



**UNIVERSITETET
I OSLO**

Computational modeling of electric brain signals across scales

From the spikes of single neurons to the local field
potentials of brain areas

Atle Eskeland Rimehaug

Supervisors:

Gaute T. Einevoll, Anders M. Dale, and Marianne H. Fyhn

© Atle Eskeland Rimehaug, 2024

*Series of dissertations submitted to the
Faculty of Mathematics and Natural Sciences, University of Oslo
No. 2791*

ISSN 1501-7710

All rights reserved. No part of this publication may be
reproduced or transmitted, in any form or by any means, without permission.

Cover: UiO.

Print production: Graphic center, University of Oslo.

I had always considered my thoughts as something abstract, but they weren't; they were as material as the heart beating in my chest. The same was true of the mind, the soul, the personality; all of it was fixed in the cells and originated as a result of the various ways in which these cells reacted with one another. All of our systems, too — communism, capitalism, religion, science — they also originated in electrochemical currents flowing through this three-pound lump of flesh encased in the skull.

- Karl Ove Knausgård, *The Terrible Beauty of Brain Surgery*

Abstract

The study of electric brain signals with biophysical simulations was born with the seminal work by Hodgkin & Huxley (1952). They characterized mathematically how ionic currents across the neuronal membrane give rise to action potentials - the electrochemical signals fundamental to most communication in the brain. In the decades following their discoveries, meticulous experimental and theoretical work has given us a fairly good understanding of how single cells work. However, it is the cooperative effort of neurons in networks that gives rise to thoughts, memories, emotions, and advanced behaviors of humans and animals. Additionally, the extracellular and extracortical signals we measure in experiments reflect the activity of large populations of neurons, sometimes spanning entire brain regions. Compared to our knowledge of single neurons, our insight into how large numbers of neurons interact and produce electric signals remains underdeveloped.

In recent years, both the amount and quality of experimental data as well as access to computational resources have increased prodigiously, enough to enable large-scale biophysically detailed modeling of brain networks. By simulating the activity of large networks composed of detailed neuron models, we can emulate the signals we observe in experiments stemming from populations, structures, or brain regions.

In this thesis, we have developed, validated, and studied a large-scale biophysically detailed model of mouse primary visual cortex (V1). In our first project, we extended an existing V1 model to quantitatively reproduce experimental brain activity across scales: from the spikes of individual cells to the local field potentials (LFPs) of the structure as a whole. We demonstrated that constraints on model architecture and parameters are enhanced by requiring a model to reproduce both signals simultaneously, and that spikes and LFP therefore represent complementary aspects of neural activity.

In the subsequent projects, we used data simulated with the V1 model to validate two algorithms for estimating LFPs from presynaptic firing rates. In both methods, presynaptic firing rates are convolved with kernels. The purpose is to approximate the process by which postsynaptic LFP is generated by presynaptic spikes. The first method, named Laminar Population Analysis (LPA), estimated the salient features in the LFP generated by external inputs to V1 and from

the recurrent activity in V1 as a whole. This indicates that the LPA method can potentially be used to disentangle LFP contributions from different presynaptic structures in experimental data.

In the second method, more detailed kernels are constructed by utilizing information about membrane potentials, synaptic parameters, and connection patterns in the model. We validated this framework on layer 2/3 (L2/3) in the V1 model, and found that we could precisely estimate the LFP contributions from external structures as well as different internal populations of L2/3. The detailed information required to construct these kernels is typically not experimentally available together with LFP recordings, which may limit its application to experimental data at present. However, the precision with which it could estimate LFP suggests that it can in some cases be used as a substitute for computationally expensive full-scale simulations.

The overarching ambition of these projects was to contribute to the foundation of a bridge between our knowledge of single neurons and our knowledge of populations of neurons and the electric signals they produce. We hope that this bridge, composed of computational building blocks, may both motivate future experiments and provide greater insight into the neural mechanisms behind already collected data. We look forward to both extensions and applications of this work in the coming years.

