vs. correlational connectivity across scales.....	33
2.5 Preliminary results: Intrinsic timescales of cortical areas.....	36
2.6 Discussion.....	37
2.6.1 Summary of findings.....	37
2.6.2 System characteristics and limitations.....	38
2.6.3 Assumptions and interpretations of causal coupling via optogenetics.....	39
2.6.4 Hierarchical timescales in the cortex.....	40
Chapter 3: Global origins of trial-to-trial variability in V1.....	41
3.1 Introduction.....	41
3.2 Methods.....	42
3.2.1 Passive stimulus presentation.....	42
3.2.2 Behavioral monitoring.....	43
3.2.3 Widefield imaging and optogenetics.....	43
3.2.4 Poisson reduced-rank regression model.....	43
3.2.5 Quantifying explained trial-to-trial variability.....	47
3.2.6 Variance partitioning.....	49
3.2.7 Manifold geometry analysis.....	49
3.2.8 Decoding analysis.....	51
3.2.9 HMM analysis.....	51

3.3 Results.....	53
3.4 Discussion.....	67
3.4.1 Summary of findings.....	67
3.4.2 Locally shared activity dominates super-Poisson V1 variability.....	67
3.4.3 Globally distributed activity accounts for the majority of locally shared variability.....	69
3.4.4 Model assumptions and limitations.....	69
Chapter 4: Discussion.....	71
4.1 Summary of findings.....	71
4.2 Implications for understanding distributed cortical dynamics.....	71
4.3 Correlational versus causal measures of cortical dynamics.....	72
4.4 Limitations.....	73
4.5 Future directions.....	73
Supplementary Figures.....	75
Chapter 2 Supplementary Figures.....	75
Chapter 3 Supplementary Figures.....	80
References.....	90

Acknowledgments

When I started graduate school in 2020, I remember thinking that the end felt eons away. Now, having defended, with the end well in sight, it's felt both like a lifetime and like no time at all. Certainly a lot of time is spent performing experiments, conducting analyses, writing, and more, but much of the change happens in your brain and often feels imperceptible. I'm immensely grateful for the privilege to have spent six years growing as a thinker and scientist.

Thank you to my advisor, Nick Steinmetz, for his support and mentorship over the years. Nick is one of the most thorough, "let's do this the right way" people I know, and inspires me to bring that principle into all aspects of my own life. In addition to thinking about the development of his trainees, he is always thinking about how to improve himself as a mentor and as a scientist. I'm glad to have spent such a formative time in my life in the Steinmetz lab.

The people in a lab reflect it as much as the PI, and in this regard I've been very fortunate as well. Thank you to Zhiwen Ye, Dan Birman, and Anna Bowen. Their guidance was instrumental especially early on, and still continues as they've moved on in their careers. Zhiwen once voluntold me to ski down a steep (to me) slope, saying, "I know you can do it." He was right, and often encouraged me this way in research as well. Kim Miller, who keeps the lab running as a team, inspires me with her curiosity and endless knowledge of all things. Pascha Matveev, now a Ph.D. student, made foundational contributions to much of this work as an undergraduate student. Finally, thank you to every single member of the lab who has made it such a fun and thoughtful place to do research.

Eric Shea-Brown and Ziyu Lu have been there since nearly the very first day of my Ph.D. Ziyu is incredibly rigorous as a scientist and mathematician, and always excited to understand more about the biology of the systems she's analyzing. Eric is always ready with an insightful question and a kind comment. My supervisory committee, which includes Eric as well as Fred Rieke and Bing Brunton, has constantly provided perceptive feedback and steady encouragement. Committee meetings can often feel like a scary evaluation, but I left every single one of mine with a renewed excitement about what I was doing next.

Outside of the lab, I'm enormously grateful to my friends and family in Seattle, across the country, and across the world. My fiancé, Kevin Ho, moved to Seattle with me when I started at UW, and has been my unwavering supporter throughout graduate school. Kevin ended up in the ECE Ph.D. program at UW himself, and also happened to do a lot of the Steinmetz lab's electrical engineering. He and our feline daughter, Rio, have kept me sane and grounded. I'm so lucky to have such a kind and brilliant person by my side.

I'd like to thank my parents, Sarah Shang and Wei Li, who didn't even blink an eye when I said I wanted to go back to school for six years. Thank you also to my bonus parents - John Ferraro, Yin-Fun Lum, and Kian Ho, who always reminded me how proud they are of me, Ph.D. or otherwise. My baby sister, Amanda, completed her undergraduate degree at UW while I was here, and is the greatest sister and friend. Love you all.

Author contributions

All chapters were written by A. Li, and describe research designed and conducted by her, with the following exceptions:

Chapter 2 describes research designed and written in collaboration with N. Steinmetz and Pascha Matveev. The optical system was designed by N. Steinmetz. The genetic strategy was designed by N. Steinmetz, in consultation with Jonathan Ting, with additional contributions from Ximena Opitz-Araya. Some data collection and preliminary analysis was performed by P. Matveev and Alexander E. Ladd.

Chapter 3 describes research designed and written in collaboration with N. Steinmetz, E. Shea-Brown, and Ziyu Lu. The modeling framework was developed by Z. Lu, in collaboration with A. Li, N. Steinmetz, and E. Shea-Brown. A subset of the optogenetic experiments were performed by A. Ladd and P. Matveev. Z. Lu conducted the geometric manifold analysis.

Chapters 1 and 4 describe ideas, interpretations, and future directions formulated in collaboration with N. Steinmetz.

Chapter 1: Introduction

Perception and action depend on computations from sensory, motor, and cognitive brain regions. Sensory input is transformed into representations of stimulus features, combined with internal state and prior experience, and read out to inform motor commands. This modular feedforward description, however, may not reflect how these computations are performed in the brain. For example, internal state may modulate the processing of sensory input even at the earliest stages (Livingstone and Hubel, 1981). Sensory representations are therefore shaped by more than the feedforward transformation of stimulus features.

Vision is a critical sensory modality for many organisms, playing a vital role in behaviors needed for survival like navigation, foraging, and social interaction. However, individual visual neurons encode stimuli with apparently low fidelity; their responses to identical stimuli vary from trial to trial, in contrast to the relative reliability of visually guided behavior. A substantial body of work has found this variability to be structured rather than random, reflecting processes that contribute to computation and are not simply a limit on signaling. Responses in visual cortex therefore reflect influences beyond the encoding of the stimulus. This variability may reflect the demands of natural behavior, in which the same stimulus can require different actions depending on context.

1.1 Distributed cortical activity

1.1.1 From localization to distributed function

The view of cortex as a collection of specialized regions has its origins in nineteenth-century neurology. Cytoarchitectural mapping of anatomically distinct regions established cortex as a collection of structurally and functionally specialized areas (Brodmann, 2007). Early single-unit recordings also suggested specialization at the neuronal level in visual cortex, somatosensory cortex, and auditory cortex (Hubel and Wiesel, 1962; Merzenich et al., 1975; Mountcastle, 1957). This view positioned cortical regions as relatively autonomous computational units, each performing a designated function within a specialized circuit.

Alongside this work, lesion studies pointed out a more distributed organization. The concept of diaschisis held that focal damage produces functional depression not only at the lesion site but at distant, anatomically connected regions, implying that a region's normal function depends on input from elsewhere (von Monakow, 1914). Lesion studies in rodents found that behavioral deficits scaled with the volume of removed cortex rather than the location, leading to the proposal that some functions are distributed across cortex rather than localized (Lashley, 1929). The neural activity underlying these phenomena, however, was difficult to determine empirically given the methodological limitations of the time.

Anatomical work through the late twentieth century revealed the extent of inter-areal connectivity. Mapping of primate visual cortex showed a hierarchy of specialized areas linked by extensive reciprocal connections, with feedback projections as prevalent as feedforward ones (Felleman and Van Essen, 1991). Specialization and interconnection were both features of cortical organization, but understanding how neural circuits integrated these properties remained an open question.

1.1.2 Evidence from large-scale recordings

Historically, cortical recordings were limited to small numbers of neurons in single regions (Siegle and Steinmetz, 2026). Recordings yielded a fraction of the population of any single area, providing little information about activity across broader circuits and networks. Inferring function from such a limited view is fundamentally constrained. Recent technologies have substantially expanded the scale of brain recording. High-density silicon probes such as Neuropixels record from hundreds of neurons along electrodes spanning multiple regions (Jun et al., 2017; Steinmetz et al., 2021), and combining multiple probes enables simultaneous recording from many distant regions in the same subject (Steinmetz et al., 2019). Such recordings have revealed that task-related signals are broadly distributed across cortex and subcortex, well beyond regions traditionally associated with their computations (Steinmetz et al., 2019). A substantial portion of this widespread activity reflects movement and behavioral state, though behavior does not account for all of it (Manley et al., 2024; Musall et al., 2019; Stringer et al., 2019b).

A complementary approach measures activity across the cortical surface rather than within a single site. Optical imaging of cortex has a long history, beginning with voltage- and calcium-sensitive

dyes (Grinvald et al., 1986; Orbach et al., 1985), but its routine use in mice followed the development of transgenic lines expressing genetically-encoded calcium indicators throughout cortex (Vanni and Murphy, 2014; Wekselblatt et al., 2016). Single-photon “widefield” imaging through the cleared or thinned skull captures activity across most of dorsal cortex simultaneously. The signal reflects bulk fluorescence from populations of neurons rather than individual cells, and its temporal resolution is limited by indicator kinetics. Therefore, widefield trades cellular and spike-level resolution for spatial breadth. This breadth has made it possible to characterize how activity is organized across cortex as a whole (MacDowell et al., 2025; Musall et al., 2019; Ye et al., 2026).

1.1.3 Anatomical substrates of distributed activity

The distributed nature of cortical activity is supported by extensive structural connectivity. Anatomical tracing has shown that cortical areas are interconnected through long-range projections, forming a densely interconnected network rather than a set of isolated processing units (Felleman and Van Essen, 1991; Markov et al., 2014). In the mouse, large-scale connectivity mapping has characterized projections among virtually all cortical areas, revealing that most areas project to and receive input from many others (Gămănuț et al., 2018; Harris et al., 2019; Oh et al., 2014; Zingg et al., 2014). Many projections are organized hierarchically, with feedforward and feedback pathways at different levels of processing, but the density of this connectivity suggests that influence is not restricted to hierarchically adjacent areas.

A prominent feature of cortical connectivity is the corpus callosum, which links the two hemispheres. Callosal projections are predominantly homotopic, connecting each cortical area most strongly to its counterpart in the opposite hemisphere (Oh et al., 2014). This anatomical arrangement has a clear functional correspondence; activity in homotopic regions is strongly correlated across hemispheres, a relationship observed across sensory, motor, and association cortices (Mohajerani et al., 2010; Shimaoka et al., 2019). Whether the bilateral functional coupling observed throughout cortex reflects direct callosal communication, shared input from subcortical or ipsilateral sources, or some combination has not been directly tested.

Cortical regions are also linked indirectly through the thalamus. Thalamic nuclei have traditionally been divided into first-order relays, which convey peripheral sensory input to primary cortical areas, and higher-order nuclei, which receive inputs from cortex and project back to other cortical regions (Sherman, 2016). Higher-order nuclei such as the pulvinar in primates and the lateral posterior nucleus (LP) in mice are positioned to relay activity between cortical areas, providing a subcortical route for inter-areal communication alongside direct cortico-cortical projections (Bennett et al., 2019; Saalman et al., 2012). Cortical regions also converge in the striatum, with patterns of convergence defining functional subdivisions, and where basal ganglia output returns to cortex by way of thalamus (Hunnicutt et al., 2016). Therefore, inter-area interactions need not travel by direct cortico-cortical projections alone, and the route underlying functional relationships between two regions may not be obvious from anatomical connectivity maps.

Among the cortical inputs to visual cortex are direct projections from frontal areas. The anterior cingulate area (ACA) projects to V1, and optogenetic activation of these neurons enhances V1 responses and discrimination performance (Zhang et al., 2014). This circuit appears to be specific to vision: circuit mapping has shown that while visual cortex is reciprocally connected with ACA, somatosensory and auditory cortices connect preferentially with primary motor cortex (Zhang et al., 2016). ACA also projects to visual cortex indirectly, through the superior colliculus (Hu et al., 2019) and through higher-order thalamus (Bennett et al., 2019). Frontal cortex is therefore positioned to convey signals arising from outside the visual system, and ACA in particular is a candidate source of distributed activity shaping V1 responses.

A final class of distributed input arises from neuromodulatory systems. Cholinergic projections from the basal forebrain (Záborszky et al., 2018), noradrenergic projections from the locus coeruleus (Aston-Jones and Cohen, 2005), and serotonergic projections from the raphe nuclei (Ren et al., 2019) reach broadly across cortex, in many cases targeting most or all cortical areas. Neuromodulatory transmission has classically been described as slow and spatially diffuse, operating over tens of seconds to minutes, although recent work has demonstrated precise spatiotemporal dynamics within single cortical areas (Rohner et al., 2026). Neuromodulatory activity is closely tied to arousal, movement, and behavioral state, and modulates cortical responses across sensory areas (McGinley et al., 2015; Pinto et al., 2013;

Reimer et al., 2016). Because neuromodulatory input reaches frontal cortex as well as sensory areas, activity recorded in any one cortical region may partly reflect this shared modulatory drive rather than signals originating in that region itself.

Such connectivity provides the anatomical substrate for distributed influence on single cortical regions. However, there is no straightforward prediction of which regions are functionally coupled. Regions connected by dense direct projections may have weak synaptic connections, and regions with sparse direct connectivity may show robust coupling. The multiple parallel routes mean that a functional relationship between two cortical areas may reflect direct communication, indirect communication, shared input, or some combination. Computation in any single cortical region is therefore shaped by a densely connected network whose influence cannot be inferred from connectivity alone.

1.2 Correlational vs. causal approaches to cortical interaction

Measuring activity across regions establishes correlational relationships but does not address causality. Optogenetics provides temporally precise, cell-type-specific control over neural activity (Boyden et al., 2005; Luo et al., 2008), improving on lesions and pharmacological inactivation in both timescale and specificity. In mouse cortex, transcranial illumination of mice expressing opsins allows activity in a targeted region to be manipulated without surgical access to the tissue (Guo et al., 2014; Li et al., 2019). Scanning the illumination site across the cortical surface extends this approach from single regions to systematic maps of where activity contributes to a given behavior (Pinto et al., 2019; Zatka-Haas et al., 2021). These experiments identify which regions are necessary, but measure the consequences of perturbation only in behavior or at a small number of recording sites.

Each of these approaches has complementary strengths and limitations. Widefield imaging measures activity across cortex but not at spike resolution; electrophysiology resolves spiking but samples few sites; optogenetics perturbs specific regions but is typically paired with limited readout. However, combining these techniques has been constrained by practical considerations around isolating optical channels and genetic strategies.

1.3 Trial-to-trial variability in sensory cortex

Extensive connectivity between cortical regions suggests that each area operates as a node in a brainwide network, carrying signals related to its own computations and those of the regions it connects to. Under this view, sensory responses depend jointly on the stimulus and on ongoing activity in connected regions. Indeed, spike trains produced by sensory cortical neurons in response to repeated presentations of identical stimuli can vary substantially from trial to trial (Bair and Koch, 1996; Tolhurst et al., 1983). This phenomenon is robust across sensory modalities, species, and recording conditions (Cohen and Kohn, 2011; Faisal et al., 2008; Shadlen and Newsome, 1998), and has been observed from early descriptions of single-neuron spiking through the era of large-scale population recordings (Siegle and Steinmetz, 2026).

1.3.1 From noise to structured fluctuations

Variability poses challenges for theories of neural computation, which must address how effective information processing occurs despite unreliable neuronal signaling. While spike generation within a neuron can be highly reliable (Mainen and Sejnowski, 1995), it remains unclear what proportion of variability is attributable to processes at the single-neuron versus network level. For example, spiking variability may arise via irreducible sources such as stochastic quantal release from synapses (Fatt and Katz, 1952) or other biophysical noise sources (Faisal et al., 2008). Additionally, local network mechanisms may give rise to shared chaotic fluctuations (Huang et al., 2019; van Vreeswijk and Sompolinsky, 1996) or other types of shared local fluctuations (Ahmadian and Miller, 2021; Renart et al., 2010), which may drive spiking variability (Pattadkal et al., 2024). Accordingly, neurons in cortical networks have been modeled as Poisson spike generators, with super-Poisson components of variability shared at the population level (Goris et al., 2014). In some cases, such locally-shared fluctuations may constrain the nature and precision of encoded information (Averbeck et al., 2006; Rumyantsev et al., 2020; Stringer et al., 2019a; Zylberberg et al., 2017). In other circumstances, shared correlations may enhance or be unrelated to the information represented in a population (Kohn et al., 2016; Moreno-Bote et al., 2014; Ruff and Cohen, 2014).

Beyond locally shared fluctuations, behavioral and arousal state contribute substantially to cortical variability. Locomotion, pupil dilation, and facial movements modulate firing rates in sensory cortex (Martin Vinck et al., 2015; McGinley et al., 2015; Niell and Stryker, 2010), and recent large-scale recordings have shown that face- and body-derived behavioral signals predict large fractions of activity throughout cortex (Musall et al., 2019; Steinmetz et al., 2019; Stringer et al., 2019b). The brain-wide footprint of these signals indicates that a substantial component of trial-to-trial variability in any single sensory region reflects ongoing fluctuations in behavioral state shared across many regions.

However, behavioral signals account for only a subset of this brain-wide shared activity, suggesting that other sources remain to be identified. One such source is globally distributed activity: patterns shared across many cortical regions. Inter-areal recordings have revealed that activity in one region is partly predictable from activity in another, with the predictable component occupying a low-dimensional “communication subspace” between the regions (Ebrahimi et al., 2022; MacDowell et al., 2025; Semedo et al., 2019; Srinath et al., 2021; Steinmetz et al., 2019; Veuthey et al., 2020). Prior work has quantified how local activity from visual cortical subregions contributes to single-neuron rate variability in primary visual cortex (Akella et al., 2025; Semedo et al., 2019). Whether more globally distributed activity patterns account for super-Poisson components of spiking in single neurons, and how their contribution compares to that of local and behavioral sources, remains unresolved. Chapter 3 addresses this question by partitioning single-neuron V1 variability across local, behavioral, and globally distributed sources.

1.3.2 Approaches to partitioning variability

A range of statistical approaches have been developed to partition trial-to-trial variability into shared and private components, each with its own assumptions and scope.

The modulated Poisson framework models single-neuron spiking as Poisson around a rate that is itself modulated by a shared latent signal (Ecker et al., 2014; Goris et al., 2014; Lin et al., 2015; Rabinowitz et al., 2015; Rupasinghe et al., 2026). In this formulation, a multiplicative gain applied across the population captures the bulk of super-Poisson variability. These models fit well to data, but the latent gain is not attributed to any specific signal.

Encoding models with externally defined regressors attribute single-neuron activity across predictors such as stimulus, choice, and movement. These approaches identify what fraction of each neuron’s activity is attributable to each named source and have been applied broadly across cortex (Musall et al., 2019; Park et al., 2014; Pillow et al., 2008). However, they typically fit each neuron independently and do not capture the low-dimensional structure that is characteristic of cortical populations (Cunningham and Yu, 2014; Ecker et al., 2014; Gao and Ganguli, 2015; Rabinowitz et al., 2015).

Reduced-rank regression extends this approach to the population level, jointly fitting all neurons under a shared low-dimensional constraint and decomposing their activity using observed predictors (Semedo et al., 2019). Given simultaneous recordings from two regions, the method identifies the low-dimensional subspace of source-region activity that best predicts target-region activity, providing a measure of inter-areal coupling. Other latent variable approaches recover shared structure from neural activity alone, without conditioning on externally specified predictors (Pandarinath et al., 2018; Yu et al., 2009). Standard formulations of reduced-rank regression, however, assume Gaussian observation noise and operate on continuous rate signals, making them less directly applicable to single-trial spike count data.

No existing framework has been used to simultaneously handle Poisson spike count observations, partition across multiple externally defined predictor sources, and capture low-dimensional shared structure across the population. Chapter 3 introduces the Poisson reduced rank-regression (p-RRR) framework, developed in collaboration with Z. Lu, to address this gap. Variants of p-RRR have previously been applied to other count-based data (Fitzgerald et al., 2022). Chapter 3 develops and applies the framework for partitioning neuronal variability.

1.4 Main objectives

In this dissertation, we determine whether shared signals between cortical regions reflect direct interactions, and the extent to which activity in a single cortical region is shared in distributed neural circuits. First, we describe a system designed to measure causal interactions across the brain. The system combines widefield imaging, scanning optogenetic manipulation, and high-density

electrophysiology in a single preparation (Chapter 2). Using it, we found that inhibiting a cortical area changed activity in the opposite hemisphere only near the midline, despite strong correlations between homotopic areas. Preliminary work comparing intrinsic timescales measured by autocorrelation and by recovery from perturbation suggests a similar divergence within local circuits. Second, we partitioned trial-to-trial variability in V1 spike counts across stimulus, behavioral, locally shared, and globally distributed sources using a Poisson reduced-rank regression framework (Chapter 3). A low-dimensional subspace of activity accounted for most super-Poisson variability, and activity measured elsewhere in cortex accounted for more than half of that locally explainable variability. Behavioral measurements captured less variability than global signals, and added no further explanatory power when combined with them. In Chapter 4, we consider the findings of both chapters together, examine the limitations and strengths of perturbation-based and correlational measures, and outline future directions.

1.5 General methods

1.5.1 Subjects

All experimental protocols were approved by the Institutional Animal Care and Use Committee at the University of Washington and conducted according to US National Institutes of Health guidelines.

Experiments were performed on male and female mice singly housed on a 12 hour light/dark cycle. For experiments with widefield imaging, multiple GCaMP-expressing genotypes were used: either TetO-GCaMP6s;Camk2a-tTa (B6;DBA-Tg(tetO-GCaMP6s)2Niell/J or TetO-GCaMP8s;Camk2a-tTA (B6;D2-Tg(tetO-GCaMP8s)1Genie/J crossed with B6.Cg-Tg(Camk2a-tTA)1Mmay/DboJ). Wild-type (C57Bl6/J) mice were used for experiments with only electrophysiology.

1.5.2 Surgeries

An initial surgery was performed under isoflurane (1-4% in O₂). Carprofen (5 mg/kg) and lidocaine (2 mg/kg) were administered subcutaneously in the shoulders and over the skull, respectively. The skin was cleared to expose the dorsal skull surface, and incision edges were adhered to the skull with cyanoacrylate (VetBond, World Precision Instruments) to protect the underlying muscle. The skull

was treated with 10% citric acid (Dentin Activator Liquid, Parkell) to improve implant adhesion, then a 3D-printed recording chamber was secured with dental cement (Metabond, Parkell). Thin layers of UV-cured optical adhesive (Norland Optical Adhesive 81, Norland Products) were applied to the exposed skull surface within. A 3D-printed titanium headpost (ProtoLabs) was then cemented to the posterior end of the chamber and allowed to cure. After surgery, mice were treated with carprofen for two days and given at least one week to recover. At least two days before experiments, mice were acclimated to handling and head fixation and placed on mild water restriction.

Before electrophysiological recordings, a brief (~20 min) surgery under light isoflurane (1-2% in O₂) was performed to open craniotomies (~2 mm in diameter) for probe insertions using a handheld dental drill. Exposed dura was sealed with silicone gel (Dow Corning 3-4680 Sil Gel), and the chamber covered with a 3D-printed cap to protect the craniotomies. Mice recovered for at least two hours before experiments.

For optogenetics experiments (Matveev et al., 2024), we retro-orbitally injected tetO-GCaMP8s mice aged P28-42 with the DLX2.0-ChrimsonR-tdTomato opsin packaged in a PHP.eB serotype AAV. Animals were anesthetized with isoflurane (4% in O₂) and injected with virus diluted into 100-150 μ L of PBS.

1.5.3 Electrophysiology and spike sorting

Recordings were made with Neuropixels 2.0 electrode arrays (Steinmetz et al., 2021), which have 384 selectable recording sites across 5120 total sites distributed along four 1 cm shanks. Shanks are 250 μ m apart and each have two vertical columns of sites. Recording configuration was across all four shanks, and spanning either a depth of 700 μ m or 1400 μ m. Probes were mounted to a dovetail rail and attached to a steel rod for micromanipulation (Sensapex). Probe trajectories were planned with Pinpoint (Birman et al., 2023). For later track localization, probes were coated in Dil (ThermoFisher Vybrant V22888) before insertion. In some experiments, probes were inserted through the silicone gel, a saline bath was used in others. Probes were lowered to their final positions at approximately 5 μ m/s. Recordings were made with 30 kHz sampling. In some cases, additional recordings were made at the same locations on subsequent days.

Spike sorting was performed with Kilosort 2.5 (Steinmetz et al., 2021) or Kilosort 4 (Pachitariu et al., 2024) using default parameters. After automatic sorting, units were checked manually in Phy (<https://github.com/cortex-lab/phy>) for any noise artifacts. We then applied a refractory period metric (Roth et al., 2026). These metrics passed 29.2% of units relative to initial yield.

1.5.4 Widefield calcium imaging and optogenetics

Images were collected with a CMOS camera (Basler acA2440-75m; 17.3 $\mu\text{m}/\text{pixel}$ resolution) at a true frame rate of 70 Hz, with an effective rate of 35 Hz due to alternating 405 nm and 480 nm excitation for hemodynamic correction. Frames were captured at 560 x 560 pixels through a 0.63x objective lens (Leica Planapo 0.63x). LED onset and offset were shaped as a 4-ms sine-wave ramp to reduce light artifacts in simultaneous Neuropixels recordings. To avoid light contamination from the stimulus screens, a 3D-printed black cone was affixed to the top of the recording chamber with Blu-Tack. The top of the cone was attached to a piece of light-blocking fabric (ThorLabs) that rested on the top edges of the screens. For simultaneous scanning optogenetics experiments, a 638 nm laser (LuxX 638nm-200mW) was steered by a two-dimensional galvanometer system (ThorLabs GVS002), with a dichroic filter series preventing cross-contamination between imaging and optogenetic wavelengths. Light path details and image preprocessing are as described in Chapter 2.2.1.

Chapter 2: Cortex-wide measurement and manipulation reveals causal organization of cortical dynamics

2.1 Introduction

A central challenge in systems neuroscience is causally linking behavior to neural activity patterns in defined brain regions or cell populations. Correlated activity across regions may reflect functional interdependence, such that disrupting one region impairs processing in connected areas. Alternatively, correlations may arise from shared inputs or parallel processing streams. Distinguishing between these possibilities requires methods that simultaneously measure neural activity across multiple regions while selectively manipulating specific areas to directly test causal relationships within distributed networks.

Focal manipulation of one brain region can produce functional effects at distant sites, a phenomenon termed “diaschisis” and described over a century ago (von Monakow, 1914). Early studies demonstrated that lesions to one hemisphere alter dynamics in the other (Feeney and Baron, 1986), implying that the resulting deficit reflects not the lesioned region alone but the distributed network whose activity it disrupts (Carrera and Tononi, 2014). Recent large-scale recordings have revealed that neural activity patterns are often distributed across brain regions rather than localized to single areas (Manley et al., 2024; Musall et al., 2019; Steinmetz et al., 2019; Stringer et al., 2019b; Zatka-Haas et al., 2021). Sensory, motor, and cognitive information can be encoded across subcortical and cortical networks (Makino et al., 2017; Ottenheimer et al., 2023; Peters et al., 2021).

For decades, lesions and pharmacological manipulations were the key techniques for establishing causal contributions of brain regions, but these have significant limitations in their temporal resolution and targeting precision (Martin and Ghez, 1999; Vaidya et al., 2019). Optogenetics now provides a method for fast-timescale, spatially precise, and cell-type specific manipulations (Boyden et al., 2005; Luo et al., 2008). However, linking effects of manipulations directly to behavioral outcomes, while powerful, has proven to be misleading in some cases (Otchy et al., 2015). Indeed, any manipulation likely produces diverse and far-ranging effects on activity around the brain (Kim et al., 2023; Logothetis et al.,

2010), effects that may vary from trial-to-trial. The ability to simultaneously read and perturb activity patterns could overcome this challenge in two ways. First, we could directly observe the outcome of the manipulation on each trial, and thus link the specific stimulus-induced changes of brain activity to the behavioral outcomes. Second, we could design closed-loop control algorithms to reliably achieve desired manipulation patterns on each trial (Shenoy and Carmena, 2014).

“Widefield” imaging, which uses a low-magnification, large field-of-view microscope and single-photon imaging of intrinsic signals or fluorescent reporters, has been used for decades to measure neural activity across large areas of cortex (Grinvald et al., 1986; Orbach et al., 1985). In recent years, application of this approach to mice has yielded imaging of all dorsal cortical areas, at finer temporal resolution than intrinsic signal imaging (Vanni and Murphy, 2014; Wekselblatt et al., 2016). Key advantages of the mouse model include thin skulls that permit transcranial imaging and the availability of transgenic lines with brain-wide expression of high signal-to-noise calcium sensors. This mesoscale imaging approach has been used to study the relationship of cortical dynamics to perception, action, and cognition (Allen et al., 2017; Couto et al., 2021; Makino et al., 2017; Musall et al., 2019; Zatzka-Haas et al., 2021). In parallel, the same characteristics of mice (transparent skulls and transgenic suitability) have enabled the application of optogenetics across cortical areas (Ambrosone et al., 2025; Guo et al., 2014; Kim et al., 2023; Kochalka et al., 2026; Pinto et al., 2019; Zatzka-Haas et al., 2021). These experiments typically use a similar preparation, expressing an excitatory opsin in inhibitory neurons to suppress cortical activity around the stimulated region (Atallah et al., 2012; Li et al., 2019). Moreover, it is possible to combine electrophysiology with both widefield imaging (Clancy et al., 2019; Peters et al., 2021; Xiao et al., 2017) and with transcranial optogenetics (Zatzka-Haas et al., 2021). Therefore, we hypothesized that with a suitable genetic strategy and microscope design, it would be possible to combine all three techniques in a single experiment.

Here we introduce a microscope design and genetic strategy for simultaneously measuring and manipulating brain activity across the whole dorsal cortical surface in awake, head-fixed mice using single photon widefield calcium imaging and optogenetics. We show that precise, graded measurement and manipulation are simultaneously possible with this strategy, and the measurement and manipulation spectra do not interfere. Additionally, this technique can be used many times within a single

subject with repeatable results over at least 52 weeks. High temporal and spatial precision allows for complex stimulation patterns, including simultaneous multi-target manipulation and generation of traveling waves. Using these techniques in tandem with electrophysiological recordings, we demonstrate a three-modality experiment with multi-timescale and whole-brain access. We employed our approach to characterize the relationship between the optogenetic stimulation paradigm, widefield imaging, and the underlying single-neuron spiking activity in cortex and subcortex. We identified striking discrepancies between correlational and causal coupling of cortical regions. This approach enables systematic brain-wide measurement of cortical perturbation effects on mesoscale neural dynamics.

2.2 Methods

2.2.1 Microscope design

We designed the microscope to have a single shared light path to the brain, with a triple-dichroic setup (see Figure S2.1 for to-scale diagram). The GCaMP excitation light path involved blue (470 nm) and violet (405 nm) LEDs combined with a 425 nm dichroic mirror and coupled into a liquid-core light guide (Cairn OptoLED system). The blue/violet excitation light was introduced to the central light path at a right angle and reflected off a 495 nm dichroic mirror downward towards the 0.63x objective lens and brain. The green GCaMP emission light passed through the 495 nm dichroic mirror and reflected off a 580 nm dichroic mirror to exit the central light path at a right angle towards the 1x condenser lens and camera. The red light for optogenetic excitation was produced by a 638 nm diode laser coupled to a Ø50 µm, 0.22 NA fiber patch cable. In some experiments, a 594 nm diode laser coupled to a Ø100 µm, 0.22 NA fiber patch cable was used for optogenetic excitation. The fiber output was collimated and coupled into a two-dimensional galvanometer mirror system to steer the beam downwards from the top of the central light path. This light passed through scan and tube lenses with long focal lengths such that the red light could travel through the long central chamber housing the 580 and 495 nm dichroics before being focused on the brain by the objective lens.

An alternative strategy to introduce the red laser light could have been by a separate light path with a separate objective lens, for example angled toward the brain from above and to the side of the subject. This system would have been simpler in the sense of avoiding the need to develop a way to couple the red light into the same light path as the imaging. However, our approach has several advantages: 1) There are no extra parts around the microscope relative to a standalone widefield microscope, preserving as much space as possible for combining the imaging and stimulation with electrophysiology; 2) The laser light naturally has access to the same field of view as the microscope and does not require a perspective transformation to map galvanometer voltages to coordinates on the surface of the brain. 3) The surface of the brain is as close to the focal plane of the laser as possible - if the light were introduced at a tilt from the side, the laser would be differently in focus at different points on the skull.

2.2.2 Genetic strategy

We used double-transgenic mice obtained by crossing CaMKII-tTA (B6.Cg-Tg(Camk2a-tTA)1Mmay/DboJ; Jax #007004) with either tetO-GCaMP6s (B6;DBA-Tg(tetO-GCaMP6s)2Niel/J; Jax #024742) or tetO-jGCaMP8s mice (B6;D2-Tg(tetO-GCaMP8s)1Genie/J; Jax #037717, provided by M. Pachitariu; (Wang et al., 2023)). To express the opsin, we retro-orbitally injected mice aged P28-42 with the DLX2.0-ChrimsonR-tdTomato opsin packaged in a PHP.eB serotype AAV. The ChrimsonR particles were produced from pAAV-Syn-ChrimsonR-tdT (Addgene #59171).

2.2.3 Optogenetic experiments

Laser stimulus design. Experiments in Figures 2.1, 2.2F-G, and 2.S5 used a 100 ms step function laser stimulus. The onset and offset of the laser ramped in power in a sinusoidal shape over 1 ms. For experiments with longer stimulus durations using a sine wave laser stimulus (Fig. 2.2C-E), the waveform was adjusted such that the trough was zero and the peak was twice the target power (such that the mean power over time matched the target power). In order to avoid unintended calcium activity from the sound of the galvanometers moving, we added a uniformly random 100-800 ms delay between the movement of the galvanometers and the laser stimulus onset. Laser power and galvo position were controlled via analog output from a NI DAQ (National Instruments USB6343). For the experiment testing low-wavelength activation of Chrimson by the imaging LEDs (Fig. 2.1G), stimuli were 1-second bursts of the alternating 470 nm and 405 nm LEDs, mimicking the standard widefield imaging protocol. The LEDs were set to their maximum power output.

For simultaneous multi-target experiments (Fig. 2.3A, B), the galvos alternated between two positions every 12.5 ms using a half-cosine drive, giving a 25 ms cycle (40 Hz alternation). At each target, the galvos paused for approximately 7 ms, during which the laser was activated for 5 ms. Each stimulus consisted of 160 pulses over 2 seconds.

For rotating wave generation experiments (Fig. 2.3C, D), the galvos scanned continuously around a circle of 1.5 mm radius at a constant angular velocity, with the laser held at constant power throughout the scan. Each stimulus consisted of four complete rotations of 250 ms each (a 4 Hz rotation frequency), for a total

stimulus duration of one second. The scan direction was set to either clockwise or counterclockwise, randomized across trials. The inter-trial interval was 4.3 ± 0.4 s (mean $\pm$ standard deviation).

For cortex-wide mapping experiments (Fig. 2.4, 2.5, 2.6), there were 52 laser targets spaced 1 mm apart. Briefly, the galvos steered to a random target, and the laser waited a random 15-40 ms before ramping to stimulation power. Laser duration was 25 ms, and the inter-trial interval was 470 ± 145 ms (mean $\pm$ standard deviation).

Laser and imaging LED calibrations. In order to measure laser power, we used a light meter (Thorlabs PM100D with S130C) placed in the sample plane to record power in milliwatts. We made measurements in 0.1 V steps along the voltage modulation range (0-5 V) of the laser. The relationship between input voltage and measured power was linear, and consistent across sessions. To compute laser intensity (power per area), we measured the beam diameter using a mirror positioned at the sample plane. The beam power, when reflected by this mirror, is high enough that the small portion reflected by the 550 nm longpass dichroic is detectable on the camera. During *in vivo* imaging, the skull preparation does not reflect enough laser power to be detected. We computed the LED intensity for blue and violet imaging using the area of the photodiode sensor (Thorlabs S130C), which has a diameter of 9.5 mm (area of 70.8 mm²).

2.2.4 Analysis

Duration and spatial extent of inhibition in widefield. We computed the duration of significant inhibition (as in Fig. 2.2C-E) via a threshold determined from a null distribution. The null distribution consisted of calcium activity at baseline (1000 ms before stimulus onset), and the threshold was mean-2*S.D of the distribution. To determine the spatial extent of activation (as in Fig. 2.2F-G, 2.3C-D), we computed the dF/F map shortly after laser onset. The inactivated area was taken as the contiguous region in which suppression reached at least 70% of the peak pixel. Areas were converted to mm² using the measured resolution (17.3 μ m/pixel, Fig. 2.S2)

For visual response experiments (Fig. 2.S5), we averaged across a 519 x 519 μ m (30 x 30 pixels) region of interest to match the spatial extent of the laser-evoked visual response.

Identifying enhanced and suppressed units. To classify units, we compared each unit's mean firing rate in a 100 ms window before laser onset with its mean firing rate in a post-onset response window, using a paired t-test across trials at the highest laser power delivered to the stimulation site. The response window was matched to the stimulus duration, with a 50 ms minimum (50-100 ms across sessions). Units with $p < 0.05$ were classified as Enhanced or Suppressed according to the sign of the firing-rate change (response - baseline); units with $p \geq 0.05$ were classified as non-significant.

Enhanced and suppressed area of effect in electrophysiology. We quantified the spatial extent of each effect with an intensity-weighted effective radius metric. First, we computed the center of mass with the magnitude of effect as weights. The root-mean-square (RMS) radius was then computed as the square root of the intensity-weighted average of squared distances from the center of mass. The effective diameter was defined as $2 \times \text{RMS}$. We used bootstrapping to determine the 95% confidence intervals.

Rotating traveling waves. We visualized the phase of traveling waves by bandpass filtering the widefield data between 2-8 Hz and performing a Hilbert transform (Ye et al., 2026).

Cortical grid stimulation maps. We performed principal component analysis on each neuron's mean firing rate across laser targets, i.e. [$n_{\text{positions}} \times n_{\text{timepoints}}$], binned at 5 ms and averaged over trials at the highest laser power delivered. Neurons with no response variance were not analyzed. PCA was applied with per-stimulus time points as samples and laser targets as features, so that each principal component decomposes into a temporal profile (timecourse of response) and a spatial loading across targets (spatial profile of response). We defined each neuron's causal map as the PC1 loading scaled by the peak-to-peak amplitude of the PC1 temporal profile. Maps were maximum-normalized for display. To test how well one component summarizes the response, we split each target's trials randomly into two equal halves and averaged within reach. A rank-k reconstruction fitted to one half was used to predict the other over 50 random splits. Reproducibility was quantified as the fraction of held-out variance (R^2) shared between the two halves. Cross-validated R^2 at a single component was 92% for responsive cortical neurons and 74% for responsive thalamic neurons. Causal maps computed independently from each half correlated at a median $r=0.90$ and $r=0.93$, respectively.

Spontaneous widefield correlation maps. We quantified each unit's coupling to widefield cortical activity during spontaneous epochs (~30 min total). Widefield activity was averaged within a 20×20 px

region of interest at each cortical location (matched to the grid stimulation protocol). Each unit's spike train was binned at 5 ms and resampled to the widefield frame times by linear interpolation, then both series were z-scored across time. The correlation map was defined as the zero-lag Pearson correlation between the unit's binned firing rate and the widefield activity at each location. For the population summary, correlations were computed at full pixel resolution. Maps were maximum-normalized for display.

Spike-triggered activity maps. To separate each unit's specific coupling from its participation in cortex-wide activity, we computed the same correlation with the global signal removed (Clancy et al., 2019). The global signal was the mean widefield timecourse across the imaged field. For each location we took the first-order partial correlation between the unit's binned firing rate and the widefield activity, controlling for the global signal. Locations were disks of radius 8 px, centered on each of the photostimulation targets for the per-unit maps, and single pixels for the population summary. Maps were maximum-normalized for display.

Predicting target regions from contralateral activity. We fit a separate ridge regression ($\alpha=100$) from each contralateral ROI to the target ROI on spontaneous frames, and tested on held-out data. We then applied the same fitted weights to the optogenetically evoked change.

Correlational and causal networks. A correlational edge was drawn between two ROIs when their mean correlation across spontaneous sessions exceeded 0.85. A causal edge was drawn from a stimulation target to an ROI where the per-trial evoked change was suppressed relative to zero (one-sample t-test).

2.3 Results: System characterization

We first characterized the optical performance of the system. Imaging resolution was 17.3 $\mu\text{m}/\text{pixel}$ at 70 frames per second (35 frames per second after hemodynamic correction, see Chapter 2.2.1), over a field of view of up to 14.2 x 11.8 mm (820 x 685 pixels after 3 x 3 binning; Supplemental Figure S2.1A). The laser spot was 0.12 mm in diameter and delivered 0.1-4.0 mW (~ 10 -350 mW/mm^2 ; Fig. S2.2C), and could be steered over a 10 x 12 mm range in the medial-lateral and anterior-posterior axes, respectively (Fig. S2.2B). The working distance of 7.5 cm left sufficient space to insert recording electrodes between the objective lens and the skull.

2.3.1 Validation with simultaneous imaging, electrophysiology, and optogenetic stimulation

Electrophysiological recordings confirmed that red-light delivery suppressed most local neurons while enhancing the activity in a small subset. We recorded cortical neurons with Neuropixels 2.0 probes (Steinmetz et al., 2021) ($n=3$ sessions, $n=2$ subjects, $n=1019$ total neurons, 339 ± 11 neurons per session; Fig. S2.4). To minimize light artifacts from both imaging and stimulation light sources, LED and laser onsets and offsets were ramped over short periods (see Methods). On average, $3.1\pm 0.7\%$ of neurons showed enhanced firing rates within 5 ms of laser onset, presumably corresponding to opsin-expression inhibitory neurons (“Enhanced” neurons, Fig. 1C-E). These neurons are not observed in widefield imaging, as the GCaMP expression is restricted to excitatory neurons. Suppression of activity in other neurons began within 10.6 ± 0.5 ms and persisted up to 129 ± 3 ms for a 100 ms laser stimulus. At sufficiently high power (0.8 mW), the laser suppressed local spiking activity by $62\pm 7\%$ (Fig. 2.1C-D), with 48% of recorded neurons significantly suppressed (“Suppressed” neurons; paired t-test, 100 ms pre- versus post-onset, $n=30$ trials; Fig. 2.1C). Direct Chrimson activation was spatially restricted, with opsin-expressing neurons responding within $664\text{ }\mu\text{m}$ of the laser target (95% CI: $564\text{--}780\text{ }\mu\text{m}$; Fig. 2.1F, left). However, downstream suppression extended up to $1168\text{ }\mu\text{m}$ (95% CI: $1029\text{--}1327\text{ }\mu\text{m}$; Figure 2.1F, right).

Widefield imaging reliably detected laser-evoked suppression of activity across cortical regions and subjects. Consistent with the rapid suppression measured in electrophysiology, inhibition was observed in widefield imaging within one frame from laser onset (~ 29 ms, Fig. 2.1G). The magnitude of inhibition scaled with laser power and was strongly correlated between electrophysiology and widefield imaging (Fig. 2.1H, Pearson’s r , $r= 0.95$, $p<1e-2$). Inhibitory responses were consistent across multiple target laser locations on the dorsal cortical surface (Fig. 2.1I). Reliable simultaneous calcium imaging and optogenetic manipulation was achieved in all seven opsin-expressing mice (Fig. 2.1J, Table 1).

The widefield imaging and optogenetic manipulation light channels were spectrally separate. In principle, this strategy could be reversed using a red-shifted calcium sensor (Dana et al., 2016) with a blue-shifted opsin (Boyden et al., 2005; Klapoetke et al., 2014); however, red-shifted calcium indicators have substantially lower signal-to-noise ratio than GCaMP (Dana et al., 2016), and fewer transgenic lines

are available. A potential concern with our strategy is that red-shifted opsins have weak short-wavelength sensitivity (Klapoetke et al., 2014), but this did not create a confound in practice. We found no response from opsin-expressing neurons when blue and violet imaging LEDs were present without laser stimulation (Fig. 2.1K, n=130 trials). This is likely due to the low light levels needed to obtain high signal-to-noise ratios in widefield imaging (0.012 mW/mm² for blue; 0.019 mW/mm² for violet, compared to 15.9 mW/mm² for laser stimulation at the power for minimum detectable response). A 580 nm longpass dichroic placed before the imaging camera prevented red light from contaminating the imaging channel. In control subjects without the Chrimson opsin, laser stimulation produced no excitatory or inhibitory effects (Fig. 2.1L, right). However, at higher laser powers, anterior laser stimulation elicited off-target visual responses, presumably via retinal photoreceptor activation (Fig. 2.S5). Prior work has shown that this effect might be reduced with retinal adaptation or shorter-wavelength laser stimulation (Danskin et al., 2015; Weiler et al., 2024).




Figure 2.1. Simultaneous mesoscopic cortical measurement and manipulation via combined widefield imaging, optogenetics, and electrophysiology.

A. Schematic of the microscope optical path (not to scale). Blue and violet LEDs provide

excitation for calcium imaging; emitted fluorescence is directed to the camera. A red laser is steered by galvanometers through the shared path through the brain. **B.** Top, schematic of mouse experimental setup. Bottom, example mean frame from widefield imaging. **C.** Single-neuron firing rate responses to laser stimulation (n=30 trials, example session). Units are categorized by significance and direction of laser response (paired t-test between pre- and post-stimulus spike counts, $p < 0.05$ per unit): enhanced (red), suppressed (blue), and non-significant (gray). **D.** Mean firing rate timecourse for each category in **C.** Shaded regions indicate s.e.m. across units. Red bar indicates laser stimulus. **E.** Trial-aligned spike rasters for an example enhanced unit (left) and an example suppressed unit (right). **F.** Spatial profile of laser effect on enhanced units (left) and suppressed units (right), showing change in firing rate as a function of distance from laser target (n=40 trials, example session). Color indicates mean change in spike rate. **G.** Timecourse of inactivation measured with electrophysiology (top) and widefield calcium imaging (bottom) at varying laser powers (colored traces). Lines indicate mean; shaded regions indicate s.e.m. across trials (n=30 trials, example session). Red bar indicates laser stimulus. **H.** Magnitude of inactivation measured by widefield imaging (y-axis, standard deviations of deconvolved signal) versus electrophysiology (x-axis, spikes/s) across laser powers. Each point represents one power. Pearson's $r = 0.95$. **I.** Mean dF/F images of calcium response at 110 ms after stimulus onset for different laser target locations. Red dot indicates laser target in each panel. **J.** Deconvolved calcium timecourse in response to 0.6 mW stimulation across six subjects (one example session per subject; bottom to top: Subject C1, n=26 trials; Subject C2, n=40 trials; Subject C3, n=50 trials, Subject C4, n=40 trials, Subject C5, n=45 trials, Subject C6, n=30 trials; example session per subject). Lines indicate mean; shaded regions indicate s.e.m. across trials. Red bar indicates laser stimulus. **K.** Top, spike rasters of example enhanced units in response to 638 nm laser (left, n=50 trials) and 405/470 nm imaging LEDs (right, n=130 trials). Bottom, mean peri-stimulus time histogram averaged across all enhanced units (n=9 units, one session). Wilcoxon signed-rank test between baseline and stimulus period, $p < 1e-5$ for laser, $p = 0.9$ for LED). **L.** Left, mean dF/F images 30 ms before (top) and 110 ms after (middle) laser onset in an opsin-expressing subject (n=45 trials), and deconvolved calcium timecourse across laser powers (bottom; lines indicate mean, shaded regions indicate s.e.m. across trials). Right, same as left for a control subject expressing GCaMP only (n=30 trials). Red dot indicates laser target.