Sammendrag

Studiet av elektriske hjernesignaler med biofysiske simuleringer begynte med det banebrytende arbeidet til Hodgkin & Huxley (1952). De karakteriserte matematisk hvordan ionestrømmer på tvers av nervemembranen gir opphav til aksjonspotensialer - de elektrokjemiske signalene som er grunnleggende for mesteparten av kommunikasjon i hjernen. I tiårene etter oppdagelsene deres har omhyggelig eksperimentelt og teoretisk arbeid gitt oss en god forståelse for hvordan enkeltnevroner fungerer. Det er imidlertid samarbeidet mellom mange nevroner i store og små nettverk som gir opphav til tanker, minner, følelser og avansert atferd hos mennesker og dyr. Videre reflekterer de ekstracellulære og ekstrakortikale signalene vi måler i eksperimenter aktiviteten i store populasjoner av nevroner som noen ganger strekker seg over hele hjerneregioner. Vår forståelse av hvordan et stort antall nevroner samhandler og produserer elektriske signaler er fortsatt underutviklet sammenlignet med vår forståelse av enkeltnevroner.

De siste årene har både mengden og kvaliteten på eksperimentelle data samt tilgang til beregningsressurser økt betraktelig, nok til å muliggjøre storskala, biofysisk detaljert modellering av hjernenettverk. Ved å simulere aktiviteten til store nettverk sammensatt av detaljerte nevronmodeller, kan vi etterligne signalene vi observerer i eksperimenter som stammer fra populasjoner, strukturer eller hjerneregioner.

I denne avhandlingen har vi utviklet, validert og studert en storskala, biofysisk detaljert modell av primær visuell cortex (V1) hos mus. I vårt første prosjekt utvidet vi en eksisterende V1-modell til den reproduserte eksperimentell hjerneaktivitet på tvers av skalaer: både aksjonspotensial fra enkeltceller og lokale feltpotensial (LFP) fra strukturen som helhet. Vi demonstrerte at modellarkitektur og parametere blir bedre begrenset ved å kreve at en modell reproduserer begge signalene samtidig, og at aksjonspotensialer og LFP derfor representerer komplementære aspekter ved nevralt aktivitet.

I de påfølgende prosjektene brukte vi data simulert med V1-modellen til å validere to algoritmer for estimering av LFP fra presynaptiske fyringsrater. I begge metodene blir presynaptiske fyringsrater konvolvert med kjerner. Hensikten er å tilnærme prosessen hvor postsynaptisk LFP genereres av presynaptiske aksjonspotensial. Den første metoden, kalt Laminar Population Analysis (LPA), estimerte de fremtredende kjennetegnene i LFP generert fra eksternt input til V1 og fra den tilbakevendende aktiviteten i hele V1. Dette

indikerer at LPA-metoden kan potensielt brukes til å skille LFP-bidrag fra forskjellige presynaptiske strukturer i eksperimentelle data.

I den andre metoden konstrueres mer detaljerte kjerner ved å bruke informasjon om membranpotensial, synaptiske parametere og koblingsmønstre fra modellen. Vi validerte dette rammeverket på lag 2/3 (L2/3) i V1-modellen, og fant at vi kunne estimere LFP-bidragene fra eksterne strukturer så vel som forskjellige intrakortikale populasjoner av L2/3 med høy nøyaktighet. Den detaljerte informasjonen som kreves for å konstruere disse kjernene er vanligvis ikke eksperimentelt tilgjengelig sammen med LFP-opptak, noe som foreløpig kan begrense dens anvendelse på eksperimentelle data. På den andre siden, presisjonen i LFP-estimatene med denne metoden tyder på at den kan brukes som substitutt for å kjøre fullskalasiluleringer, og dermed redusere store mengder kjøretimer.

Den overordnede målet med disse prosjektene var å bidra til å bygge en bro mellom vår kunnskap om enkeltnevroner og vår kunnskap om populasjoner av nerveceller og de elektriske signalene de produserer. Vi håper at denne broen, bestående av komputasjonelle byggematerialer, kan både motivere fremtidige eksperimenter og gi større innsikt i nevralk mekanismer bak de eksperimentelle dataene vi allerede har. Vi ser frem til både anvendelser og utvidelser av dette arbeidet i de kommende årene.

List of papers

Paper I

Rimehaug, A.E., Stasik, A.J., Hagen, E., Billeh, Y.N., Siegle, J.H., Dai, K., Olsen, S.R., Koch, C., Einevoll, G.T. and Arkhipov, A., 2023. Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex. *elife*, 12, p.e87169, doi: <https://doi.org/10.7554/eLife.87169>

Paper II

Rimehaug, A.E., Dale, A.M., Arkhipov, A. and Einevoll, G.T., 2024. Uncovering population contributions to the extracellular potential in the mouse visual system using Laminar Population Analysis. *bioRxiv*, pp.2024-01, doi: <https://doi.org/10.1101/2024.01.15.575805>. *Manuscript under review*.

Paper III

Meneghetti, N., Rimehaug, A.E., Einevoll, G.T., Mazzoni, A., Ness, T.V., 2024. Estimating simulated local field potentials from presynaptic firing rates and network properties. *Manuscript in preparation*.

Additional papers not included in the thesis

Dai, K., Gratiy, S.L., Billeh, Y.N., Xu, R., Cai, B., Cain, N., Rimehaug, A.E., Stasik, A.J., Einevoll, G.T., Mihalas, S. and Koch, C., 2020. Brain Modeling ToolKit: An open source software suite for multiscale modeling of brain circuits. *PLOS Computational Biology*, 16(11), p.e1008386, doi: <https://doi.org/10.1371/journal.pcbi.1008386>.

Contents

Abstract	iv
List of papers	vii
Acknowledgements	xiv
I Introduction	1
1 Background and motivation	3
1.1 The local field potential and the current source density	6
1.2 Biophysical modeling	11
2 Methods	13
2.1 Modeling at different levels of detail	13
2.1.1 Biophysically detailed modeling	14
2.1.2 Point-neurons.	20
2.2 Software	21
II Results and Discussion	23
3 Summary of papers	25
3.1 Paper I: Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex.	25
3.2 Paper II: Uncovering population contributions to the extracellular potential in the mouse visual system using Laminar Population Analysis	27
3.3 Paper III: Estimating simulated local field potentials from presynaptic firing rates and network properties	28
4 Discussion	31
4.1 Validating large-scale models and making testable predictions	32
4.2 Inferring population contributions and the validity of assuming linear LFP generation	35
4.3 Outlook.	38

I Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex	41
II Uncovering population contributions to the extracellular potential in the mouse visual system using Laminar Population Analysis	97
III Estimating simulated local field potentials from presynaptic firing rates and network properties	147

List of Figures

1.1	Recording extracellular potential to obtain MUA and LFP.	5
1.2	CSD illustrations.	9
1.3	LFP from excitatory and inhibitory synapses at different positions on the Hay model of an L5 neuron[41].	10
2.4	Membrane patch as RC circuit.	15
2.5	Multicompartment neurite as RC circuits in series.	16
2.6	Potential measured by extracellular electrode arising from transmembrane currents can be calculated with volume conductor theory.	20

List of Figures

Acknowledgements

I want to begin by thanking my main supervisor, Gaute T. Einevoll, for all the support, guidance, and fascinating conversations and discussions over the past four years. It has been both stimulating and motivating to have a supervisor who shares my own views on the importance of scientific thinking and values outside of labs and journals. Gaute has been a role model for me not only in our scientific endeavors, but also in his efforts to make our debates on social and political issues more informed by communicating and spreading enthusiasm for science to the public.