2.3.2 Long-term, high-resolution optogenetic manipulation with flexible spatial control

The effect of optogenetic stimulation was consistent within subjects over many weeks (Figure 2.2A, B), with stable inactivation magnitudes observed up to 52 weeks after injection. In the first subject (Subject C1, not shown), the opsin could no longer be driven at 52 weeks after injection, though standardized quantitative data were not collected. In subsequent subjects, we reduced the viral titer to avoid such stability issues and implemented a repeated test over weeks. There was no quantitative decline in opsin performance for these subjects.

The laser stimulus could be controlled to produce different temporal and spatial patterns of inactivation. We reliably controlled the duration of cortical inactivation using a 40 Hz sinusoidal stimulus (Figure 2.2C, D). The sinusoidal light waveform reduced adaptation of photosensitive ion channels relative to constant-amplitude stimulation (Li et al., 2019), producing stable steady-state inactivation over seconds (Figure 2.2D). The duration of inactivation scaled linearly with stimulus duration (slope=0.94,

$p < 1e-5$; Figure 2.2E). We also observed that higher laser powers produced a larger area of significant inactivation (Figure 2.2F, G). This effect is likely mediated by stronger synaptic effects at higher laser powers (rather than increased laser scattering, see Figure 2.1F).




Figure 2.2. Temporal and spatial resolution of optogenetic manipulation.

A. Calcium response to laser (130 ms after laser onset) as a function of weeks from injection. **B.** Example imaging field of view and spatiotemporal laser response, 41 weeks apart, from Subject C2. **C.** Deconvolved calcium responses to two example stimulus durations for an example session (solid line, mean; shaded region, S.E.M. across trials; $n=35$ trials, 0.7 mW laser power). Red line indicates laser stimulus shape and duration. Bracket indicates time of significant suppression ($p < 0.05$ based on a null distribution of baseline activity, see Chapter 2.2.4). **D.** Deconvolved calcium responses for all stimulus duration conditions. **E.** Inactivation duration (observed in widefield, significance test as in C) as a function of stimulus duration ($n=2$ subjects, $n=180$ trials per condition per subject, 0.7 mW laser power). **F.** Top, example images of calcium activity in response to different laser powers (example session, $n=40$ trials per condition, 110 ms after laser onset). Bottom, masked images showing significantly inactivated area (contour where suppression reached at least 70% of maximum, see Chapter 2.2.4). **G.** Average cortical inactivation area as a function of laser power across subjects ($n=120$ trials per condition per subject). Black line indicates mean and S.E.M. across all trials.

The scanning galvanometers and laser power were controllable with sufficient speed and precision to design a wide range of custom stimuli. We inactivated two target areas simultaneously by rapidly scanning the galvos between targets, pausing approximately 7 ms at each target (Fig. 2.3A). The laser was activated for 5 ms during each pause and turned off during galvo scanning. Using this approach, we simultaneously inhibited both left and right somatosensory cortices (Fig. 2.3B). We also designed stimuli to elicit complex spatiotemporal patterns, such as traveling rotating waves (Ye et al., 2026), by scanning the laser in a circular pattern (Fig. 2.3C). Rotating waves were generated with direction and center positions matching the designed laser stimulus (Fig. 2.3D).




Figure 2.3: Flexible spatiotemporal control of neural activity

A. Schematic of hardware control for simultaneous multi-target optogenetic manipulation. Top, galvos move rapidly between two target spots, with 6 ms drive time between targets and pausing 7 ms at each target. Bottom, laser pulses are delivered for 5 ms while the galvos are paused at each target. **B.** Top left, calcium activity response to multi-target inhibition. Top right, calcium activity response to single-target inhibition. Bottom, timecourses of deconvolved calcium responses during multi-target inhibition. **C.** Rotating wave generation via optogenetic inhibition. Red dashed circles indicate laser path; black dashed box indicate ROIs plotted in D. Contours indicate significantly inhibited areas ($2 \times \text{s.t.d.}$ below baseline), colored by time during laser rotation cycle ($n=4$ cycles per trial, $n=50$ trials). **D.** Top, example frames from calcium imaging and phasemaps during clockwise rotating wave generation. Bottom, as top but for counterclockwise stimuli. Circles indicate laser path, dots indicate laser position during frame. Frames are averaged across cycles.

2.3.3 Compatibility with subcortical electrophysiology

Cortical optogenetic manipulation also affected subcortical areas beyond the reach of widefield imaging. To measure these effects, we designed an experiment combining optogenetic inhibition at 52 cortical targets (Guo et al., 2014; Zatzka-Haas et al., 2021) and thalamic recording with a Neuropixels probe. Principal component analysis on each neuron's mean firing rate across targets (Fig. 2.4B) revealed the primary temporal response to stimulation (Fig. 2.4C, top) and its spatial distribution across cortex (Fig. 2.4C, bottom). The top component captured 74% of the reproducible variance in the trial-averaged response ($n=31$ thalamic units). Clustering of these responses identified diverse temporal properties (Fig. 2.4D). The timescale of response, defined as the interval from inhibitory response to excitatory rebound, was bimodally distributed with respect to recording depth (Fig. 2.4E). Dorsal neurons had shorter timescales (93 ± 3 ms, median $\pm$ median absolute deviation; $n=30$ neurons), while ventral neurons had longer timescales (175 ± 10 ms, median $\pm$ median absolute deviation; $n=17$ neurons). We fit a Gaussian mixture model to estimate category boundaries. Spatial response profiles were more distributed and heterogeneous across neurons (Fig. 2.4F). Together, these results demonstrate the ability to measure detailed temporal and spatial features of causal cortico-subcortical relationships at single-neuron resolution.




Figure 2.4: Causal connectivity between cortex and thalamus

A. Top, schematic of cortical imaging and optogenetic manipulation targets (left) and subcortical electrophysiology (right). Bottom, heatmap of electrophysiology data. **B.** Example subcortical neuron response to cortical inhibition at different targets. **C.** Top, first principal component of optogenetic response for example neuron in B. Bottom, weight on PC1 across cortical inhibition targets. **D.** Heatmap of PC1s for thalamic neurons. Neurons are sorted by response similarity via Rastermap. **E.** Timescale of neuron responses to optogenetic stimulation as a function of depth. Ellipses show boundaries (two standard deviations from mean) of clusters found by Gaussian mixture model. Insets show PC1s from two example neurons, indicated by color. Gray shading in the inset indicates timescale value plotted for the main panel (y-axis). **F.** Spatial distribution of weight on PC1 for nine example thalamic neurons.

2.4 Results: Causal vs. correlational connectivity across scales

Optogenetic mapping revealed cortical coupling structure not identifiable from imaging alone. We recorded well-isolated cortical units (Fig. 2.5A, B) during simultaneous widefield imaging and cortex-wide optogenetic perturbation. Although spiking was highly correlated with multiple sites in the widefield imaging, the spatial range of the optogenetic perturbation was more limited (Fig. 2.5C-E). We compared three measures of spatial coupling between single neurons and global cortical activity. The correlation map was the Pearson correlation between binned spiking and the calcium activity at each pixel. The spike-triggered activity map was the mean calcium activity surrounding each spike, computed as a partial

correlation to remove the shared cortex-wide signal (Clancy et al., 2019). For the example unit, correlation and spike-triggered maps were high across most of the target grid, while the causal map was confined to a small number of adjacent targets (Fig. 2.5F). Across the recorded population, the causal map had a localized, unilateral peak, while the correlation and spike-triggered maps were broader and bilateral (Fig. 2.5G). Our data suggest that correlational and spike-triggered coupling reflect a unit's participation in distributed cortical activity rather than the sources of causal influence on its firing, although the perturbation may constrain which forms of influence are detectable (Chapter 2.6.3).




Figure 2.5: Three-modality experiment reveals correlational and causal connectivity at single-neuron resolution

A. Left, schematic of the cortical optogenetic stimulation targets. Right, schematic of Neuropixels recording during simultaneous widefield imaging and cortex-wide stimulation. **B.** Spike waveforms (left, 200 randomly selected spikes) and autocorrelogram (right) for an example single unit. **C.** Stimulation target grid with two example targets highlighted (purple, ML -1.5 mm; orange, ML +2.5 mm; AP +0 mm from bregma). **D.** Spike times of the unit in B (black) with simultaneous widefield activity from ROIs centered on the two locations in C. Values indicate the Pearson correlation between binned spiking and the widefield timecourse. **E.** Spike rasters of the unit in B aligned to laser onset at each of the two targets in C (n=50 trials per target). Red bar indicates laser stimulus. **F.** Causal, correlation, and spike-triggered activity maps for the unit in B, evaluated at each of the 52 target locations. Each map is normalized to its maximum absolute value. Circled positions indicate the two locations in C. **G.** Number of units whose causal (left), correlation (middle), and spike-triggered activity (right) map exceeded 70% of its peak at each pixel (n=41 units, one session, one subject).




Figure 2.6: Cortical networks are ipsilaterally dominant

A. Correlation maps between target areas (black circles) and other cortical regions. **B.** Cross-validated R^2 scores of target areas (x markers) from contralateral predictors, trained and tested on spontaneous activity. **C.** Cross-validated R^2 scores of target areas (x markers), using model trained in **B**, tested on data with direct perturbations (time of perturbation response, 50-200 ms after stimulus onset). **D.** Measured and predicted timeseries of activity from target regions in **C**. **E.** Top, network graph based on correlations between regions. An edge is drawn if the correlation is greater than 0.85. Bottom, network graph based on causal perturbation. An edge is drawn if there is significant response to optogenetic stimulation at the node location. Colors indicate distance from midline.

A discrepancy between correlational and causal coupling was also present at the population level. Homotopic cortical areas exhibited strongly correlated spontaneous activity (Fig. 2.6A). These bilateral correlations could reflect the influence of direct cross-hemispheric projections or of shared upstream inputs (Oh et al., 2014; Shimaoka et al., 2019). To test this, we probed causal coupling with cortex-wide optogenetic manipulation, and found that perturbations were not reflected bilaterally except

near the midline. We quantified this by fitting a linear regression model to predict activity in a target region from spontaneous contralateral activity. The model performed well on held-out spontaneous data (Fig. 2.6B) but failed to predict optogenetically driven changes (Fig. 2.6C, D). Thus, correlational coupling between hemispheres appears strong (Fig. 2.6E, top), but causal coupling is predominantly local (Fig. 2.6E, bottom).

2.5 Preliminary results: Intrinsic timescales of cortical areas

Prior work has found that “intrinsic” timescales, estimated from the width of the spontaneous activity autocorrelation, increase across the cortical hierarchy, from primary sensory areas to frontal and association areas (Murray et al., 2014; Song et al., 2024; Zeraati et al., 2026). One interpretation of this is that cognitive functions requiring persistent activity, such as sensorimotor integration or working memory, depend on slower local circuit dynamics. However, an autocorrelation measured from spontaneous activity reflects both the intrinsic properties of the local network and the timescales of its inputs. Distinguishing these contributions requires a measurement that isolates the local circuit.




Figure 2.7: Hierarchy of intrinsic timescales in mouse cortex

A. Left, autocorrelogram of spontaneous calcium widefield activity at location indicated in cyan. Tau indicates time constant of exponential fit to autocorrelogram decay. Right, map of cortical locations colored by time constant of exponential fit. **B.** Left, timecourse of calcium activity recovery (blue) after brief optogenetic ping (red bar), and corresponding exponential recovery fit (orange). Right, map of cortical locations colored by time constant of exponential fit. **C.** Time constant measured by autocorrelation versus time constant measured by response recovery.

We addressed this by comparing two estimates of timescale in the same cortical areas: a correlational estimate from spontaneous activity, and a causal estimate from the recovery of activity following a brief local perturbation. We measured spontaneous activity and optogenetic responses at 52

cortical areas with simultaneous widefield calcium imaging and scanning optogenetics. The correlational timescale, defined as the autocorrelation width of the spontaneous calcium signal at each area (Song et al., 2024), varied systematically across cortex (Fig. 2.7A). Consistent with prior work, the shortest timescales were in sensory regions (e.g. VISp) and the longest were in frontal and association regions (e.g. MOs, RSP). In principle, this pattern may arise from the local circuit properties in each area, or instead from differences in the timescales of the inputs each area receives. To estimate the timescale of the local circuit directly, we measured the recovery latency of each area following a brief optogenetically-driven pulse, defined as the time constant of an exponential fit to the activity recovery to baseline (Fig 2.7B). Surprisingly, the causal timescale showed the opposite spatial pattern: sensory regions recovered more slowly than frontal and association regions (Fig. 2.7C). Across areas, correlational and causal timescales were negatively correlated. The divergence between the two measures indicates that the hierarchy of intrinsic timescales measured from spontaneous activity does not reflect a corresponding hierarchy of local circuit dynamics. Slow spontaneous fluctuations in frontal and association areas are therefore more consistent with slow input to those areas than with slow local recurrent dynamics.

2.6 Discussion

2.6.1 Summary of findings

Here we developed a system for simultaneous mesoscopic measurement and manipulation of mouse cortical activity in combination with large-scale electrophysiology. We observed reliable inactivation of neural populations across the dorsal cortical surface in awake mice, with these effects in both the widefield calcium imaging and electrophysiological recordings. The spectra of the two modalities were effectively separable: the blue and violet lights used for fluorescence imaging did not activate the opsin, and the red light used to drive the opsin did not interfere with the green imaging channel. The inactivation magnitude depended on both laser intensity and laser stimulus duration; higher laser intensities caused larger spatial inactivation across the cortex. The viral strategy remained stable over many weeks within mice, and though the magnitude of inactivation by laser power varied across subjects,

the slope of the relationship between laser power and inactivation magnitude was similar. We measured causal cortico-subcortical interactions at single-neuron resolution, resolving the temporal and spatial structure of thalamic responses to inactivation of specific cortical targets. Finally, we compared correlational coupling, estimated from spontaneous calcium activity, with causal coupling, estimated from responses to optogenetic perturbation, and found that the two could diverge in single units. At the population scale, homotopic cortical areas were strongly correlated during spontaneous activity, yet perturbations propagated largely ipsilaterally.

2.6.2 System characteristics and limitations

We designed the system to be cost-effective and compatible with widely available genetic tools. For beam steering, digital micromirror devices (DMDs) are an alternative approach to the galvo-targeted stimulation used here (Kim et al., 2023; Kochalka et al., 2026). Whereas DMD-based patterning requires a multi-watt source to compensate for distributing light across the full micromirror array, galvanometer steering concentrates the beam into a single focal spot, allowing a significantly lower-power laser to achieve equivalent intensities. An advantage of DMD is the ability to design arbitrary spatial patterns. However, the speed of the galvo system allows for pseudo-simultaneous targeting at multiple spots (Figure 2.3A). For imaging, a red-shifted calcium sensor (Dana et al., 2016) could be used instead of GCaMP, paired with a blue-shifted opsin (Boyden et al., 2005; Klapoetke et al., 2014). However, red-shifted calcium indicators have substantially lower signal-to-noise ratio than GCaMP (Dana et al., 2016), and fewer transgenic lines are available.

This system is constrained to head-fixed mice, rather than freely moving and naturally behaving subjects. Though head-mounted widefield imaging microscopes have been developed (Rynes et al., 2021; Scott et al., 2018), additional modifications of those systems would be necessary to incorporate targetable, precise stimulation for optogenetic control. Furthermore, combining electrophysiological measurements, as we have demonstrated here, may be challenging. The present system is compatible with the large range of behaviors that have been developed for head-fixed mice.

The spatial spread of inactivation that we observed was unexpectedly large, with significant suppression detected several millimeters from the laser target (Figure 2.2F) despite the small size of the

focused laser beam (0.12 mm diameter, Supplemental Figure 2.1C). We considered that light scattering might compromise spatial precision, however, our electrophysiological observations suggest this is not the case: neurons activated by the laser (i.e. Chrimson-expressing inhibitory interneurons, which do not express GCaMP) were only activated near the laser location (Figure 2.1F), though laser stimulation further away did suppress many of the other recorded neurons (presumably mostly those expressing GCaMP under the CaMKII driver; Figure 2.1F). The local effect of the laser on inhibitory interneurons but longer-range suppression of excitatory neurons, as reported previously with similar optogenetic strategies (Li et al., 2019), suggests that the range of suppression is determined by network propagation of the suppressive effect.

2.6.3 Assumptions and interpretations of causal coupling via optogenetics

Correlational and causal coupling between cortical regions diverged at both single-unit and population scales. At the population scale, homotopic areas were among the most strongly correlated pairs in cortex, yet suppression of one area produced little change in its homotopic partner, despite dense callosal projections (Oh et al., 2014). This result may arise from relatively low-weight homotopic connections not measurable from bulk calcium imaging; however, we did not find individual neurons that responded reliably to contralateral stimulation, which would be expected if the effect were present but too sparse to appear in the widefield signal. For single units, correlational and spike-triggered maps were spatially broad, while responses to perturbation were confined to nearby targets (Figure 2.5G). These data are most consistent with common input driving both hemispheres, for example subcortical or neuromodulatory signals.

Another possibility is that the perturbation used here does not functionally engage the callosal circuits underlying homotopic correlations. Optogenetic suppression produces focal activity patterns outside the range observed spontaneously (MacDowell et al., 2025), and removing activity uniformly within a single area may not reproduce the patterns through which homotopic influence operates spontaneously. A systematic survey of multi-site and temporally patterned perturbations may establish whether homotopic influence is engaged under different conditions.

2.6.4 Hierarchical timescales in the cortex

The hierarchy of intrinsic timescales measured from spontaneous activity was not reproduced by a causal measure of the same areas. Autocorrelation timescales were shortest in sensory regions and longest in association and frontal regions, consistent with prior reports (Murray et al., 2014; Song et al., 2024), while recovery from local suppression was slowest in sensory areas. Autocorrelation reflects both properties of the local network and the timescales of activity arriving to that network, which may be unrelated. One interpretation is that frontal and association areas receive input on slower timescales than sensory areas, while their local dynamics are comparatively fast. However, this interpretation relies on several assumptions.

First, that the recovery from suppression is determined by the same dynamics that shape ongoing activity. Our measurement, recovery from a strong exogenous perturbation, may not necessarily reflect the circuit's response to input in an ordinary regime.

Second, that the perturbation was equivalent across targets. Laser power was matched at each location, and suppression appeared comparable across target locations in widefield (Fig 2.1I). However, suppression is produced by network propagation from locally activated interneurons, so its extent may vary depending on the cell type composition of the local circuit, as well as with opsin expression and light attenuation through the skull.

Third, that both autocorrelation and perturbation measures identify the timescales of the underlying circuit rather than those of the calcium indicator. Widefield fluorescence also reflects bulk activity pooled across many neurons, and so the timescales measured here need not correspond to those of individual neurons, which may be heterogeneous. Repeating both measurements with electrophysiology would resolve whether these observations are consistent across modalities, and whether the perturbation was equivalent across-targets at single neuron resolution.

Chapter 3: Global origins of trial-to-trial variability in V1

3.1 Introduction

Trial-to-trial variability in cortical spike counts can arise from multiple sources, including cell-autonomous processes, locally shared network fluctuations, behavioral state, and globally distributed activity. Although prior work has partitioned variability into single-neuron and locally shared components within individual cortical regions (Goris et al., 2014; Lin et al., 2015), it remains unclear how much of the variability observed in any single region reflects activity that is shared with the rest of cortex. A common framework for organizing these sources treats individual neurons as Poisson spike generators whose rates are modulated by shared latent signals, with super-Poisson variability arising from fluctuations in those latent signals (Ecker et al., 2014; Goris et al., 2014; Lin et al., 2015; Rabinowitz et al., 2015; Rupasinghe et al., 2026)(cite). Prior work has identified some of these shared signals, including locally shared network activity (Akella et al., 2025; Semedo et al., 2019) and spontaneous behaviors (Ährlund-Richter et al., 2026; Manley et al., 2024; Musall et al., 2019; Raut et al., 2025; Steinmetz et al., 2019; Stringer et al., 2019b; Yin et al., 2025). However, behaviors account for only a subset of distributed cortical activity, and prior measurements of behavior-related fluctuations have largely relied on calcium imaging, precluding the quantification of spike count variability. Therefore, the relative contributions of locally shared, behaviorally correlated, and other globally distributed sources to single-neuron spike count variability in V1 have not been quantitatively resolved. Here, we used simultaneous multi-region electrophysiology, mesoscale widefield imaging, and behavioral monitoring to partition trial-to-trial variability in V1 spike counts.

Resolving the relative contributions of these sources requires simultaneous measurements of multiple potential sources, single-neuron spiking resolution, and an analytical framework that can attribute variability to multiple, potentially overlapping sources under a Poisson observation likelihood. Prior studies have typically achieved only a subset of these. Linear reduced-rank regression has been used to partition shared activity across cortical regions but assumes Gaussian observations and does not handle Poisson spike counts (Semedo et al., 2019). Modulated count models

capture super-Poisson variability in spiking but typically use a single latent gain term rather than partitioning across multiple, externally defined predictors (Goris et al., 2014; Rabinowitz et al., 2015). To address these gaps, we developed a Poisson reduced-rank regression (p-RRR) framework that fits multiple predictors jointly under a Poisson likelihood and enables quantification of both unique and shared variability attributable to each source.

To partition variability across these sources, we recorded spikes from hundreds of neurons in mouse primary visual cortex (V1) with four-shank Neuropixels 2.0 probes (Steinmetz et al., 2021). We employed a method not previously used for partitioning neuronal variability, Poisson reduced-rank regression, which models low-dimensional shared structure between neurons. We identified a reliable low-dimensional subspace of activity locally shared among V1 neurons. This subspace represented a large fraction of trial-to-trial spike count variability, and in particular, explained the shared patterns of activity underlying most super-Poisson variability at the population level. Globally distributed patterns of cortical activity, recorded simultaneously from either widefield calcium imaging or with large-scale electrophysiology in frontal cortex, accounted for more than half of this locally explainable variability, significantly more than that attributable to behavior. Geometric analysis of the predicted activity subspaces revealed that globally-predicted activity encompassed behavior-related variability, but extended into additional dimensions unrelated to behavior. Optogenetic perturbations suggested a possible causal pathway for these shared fluctuations from the anterior cingulate area. Our results demonstrate that V1 neurons have responses that can be understood as modulated Poisson processes with at least half of the modulations shared beyond V1. These findings reframe single-neuron variability in sensory cortex as substantially reflecting distributed rather than local processes, and provide a framework for partitioning variability across spatial scales of cortical organization.

3.2 Methods

3.2.1 Passive stimulus presentation

During each experimental session, mice were presented with visual stimuli on three 60 Hz refresh rate screens (Adafruit, LG LP097QX1). Experiment hardware was controlled with Rigbox (Bhagat et al.,

2020). Screens were placed 10 cm in front of the mouse, to its left, and to its right. Stimuli were ten natural scenes from the Allen Brain Observatory Visual Coding dataset (de Vries et al., 2020). Images were randomly interleaved, and each image was shown for 150 repetitions. In most sessions ($n=10$), images were shown for 300 ms. In two sessions, images were shown for 230 ms. Inter-trial intervals were drawn from an exponential distribution (mean: 700 ms, min: 400 ms, max: 3000 ms). To maintain arousal throughout visual stimulus presentations, mice were randomly rewarded (mean: 5 seconds) with drops of 10% sucrose delivered through a small metal tube positioned near the mouth.

3.2.2 Behavioral monitoring

Mice were able to turn a wheel freely during recording experiments. The wheel was attached to a rotary encoder to record movements. A camera and 16 mm focal length lens (ThorLabs MVL16M23) with IR filter was used to record the subject's face and forelimb movements during experiments. The camera acquired at a 70 Hz frame rate (~ 14.3 ms), using the same hardware trigger as for the widefield imaging camera.

3.2.3 Widefield imaging and optogenetics

Briefly, widefield calcium images were collected at 35 Hz (after hemodynamic correction). See Chapter 1.5.4 for technical details. For cortex-wide optogenetic mapping (Fig 3.3A, B), 46 laser targets spaced 1 mm apart were stimulated in randomized order with 150-650 ms inter-trial intervals. Each stimulus was 25 ms in duration at 0.75 mW. For ACA-targeted optogenetic experiments (Fig 3.3C), the laser was targeted to ± 0.75 mm ML, +0.5 mm AP relative to bregma. Each stimulus was 400 ms in duration at 1 mW.

3.2.4 Poisson reduced-rank regression model

Model fitting and relationship to Poisson GLM. Let $y^{(i)} = [y_1^{(i)}, y_2^{(i)}, \dots, y_T^{(i)}]^T \in R^{T \times 1}$ denote the binned spike train of neuron i , and $X = [x_1, x_2, \dots, x_T]^T \in R^{T \times M}$ denote the M -dimensional measurements (e.g. sensory inputs, behavioral or neural features) recorded at time bins 1, 2, ..., T . p-RRR assumes that

$y_t^{(i)} \sim \text{Poisson}(\lambda_t^{(i)})$, where $\lambda_t^{(i)} = f(x_t \cdot h^{(i)} + b^{(i)})$. $b^{(i)}$ is an intercept term accounting for the mean activity, and f is a nonnegative nonlinearity taken to be the softplus function (Keeley et al., 2020). Once learned, the model parameter, $h^{(i)} \in R^{M \times 1}$ specifies the relationship between observed spike counts and the other measurements.

To learn $h^{(i)}, b^{(i)}$ from data, we can write down the log-likelihood of the spike train

$$\text{logp}(y^{(i)}|X, h^{(i)}, b^{(i)}) = c + \sum_{t=1}^T y_t^{(i)} \log f(x_t \cdot h^{(i)} + b^{(i)}) - f(x_t \cdot h^{(i)} + b^{(i)}), \quad c = - \sum_{t=1}^T \log(y_t^{(i)}!),$$

and the optimal $h^{(i)}, b^{(i)}$ are those that maximize the likelihood of the observed spike train.

When spikes from multiple neurons are recorded, for example, $i = 1, \dots, N$, the joint log-likelihood for observing all spike trains is

$$\text{logp}(y^{(1)}, \dots, y^{(N)}|X, h^{(1)}, \dots, h^{(N)}, b^{(1)}, \dots, b^{(N)}) = C + \sum_{i=1}^N \sum_{t=1}^T y_t^{(i)} \log f(x_t \cdot h^{(i)} + b^{(i)}) - f(x_t \cdot h^{(i)} + b^{(i)}).$$

Assuming that $\{h^{(i)}\}_{i=1}^N$ are independent of each other, previous approaches maximize the joint log-likelihood by maximizing $\text{logp}(y^{(i)}|X, h^{(i)}, b^{(i)})$ for each i separately. However, given the significance of low-dimensional structure in neural populations, instead of independence, it may be more appropriate to assume $H = [h^{(1)}, \dots, h^{(N)}] \in R^{M \times N}$ to be low-rank. Rank constraint is a popular regularization method (Candès and Recht, 2012) that can improve the model's robustness to noise in individual neuron spike trains. Therefore, we propose p-RRR, which seeks to maximize

$$\text{logp}(y^{(1)}, \dots, y^{(N)}|X, h^{(1)}, \dots, h^{(N)}, b^{(1)}, \dots, b^{(N)}) + \alpha \|H\|_F^2, \quad \text{subject to } \text{rank}(H) \leq r,$$

where $\alpha \|H\|_F^2$ is an additional L2 regularization penalty to prevent overfitting (Paninski, 2004). α, r are hyperparameters which can be determined by cross-validation. r is typically taken to be less than N (the number of neurons), and when $r = N$, p-RRR is equivalent to N independent Poisson GLMs. When $r < N$, $\{h^{(i)}\}_{i=1}^N$ have to be fitted simultaneously in one matrix H . The optimal H does not have a closed-form solution, but nonetheless can be solved numerically. In practice, we implemented p-RRR as a simple feedforward network in Pytorch (Paszke et al., 2019), and optimized it with the limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm (Liu and Nocedal, 1989).

Model selection. The two major hyperparameters of p-RRR are the rank constraint (r) and the weight for L_2 regularization penalty (α). Hyperparameters were optimized independently for each session and predictor combination. When X contains multiple types of measurements, since different parts of $h^{(i)}$ describe relationships with different measurements, they are given separate L_2 regularization weights. For example, if $x_t = \begin{bmatrix} x_t^{drift} \\ x_t^{behav} \end{bmatrix}$, where $x_t^{drift} \in R^{M_1}$, $x_t^{behav} \in R^{M_2}$, $M_1 + M_2 = M$, then $h_{1:M_1}^{(i)}$ will correspond to x_t^{drift} , and $h_{M_1:M}^{(i)}$ will correspond to x_t^{behav} . So the L_2 penalty will be $\alpha_1 \|H_{1:M_1}\|_F^2 + \alpha_2 \|H_{M_1:M}\|_F^2$, where

$$H_{1:M_1} = \begin{bmatrix} h_{1:M_1}^{(1)} & \dots & h_{1:M_1}^{(N)} \end{bmatrix} \in R^{M_1 \times N}, H_{M_1:M} = \begin{bmatrix} h_{M_1:M}^{(1)} & \dots & h_{M_1:M}^{(N)} \end{bmatrix} \in R^{M_2 \times N}, \text{ and } \alpha_1, \alpha_2 \text{ are determined individually.}$$

We used grid search to find the optimal values of α and r . The typical values we tried for α are 0, 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , and for r are integers between 1 and 60 (typically as multiples of five). The optimal values of α and r were determined through a two-fold cross-validation. We split the recorded timestamps (53581 ± 632 per session) into ten even blocks. Alternating blocks were used as training or testing sets. This method was chosen over random sampling to avoid autocorrelations between training and testing data. We evaluated model performance using D^2 (percentage of Tweedie deviance explained), which is similar to variance explained (R^2) but for Poisson-distributed data (Cameron and Windmeijer, 1996). α, r were chosen based on c.v. D^2 : we first computed a testing D^2 for each neuron, and then obtained a population average D^2 (weighted by each neuron's variance on the test set) for each fold. c.v. D^2 is taken as the average population D^2 across two folds. For each r , the optimal α was chosen to achieve the highest c.v. D^2 across all values of α . The optimal r was then chosen to be the minimal r achieving at least 99% of the maximum c.v. D^2 across all r . To identify the rank for 95% performance (Fig.3.3.1L), we first created a linear interpolation on the c.v. D^2 from ranks we experimented, and found the minimal rank achieving at least 95% of the maximum c.v. D^2 across all ranks.

Null predictors. To account for slow fluctuations in electrophysiology (see Fig 3.1), binned spikes were low-pass filtered at 0.1 Hz. The resulting “infraslow” spike matrix was decomposed using PCA, and the first 10 PCA components served as the null term. This term was included in all models; the Null model contained only this term and an intercept.

Stimulus predictors. Each unique image i ($i = 1, \dots, 10$) was represented by a predictor matrix X_i of size $T \times L$, where T is the total number of time points and $L = 31$ is the number of temporal lags spanning 0-857 ms relative to stimulus onset. Each column of X_i is a binary vector corresponding to a single lag, with entries $X_i(t, l) = 1$ if $t - l \in S_i$ and 0 otherwise, where S_i is the set of onset time points for image i , $l = 0, \dots, 30$. The 10 per-image predictor matrices were horizontally concatenated into a stimulus predictor matrix X of size $T \times 310$. See Fig. 3.S1B for schematic.

Behavior predictors. Face video was analyzed using FaceMap (Syeda et al., 2024); motion energy was computed between frames and decomposed with SVD. The video was cropped to exclude the reward spout while retaining visibility of the tongue for licking behavior. For analyses with neural data, motion energy temporal components were linearly downsampled to 35 Hz to match the binning resolution for electrophysiology and widefield imaging (35 Hz). The top 30 temporal components, each with 30 time-shifted copies, were used as behavior predictors, capturing both current behavior and recent behavior history.

Global predictors. For widefield data, frames were decomposed into 500 orthogonal spatial and temporal components using SVD, then first-order differencing was applied to the temporal components. The top 30 components and their 30 time-shifted copies were included as predictors. For spiking activity from ACA, neurons with average firing rate below 0.2 spikes/s were excluded. Other quality control metrics were not applied, as only model targets required stringent identification of single units. Spikes were binned at 35 Hz, and 30 time-shifted copies per neuron were included.

Local predictors and targets. Following initial quality control (see Chapter 1.5.3, Electrophysiology), neurons with mean firing rate below 0.2 spikes/sec were excluded. Spikes from the remaining single units were binned at 35 Hz. When used as predictors, 30 time-shifted copies per neuron were also included. For analyses using different counting windows (Fig. 3.S7), binned spikes were convolved with a boxcar filter such that each timepoint represented sum over its recent history (e.g., for a 5-bin counting window, each timepoint summed the current bin and the four preceding bins).

3.2.5 Quantifying explained trial-to-trial variability

Residual Fano factor. Let Y, S, O be random variables representing the binned spike count of some neuron, the stimulus, and non-stimulus experimental measurements, such as other behavioral or neural features. p-RRR is fitted under the assumption $Y|S, O \sim \text{Poisson}(g(S, O))$, where g is some function to be modeled. When fitting p-RRR, observations of Y , i.e., $y_1, \dots, y_T$, are organized into the response vector y of some neuron; observations of S, O are organized into the design matrix X .

We devised a novel measure, called the residual Fano factor, to quantify the amount of trial-to-trial variability $\text{Var}(Y|S)$ explainable by O :

$$\text{residual Fano factor} = \frac{\text{Var}(Y - g(S, O)|S)}{E[Y|S]}$$

In the following, we show that when the p-RRR assumption $Y|S, O \sim \text{Poisson}(g(S, O))$ is true, the residual Fano factor equals 1.

Rewrite the numerator as

$$\text{Var}(Y - g(S, O)|S) = \text{Var}(Y|S) + \text{Var}(g(S, O)|S) - 2\text{Cov}(Y, g(S, O)|S)$$

where

$$\text{Var}(Y|S) = E[Y^2|S] - E[Y|S]^2$$

$$\text{Var}(g(S, O)|S) = E[g^2(S, O)|S] - E[g(S, O)|S]^2$$

$$= E[g^2(S, O)|S] - E[E[Y|S, O]|S]^2$$

$$= E[g^2(S, O)|S] - E[Y|S]^2$$

$$\text{Cov}(Y, g(S, O)|S) = E[Yg(S, O)|S] - E[Y|S]E[g(S, O)|S]$$

$$= E[Yg(S, O)|S] - E[Y|S]^2$$

Combine (1)(2)(3),

$$(\star) = E[Y^2|S] - E[Y|S]^2 + E[g^2(S, O)|S] - E[Y|S]^2 - 2E[Yg(S, O)|S] + 2E[Y|S]^2$$

$$= E[(Y - g(S, O))^2|S]$$

$$= E[E[(Y - g(S, O))^2|S, O]|S]$$

$$= E[E[(Y - E[Y|S, O])^2|S, O]|S]$$

$$= E[\text{Var}(Y|S, O)|S]$$

$$= E[E[Y|S, O]|S]$$

$$= E[Y|S],$$

where the equality between (4) and (5) follows from the property of the Poisson distribution. So

$$\text{residual Fano factor} = \frac{E[Y|S]}{E[Y|S]} = 1.$$

In other words, neural covariates represented by O can account for all trial-to-trial variability down to Poisson noise. If the Poisson assumption does not hold, that is, if $Y|S, O \sim D(g(S, O))$, where D is different from the Poisson distribution, then p-RRR may not be the best model for g . Nonetheless, if using g modeled by p-RRR we still have $E[g(S, O)|S] = E[Y|S]$, then we can still follow the calculation to (4). If D has sub-Poisson variability, then (4) < (5), and *residual Fano factor* < 1. If D has super-Poisson variability, then (4) > (5), and *residual Fano factor* > 1.

We estimated the residual Fano factor using trial-aligned data and p-RRR prediction on the test set (which was concatenated between the two folds). For each stimulus, to estimate the numerator, we took the difference between the observed spikes and p-RRR prediction, and computed its variance across trials. An estimation of the denominator was obtained by computing the average spikes across trials. The numerators from all stimuli were scattered against the denominators, and the residual Fano factor was taken as the slope k of the best-fit linear model $y = kx$ on these scattered points. In Fig. 2, each trial time point had a separate linear fit. In Extended Data Fig. 5, all trial time points were included in a single linear fit.

We also verified the condition $E[g(S, O)|S] = E[Y|S]$ for p-RRR models where O consists of the null predictor together with one of the following: the behavior predictor, the global predictor, both the behavior and global predictors, or the local predictor. For every neuron and every stimulus condition, we computed the total observed spikes and p-RRR prediction in the -143ms to 571ms window around stimulus onset for every trial, and tested whether the distributions of observed spikes and p-RRR prediction across trials have the same mean. For all sessions and all choices of O , over 90% of p-RRR predictions across all neurons and all stimulus conditions satisfy the condition (two sample t-test, $\alpha = 0.01$).

Poisson simulation analysis. For each model, test predictions from both folds were concatenated to reconstruct the whole time series. Each neuron's model prediction was treated as the

underlying inhomogeneous rate, from which 100 Poisson-random draws were taken and scored against the prediction using D^2 . A z-score was then computed using the simulated D^2 distribution as the null and the observed D^2 (scored against the data) as the test statistic.

3.2.6 Variance partitioning

To partition D^2 into unique and shared components across predictors, we compared the performance of models containing different predictor combinations. All terms were computed after subtracting the contribution of the Null model. Unique D^2 of each predictor P from {Stimulus, Behavior, Global} (Fig 3.3I) is computed relative to other predictors as $D^2(P \text{ unique}) = D^2(\text{Full}) - D^2(\text{Full without P})$, where Full denotes the model containing all three predictors plus Null, and Full without P omits predictor P. Fig. 3.S8G further partitions the variance explained by all predictors beyond Null into seven non-overlapping components: three unique terms, three pairwise-shared terms, and one triple-shared term. The unique terms were computed in the same way as above. Pairwise-shared terms were computed by inclusion-exclusion; for example, $D^2(\text{Global}\&\text{Behav}) = [D^2(\text{Global}+\text{Stim}+\text{Null}) - D^2(\text{Stim}+\text{Null})] + [D^2(\text{Behav}+\text{Stim}+\text{Null}) - D^2(\text{Stim}+\text{Null})] - [D^2(\text{Full}) - D^2(\text{Stim}+\text{Null})]$. The triple-shared term was defined as the remainder: $D^2(\text{Global}\&\text{Behav}\&\text{Stim}) = D^2(\text{Full}) - D^2(\text{Null}) - (\text{sum of the six other components})$.

For two predictors (e.g., Stimulus and Behavior; Fig 3.4G), individual and shared D^2 were defined as: $D^2(\text{Stim}) = D^2(\text{Null}+\text{Stim}) - D^2(\text{Null})$, $D^2(\text{Behav}) = D^2(\text{Null}+\text{Behav}) - D^2(\text{Null})$, $D^2(\text{Stim}+\text{Behav}) = D^2(\text{Null}+\text{Stim}+\text{Behav}) - D^2(\text{Null})$, and shared $D^2(\text{Stim}, \text{Behav}) = D^2(\text{Stim}) + D^2(\text{Behav}) - D^2(\text{Stim}+\text{Behav})$. Shared D^2 between Stimulus and Global in Fig 3.4G was computed analogously.

3.2.7 Manifold geometry analysis

We used the “half-half” setup (Fig. 3.S4A) to allow for identification of the Local manifold. Half of the neurons in the V1 population were used as predictors and the other half as targets. Stimulus, Behavior, and Global models were also fit to predict the target population. To obtain prediction manifolds, we performed PCA on p-RRR predictions, and took the top PCs accounting for over 99% variance as the predicted manifold.

The alignment between two manifolds is quantified by the amount of activity variance in one manifold captured by the other, with more variance captured indicating a higher degree of alignment. As an example, consider the Stimulus and Behavior manifolds. Take the Behavior model prediction and perform PCA. Suppose the top R_{Behav} PCs account for over 99% variance, then the Behavior manifold is R_{Behav} -dimensional, and the PCs form an orthonormal basis $B_{Behav} \in R^{N \times R_{Behav}}$ of it, where N is the number of V1 neurons. Let $\Lambda_{Stim} \in R^{T \times N}$ be the Stimulus model prediction, where T is the number of timestamps. Following , we can find the dimensions inside Behavior manifold that contain information about Stimulus prediction by identifying a set of $e_i \in R^N, i = 1, \dots, R_{Behav}$, such that $e_i = \underset{e}{\operatorname{argmax}} \left(\|\Lambda_{Stim} e\|_2^2 \right)$, subject to $\|e\|_2^2 = 1, e \in \operatorname{col}(B_{Behav}), e \perp e_1, \dots, e_{i-1}$. e_i are given by the columns of $B_{Behav} V$, where $V \in R^{R_{Behav} \times R_{Behav}}$ contains the right singular vectors of $\Lambda_{Stim} B_{Behav}$. The percentage of variance on the Stimulus manifold captured by Behavior manifold is $100 \times \frac{\|\Lambda_{Stim} E\|_F^2}{\|\Lambda_{Stim}\|_F^2}$, where $E = \begin{bmatrix} e_1 & \dots & e_{R_{Behav}} \end{bmatrix} \in R^{N \times R_{Behav}}$.

Our manifold alignment metric is closely related to those used in previous studies. Our metric is equivalent to the metric neural variance accounted for (VAF), defined in equation 10 in . Both metrics share the same denominator, and using the fact that E has orthonormal columns, we can establish the equivalence of the numerators:

$$\|\Lambda_{Stim} E\|_F^2 = \|E^\top \Lambda_{Stim}\|_F^2 = \|EE^\top \Lambda_{Stim}\|_F^2 = \|\Lambda_{Stim}\|_F^2 - \|\Lambda_{Stim} - EE^\top \Lambda_{Stim}\|_F^2.$$

Also, note that $\|\Lambda_{Stim} - EE^\top \Lambda_{Stim}\|_F^2$ is the squared numerator of the normalized subspace error (e_{sub}), defined by equation 9 in , so our metric also equals to $1 - e_{sub}^2$.

We call E the shared dimensions between the Stimulus and Behavior manifolds, with the amount of stimulus-related variance accounted for by individual dimensions decreasing from its first to the last column. The alignment metric is between 0 and 100 (inclusive), where 0 means the Stimulus manifold is orthogonal to the Behavior manifold, and 100 means all modes in the Stimulus manifold can also be found in the Behavior manifold. If the Stimulus manifold has both shared and private modes relative to the

Behavior manifold, then the metric is between 0 and 100 (exclusive), with higher values indicating that more stimulus-predicted variance was captured by shared modes.

The reported alignment score was cross-validated using our two-fold cross-validation scheme. Following the previous example, on each fold, we first used Λ_{Stim} and B_{Behav} derived from training set predictions of Stimulus and Behavior models to identify e_i . The amount of information captured by each e_i was then evaluated using Λ_{Stim} from the test set. Lastly, the average score of the metric across two folds was reported.

Because fitting p-RRR involves solving a non-convex optimization problem numerically, learned manifolds may vary across random initializations. To assess robustness, we retrained each model 5 times from different random seeds after fixing hyperparameters at their cross-validated optima, yielding P_i and $Local_j$ fits ($i, j = 1, \dots, 5$). For each pair (i, j), we repeated the manifold projection analyses described above (Fig. 4D–F), producing 25 estimates of variance captured per session. We report the mean and standard deviation across these 25 values (Fig. 3.S10D–F).

3.2.8 Decoding analysis

Image identity was decoded from recorded and predicted V1 peri-stimulus time histograms (PSTHs) using multiclass linear discriminant analysis. For each image, even- and odd-numbered trials were used for training and testing, respectively. PSTH matrices were flattened to shape $[n_trials, n_neurons \times n_trial_timepoints]$, and PCA was fitted on the flattened training data. The top 200 principal components were used to fit the LDA model. Decoding performance was evaluated as the posterior probability assigned to the correct image on each trial.

3.2.9 HMM analysis

For each session, a Poisson Hidden Markov Model (Poisson-HMMs) was fitted to the V1 spike matrix using the dynamax (Linderman et al., 2025) package. The number of Poisson-HMM states was treated as a hyperparameter and selected via the Bayesian Information Criterion. Poisson-HMM states were decoded with an L2-regularized logistic regression model with the same predictors as in p-RRR. Timepoints were split into contiguous training, validation, and test blocks of equal size. Decoding models

were fitted on the training set, with optimal regularization weight selected on the validation set. Decoding accuracy on the test set is reported in Fig. 3.S3.

3.3 Results

To determine the magnitude and origin of explainable variability in primary visual cortex (V1), we recorded with a Neuropixels 2.0 probe in awake mice ($n=12$ sessions, $n=10$ mice, $n=2722$ total neurons, 227 ± 64 neurons per session, Fig 3.1A, B). Subjects passively viewed natural scenes (10 images, $n=150$ trials per image, Fig. 3.S1A) while their face movements were monitored with video cameras (Fig 3.1C, D). The presentation of images reliably drove responses in V1 neurons. Individual neurons showed tuning to different images (Fig 3.1E) and trial-to-trial variability, measured by Fano factor (trial-aligned spike count variance to mean ratio), in their responses to identical images (Fig 3.1F).

To capture the influence of factors that may explain trial-to-trial variability in V1 spike counts, we developed a Poisson reduced-rank regression (“p-RRR”) model (see Chapter 3.2.4). The p-RRR model is a hybrid of regression-based and latent variable models, with a nonlinearity and Poisson “output” step to model observed spike trains. It models a low-dimensional shared latent structure between predictors and targets through its rank constraint, guided by the observation that neural populations often have low-dimensional shared structure (Cunningham and Yu, 2014; Ecker et al., 2014; Gao and Ganguli, 2015; Rabinowitz et al., 2015). A rank constraint on stimulus and task variables has been shown to be useful in a previous model for neuron spiking response (Latimer and Freedman, 2023), and p-RRR extends it to model correlations between simultaneously recorded neurons. The use of a Poisson output nonlinearity describes the distributions of spike counts more accurately than a linear model (Semedo et al., 2019), and successfully captured ground-truth latent factors in simulated data (Fig. 3.S2A, B). The p-RRR model performed better than the linear model on our dataset (Fig. 3.S2E) and captured both fast, multi-dimensional fluctuations (Fig 3.1L, M) and slower, coordinated state changes in population activity (Fig. 3.S3). When used to predict a one-dimensional target, the p-RRR model behaves as a Poisson generalized linear model (GLM) (Paninski, 2004; Truccolo et al., 2005) (Fig 3.1G). When non-stimulus factors are included alongside stimulus information as predictors, p-RRR models the unobserved source of trial-to-trial variability as a function of these factors and serves as a flexible generalization of the modulated Poisson model (Goris et al., 2014). Specifically, in prior work, shared super-Poisson variability

(i.e., variability in excess of that expected from a Poisson process) was modeled as a time-varying quantity of unknown origin, and this model enables identification of the correlates of that shared variability. Note that at the firing rates observed in our data, the output nonlinearity can capture a relationship between variance and mean that is consistent with the multiplicative gain factor in the modulated Poisson model (Fig. 3.S2I). Besides neural activity, the p-RRR model may also be applicable to other count-based data; a variant of it has been used for analyzing sequencing data (Fitzgerald et al., 2022).




Figure 3.1: Shared local activity patterns in primary visual cortex explain more spike count variability than behavioral or sensory factors

A. Schematic of Neuropixels 2.0 experimental setup. **B.** Left, example time series of simultaneously recorded data: stimulus presentation times (gray), face motion SV1 (magenta), and

mean V1 population spiking (blue). Right, spike matrix from example session binned at 35 Hz (~28.6 ms), $n=220$ individual V1 units passing quality control. **C.** Mean video frame of the subject's face, and spatial masks of the top 3 SVD components of the movie, emphasizing parts of the subject's snout, whisker pad, and eyes. **D.** Stimulus presentation timing. Ten natural images were randomly interleaved, each shown for 300 ms (or 230 ms in two sessions), with 400-3000 ms inter-stimulus intervals (median: 680 ms), and 150 presentations per image. **E.** Normalized tuning curves across images for five example V1 neurons (columns) from one session. Blue point indicates the image used in F. Error bars indicate s.e.m. **F.** For each example neuron in E: top, spike rasters for all trials ($n=150$) of the image indicated in blue; middle, peri-stimulus time histogram; bottom, Fano factor computed across all ten images as a function of time relative to stimulus onset. Gray shading indicates stimulus duration. **G.** Schematic of the Poisson reduced-rank regression (p-RRR) model. Top, multi-target model with rank constraint. Bottom, single-target model without rank constraint (equivalent to a Poisson-GLM). Nodes from left to right: predictors (stimulus, gray; local peers, pink; behavior, purple; null, dark gray), rank constraint, softplus nonlinearity, Poisson draws. **H.** Example predictions from each model in I (colored) overlaid on recorded binned spikes (blue) for one neuron. Cross-validated D^2 shown for each model. **I.** Distribution of cross-validated D^2 across all neurons in an example session for four models: Local, Behavior, Behavior+Stimulus, and Null. Dashed vertical lines indicate means. **J.** Mean $\pm$ s.e.m. D^2 across sessions ($n=12$, $n=2722$ neurons total). Asterisks indicate significant difference from the Null model (two-sided paired t-test across sessions). **K.** Cross-validated D^2 of the single-target Local model as a function of the number of peer predictors. Each line represents the median across subsamples of predictors and targets in one session. **L.** Cross-validated D^2 of the multi-target Local model (normalized to maximum performance per session, mean $\pm$ s.e.m. over $n=12$ sessions) as a function of rank constraint. Shaded region indicates s.d. across sessions. Arrows indicate ranks at 95% and 100% of maximum performance. **M.** Dimensionality (rank at 95% of maximum performance) for each model. Asterisk indicates significant difference between Local and Behavior models (two-sided paired t-test across sessions).

A “peer prediction” of V1 neurons captured a major fraction of the V1 variability via locally shared factors in the neural population. We first used the p-RRR as a single-target GLM, predicting each V1 neuron from the spiking activity of other single neurons in the simultaneously recorded population. We used a D^2 score (Tweedie deviance, similar to R^2 but for Poisson-distributed model outputs; see Chapter 3.2.4) on held-out data to score each model. This “Local” model explained $26.7 \pm 1.7\%$ (mean $\pm$ s.e.m. across $n=12$ sessions) of the total variance in held-out test data (Fig 3.1H-J). The activity fluctuations captured by this model can in principle be driven by inputs to V1 or by local recurrent dynamics, but we refer to it here as “Local” as a description of the predictors included in the model. The optimal model performance was similar whether we used this metric or R^2 (Fig. 3.S2E), and whether we structured the model to predict single or multiple targets (Fig. 3.S4A, B). As a baseline against which to compare model performance, we constructed a “Null” model whose predictors were the slow (<0.1 Hz) components of the population activity (see Chapter 3.2.4). The Null model captures slow processes that can result in nonsense correlations (Harris, 2021) and are not the focus of this study, such as representational drift (Deitch et al., 2021), “infraslow” population activity (Okun et al., 2019), or artifactual relative motion

between electrodes and the tissue. Null predictors were included in all models. The Local model predicted significantly more variance than the Null model ($p < 0.001$, two-sided paired t-test across session mean D2). Previous work reported that a significant fraction of V1 neural variability is explained by overt movements observable in facial videos (Musall et al., 2019; Stringer et al., 2019b). These factors also explain some V1 spike count variability in our dataset (“Behavior” model; $p < 0.001$ relative to Null), but leave a major fraction unexplained (Fig 3.1J; $p < 0.001$ relative to Local). The Stimulus model, which predicts each neuron’s time-varying mean response to each unique image (see Chapter 3.2.4), also explained significantly less variability than the Local model ($p < 0.001$). The combined Behavior+Stimulus model explained less than the Local model ($p < 0.001$), indicating that the Local model’s high performance was not solely due to capturing variability related to behavior, stimulus timing and identity, or both.

Spike count variability was accounted for by a low-dimensional set of shared local factors. While performance did not saturate with increasing numbers of neurons used to make predictions (Fig 3.1K), this lack of saturation can arise due to spiking noise even in low-dimensional simulated populations (Fig. 3.S4C, D). In other words, improved performance with more neurons can arise from a better estimation of low-dimensional shared variables, and therefore does not necessarily indicate a population dimensionality higher than the number of recorded neurons. We more directly estimated the dimensionality via the reduced rank structure of our multi-target model, determining the rank needed to reach at least 95% of maximum performance (Fig 3.1L). Measured in this way, the dimensionality of spike count variability of local V1 activity was 11.0 ± 0.4 (mean $\pm$ S.E.M. across $n=12$ sessions), significantly greater than the dimensionality explainable by correlations with behavior (7.6 ± 0.5 ; Fig 3.1M; $p < 0.001$, two-sided paired t-test) indicating that the full space of V1 activity explores dimensions not associated with behavioral factors.




Figure 3.2: Shared factors within local populations explain super-Poisson components of trial-to-trial variability

A. Fano factor calculation for an example neuron. Left, spike rasters across trials of a single repeated stimulus. Middle, mean (black) and variance (gray) of spike counts computed across trials in 35 Hz (~28.6 ms) bins. Right, Fano factor (variance-to-mean ratio). Gray shading indicates stimulus duration. Error bars indicate 95% confidence intervals from bootstrapping across trials. **B.** Left, mean spike rate (black) and variance (gray) averaged across ten image conditions for $n=2722$ neurons from $n=12$ sessions. Right, median Fano factor across neurons. Gray shading indicates stimulus duration. Error bars indicate 95% confidence intervals from bootstrapping within sessions. Asterisk indicates significant change in Fano factor during versus outside the stimulus period (two-sided paired t-test, $p=0.013$). **C.** Left, single-trial responses of an example V1 neuron to one stimulus. Colors indicate individual trials; black line indicates the trial-averaged response. Right, total predicted versus recorded spike counts per trial. Colors indicate the sum of predicted spikes. Pearson correlation coefficient shown ($p<1e-8$). **D.** Left, recorded spike count variance (blue) and residual variance after subtracting the Local model prediction (green) for the neuron in C. Right, corresponding recorded and residual Fano factors (variance divided by recorded mean). **E.** Recorded population activity (left) and Local+Stimulus model prediction (right) for one session. Gray bars indicate stimulus presentations. **F.** Recorded Fano factor (blue) and residual Fano factor after accounting for Behavior (cyan) or Local (pink) model predictions, as a function of time relative to stimulus onset, across sessions. Gray shading indicates stimulus duration. Error bars indicate 95% confidence intervals from bootstrapping within sessions. **G.** Distribution of recorded (blue) and Local model residual (pink) Fano factors across neurons (median $\pm$ m.a.d.: recorded, 1.06 ± 0.14 ; residual, 0.98 ± 0.12 ; $p<0.001$, Wilcoxon signed-rank test). **H.** Recorded (blue) and Local model residual (pink) Fano factor as a function of spike counting window size. **I.** Schematic of Poisson simulation test. Poisson-random time series were taken from the

model prediction (i.e., the predicted value at each time point was used as the mean parameter for a draw from a Poisson distribution). **J.** D^2 scores between the Local model prediction and each of the 100 Poisson draws (purple histogram) for an example neuron. The D^2 score between the prediction and recorded data (pink line) falls above the simulated distribution, yielding a simulation z-score of 7. **K.** Distribution of simulation z-scores across neurons for the Local model (pink) and Null model (gray). Positive z-scores indicate that the model predicts the data better than expected from Poisson variability alone ($p < 0.001$, Wilcoxon signed-rank test between Local and Null across sessions, on the fraction of neurons with $z < -3$).

Accounting for locally-shared variability revealed that, on average, V1 neurons fired with variability close to a Poisson process. We quantified the total trial-to-trial spiking variability with Fano factor (Fig 3.2A). Consistent with prior findings in other species (Churchland et al., 2010), we observed a decrease in Fano factor immediately after stimulus onset (Fig 3.2A, B; paired t-test $p = 0.013$). The median Fano factor of V1 neurons across neurons was 1.06 ± 0.14 (median $\pm$ m.a.d. across neurons), with a bin-size dependence consistent with prior work (Rupasinghe et al., 2026) (Fig 3.2H, Fig. 3.S7). The Local model captured much of this trial-to-trial variability, correctly predicting differences in amplitude of responses around the time of stimulus presentations (Fig 3.2C, E). We noted that individual neurons were heterogeneous in their spiking statistics, with some neurons far above or below Poisson expectation, despite the overall population average being close to the variability expected from a Poisson process. The level of individual neuron spiking variability was only weakly correlated with spike rate, spike waveform width, and unit depth (Fig. 3.S6A).

To assess how close the remaining unexplained variability was to that expected from a Poisson process, we computed a “residual” Fano factor (variance of the difference between observed and predicted single-trial spike count data, divided by the observed trial-average spike count, Fig. 3.2D, see Chapter 3.2.5). This “floor” of unexplainable variability has a Fano factor of 0.98 ± 0.12 (median $\pm$ m.a.d. across neurons; Fig. 3.2F, G), significantly lower than the raw Fano factor of 1.06 ± 0.14 (median $\pm$ m.a.d. across neurons; Fig 3.2F, G, Wilcoxon $p < 0.001$). We refer to this as the “floor” since, given that most reliable V1 variance is captured by a small number of dimensions, most explainable variability should be shared across neurons and thus captured by this model. We observed that this near-Poisson floor was lowest at the timescale of our analysis (~ 29 ms bins); the residual Fano factor increased at coarser timescales (Fig 3.2H, Fig. 3.S7), consistent with the accumulation of small residual autocorrelations expected within larger counting windows (Rupasinghe et al., 2026). As with the total predicted variance

(CV D^2 in Fig 3.1L), the residual Fano factor did not saturate as a function of the number of local peer predictors, but nevertheless remained in a similar range even after extrapolating to much larger population sizes (Fig. 3.S5D, F).

Despite its widespread use, the Fano factor also has challenges as a reliable indicator of Poisson-like spiking variability (Charles et al., 2018; Churchland et al., 2011), including its limitation to trial-aligned activity. We therefore developed a complementary, alternative analysis using simulated Poisson counts. This method directly examines the relationship between the residual variability observed for each individual neuron and the variability expected from a Poisson process. For each neuron, we generated 100 Poisson-random samples from the Local or Null model predictions (Fig 3.2I). We then scored these samples against the prediction, creating a distribution of D^2 scores expected under conditions of perfectly accurate mean rate predictions (by construction) combined with exactly Poisson spike generation (Fig 3.2J). Finally, we computed a z-score for each neuron's actual D^2 score relative to this distribution. A z-score of zero indicates that the model predicted the same amount of variability that would be expected if the neuron were a Poisson spike generator with the predicted mean rate. Positive z-scores indicate more variability predicted, thus implying that the neuron fires more regularly than Poisson given the model prediction, and vice-versa for negative z-scores. The results from this method were robust to cross-validation and imprecision in firing rate estimation (Fig. 3.S6B-D). As expected, when the Null model was used for prediction, most neurons had significantly ($z < -3$, 99.7% confidence interval) negative z-scores indicating their super-Poisson variability (61% of neurons with $z < -3$). However, when the Local model was used for prediction, a substantial fraction of the population had responses that were at or below Poisson variability (84% of neurons with $z > -3$). These simulations demonstrate that most neurons fired similarly to a Poisson process or with somewhat more regular firing, but only after accounting for locally-shared variability (Fig 3.2K). These findings were robust to particular choices about the bin size used for counting spikes (Fig. 3.S7).

In principle, the locally-shared activity patterns that we identified could be partially or completely isolated in V1, or may be shared with fluctuations in other parts of cortex. In order to identify a potential source for globally-shared activity, we mapped causal cortical inputs to V1 using simultaneous widefield imaging and scanning optogenetics (Matveev et al., 2024). The anterior cingulate area (ACA) in mice

projects directly to V1 (Zhang et al., 2014), and indeed we observe that suppressing activity there directly resulted in suppression of V1 responses (Fig 3.3A-C). Therefore, we hypothesized that ACA might be a conduit of globally shared activity underlying trial-to-trial variability V1. While recording in V1, we measured nonlocal (“Global”) cortical activity using either a second Neuropixels 2.0 probe in anterior cingulate area (ACA; n=6 sessions; Fig 3.3D-E), or widefield calcium imaging of the cortical surface (n=5 sessions; Fig 3.3F-G), or both methods simultaneously (n=1 session). We fit the p-RRR model to predict V1 using this Global activity, face movements, stimulus kernels, and the null term. This “Full” model achieved over half of the predictive power of the Local model ($15.9 \pm 1.1\%$ for Full vs $26.7 \pm 1.7\%$ for Local, Fig 3.3H, I). Adding Global predictors to the Local model did not improve performance (Fig 3.3I), consistent with the assumption that all V1 signals arising from external sources are also shared locally among neurons within the V1 population.




Figure 3.3: Globally-shared activity patterns account for a majority of local spike count variability and subsume nearly all behavioral contributions

A. Left, schematic of cortex-wide optogenetic inhibition (via injected DLX2.0-Chrimson) and widefield calcium imaging (via transgenic GCaMP8s). Colored circles indicate laser target locations. Right, map of V1 responses to inhibition at each target (25 ms pulses, 0.75 mW). Circle color indicates mean change in V1 fluorescence following inhibition ($n=620$ trials, $n=8$ sessions). Circle size indicates statistical significance (two-sided one-sample t-test on per-session means of baseline versus response, Benjamini-Hochberg corrected across targets). Targets adjacent to V1 are excluded for clarity (see **B**).

B. V1 response magnitude as a function of distance from V1 to the optogenetic target. Circle size indicates significance as in **A**. Colors correspond to target locations in **A**.

C. Time series of V1 (blue) and ACA (red) calcium responses during ACA-targeted optogenetic inhibition (± 0.75 mm ML, $+0.5$ mm AP; 400 ms, 1 mW; $n=400$ trials, example session). Red bar indicates laser stimulation period.

D. Left, schematic of dual Neuropixels 2.0 probe recording in V1 and ACA. Right, example simultaneously recorded time series: face motion SVD component 1 (magenta), mean V1 spiking (blue), and mean ACA spiking (red). Gray bars indicate stimulus presentations.

E. Spike count matrices from simultaneously recorded V1 (blue, top) and ACA (red, bottom) populations, binned at 35 Hz (~ 28.6 ms), from one session.

F. Left, schematic of widefield calcium imaging paired with Neuropixels 2.0 recording. Right, example simultaneously recorded time series: face motion SVD component 1 (magenta), mean V1 spiking (blue), and mean widefield ΔSV (green).

G. Avg frame and four regions of interest (SV1, SV2, SV3, SV4) showing calcium responses.

H. Recorded (blue), Stimuli (gray), Behav+Stim pred (magenta), and ACA+Stim pred (red) time series for 50 neurons over 2 seconds.

I. Scatter plot of % CV D² for different models: Local, Local+Full, Full, Stim, Behav, ACA, Widefield, Global. Legend: Total var expl (black circle), Unique var expl (colored circle).

J. Hypothesis 1 and Hypothesis 2 scatter plots showing $D^2(\text{Global}) - D^2(\text{Null})$ vs $D^2(\text{Behav}) - D^2(\text{Null})$. Legend: Behavior (purple), No behavior (pink).

K. ACA model and WF model scatter plots showing % $D^2(\text{Global}) - \% D^2(\text{Null})$ vs % $D^2(\text{Behav}) - \% D^2(\text{Null})$. Legend: Single trial (white circle).

in V1. Right, example simultaneously recorded time series: face motion SVD component 1 (magenta), mean V1 spiking (blue), and the first-order difference of widefield SVD component 1 (green). Gray bars indicate stimulus presentations. **G.** Average frame from widefield calcium imaging with cortical area boundaries overlaid (left), and spatial masks of the top four SVD components (right). **H.** Recorded V1 population activity (left), Behavior+Stimulus model prediction (middle), and ACA Global+Stimulus model prediction (right) for one session. Gray bars indicate stimulus presentations. **I.** Total cross-validated D^2 (circles) and unique D^2 (squares) for each model, averaged across sessions. Unique D^2 is computed relative to the Full model (containing all predictors). **J.** Schematics of two hypotheses for single-trial Global model performance. Hypothesis 1: global signals carry only behavioral information, predicting that Global and Behavior model performance are correlated across trials. Hypothesis 2: global signals carry additional non-behavioral information, predicting that the Global model outperforms Behavior on trials with and without overt behavior. **K.** Single-trial D^2 from the ACA Global model (left) or widefield Global model (right) versus the Behavior model, both relative to the Null model. Each point is one trial; colors indicate local point density. Red contour lines indicate density.

The “Global” model, which included nonlocal activity and the null term but not behavior or stimulus kernels as predictors, was significantly better than the Behavior model in predicting V1 spike counts ($8.3 \pm 0.7\%$ for behavior, $11.7 \pm 1.2\%$ for ACA (two-sided paired t-test $p < 0.01$), $11.7 \pm 0.9\%$ for widefield (two-sided paired t-test $p < 0.05$), Fig 3.3I). Some of the widefield Global prediction, but not the ACA Global prediction, included stimulus-related variability (Fig. 3.S8). Moreover, when Global predictors were included in the model, Behavior predictors did not add explanatory power. We computed the unique contribution of each predictor by removing it from the Full model and quantifying the reduction in performance (see Chapter 3.2.5). The unique contribution of behavior was nearly zero ($0.5 \pm 0.2\%$). The lack of unique explained variability for the Behavior model relative to the Global model was unlikely due to undersampling of the behavioral space. In prior work, it had been observed that cortical population activity predictions were minimally improved with additional videos of the mouse beyond the face, due to the high correlation of bodily with facial movements (Manley et al., 2024). In principle, both ACA and V1 could receive copies of the same behavioral signal that give rise to the ability of ACA to predict V1. Our optogenetic experiments demonstrate that arbitrary non-behavioral signals injected into ACA do transmit to V1, confirming the possibility that activity patterns in ACA unrelated to behavior may directly contribute to variability in V1. We recorded both the subject’s body and face in a subset of sessions, and also found that body movements were highly correlated with face movements (Fig. 3.S9B, C). To test this another way, we considered that if Global model predictions of V1 truly only reflected behavioral signatures, then this information should not be predictive of V1 activity on trials when mice are sitting very still with no overt behaviors (Fig 3.3J, Hypothesis 1). Under this hypothesis, higher predictability with the Global model reflects better predictions than the Behavior model on the same trials, potentially arising from a

higher-resolution representation of the same behavioral information. If instead the Global activity carried both behavioral and other, non-behavioral predictive signals, then the trials without behavior would still be predictable from the Global model (Fig. 3.3J, Hypothesis 2). We found that trials in which behavior held no predictive power were nevertheless predictable from global activity, consistent with Hypothesis 2 (Fig. 3.3K; ACA). The superior predictability of V1 activity by the Global model relative to the Behavior model was consistent across synchronized and desynchronized cortical states (Fig. 3.S6E, F). These results demonstrate that the predictions of V1 activity by global signals include fundamentally distinct factors rather than being a more detailed representation of behavior alone.




Figure 3.4: The V1-Behavior manifold is a subspace of the Global manifold, and shares

dimensions but not temporal structure with the Stimulus manifold.

A. Schematic of predicted low-dimensional manifolds in the V1 neural space. Each manifold represents the subspace of V1 activity patterns predicted by a given model. **B.** Distribution of manifold dimensionality across all models (Null, Stimulus, Behavior, Global, Local, and their combinations), cross-validation folds, and sessions ($n=298$ total). Dashed line indicates the minimum number of predicted V1 neurons. **C.** Manifold dimensionality versus p-RRR model rank (same models as **B**). Points generally fall near the diagonal; slight deviations arise from the softplus nonlinearity in p-RRR. **D.** Percentage of nonlocally predicted variance (from Stimulus, Behavior, or Global models; colors) captured by the Local manifold alone (x-axis) versus by the combined Local+nonlocal manifold (y-axis). Points near the diagonal indicate that the Local manifold already captures the nonlocal prediction, meaning no additional dimensions are needed. Insets illustrate the two possible outcomes: nonlocal activity fully contained within (right) or extending beyond (left) the Local manifold. **E.** Percentage of combined Global+Behavior predicted variance captured by the Local, Global, or Behavior manifold individually. If the Global manifold fully contains the Behavior manifold, it should capture all combined variance. Local and Global do not differ significantly ($p=0.99$, two-sided paired t-test). **F.** Percentage of Stimulus-predicted variance captured by the Behavior (purple) or Global (red) manifold (y-axis) versus by the Null manifold (x-axis). Points above the diagonal indicate that the manifold captures stimulus-related variance beyond what is explained by slow drift alone. **G.** Individual and shared D^2 of Stimulus and Behavior models (left bar) and Stimulus and Global models (right bar; one-sided t-test against zero across sessions, $p<0.001$). Individual D^2 quantifies variance explained by each model; shared D^2 quantifies variance jointly explained by both models. All terms are computed after subtracting the Null model contribution (see Chapter 3.2.4). **H.** Temporal correlation between Stimulus-Behavior or Stimulus-Global model predictions projected onto their shared neural dimensions (see Chapter 3.2.7). Higher values indicate that the two models not only share subspace geometry but also co-fluctuate in time along those dimensions. Correlation is significantly higher for Stimulus-Global than Stimulus-Behavior ($n=12$, $p<0.001$, one-sided paired t-test). **I.** Left, temporal correlation between pairs of projected activities along the top Stimulus-Behavior shared dimension (the dimension within the Behavior manifold capturing the most stimulus-related variance) across sessions. Right, projected time series for an example session: recorded V1 activity, Local, Stimulus, and Behavior model predictions. Shaded regions indicate stimulus presentations. **J.** Same as I, with the Global model substituted for the Behavior model.

The finding that behavioral variables predicted no unique variance relative to global neural measurements could be consistent with multiple different organizations of activity in the population space. To disambiguate, we conducted a geometric analysis, which revealed that Stimulus, Behavior, and Global predictors act within the subspace spanned by local peer activity, and that Global and Behavior predictors occupy nested subspaces within it. Following previous studies (Cunningham and Yu, 2014; Gallego et al., 2018, 2017; Sadtler et al., 2014), we took V1 neural space as a high-dimensional space where each axis represents the activity of one neuron. For all models, both the optimal model rank and prediction dimensionality were significantly lower than the number of simultaneously recorded neurons (Fig 3.4A-C), suggesting that the predictable activity mainly evolves on low-dimensional manifolds of the neural space. To identify the manifold for each model, we applied principal component analysis (PCA) to the model prediction, and defined its manifold as the r -dimensional subspace spanned by the top r PCs capturing

over 99% variance (Gallego et al., 2018). To quantify the alignment between the Local and Global manifolds, for example, we computed the percentage of variance of the Global model prediction that was captured within the Local manifold (see Chapter 3.2.7). Larger values indicate that more variance is captured by shared neural modes between the two manifolds, and thus a higher degree of alignment.

Stimulus, Behavior, and Global predictors did not drive activity along dimensions beyond those captured by local peer activity. Consistent with results from the unique variance analysis (Fig 3.3I, Fig. 3.S8G), we found that the Stimulus, Behavior, and Global manifolds were subspaces of the Local manifold. Combined Local+Global, Local+Behavior, and Local+Stimulus manifolds did not explain significantly more variance than the Local manifold alone (Fig. 3.4D; two-sided paired t-tests, $p > 0.7$). Note that a subset of points falling below the diagonal (i.e., the Local+P manifold capturing slightly less variance than the Local manifold alone) may reflect estimation variability introduced by cross-validation and model fitting (Fig. 3.S10D-F).

Next, we showed that the Global manifold contained the Behavior manifold, and extended into unique dimensions. To test this, we asked how much of the combined Global+Behavior prediction was captured by each manifold individually. If the Global manifold contains the Behavior manifold, it should capture all of the combined variance; if not, the combined manifold should extend beyond either one alone. Nearly all of the combined model variance was captured within the Global manifold (98%), which did not differ significantly from captured by the Local manifold (Fig. 3.4E; two-sided paired t-test, $p = 0.99$). In contrast, the Behavior manifold captured significantly less ($p < 0.001$, one-sided paired sample t-test; Fig. 3.4E), indicating that the Global model predicts along dimensions not predicted by Behavior, but not vice versa.

The Stimulus manifold shared dimensions with both the Behavior and Global manifolds, but only carried correlated temporal information with the Global manifold. Both the Behavior and Global manifolds captured more stimulus-predicted variance than the Null manifold alone ($p < 0.001$ for both, one-sided paired sample t-test; Fig. 3.4F), indicating shared dimensions with the Stimulus manifold in both cases. However, shared geometry does not necessarily imply shared temporal structure. The Stimulus and Behavior models explained nearly non-overlapping portions of D^2 (Fig. 3.4G; not significantly greater than zero, one-sided t-test $p > 0.5$), whereas the Stimulus and Global models shared significant D^2 (Fig 3.4G;

shared $D^2 = 1.9\%$, $p < 0.001$, one-sided t-test against zero). Projecting predictions onto stimulus-behavior shared dimensions supported this dissociation: the two predictions showed little temporal correlation (Fig 3.4H, I). In contrast, stimulus and global predictions were temporally correlated along their shared dimensions (Fig 3.4H, J). These results indicate that behavior-related predictions share subspace geometry but not temporal structure with stimulus-driven activity, while global signals share both. Together with results from the explained variance comparison (Fig 3.3), the geometric manifold analysis suggests that most of the behavior-related V1 variability is embedded within shared fluctuations coordinated across the brain.

3.4 Discussion

3.4.1 Summary of findings

We used Poisson reduced-rank regression to partition trial-to-trial variability in V1 spike counts across stimulus, behavioral, locally shared, and globally distributed sources. A low-dimensional subspace of locally shared V1 activity accounted for most super-Poisson variability at the population level. Globally distributed activity, captured either by ACA spiking or mesoscale widefield imaging, accounted for more than half of this locally explainable variability. Behavioral measurements alone captured significantly less variability than global signals, and added no further explanatory power when combined with them. Consistent with this finding, geometric analysis revealed that globally predicted activity contained behavior-related variability as a subspace, and extended into additional dimensions unrelated to behavior. Together, these results indicate that single-neuron variability in V1 substantially reflects activity distributed across cortex, rather than fluctuations private to V1.

3.4.2 Locally shared activity dominates super-Poisson V1 variability

We adopted a Poisson noise model as a common and simple assumption for spike count distributions (Goris et al., 2014; Pandarinath et al., 2018; Yu et al., 2009). Although the population-average residual variability was near-Poisson at fast timescales after accounting for locally shared modulation, individual neurons were heterogeneous: many had sub-Poisson residuals (Fig. 3.2G,

K), consistent with biophysical constraints such as refractory periods (Berry and Meister, 1998; Dayan and Abbott, 2005; Ferrari et al., 2018; Kara et al., 2000). The scaling of residual Fano factor with counting window size further indicated that not all non-Poisson temporal structure was captured by the model (Fig. 3.S7). The remaining autocorrelation may reflect cell-autonomous biophysical processes such as adaptation, bursting, or subthreshold oscillations. These processes, when private to individual neurons, would be unrelated to the shared population patterns captured here.

A subtlety of our results represents a key deviation from the Poisson model: while the population average of unexplainable variance was close to Poisson, this property was highly heterogeneous across neurons, with some neurons exhibiting reliably less unexplainable variance and others reliably more (Fig. 3.2G, J). One possibility to explain this outcome is that different cell types have different levels of explainable spike count variability. We separately analyzed neurons by waveform duration, and observed only weak differences in Fano factor (Fig. 3.S6A). A finer division by cell type identity in the future, using opto-tagging, could resolve this question further. A more intriguing possibility is that approximate Poisson variability is enforced by mechanisms at the population level, with individual neurons deviating from this pattern to reflect specific roles in the network.

Our results position V1 spike count variability within the modulated Poisson framework explored in prior work (Goris et al., 2014; Rabinowitz et al., 2015; Rupasinghe et al., 2026), while extending it in two ways. First, the modulating signal is multi-dimensional rather than scalar, and its dominant components are identifiable from simultaneously recorded neural and behavioral signals. Second, in prior approaches rate inhomogeneity has been captured with a multiplicative gain factor, whereas our model captures rate fluctuations using a nonlinear transformation of additive contributions from identified sources. The particular nonlinearity used here implements a form of multiplicative interaction between the additive predictors (Fig. 3.S2I), and prior work has shown that additive and multiplicative formulations can yield similar quality of fit (Lin et al., 2015). Although these formulations may perform similarly in accounting for the data, they imply different mechanisms by which shared factors influence neuronal activity, and disambiguating them will require more mechanistic approaches.

3.4.3 Globally distributed activity accounts for the majority of locally shared variability

Globally distributed activity, captured by either ACA spiking or mesoscale widefield imaging, accounted for more than half of the locally explainable variability in V1 (Fig. 3.3I). Behavior-related signals have been widely reported to predict cortical activity (Manley et al., 2024; Musall et al., 2019; Steinmetz et al., 2019; Stringer et al., 2019b). Our results show that, in V1, this behavior-related variability is a subset of a broader globally distributed signal: behavior captured part but not all of the global modulation, and added nothing beyond what was already captured by ACA spiking or widefield activity. Geometric analysis of the predicted activity subspaces was consistent with this partition, showing that globally predicted activity contained behavior-related variability as a subspace and extended into additional dimensions unrelated to behavior (Fig. 3.4). These additional dimensions may reflect internal cortical states, top-down signals, or other distributed processes not directly observable in overt behavior.

3.4.4 Model assumptions and limitations

Our model makes several key assumptions. First, we focus primarily on capturing variability on faster timescales of milliseconds to seconds. The baseline Null model accounts for activity slower than 0.1 Hz, for example fluctuations with slow temporal dynamics (Goris et al., 2014; Hénaff et al., 2020; Rabinowitz et al., 2015). Second, the model assumes that neuron spike counts are described by a linear combination of predictive factors (Musall et al., 2019; Semedo et al., 2019) with a nonlinearity after summation (Park et al., 2014; Pillow et al., 2008). While such models have provided good fits to spike count data in a variety of contexts, we have not explored the extent to which model fits could be improved with, for example, nonlinear interactions between predictors (Ecker et al., 2014; Goris et al., 2014; Rabinowitz et al., 2015; Rupasinghe et al., 2026; Syeda et al., 2024). Third, we assume that predictive relationships are static over time. However, V1 stimulus representations may drift on slow timescales (Deitch et al., 2021) and cortical dynamics may be better accounted for by HMM-derived states (Akella et al., 2025; Ashwood et al., 2022; Engel et al., 2016; Escola et al., 2011), switching dynamical systems (Linderman et al., 2017), or other kinds of non-stationary models. In the future, models that allow for non-stationarity in predictive relationships may improve fits. Due to these assumptions, our model represents only a lower bound of the possible predictive performance.

Considering the Global model, if a version were devised with substantially superior performance, either due to improved model design or extended measurements of global neural activity, this would only strengthen our conclusion that a majority of spike count variability is shared globally.

Chapter 4: Discussion

4.1 Summary of findings

Chapter 2 introduced a system for simultaneous mesoscale imaging and optogenetic manipulation of dorsal cortex, combining spectrally separate calcium indicator and opsin with scanning illumination through a shared optical path. Using this system, we found that activity in homotopic regions was strongly correlated across hemispheres, but optogenetic perturbation of one hemisphere produced little change in the other except near the midline. This discrepancy was also reflected at the level of single neurons. Therefore, bilateral correlation may not reflect direct interhemispheric communication in most of dorsal cortex. Preliminary work suggested that a hierarchy of intrinsic timescales defined by local circuit manipulation may be inverted relative to the hierarchy defined by autocorrelation.

Chapter 3 introduced a Poisson reduced-rank regression framework to partition trial-to-trial variability in V1 spike counts across local, behavioral, and globally distributed sources. Locally shared V1 activity accounted for most super-Poisson variability at the population level. Globally distributed activity, measured by simultaneous brain-wide recordings, accounted for more than half of this variability and added no variance beyond what behavior captured. Together, these results indicate that single-neuron variability in V1 substantially reflects activity shared with the rest of cortex, rather than fluctuations private to V1.

4.2 Implications for understanding distributed cortical dynamics

The findings in this dissertation reframe trial-to-trial variability in single cortical regions as a signature of distributed organization rather than a property of local circuits. The dominant frameworks for thinking about V1 variability have emphasized within-region modulation (Ecker et al., 2014; Goris et al., 2014; Lin et al., 2015; Rabinowitz et al., 2015), local inter-areal coupling among visual cortical regions (Akella et al., 2025; Semedo et al., 2019), and behavior-related modulation distributed broadly across cortex (Musall et al., 2019; Stringer et al., 2019b). Chapter 3 extends each of these frameworks by

showing that the relevant shared signals are multi-dimensional, linked to identifiable sources, and extend beyond local visual circuits and overt behavior.

The mechanistic and functional nature of these global signals is unknown. In our analyses of locally-shared spiking variability, the activity patterns we term “Local” may not be truly local in origin in the sense of arising purely from local recurrent activity. They may be instead driven by inputs to V1 that we do not measure in Chapter 3, such as neuromodulatory (Pinto et al., 2013) or subcortical activity (Hu et al., 2019; Peters et al., 2021; Ye et al., 2026), and determining their origins would require extended measurements and causal manipulation of potential sources. However, dissecting causality via bulk stimulation (e.g., optogenetics or pharmacological manipulation) is challenging due to state changes and off-target effects potentially elicited by these techniques (Li et al., 2019; Otchy et al., 2015; Yizhar et al., 2011). Manipulation of specific patterns of activity at single-cell resolution, for example with two-photon holographic stimulation (Packer et al., 2015), may be required.

In addition to the anatomical mechanisms by which global signals are distributed across the brain, our work raises the question of what information they carry beyond behavioral movements. Our data suggest that they are not simply a richer readout of behavior; the Global model was predictive even on trials without overt behaviors (Fig. 3.3K), and global signals contain predictive dimensions orthogonal to behavior (Fig. 3.4). This asymmetry implies that global signals carry information between cortical areas that is not captured by overt movements. These globally shared fluctuations may reflect internal cognitive processes, such as memory consolidation (Nguyen et al., 2024), long-term behavioral planning, or covert attentional modulation (Speed and Haider, 2021).

4.3 Correlational versus causal measures of cortical dynamics

The experiments in Chapter 2 directly tested whether correlated activity between cortical regions reflects direct communication between them. Homotopic areas were among the most strongly correlated pairs and are linked by dense callosal projections (Oh et al., 2014). However, optogenetic suppression of one area generally produced little change in its homotopic partner. The same divergence appeared at the level of single neurons: correlational maps were spatially broad, while causal maps were more localized.

Therefore, correlated activity across hemispheres likely reflects common inputs, perhaps from subcortical structures or neuromodulatory activity.

Similar considerations extend to the Global predictors in Chapter 3. A majority of explainable trial-to-trial variability in V1 is predictable from activity in ACA and across dorsal cortex. The mechanism may be common input, direct projections, or a combination. Optogenetic perturbation of ACA did produce changes in V1 activity, establishing a causal pathway between them. However, whether this pathway carries the patterns of activity for trial-to-trial variability is an open question, as Chapter 2 demonstrates that common input may dominate the correlation between two regions even with direct projections.

4.4 Limitations

The manipulations in Chapter 2 and the observations of global activity in Chapter 3 were limited to dorsal cortex. The partition of global activity therefore represents a lower bound on possible model performance. Although we perturbed ACA and observed a change in V1 (Chapter 3.3), ACA also reaches V1 indirectly through superior colliculus and thalamus (Bennett et al., 2019; Hu et al., 2019; Zhang et al., 2014), and thus our experiments do not identify which of these routes carried the effect. More broadly, thalamic and other subcortical structures may participate in distributed cortical fluctuations as strongly as cortex itself. Inactivating cortical targets produced structured responses in thalamic neurons, with a range of timescales and temporal profiles (Chapter 2.3.3). Broader coverage would be expected to increase the estimated extent of distributed influence.

The experiments described used head-fixed mice passively viewing visual stimuli or resting spontaneously. This preparation supports stable imaging and recording, but may span a narrower range of behavioral and arousal states relative to natural behavior. For example, the global activity identified in Chapter 3 may differ in magnitude or structure during tasks that impose greater cognitive demand. Similarly, the causal influence of one cortical region on another may depend on behavioral state.

4.5 Future directions

The system described in Chapter 2 opens several directions beyond the experiments reported here. First, mapping the zone of influence of each cortical location, i.e. the specific set of areas affected

by its inactivation, would characterize the causal architecture of dorsal cortex systematically. We expect this influence may be asymmetrical between any pair of areas. Second, future studies can leverage this technique to further investigate neural correlates of behavior. Rather than inferring manipulation targets, areas engaged during specific behaviors can be identified by imaging and then perturbed directly to test their causal contribution.


Extensions of the system allow for a wide range of experimental designs. For example, any combination of promoters may be used to drive the GCaMP sensor and opsin expression, which facilitates the study of cell type function in distributed neural circuits. The addition of high-density electrophysiology provides additional resolution and allows for linking the activity of deeper brain structures to causal perturbations and mesoscale dynamics in cortex (Huo et al., 2020). The ability to record and manipulate across the entire dorsal cortical surface also provides the opportunity to perform closed-loop perturbations at the mesoscale.

Chapter 3 demonstrated that global activity predicts a majority of explainable trial-to-trial variability in V1. Optogenetic experiments showed a potential causal pathway from ACA to V1, consistent with prior work. However, these experiments do not establish that the predictive relationship is mediated by such a pathway. Testing this hypothesis would require simultaneous ACA and V1 recording during ACA stimulation, asking whether the perturbation changes V1 activity along the dimensions identified by the Global model. However, as described in Chapter 4.2, bulk manipulation induces state changes and off-target effects, and the perturbation would itself alter the global activity whose influence is being tested. Manipulations at single-cell resolution may be required to induce specific patterns of activity rather than suppressing an entire region.

Supplementary Figures


Chapter 2 Supplementary Figures

Figure 2.S1: Scale diagram of microscope with combined imaging and optogenetic light paths.



A. To-scale schematic of microscope system described in Chapter 2.


Figure 2.S2: Spatial calibration of imaging field of view and galvo-directed targeting.



A. Frames captured by widefield microscope with ruler placed in the field of view. Each square is 1 x 1 mm. Left, 560 x 560 px FOV (after 3x3 binning) as used in experiments. Middle, maximum FOV size after 3x3 binning; camera sensor is 2464 x 2056. Right, field of view with implant well. **B.** Diameter of the laser beam via mirror placed at sample plane. Each pixel is 17.3 μm . **C.** Distance of galvo-directed laser movement with analog voltage control. Each square is 1 x 1 mm. Left, coordinates indicate (X, Y) galvo position in volts. Middle, working range of galvo-directed laser movement. X and Y galvo targets span -2 to 5 V in 1 V steps. Right, range of movement with implant well. **D.** Spatial profile of laser transmission as a function of galvo position. Reduced transmission is caused by eclipsing of the beam at the edges of optical components.


Figure 2.S3: Fitting and across-subject generalization of widefield-to-spiking deconvolution

kernel.




A. Schematic of deconvolution kernel fitting. A linear regression predicts measured population spiking (y , blue) from time-lagged copies of the simultaneously recorded widefield calcium signal (design matrix X , green). The dark green trace (C_t) is the unshifted signal; light green traces are copies shifted by lag, from C_{t-n} to C_{t+n} . The fitted weights (k , purple) form the deconvolution kernel, plotted on the right as a function of lag. Lines indicate kernels from individual subjects. Calcium and spiking were recorded at the same cortical site with simultaneous widefield imaging and Neuropixels electrophysiology. **B.** Generalization across subjects. Measured population spiking in subject B (light blue) overlaid with deconvolved activity from the same subject, using a kernel fit to subject B's own data (dark gray) and a kernel fit to a different mouse, subject A (light gray). All traces z-scored.

Figure 2.S4: Electrophysiological responses to laser stimulation across sessions and cortical depth.



A. Heatmap of binned spikes relative to laser stimulus for four sessions across three subjects. Units are categorized by significance and direction of laser response (paired t-test between pre- and post-stimulus spike counts, $p < 0.05$ per unit): enhanced (red), suppressed (blue), and non-significant (gray). **B.** Mean firing rate timecourse for each category in C. Shaded regions indicate s.e.m. across units. Red bar indicates laser stimulus. **C.** Spike rate change during laser stimulus as a function of unit depth on probe. Neurons are colored by their response category.


Figure 2.S5: Visual cortex fluorescence during red-laser photostimulation and during spontaneous activity in GCaMP-only control mice.



A. Continuous visual cortex fluorescence in GCaMP-only control mice during 638 nm photostimulation of anterior/lateral targets near the eye (left) and during a spontaneous epoch of the same session (right). Vertical red bands are laser pulses, colored by power. Grey band is ± 1 s.d. of that subject's spontaneous activity. **B.** As A, for midline and posterior targets.


Chapter 3 Supplementary Figures

Figure 3.S1: Natural image stimuli and validation of image-specific encoding scheme.




A. Natural image stimuli shown in experiments. **B.** Schematics of stimulus encoding schemes. Left, each unique image is represented by 31 binary vectors, with each row corresponding to a specific temporal offset for a specific image (block-diagonal structure). This encoding scheme is used throughout the work. Right, an alternative encoding scheme that does not uniquely identify images. Each row corresponds to a specific temporal offset shared across all images. **C.** Trial-averaged recorded responses (blue) and Stimulus model predictions using the image-specific (pink) and image-agnostic (green) encoding schemes from B. Each subplot shows the response to one image. Shaded regions indicate stimulus presentations. **D.** Same as C, for a second example neuron. **E.** Distribution of the difference in cross-validated D^2 score between the image-specific and image-agnostic encoding schemes, across all neurons. Dashed lines indicate example neurons from C (cyan) and D (magenta).

Figure 3.S2: Poisson-RRR recovers ground-truth structure and outperforms linear RRR.




A. Cross-validated D^2 performance on ground-truth simulated data (see Chapter 3.2.4), as a function of rank and L_2 regularization weight. The optimal rank was 10, matching the rank used to generate the ground-truth weights. **B.** p-RRR recovered weights (blue) and bias (red) versus ground truth weights and bias. **C.** Distribution of p-RRR ranks selected via cross-validation across all sessions and predictor types ($n=88$ models). **D.** Distribution of L^2 regularization weights selected via cross-validation, shown as a heatmap across predictor types and weight values (all sessions combined; see Chapter 3.2.4). **E.** Cross-validated R^2 (top) and D^2 (bottom) for individual neurons in an example session, comparing p-RRR (x-axis) and linear RRR (y-axis). For D^2 computation, negative RRR predictions were clipped to machine epsilon ($\sim 2 \times 10^{-16}$). Colored letters indicate example neurons shown in F-H. **F-H.** Recorded spiking (black), p-RRR predictions (pink), and linear RRR predictions (blue) for three example neurons indicated in E. Below each, the Poisson log-likelihood of each model's prediction and the squared error difference (RRR minus p-RRR) are shown; pink shading indicates where p-RRR achieves lower error, blue where RRR does. Gray bars indicate stimulus presentations. **I.** Simulated variance-mean relationships under additive, multiplicative (Goris et al., 2014), and softplus (p-RRR) models of trial-to-trial variability, for three levels of input noise variance.

Figure 3.S3: Full model predictors track transient cortical state changes in V1.




A. Normalized Bayesian information criterion (BIC) as a function of the number of HMM states. BIC scores for each session were normalized by the 1-state model score to make values comparable across sessions. Error bars indicate s.e.m. across sessions. **B.** Mean firing rate of all neurons ($n=2722$) in the High state versus the Low state of the 2-state HMM (92.8% of neurons fired faster in the High state; High-Low = 5.4 ± 0.5 spikes/sec, mean $\pm$ s.e.m. across neurons; two-sided paired t-test across sessions, $p < 0.001$). **C.** Binned spike rates for simultaneously recorded V1 neurons (rows) over time (columns) in an example session. Gray bars indicate stimulus presentations. **D.** HMM states fit to V1 population activity (magenta), decoded from Full model predictors (cyan), and decoded from Null model predictors (black) for the same time period as C. Decoded states are from held-out time points. **E.** Cross-validated decoding accuracy of HMM states from Full model predictors as a function of minimum state duration. Error bars indicate standard deviation across sessions. Dashed lines indicate chance level (50%) and 90% accuracy.

Figure 3.S4: Local model variants and low-dimensional local population structure.




A. Schematics of two Local model variants. Left ("half-half"), recorded V1 neurons are split into random halves serving as predictors and targets, and p-RRR is fit with rank constraint (as in Fig. 1L, M). Right ("leave-one-out"), a single-target model is fit for each V1 neuron using all other simultaneously recorded V1 neurons as predictors; p-RRR reduces to a Poisson-GLM in this case (as in Fig. 1J, K). **B.** Cross-validated D^2 for the half-half Local model as a function of the number of random splits, compared to the leave-one-out model (dashed pink) and Null model (dashed gray). Performance improves with more splits because averaging reduces cases where a target neuron's most predictive peers happen to fall in the target set. Even a single split performs only slightly below leave-one-out, suggesting that most V1 neurons are driven by shared low-dimensional fluctuations rather than strong pairwise correlations with specific peers. **C.** Schematic of rank-constrained population simulation. Spike trains were generated from a simulated population with a known rank constraint (see Chapter 3.2.4). **D.** Cross-validated D^2 as a function of the number of local predictors for simulated populations with varying rank constraints (colors) and experimental recordings (gray, averaged across sessions). The observed scaling from data most closely resembled simulated populations with low rank (e.g., $r=4$).

Figure 3.S5: The Local model represents the majority of shared fluctuations in V1.




A. Cross-validated D^2 of each model normalized by the Local model D^2 , for each session (lines). Values above 1 indicate better performance than the Local model alone. **B.** Poisson simulation z-score (see Chapter 3.2.5) as a function of the number of local peer predictors. Points and error bars represent mean and s.e.m. across neurons. **C.** Fraction of neurons with Poisson simulation z-score greater than -3 as a function of the number of local peer predictors. Points and error bars represent mean and s.e.m. across neurons. **D.** Residual Fano factor as a function of the number of local peer predictors in the single-target Local model. Each line represents one session. Colors indicate example sessions shown in E. **E.** Distribution of residual Fano factor for individual neurons after accounting for the Local model predictions, using the maximum number of available peers. Colors correspond to sessions in D. **F.** Residual Fano factor as a function of number of local peer predictors, with linear (blue) and quadratic (orange) extrapolations to larger population sizes. Annotated values indicate the extrapolated residual Fano factor at 2^{12} and 2^{16} neurons.

Figure 3.S6: The Poisson measures are reliable, state-dependent, and weakly correlated with basic neuron properties.




A. Fano factor of individual neurons as a function of mean spike rate (left; Spearman $\rho = -0.22$, $p < 0.001$; $n = 2722$ neurons), spike waveform width (middle; Spearman $\rho = -0.19$, $p < 0.01$; $n = 1871$ neurons, units with inverted waveforms excluded), and depth on probe (right; Spearman $\rho = -0.11$, $p = 0.02$; $n = 2722$ neurons). Waveform width is the time between peak and trough. **B.** Poisson simulation z-scores computed from different cross-validation blocks. **C.** Null model-predicted mean firing rate versus recorded mean firing rate across neurons. **D.** Left, Poisson simulation z-score with adjusted mean (y-axis) versus unadjusted mean (x-axis), where "adjusted" means the Null model prediction was additively offset to exactly match the recorded mean firing rate on held-out data. Right, distribution of z-scores before (blue) and after (orange) mean adjustment. **E.** Example trials from an example session, sorted by synchronization category (desynchronized to synchronized, left to right). Rows are neurons, color indicates spike rate. **F.** Fano factor (blue) and residual Fano factor after accounting for Behavior (purple), Global (red), or Local (pink) model predictions, across synchronization categories. Error bars indicate s.e.m. across neurons.

Figure 3.S7: Fano factor increases with integration window.



A. Tuning curves across images for five example V1 neurons (columns) from one session. Blue arrow and point indicate the image used for panels B–D. Error bars indicate s.e.m. **B.** Spike rasters for all trials ($n=150$) of the image indicated in A. Gray shading indicates stimulus duration. **C.** Peri-stimulus time histograms for each example neuron, computed with 1-bin (28 ms), 5-bin (143 ms), and 10-bin (286 ms) counting windows. **D.** Fano factor computed across all ten images as a function of time relative to stimulus onset, for each counting window size. **E.** Residual Fano factor (left of each pair) and Fano factor distributions (right of each pair) computed with 5-bin (left) and 10-bin (right) counting windows. Format as in Fig. 2F, G. **F.** Poisson simulation z-score distributions computed with 5-bin (left) and 10-bin (right) counting windows. Format as in Fig. 2K. **G.** Cross-validated total variance explained (circles) and unique variance explained (squares) for each model, computed with 5-bin (left) and 10-bin (right) counting windows.

Figure 3.S8: Stimulus-related variability in the Global model.



A. Mean V1 population response to each natural image stimulus in an example session. Each row shows one image; colors correspond to image identity. Gray shading indicates stimulus duration.

B. Schematic of the decoding pipeline. Left, single-trial responses are flattened so that each row is one trial and columns index neuron × timepoint combinations. Right, PCA is applied to training data, and held-out trials are projected onto the top 200 components as inputs to a linear discriminant decoder.

C. Held-out single trials projected onto the top two linear discriminant axes. Colors correspond to image identity, as in A.


D. Distribution of posterior probability assigned to the correct image on each held-out trial. Colors correspond to image identity, as in A.

E. Cross-validated decoding performance (mean posterior probability of the correct image) as a function of the number of V1 neurons used for decoding. Each line represents one session.

F. Cross-validated decoding performance using recorded V1 neurons or model-predicted V1 activity from Stimulus, Behavior, ACA Global, and widefield Global models.


G. Partition of mean D^2 into unique and shared components across Stimulus, Behavior, and Global predictors. Unique contributions (Stim, Global, Behav) reflect D^2 exclusively attributable to one predictor. Shared components reflect D^2 jointly attributable to predictor pairs (Global&Stim, Global&Behav, Behav&Stim) or all three (Global&Behav&Stim). All terms are computed after subtracting the Null model contribution.

Figure 3.S9: Considerations for the Behavior model.



A. Cross-validated D^2 from the Behavior model versus the Null model for individual neurons across all sessions (paired t-test across sessions, $p < 0.001$). **B.** Left, example frames from face (top) and body (bottom) video cameras. Right, example time series of the top 5 temporal SVD components from each video. **C.** Overlaid time series of face (blue) and body (red) motion energy from an example session. Correlation coefficient is computed across the entire session. **D.** Example frame from the face video with the spatial loading of the top SVD component overlaid (red). **E.** Left, time series of the temporal SVD component corresponding to D. Dashed vertical lines indicate reward deliveries. Right, expanded view of the region indicated by the black box in the left panel. **F.** Cross-validated D^2 using reward times alone, face video alone ("Behavior" model), or reward times and face video together (Wilcoxon signed-rank test across sessions, Reward-Face, $p < 0.001$, Face-Face+Reward, $p = 0.08$).

Figure 3.S10: Manifold geometry dependence on the presence of behavior and random initialization.



A. Cross-validated D^2 score of the Behavior model versus the Global model (both relative to the Null model) on single trials. Trials are categorized as "Behavior" or "No behavior" based on whether the Behavior model exceeds the Null model (position on the x-axis). Histograms show distributions of the trials on each axis. **B.** Left, for "no behavior" trials, percentage of variance predicted by combined Behavior and Global models accounted for by Local, Global, or Behavior manifolds. Right, as left but for "behavior" trials. **C.** Percentage of variance predicted by combined Behavior and Global models accounted for by the Behavior manifold, compared between trials with and without behavior. **D.** Global-, Behavior-, or Stimulus-predicted variance accounted for by manifolds predicted by combined Local and nonlocal predictors (y-axis), versus by local peer predictors alone (x-axis), as in Figure 4D, repeated across different random initializations (see Chapter 3.2.7). Points and error bars show mean $\pm$ s.d. across initialization pairs. **E.** Stimulus-predicted variance accounted for by Behavior or Global manifolds (y-axis), versus by Null manifold (x-axis), as in Figure 4F, repeated across different random initializations. Points and error bars show mean $\pm$ s.d. across initialization pairs. **F.** Variance predicted by combined Behavior and Global models accounted for by Local, Global, or Behavior manifolds, as in Figure 4E, repeated across different random initializations. Points and error bars show mean $\pm$ s.d. across initialization pairs.

References

- Ahmadian Y, Miller KD. 2021. What is the dynamical regime of cerebral cortex? *Neuron* **109**:3373–3391. DOI: <https://doi.org/10.1016/j.neuron.2021.07.031>, PMID: 34464597
- Ährlund-Richter S, Osako Y, Jenks KR, Odom E, Huang H, Arnold DB, Sur M. 2026. Distinct roles of prefrontal subregion feedback to the primary visual cortex across behavioral states. *Neuron* **114**:492–506.e6. DOI: <https://doi.org/10.1016/j.neuron.2025.10.037>
- Akella S, Ledochowitsch P, Siegle JH, Belski H, Denman DD, Buice MA, Durand S, Koch C, Olsen SR, Jia X. 2025. Deciphering neuronal variability across states reveals dynamic sensory encoding. *Nature Communications* **16**:1768. DOI: <https://doi.org/10.1038/s41467-025-56733-w>
- Allen WE, Kauvar IV, Chen MZ, Richman EB, Yang SJ, Chan K, Gradinaru V, Deverman BE, Luo L, Deisseroth K. 2017. Global Representations of Goal-Directed Behavior in Distinct Cell Types of Mouse Neocortex. *Neuron* **94**:891–907.e6. DOI: <https://doi.org/10.1016/j.neuron.2017.04.017>, PMID: 28521139
- Ambrosone M, Montagni E, Resta F, Ulivi T, Curti L, Polverini F, Mazzamuto G, Mannaioni G, Masi A, Pavone FS, Allegra Mascaro AL. 2025. All-optical mapping reveals distributed suppression of cortical sensory responses after optogenetic silencing. *Brain Stimulation* **18**:1514–1522. DOI: <https://doi.org/10.1016/j.brs.2025.07.003>
- Ashwood ZC, Roy NA, Stone IR, Urai AE, Churchland AK, Pouget A, Pillow JW. 2022. Mice alternate between discrete strategies during perceptual decision-making. *Nature Neuroscience* **25**:201–212. DOI: <https://doi.org/10.1038/s41593-021-01007-z>
- Aston-Jones G, Cohen JD. 2005. AN INTEGRATIVE THEORY OF LOCUS COERULEUS-NOREPINEPHRINE FUNCTION: Adaptive Gain and Optimal Performance. *Annual Review of Neuroscience* **28**:403–450. DOI: <https://doi.org/10.1146/annurev.neuro.28.061604.135709>
- Atallah BV, Bruns W, Carandini M, Scanziani M. 2012. Parvalbumin-Expressing Interneurons Linearly Transform Cortical Responses to Visual Stimuli. *Neuron* **73**:159–170. DOI: <https://doi.org/10.1016/j.neuron.2011.12.013>, PMID: 22243754
- Averbeck BB, Latham PE, Pouget A. 2006. Neural correlations, population coding and computation. *Nature Reviews Neuroscience* **7**:358–366. DOI: <https://doi.org/10.1038/nrn1888>
- Bair W, Koch C. 1996. Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Computation* **8**:1185–1202. DOI: <https://doi.org/10.1162/neco.1996.8.6.1185>, PMID: 8768391
- Bennett C, Gale SD, Garrett ME, Newton ML, Callaway EM, Murphy GJ, Olsen SR. 2019. Higher-Order Thalamic Circuits Channel Parallel Streams of Visual Information in Mice. *Neuron* **102**:477–492.e5. DOI: <https://doi.org/10.1016/j.neuron.2019.02.010>, PMID: 30850257
- Berry MJ, Meister M. 1998. Refractoriness and Neural Precision. *Journal of Neuroscience* **18**:2200–2211. DOI: <https://doi.org/10.1523/JNEUROSCI.18-06-02200.1998>, PMID: 9482804
- Bhagat J, Wells MJ, Harris KD, Carandini M, Burgess CP. 2020. Rigbox: An Open-Source Toolbox for Probing Neurons and Behavior. *eNeuro* **7**. DOI: <https://doi.org/10.1523/ENEURO.0406-19.2020>, PMID: 32493756
- Boyden ES, Zhang F, Bamberg E, Nagel G, Deisseroth K. 2005. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience* **8**:1263–1268. DOI: <https://doi.org/10.1038/nn1525>
- Brodmann K. 2007. Brodmann's: Localisation in the Cerebral Cortex. Springer Science & Business Media.
- Cameron AC, Windmeijer FAG. 1996. R-Squared Measures for Count Data Regression Models with Applications to Health-Care Utilization. *Journal of Business & Economic Statistics* **14**:209–220. DOI: <https://doi.org/10.2307/1392433>
- Candès E, Recht B. 2012. Exact matrix completion via convex optimization. *Commun. ACM* **55**:111–119. DOI: <https://doi.org/10.1145/2184319.2184343>
- Carrera E, Tononi G. 2014. Diaschisis: past, present, future. *Brain* **137**:2408–2422. DOI: <https://doi.org/10.1093/brain/awu101>
- Charles AS, Park M, Weller JP, Horwitz GD, Pillow JW. 2018. Dethroning the Fano Factor: A Flexible, Model-Based Approach to Partitioning Neural Variability. *Neural Computation* **30**:1012–1045. DOI: https://doi.org/10.1162/neco_a_01062

Churchland AnneK, Kiani R, Chaudhuri R, Wang X-J, Pouget A, Shadlen MN. 2011. Variance as a Signature of Neural Computations during Decision Making. *Neuron* **69**:818–831. DOI: <https://doi.org/10.1016/j.neuron.2010.12.037>

Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, Finn IM, Ferster D, Ryu SI, Santhanam G, Sahani M, Shenoy KV. 2010. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience* **13**:369–378. DOI: <https://doi.org/10.1038/nn.2501>

Clancy KB, Orsolic I, Masic-Flogel TD. 2019. Locomotion-dependent remapping of distributed cortical networks. *Nature Neuroscience* **22**:778–786. DOI: <https://doi.org/10.1038/s41593-019-0357-8>

Cohen MR, Kohn A. 2011. Measuring and interpreting neuronal correlations. *Nature Neuroscience* **14**:811–819. DOI: <https://doi.org/10.1038/nn.2842>

Couto J, Musall S, Sun XR, Khanal A, Gluf S, Saxena S, Kinsella I, Abe T, Cunningham JP, Paninski L, Churchland AK. 2021. Chronic, cortex-wide imaging of specific cell populations during behavior. *Nature Protocols* **16**:3241–3263. DOI: <https://doi.org/10.1038/s41596-021-00527-z>

Cunningham JP, Yu BM. 2014. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience* **17**:1500–1509. DOI: <https://doi.org/10.1038/nn.3776>

Dana H, Mohar B, Sun Y, Narayan S, Gordus A, Hasseman JP, Tsegaye G, Holt GT, Hu A, Walpita D, Patel R, Macklin JJ, Bargmann CI, Ahrens MB, Schreier ER, Jayaraman V, Looger LL, Svoboda K, Kim DS. 2016. Sensitive red protein calcium indicators for imaging neural activity. *eLife* **5**:e12727. DOI: <https://doi.org/10.7554/eLife.12727>

Danskin B, Denman D, Valley M, Ollerenshaw D, Williams D, Groblewski P, Reid C, Olsen S, Waters J. 2015. Optogenetics in mice performing a visual discrimination task: measurement and suppression of retinal activation and the resulting behavioral artifact. *PLoS One* **10**:e0144760.

Dayan P, Abbott LF. 2005. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Computational Neuroscience Series MIT Press.

de Vries SEJ, Lecoq JA, Buice MA, Groblewski PA, Ocker GK, Oliver M, Feng D, Cain N, Ledochowitsch P, Millman D, Roll K, Garrett M, Keenan T, Kuan L, Mihalas S, Olsen S, Thompson C, Wakeman W, Waters J, Williams D, Barber C, Berbesque N, Blanchard B, Bowles N, Caldejon SD, Casal L, Cho A, Cross S, Dang C, Dolbeare T, Edwards M, Galbraith J, Gaudreault N, Gilbert TL, Griffin F, Hargrave P, Howard R, Huang L, Jewell S, Keller N, Knoblich U, Larkin JD, Larsen R, Lau C, Lee E, Lee F, Leon A, Li L, Long F, Luviano J, Mace K, Nguyen T, Perkins J, Robertson M, Seid S, Shea-Brown E, Shi J, Sjoquist N, Slaughterbeck C, Sullivan D, Valenza R, White C, Williford A, Witten DM, Zhuang J, Zeng H, Farrell C, Ng L, Bernard A, Phillips JW, Reid RC, Koch C. 2020. A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature Neuroscience* **23**:138–151. DOI: <https://doi.org/10.1038/s41593-019-0550-9>

Deitch D, Rubin A, Ziv Y. 2021. Representational drift in the mouse visual cortex. *Current Biology* **31**:4327–4339.e6. DOI: <https://doi.org/10.1016/j.cub.2021.07.062>, PMID: 34433077

Ebrahimi S, Lecoq J, Romyantsev O, Tasci T, Zhang Y, Irimia C, Li J, Ganguli S, Schnitzer MJ. 2022. Emergent reliability in sensory cortical coding and inter-area communication. *Nature* **605**:713–721. DOI: <https://doi.org/10.1038/s41586-022-04724-y>

Ecker AS, Berens P, Cotton RJ, Subramaniyan M, Denfield GH, Cadwell CR, Smirnakis SM, Bethge M, Tolias AS. 2014. State Dependence of Noise Correlations in Macaque Primary Visual Cortex. *Neuron* **82**:235–248. DOI: <https://doi.org/10.1016/j.neuron.2014.02.006>

Engel TA, Steinmetz NA, Gieselmann MA, Thiele A, Moore T, Boahen K. 2016. Selective modulation of cortical state during spatial attention. *Science* **354**:1140–1144. DOI: <https://doi.org/10.1126/science.aag1420>, PMID: 27934763

Escola S, Fontanini A, Katz D, Paninski L. 2011. Hidden Markov Models for the Stimulus-Response Relationships of Multistate Neural Systems. *Neural Computation* **23**:1071–1132. DOI: https://doi.org/10.1162/NECO_a_00118

Faisal AA, Selen LPJ, Wolpert DM. 2008. Noise in the nervous system. *Nature Reviews Neuroscience* **9**:292–303. DOI: <https://doi.org/10.1038/nrn2258>

Fatt P, Katz B. 1952. Spontaneous subthreshold activity at motor nerve endings. *The Journal of Physiology* **117**:109–128. PMID: 14946732

Feeney DM, Baron JC. 1986. Diaschisis. *Stroke* **17**:817–830. DOI: <https://doi.org/10.1161/01.str.17.5.817>, PMID: 3532434

Felleman DJ, Van Essen DCV. 1991. Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1–47.

Ferrari U, Deny S, Marre O, Mora T. 2018. A Simple Model for Low Variability in Neural Spike Trains. *Neural Computation* **30**:3009–3036. DOI: https://doi.org/10.1162/neco_a_01125

Fitzgerald T, Jones A, Engelhardt BE. 2022. A Poisson reduced-rank regression model for association mapping in sequencing data. *BMC Bioinformatics* **23**:529. DOI: <https://doi.org/10.1186/s12859-022-05054-6>

Gallego JA, Perich MG, Miller LE, Solla SA. 2017. Neural Manifolds for the Control of Movement. *Neuron* **94**:978–984. DOI: <https://doi.org/10.1016/j.neuron.2017.05.025>, PMID: 28595054

Gallego JA, Perich MG, Naufel SN, Ethier C, Solla SA, Miller LE. 2018. Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications* **9**:4233. DOI: <https://doi.org/10.1038/s41467-018-06560-z>

Gămănuț R, Kennedy H, Toroczka Z, Ercsey-Ravasz M, Essen DCV, Knoblauch K, Burkhalter A. 2018. The Mouse Cortical Connectome, Characterized by an Ultra-Dense Cortical Graph, Maintains Specificity by Distinct Connectivity Profiles. *Neuron* **97**:698–715.e10. DOI: <https://doi.org/10.1016/j.neuron.2017.12.037>, PMID: 29420935

Gao P, Ganguli S. 2015. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology* **32**:148–155. DOI: <https://doi.org/10.1016/j.conb.2015.04.003>

Goris RLT, Movshon JA, Simoncelli EP. 2014. Partitioning neuronal variability. *Nature Neuroscience* **17**:858–865. DOI: <https://doi.org/10.1038/nn.3711>

Grinvald A, Lieke E, Frostig RD, Gilbert CD, Wiesel TN. 1986. Functional architecture of cortex revealed by optical imaging of intrinsic signals. *Nature* **324**:361–364. DOI: <https://doi.org/10.1038/324361a0>, PMID: 3785405

Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K. 2014. Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* **81**:179–194. DOI: <https://doi.org/10.1016/j.neuron.2013.10.020>, PMID: 24361077

Harris JA, Mihalas S, Hirokawa KE, Whitesell JD, Choi H, Bernard A, Bohn P, Caldejon S, Casal L, Cho A, Feiner A, Feng D, Gaudreault N, Gerfen CR, Graddis N, Groblewski PA, Henry AM, Ho A, Howard R, Knox JE, Kuan L, Kuang X, Lecoq J, Lesnar P, Li Y, Luviano J, McConoughey S, Mortrud MT, Naeemi M, Ng L, Oh SW, Ouellette B, Shen E, Sorensen SA, Wakeman W, Wang Q, Wang Y, Williford A, Phillips JW, Jones AR, Koch C, Zeng H. 2019. Hierarchical organization of cortical and thalamic connectivity. *Nature* **575**:195–202. DOI: <https://doi.org/10.1038/s41586-019-1716-z>

Harris KD. 2021. Nonsense correlations in neuroscience. DOI: <https://doi.org/10.1101/2020.11.29.402719>

Hénaff OJ, Boundy-Singer ZM, Meding K, Ziemba CM, Goris RLT. 2020. Representation of visual uncertainty through neural gain variability. *Nature Communications* **11**:2513. DOI: <https://doi.org/10.1038/s41467-020-15533-0>

Hu F, Kamigaki T, Zhang Z, Zhang S, Dan U, Dan Y. 2019. Prefrontal Corticotectal Neurons Enhance Visual Processing through the Superior Colliculus and Pulvinar Thalamus. *Neuron* **104**:1141–1152.e4. DOI: <https://doi.org/10.1016/j.neuron.2019.09.019>

Huang C, Ruff DA, Pyle R, Rosenbaum R, Cohen MR, Doiron B. 2019. Circuit Models of Low-Dimensional Shared Variability in Cortical Networks. *Neuron* **101**:337–348.e4. DOI: <https://doi.org/10.1016/j.neuron.2018.11.034>, PMID: 30581012

Hubel DH, Wiesel TN. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology* **160**:106–154. DOI: <https://doi.org/10.1113/jphysiol.1962.sp006837>, PMID: 14449617

Hunnicutt BJ, Jongbloets BC, Birdsong WT, Gertz KJ, Zhong H, Mao T. 2016. A comprehensive excitatory input map of the striatum reveals novel functional organization. *eLife* **5**:e19103. DOI: <https://doi.org/10.7554/eLife.19103>

Huo Y, Chen H, Guo ZV. 2020. Mapping Functional Connectivity from the Dorsal Cortex to the Thalamus. *Neuron* **107**:1080–1094.e5. DOI: <https://doi.org/10.1016/j.neuron.2020.06.038>, PMID: 32702287

Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, Lee AK, Anastassiou CA, Andrei A, Aydın Ç, Barbic M, Blanche TJ, Bonin V, Couto J, Dutta B, Gratiy SL, Gutnisky DA, Häusser M, Karsh B, Ledochowitsch P, Lopez CM, Mitelut C, Musa S, Okun M, Pachitariu M, Putzeys J, Rich PD, Rossant C, Sun W, Svoboda K, Carandini M, Harris KD, Koch C, O’Keefe J, Harris TD. 2017. Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**:232–236. DOI: <https://doi.org/10.1038/nature24636>

Kara P, Reinagel P, Reid RC. 2000. Low Response Variability in Simultaneously Recorded Retinal, Thalamic, and Cortical Neurons. *Neuron* **27**:635–646. DOI: [https://doi.org/10.1016/S0896-6273\(00\)00072-6](https://doi.org/10.1016/S0896-6273(00)00072-6)

Keeley SL, Zoltowski DM, Aoi MC, Pillow JW. 2020. Modeling statistical dependencies in multi-region spike train data. *Current Opinion in Neurobiology* **65**:194–202. DOI: <https://doi.org/10.1016/j.conb.2020.11.005>

Kim S, Moon HS, Vo TT, Kim C-H, Im GH, Lee S, Choi M, Kim S-G. 2023. Whole-brain mapping of effective connectivity by fMRI with cortex-wide patterned optogenetics. *Neuron* **111**:1732–1747.e6. DOI: <https://doi.org/10.1016/j.neuron.2023.03.002>

Klapoetke NC, Murata Y, Kim SS, Pulver SR, Birdsey-Benson A, Cho YK, Morimoto TK, Chuong AS, Carpenter EJ, Tian Z, Wang J, Xie Y, Yan Z, Zhang Y, Chow BY, Surek B, Melkonian M, Jayaraman V, Constantine-Paton M, Wong GK-S, Boyden ES. 2014. Independent optical excitation of distinct neural populations. *Nature Methods* **11**:338–346. DOI: <https://doi.org/10.1038/nmeth.2836>

Kochalka J, Mitra A, White AD, Wang Y, Sheng K, Hao YA, Kadur C, Ramakrishnan C, Hernandez M, Wenderski W, Pyo KE, Andini N, Clandinin TR, Lin MZ, Druckmann S, Deisseroth K. 2026. Causal and directional elements of global brain dynamics. DOI: <https://doi.org/10.64898/2026.05.27.726039>

Kohn A, Coen-Cagli R, Kanitscheider I, Pouget A. 2016. Correlations and Neuronal Population Information. *Annual Review of Neuroscience* **39**:237–256. DOI: <https://doi.org/10.1146/annurev-neuro-070815-013851>

Lashley KS. 1929. Brain mechanisms and intelligence: A quantitative study of injuries to the brain. *Brain mechanisms and intelligence: A quantitative study of injuries to the brain* University of Chicago Press. DOI: <https://doi.org/10.1037/10017-000>

Latimer KW, Freedman DJ. 2023. Low-dimensional encoding of decisions in parietal cortex reflects long-term training history. *Nature Communications* **14**:1010. DOI: <https://doi.org/10.1038/s41467-023-36554-5>

Li N, Chen S, Guo ZV, Chen H, Huo Y, Inagaki HK, Chen G, Davis C, Hansel D, Guo C, Svoboda K. 2019. Spatiotemporal constraints on optogenetic inactivation in cortical circuits. *eLife* **8**:e48622. DOI: <https://doi.org/10.7554/eLife.48622>

Lin I-C, Okun M, Carandini M, Harris KD. 2015. The Nature of Shared Cortical Variability. *Neuron* **87**:644–656. DOI: <https://doi.org/10.1016/j.neuron.2015.06.035>, PMID: 26212710

Linderman S, Johnson M, Miller A, Adams R, Blei D, Paninski L. 2017. Bayesian Learning and Inference in Recurrent Switching Linear Dynamical Systems. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. Presented at the Artificial Intelligence and Statistics. PMLR. p. 914–922.

Linderman SW, Chang P, Harper-Donnelly G, Kara A, Li X, Duran-Martin G, Murphy K. 2025. Dynamax: A Python package for probabilistic state space modeling with JAX. *Journal of Open Source Software* **10**:7069. DOI: <https://doi.org/10.21105/joss.07069>

Liu DC, Nocedal J. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical Programming* **45**:503–528. DOI: <https://doi.org/10.1007/BF01589116>

Livingstone MS, Hubel DH. 1981. Effects of sleep and arousal on the processing of visual information in the cat. *Nature* **291**:554–561. DOI: <https://doi.org/10.1038/291554a0>

Logothetis NK, Augath M, Murayama Y, Rauch A, Sultan F, Goense J, Oeltermann A, Merkle H. 2010. The effects of electrical microstimulation on cortical signal propagation. *Nature Neuroscience* **13**:1283–1291. DOI: <https://doi.org/10.1038/nn.2631>

Luo L, Callaway EM, Svoboda K. 2008. Genetic Dissection of Neural Circuits. *Neuron* **57**:634–660. DOI: <https://doi.org/10.1016/j.neuron.2008.01.002>

MacDowell CJ, Libby A, Jahn CI, Tafazoli S, Ardalan A, Buschman TJ. 2025. Multiplexed subspaces route neural activity across brain-wide networks. *Nature Communications* **16**:3359. DOI: <https://doi.org/10.1038/s41467-025-58698-2>

Mainen ZF, Sejnowski TJ. 1995. Reliability of spike timing in neocortical neurons. *Science (New York, N.Y.)* **268**:1503–1506. DOI: <https://doi.org/10.1126/science.7770778>, PMID: 7770778

Makino H, Ren C, Liu H, Kim AN, Kondapaneni N, Liu X, Kuzum D, Komiyama T. 2017. Transformation of Cortex-wide Emergent Properties during Motor Learning. *Neuron* **94**:880–890.e8. DOI: <https://doi.org/10.1016/j.neuron.2017.04.015>, PMID: 28521138

Manley J, Lu S, Barber K, Demas J, Kim H, Meyer D, Traub FM, Vaziri A. 2024. Simultaneous, cortex-wide dynamics of up to 1 million neurons reveal unbounded scaling of dimensionality with neuron number. *Neuron* **112**:1694–1709.e5. DOI: <https://doi.org/10.1016/j.neuron.2024.02.011>, PMID: 38452763

Markov NT, Ercsey-Ravasz MM, Ribeiro Gomes AR, Lamy C, Magrou L, Vezoli J, Misery P, Falchier A, Quilodran R, Gariel MA, Sallet J, Gamanut R, Huissoud C, Clavagnier S, Giroud P, Sappey-Marinier D, Barone P, Dehay C, Toroczkai Z, Knoblauch K, Van Essen DC, Kennedy H. 2014. A Weighted and Directed Interareal Connectivity Matrix for Macaque Cerebral Cortex. *Cerebral Cortex* **24**:17–36. DOI: <https://doi.org/10.1093/cercor/bhs270>

Martin JH, Ghez C. 1999. Pharmacological inactivation in the analysis of the central control of movement. *Journal of Neuroscience Methods* **86**:145–159. DOI: [https://doi.org/10.1016/S0165-0270\(98\)00163-0](https://doi.org/10.1016/S0165-0270(98)00163-0)

Martin Vinck, Renata A. Batista-Brito, Ulf Knoblich, Jessica A. Cardin. 2015. Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding. *Neuron*. DOI: <https://doi.org/10.1016/j.neuron.2015.03.028>, PMID: 25892300

Matveev P, Li AJ, Ye Z, Bowen AJ, Opitz-Araya X, Ting JT, Steinmetz NA. 2024. Simultaneous mesoscopic measurement and manipulation of mouse cortical activity. DOI: <https://doi.org/10.1101/2024.11.01.621418>

McGinley MJ, Vinck M, Reimer J, Batista-Brito R, Zagha E, Cadwell CR, Tolias AS, Cardin JA, McCormick DA. 2015. Waking State: Rapid Variations Modulate Neural and Behavioral Responses. *Neuron* **87**:1143–1161. DOI: <https://doi.org/10.1016/j.neuron.2015.09.012>

Merzenich MM, Knight PL, Roth GL. 1975. Representation of cochlea within primary auditory cortex in the cat. *Journal of Neurophysiology* **38**:231–249. DOI: <https://doi.org/10.1152/jn.1975.38.2.231>, PMID: 1092814

Mohajerani MH, McVea DA, Fingas M, Murphy TH. 2010. Mirrored Bilateral Slow-Wave Cortical Activity within Local Circuits Revealed by Fast Bihemispheric Voltage-Sensitive Dye Imaging in Anesthetized and Awake Mice. *Journal of Neuroscience* **30**:3745–3751. DOI: <https://doi.org/10.1523/JNEUROSCI.6437-09.2010>, PMID: 20220008

Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A. 2014. Information-limiting correlations. *Nature Neuroscience* **17**:1410–1417. DOI: <https://doi.org/10.1038/nn.3807>

Mountcastle VB. 1957. Modality and topographic properties of single neurons of cat's somatic sensory cortex. *Journal of Neurophysiology* **20**:408–434. DOI: <https://doi.org/10.1152/jn.1957.20.4.408>, PMID: 13439410

Murray JD, Bernacchia A, Freedman DJ, Romo R, Wallis JD, Cai X, Padoa-Schioppa C, Pasternak T, Seo H, Lee D, Wang X-J. 2014. A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience* **17**:1661–1663. DOI: <https://doi.org/10.1038/nn.3862>

Musall S, Kaufman MT, Juavinett AL, Gluf S, Churchland AK. 2019. Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience* **22**:1677–1686. DOI: <https://doi.org/10.1038/s41593-019-0502-4>

Nguyen ND, Lutas A, Amsalem O, Fernando J, Ahn AY-E, Hakim R, Vergara J, McMahon J, Dimidschstein J, Sabatini BL, Andermann ML. 2024. Cortical reactivations predict future sensory responses. *Nature* **625**:110–118. DOI: <https://doi.org/10.1038/s41586-023-06810-1>

Niell CM, Stryker MP. 2010. Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. *Neuron* **65**:472–479. DOI: <https://doi.org/10.1016/j.neuron.2010.01.033>

Oh SW, Harris JA, Ng L, Winslow B, Cain N, Mihalas S, Wang Q, Lau C, Kuan L, Henry AM, Mortrud MT, Ouellette B, Nguyen TN, Sorensen SA, Slaughterbeck CR, Wakeman W, Li Y, Feng

D, Ho A, Nicholas E, Hirokawa KE, Bohn P, Joines KM, Peng H, Hawrylycz MJ, Phillips JW, Hohmann JG, Wohnoutka P, Gerfen CR, Koch C, Bernard A, Dang C, Jones AR, Zeng H. 2014. A mesoscale connectome of the mouse brain. *Nature* **508**:207–214. DOI: <https://doi.org/10.1038/nature13186>

Okun M, Steinmetz NA, Lak A, Dervinis M, Harris KD. 2019. Distinct Structure of Cortical Population Activity on Fast and Infralow Timescales. *Cerebral Cortex* **29**:2196–2210. DOI: <https://doi.org/10.1093/cercor/bhz023>

Orbach HS, Cohen LB, Grinvald A. 1985. Optical mapping of electrical activity in rat somatosensory and visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* **5**:1886–1895. DOI: <https://doi.org/10.1523/JNEUROSCI.05-07-01886.1985>, PMID: 4020423

Otchy TM, Wolff SBE, Rhee JY, Pehlevan C, Kawai R, Kempf A, Gobes SMH, Ölveczky BP. 2015. Acute off-target effects of neural circuit manipulations. *Nature* **528**:358–363. DOI: <https://doi.org/10.1038/nature16442>

Ottenheimer DJ, Hjort MM, Bowen AJ, Steinmetz NA, Stuber GD. 2023. A stable, distributed code for cue value in mouse cortex during reward learning. *eLife* **12**. DOI: <https://doi.org/10.7554/eLife.84604.2>

Pachitariu M, Sridhar S, Pennington J, Stringer C. 2024. Spike sorting with Kilosort4. *Nature Methods* **21**:914–921. DOI: <https://doi.org/10.1038/s41592-024-02232-7>

Packer AM, Russell LE, Dalgleish HWP, Häusser M. 2015. Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo. *Nature Methods* **12**:140–146. DOI: <https://doi.org/10.1038/nmeth.3217>

Pandarínath C, O'Shea DJ, Collins J, Jozefowicz R, Stavisky SD, Kao JC, Trautmann EM, Kaufman MT, Ryu SI, Hochberg LR, Henderson JM, Shenoy KV, Abbott LF, Sussillo D. 2018. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods* **15**:805–815. DOI: <https://doi.org/10.1038/s41592-018-0109-9>

Paninski L. 2004. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems* **15**:243. DOI: <https://doi.org/10.1088/0954-898X/15/4/002>

Park IM, Meister MLR, Huk AC, Pillow JW. 2014. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature Neuroscience* **17**:1395–1403. DOI: <https://doi.org/10.1038/nn.3800>

Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A, Köpf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J, Chintala S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. DOI: <https://doi.org/10.48550/arXiv.1912.01703>

Pattadkal JJ, O'Shea RT, Hansel D, Taillefumier T, Brager D, Priebe NJ. 2024. Synchrony dynamics underlie irregular neocortical spiking. DOI: <https://doi.org/10.1101/2024.10.15.618398>

Peters AJ, Fabre JMJ, Steinmetz NA, Harris KD, Carandini M. 2021. Striatal activity topographically reflects cortical activity. *Nature* **591**:420–425. DOI: <https://doi.org/10.1038/s41586-020-03166-8>

Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. 2008. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**:995–999. DOI: <https://doi.org/10.1038/nature07140>

Pinto L, Goard MJ, Estandian D, Xu M, Kwan AC, Lee S-H, Harrison TC, Feng G, Dan Y. 2013. Fast modulation of visual perception by basal forebrain cholinergic neurons. *Nature Neuroscience* **16**:1857–1863. DOI: <https://doi.org/10.1038/nn.3552>

Pinto L, Rajan K, DePasquale B, Thiberge SY, Tank DW, Brody CD. 2019. Task-Dependent Changes in the Large-Scale Dynamics and Necessity of Cortical Regions. *Neuron* **104**:810–824.e9. DOI: <https://doi.org/10.1016/j.neuron.2019.08.025>, PMID: 31564591

Rabinowitz NC, Goris RL, Cohen M, Simoncelli EP. 2015. Attention stabilizes the shared gain of V4 populations. *eLife* **4**:e08998. DOI: <https://doi.org/10.7554/eLife.08998>

Raut RV, Rosenthal ZP, Wang X, Miao H, Zhang Z, Lee J-M, Raichle ME, Bauer AQ, Brunton SL, Brunton BW, Kutz JN. 2025. Arousal as a universal embedding for spatiotemporal brain dynamics. *Nature* **647**:454–461. DOI: <https://doi.org/10.1038/s41586-025-09544-4>

Reimer J, McGinley MJ, Liu Y, Rodenkirch C, Wang Q, McCormick DA, Tolias AS. 2016. Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. *Nature Communications* **7**:13289. DOI: <https://doi.org/10.1038/ncomms13289>

Ren J, Isakova A, Friedmann D, Zeng J, Grutzner SM, Pun A, Zhao GQ, Kolluru SS, Wang R, Lin R, Li P, Li A, Raymond JL, Luo Q, Luo M, Quake SR, Luo L. 2019. Single-cell transcriptomes and whole-brain projections of serotonin neurons in the mouse dorsal and median raphe nuclei. *eLife* **8**:e49424. DOI: <https://doi.org/10.7554/eLife.49424>

Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, Harris KD. 2010. The Asynchronous State in Cortical Circuits. *Science* **327**:587–590. DOI: <https://doi.org/10.1126/science.1179850>

Rohner VL, Curreli S, Lamothe-Molina PJ, Kagiampaki Z, Yee AG, Nardin C, Eschholz L, Dieter A, Foustoukos G, Milanese P, Banterle L, Childs T, Canziani A, Dernic J, Ravotto L, Bhat MA, Sönmez L, Wasielewski LM, Ziebarth T, Molnar MI, Huang G, Maseck OA, Weber B, Reiner A, Lüthi A, Ford CP, Wiegert JS, Ruediger S, Fellin T, Patriarchi T. 2026. Next-generation multicolor indicators for in vivo imaging of norepinephrine. *Nature Methods* **23**:636–652. DOI: <https://doi.org/10.1038/s41592-026-03006-z>

Roth N, Chapuis G, Winter O, Laboratory IB, Ressmeyer RA, Bun LM, Canfield RA, Horwitz GD, Steinmetz NA. 2026. A flexible quality metric for electrophysiological recordings across brain regions and species. DOI: <https://doi.org/10.64898/2026.03.06.710130>

Ruff DA, Cohen MR. 2014. Attention can either increase or decrease spike count correlations in visual cortex. *Nature Neuroscience* **17**:1591–1597. DOI: <https://doi.org/10.1038/nn.3835>

Rumyantsev OI, Lecoq JA, Hernandez O, Zhang Y, Savall J, Chrapkiewicz R, Li J, Zeng H, Ganguli S, Schnitzer MJ. 2020. Fundamental bounds on the fidelity of sensory cortical coding. *Nature* 1–6. DOI: <https://doi.org/10.1038/s41586-020-2130-2>

Rupasinghe A, Charles AS, Pillow JW. 2026. Continuous partitioning of neuronal variability. *eLife* **15**. DOI: <https://doi.org/10.7554/eLife.109719.1>

Rynes ML, Surinach DA, Linn S, Laroque M, Rajendran V, Dominguez J, Hadjistamoulou O, Navabi ZS, Ghanbari L, Johnson GW. 2021. Miniaturized head-mounted microscope for whole-cortex mesoscale imaging in freely behaving mice. *Nature Methods* **18**:417–425.

Saalmann YB, Pinsk MA, Wang L, Li X, Kastner S. 2012. The Pulvinar Regulates Information Transmission Between Cortical Areas Based on Attention Demands. *Science* **337**:753–756. DOI: <https://doi.org/10.1126/science.1223082>

Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP. 2014. Neural constraints on learning. *Nature* **512**:423–426. DOI: <https://doi.org/10.1038/nature13665>

Scott BB, Thiberge SY, Guo C, Tervo DGR, Brody CD, Karpova AY, Tank DW. 2018. Imaging cortical dynamics in GCaMP transgenic rats with a head-mounted widefield macroscope. *Neuron* **100**:1045–1058.

Semedo JD, Zandvakili A, Machens CK, Yu BM, Kohn A. 2019. Cortical Areas Interact through a Communication Subspace. *Neuron* **102**:249–259.e4. DOI: <https://doi.org/10.1016/j.neuron.2019.01.026>

Shadlen MN, Newsome WT. 1998. The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *Journal of Neuroscience* **18**:3870–3896. DOI: <https://doi.org/10.1523/JNEUROSCI.18-10-03870.1998>, PMID: 9570816

Shenoy KV, Carmena JM. 2014. Combining Decoder Design and Neural Adaptation in Brain-Machine Interfaces. *Neuron* **84**:665–680. DOI: <https://doi.org/10.1016/j.neuron.2014.08.038>, PMID: 25459407

Sherman SM. 2016. Thalamus plays a central role in ongoing cortical functioning. *Nature Neuroscience* **19**:533–541. DOI: <https://doi.org/10.1038/nn.4269>

Shimaoka D, Steinmetz NA, Harris KD, Carandini M. 2019. The impact of bilateral ongoing activity on evoked responses in mouse cortex. *eLife* **8**:e43533. DOI: <https://doi.org/10.7554/eLife.43533>

Siegle JH, Steinmetz NA. 2026. Large-scale electrophysiology at single-spike resolution. *Nature Reviews Neuroscience* 1–21. DOI: <https://doi.org/10.1038/s41583-026-01042-4>

Song M, Shin EJ, Seo H, Soltani A, Steinmetz NA, Lee D, Jung MW, Paik S-B. 2024. Hierarchical gradients of multiple timescales in the mammalian forebrain. *Proceedings of the National Academy of Sciences* **121**:e2415695121. DOI: <https://doi.org/10.1073/pnas.2415695121>

Speed A, Haider B. 2021. Probing mechanisms of visual spatial attention in mice. *Trends in Neurosciences* **44**:822–836. DOI: <https://doi.org/10.1016/j.tins.2021.07.009>, PMID: 34446296

Srinath R, Ruff DA, Cohen MR. 2021. Attention improves information flow between neuronal populations without changing the communication subspace. *Current biology: CB* **31**:5299–5313.e4. DOI: <https://doi.org/10.1016/j.cub.2021.09.076>, PMID: 34699782

Steinmetz NA, Aydin C, Lebedeva A, Okun M, Pachitariu M, Bauza M, Beau M, Bhagat J, Böhm C, Broux M, Chen S, Colonell J, Gardner RJ, Karsh B, Kloosterman F, Kostadinov D, Mora-Lopez C, O'Callaghan J, Park J, Putzeys J, Sauerbrei B, van Daal RJJ, Vollen AZ, Wang S, Welkenhuysen M, Ye Z, Dudman JT, Dutta B, Hantman AW, Harris KD, Lee AK, Moser EI, O'Keefe J, Renart A, Svoboda K, Häusser M, Haesler S, Carandini M, Harris TD. 2021. Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* **372**:eabf4588. DOI: <https://doi.org/10.1126/science.abf4588>

Steinmetz NA, Zatzka-Haas P, Carandini M, Harris KD. 2019. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**:266–273. DOI: <https://doi.org/10.1038/s41586-019-1787-x>

Stringer C, Pachitariu M, Steinmetz N, Carandini M, Harris KD. 2019a. High-dimensional geometry of population responses in visual cortex. *Nature* **571**:361–365. DOI: <https://doi.org/10.1038/s41586-019-1346-5>

Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD. 2019b. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**. DOI: <https://doi.org/10.1126/science.aav7893>, PMID: 31000656

Syeda A, Zhong L, Tung R, Long W, Pachitariu M, Stringer C. 2024. Facemap: a framework for modeling neural activity based on orofacial tracking. *Nature Neuroscience* **27**:187–195. DOI: <https://doi.org/10.1038/s41593-023-01490-6>

Tolhurst DJ, Movshon JA, Dean AF. 1983. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision research* **23**:775–785.

Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. 2005. A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects. *Journal of Neurophysiology* **93**:1074–1089. DOI: <https://doi.org/10.1152/jn.00697.2004>

Vaidya AR, Pujara MS, Petrides M, Murray EA, Fellows LK. 2019. Lesion Studies in Contemporary Neuroscience. *Trends in Cognitive Sciences* **23**:653–671. DOI: <https://doi.org/10.1016/j.tics.2019.05.009>

van Vreeswijk C, Sompolinsky H. 1996. Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity. *Science* **274**:1724–1726. DOI: <https://doi.org/10.1126/science.274.5293.1724>

Vanni MP, Murphy TH. 2014. Mesoscale Transcranial Spontaneous Activity Mapping in GCaMP3 Transgenic Mice Reveals Extensive Reciprocal Connections between Areas of Somatomotor Cortex. *Journal of Neuroscience* **34**:15931–15946. DOI: <https://doi.org/10.1523/JNEUROSCI.1818-14.2014>, PMID: 25429135

Veuthay TL, Derosier K, Kondapavulur S, Ganguly K. 2020. Single-trial cross-area neural population dynamics during long-term skill learning. *Nature Communications* **11**:4057. DOI: <https://doi.org/10.1038/s41467-020-17902-1>

von Monakow C. 1914. Die lokalisation im grosshirn und der abbau der funktion durch kortikale herde. J. F. Bergmann.

Wang T, Pachitariu M, Guo C, Lee AK, Lee JS, Zhang Y, Rozsa M, Hasseman J, Kolb I, Cox JD, Lopez C, Looger L, Svoboda K. 2023. **jGCaMP8 transgenic mice**. DOI: <https://doi.org/10.25378/JANELIA.24565837.V3>

Weiler S, Velez-Fort M, Margrie TW. 2024. Overcoming off-target optical stimulation-evoked cortical activity in the mouse brain *in vivo*. *iScience* **27**:111152. DOI: <https://doi.org/10.1016/j.isci.2024.111152>

Wekselblatt JB, Flister ED, Piscopo DM, Niell CM. 2016. Large-scale imaging of cortical dynamics during sensory perception and behavior. *APSselect* **3**:2852–2866. DOI: <https://doi.org/10.1152/jn.01056.2015@apsselect.2016.3.issue-4>

Xiao D, Vanni MP, Mitelut CC, Chan AW, LeDue JM, Xie Y, Chen AC, Swindale NV, Murphy TH. 2017. Mapping cortical mesoscopic networks of single spiking cortical or sub-cortical neurons. *eLife* **6**:e19976. DOI: <https://doi.org/10.7554/eLife.19976>

Ye Z, Ladd AE, MacKenzie N, Kolich L, Li AJ, Birman D, Bull MS, Daigle TL, Tasic B, Zeng H, Steinmetz NA. 2026. Brain-wide topographic coordination of rotating waves. *Science* **392**:eadx1369. DOI: <https://doi.org/10.1126/science.adx1369>

Yin C, Melin MD, Rojas-Bowe G, Sun XR, Couto J, Gluf S, Kostiuk A, Musall S, Churchland AK. 2025. Spontaneous movements and their relationship to neural activity fluctuate with latent engagement states. *Neuron* **0**. DOI: <https://doi.org/10.1016/j.neuron.2025.06.001>, PMID: 40602403

Yizhar O, Fenno LE, Prigge M, Schneider F, Davidson TJ, O'Shea DJ, Sohal VS, Goshen I, Finkelstein J, Paz JT, Stehfest K, Fudim R, Ramakrishnan C, Huguenard JR, Hegemann P, Deisseroth K. 2011. Neocortical excitation/inhibition balance in information processing and social dysfunction. *Nature* **477**:171–178. DOI: <https://doi.org/10.1038/nature10360>

Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. 2009. Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity. *Journal of Neurophysiology* **102**:614–635. DOI: <https://doi.org/10.1152/jn.90941.2008>

Záborszky L, Gombkoto P, Varsanyi P, Gielow MR, Poe G, Role LW, Ananth M, Rajebhosale P, Talmage DA, Hasselmo ME, Dannenberg H, Mincses VH, Chiba AA. 2018. Specific Basal Forebrain–Cortical Cholinergic Circuits Coordinate Cognitive Operations. *Journal of Neuroscience* **38**:9446–9458. DOI: <https://doi.org/10.1523/JNEUROSCI.1676-18.2018>, PMID: 30381436

Zatka-Haas P, Steinmetz NA, Carandini M, Harris KD. 2021. Sensory coding and the causal impact of mouse cortex in a visual decision. *eLife* **10**:e63163. DOI: <https://doi.org/10.7554/eLife.63163>

Zeraati R, Levina A, Macke JH, Gao R. 2026. Neural timescales from a computational perspective. *Nature Neuroscience* **29**:1534–1547. DOI: <https://doi.org/10.1038/s41593-026-02343-8>

Zhang S, Xu M, Chang W-C, Ma C, Hoang Do JP, Jeong D, Lei T, Fan JL, Dan Y. 2016. Organization of long-range inputs and outputs of frontal cortex for top-down control. *Nature Neuroscience* **19**:1733–1742. DOI: <https://doi.org/10.1038/nn.4417>

Zhang S, Xu M, Kamigaki T, Hoang Do JP, Chang W-C, Jenvay S, Miyamichi K, Luo L, Dan Y. 2014. Long-range and local circuits for top-down modulation of visual cortex processing. *Science* **345**:660–665. DOI: <https://doi.org/10.1126/science.1254126>

Zingg B, Hintiryan H, Gou L, Song MY, Bay M, Bienkowski MS, Foster NN, Yamashita S, Bowman I, Toga AW, Dong H-W. 2014. Neural Networks of the Mouse Neocortex. *Cell* **156**:1096–1111. DOI: <https://doi.org/10.1016/j.cell.2014.02.023>

Zylberberg J, Pouget A, Latham PE, Shea-Brown E. 2017. Robust information propagation through noisy neural circuits. *PLOS Computational Biology* **13**:e1005497. DOI: <https://doi.org/10.1371/journal.pcbi.1005497>