I would also like to extend my thanks to my co-supervisors Anders M. Dale and Marianne H. Fyhn. Anders for two educational stays at UCSD in California which broadened my horizon and offered formative experiences with both the academic and broader culture in the US. Marianne for granting me flexibility in how and where my teaching duties were carried out. This allowed me to teach interesting topics relevant to my research and helped to make my teaching experience both rewarding and thoroughly enjoyable.

I want to express my deep gratitude to Anton Arkhipov at the Allen Institute for all his advice, scientific insight, and incisive comments. Anton has been a co-supervisor in all but the name throughout my PhD, and has taught me rigor, attention to detail, and how to reason clearly and precisely about ideas and evidence; skills that have made me a better scientist. I also want to thank Christof Koch, Joshua H. Siegle, Yazan N. Billeh, Kael Dai, and others at the Allen Institute for including me in their ambitious scientific pursuits. This thesis would not have been attainable without their ideas, suggestions, assistance, prompt answers to my questions and inquiries, and numerous other contributions.

I want to thank my parents, my siblings, and all my nephews and nieces for their unconditional love and support. Their kindness has been a source of security and confidence and has given me the strength to carry on in challenging periods. They have all, time and again, brought me back to the world outside of simulations, data, papers, and review processes and enriched my life in ways that science, for all its fascinating facts and exhilaration of discovery, can never do.

A special thanks to Kirubel Danieli, Kosio Beshkov, Frederik Rogge, Maria Perona Fjeldstad, Jonas Verhellen, Elise Reppe Olsen, my dad, Nicolò Meneghetti, and Torbjørn V. Ness for proofreading a part of my thesis each. The blame for any remaining typos, grammatical errors, or other gaffes should now be laid at their feet, not mine. I am also very grateful to Nicolò and Torbjørn for involving me in their project and our engaging discussions in the last few months of the PhD. Jonas, Elise, Sara Kriebel Thon, and Maria have been excellent office mates these four

List of Figures

years, sharing ups and downs, fortunes and mishaps, and many amusing moments of everyday life. I want to thank all other members of CINPLA for creating an inclusive and friendly workplace where it was always pleasant to be.

Thanks to Kimberly McGabe for all assistance and encouragement, both as the coordinator for the SUURPh-program and as a recent PhD student who could see things from our perspective and help us overcome the inevitable problems, large and small, that come with doing a PhD. Her organizational skills have alleviated many of our burdens and made both the SUURPh-experience and the PhD itself easier and more comfortable.

Thanks to my fellow SUURPh-students for all the hikes, trips, 17th of May celebrations, dinners, and zoom-socials during the pandemic. The stay in San Diego was made many times more fun and exciting with you guys. Thanks to all my other friends who have filled the last few years with great conversations and experiences that have brought joy, comfort, ideas and perspectives, fascination, and, sometimes most importantly, laughter.

Lastly, I want to highlight that this thesis was made possible by the generous financial support of the SUURPh program through Simula and of the Department of Informatics at the University of Oslo. The research was conducted at the Center for Integrative Neuroplasticity (CINPLA) at the University of Oslo and at the University of California San Diego.

Part I

Introduction

Chapter 1

Background and motivation

The quote from Norwegian author Karl Ove Knausgård prefacing this thesis beautifully expresses the idea that it is not only the world around us that arises from material, physical processes, so does the world within us. There is, however, another idea beyond the one that meets the reader first captured in this quote: The idea that there is a scaling process in nature where the events at the small scale dictate what happens at a bigger scale, which in turn dictate what happens on the largest scale. There is a thread weaved by the laws of nature connecting the world of atoms and elementary particles to the world of animals, humans and their minds, and finally the planets, stars, and galaxies.

To make problems more tractable, scientific studies are often limited to a single scale. For many research questions, the information at a single scale is indeed sufficient to discover the relevant explanatory mechanisms; but there are also problems for which an understanding of the microscopic is needed to comprehend the mesoscopic or macroscopic. To build the modern phone, we first had to understand how electrons move in electrical circuits, and then how electrical circuits can be constructed to transmit signals that are coded, received, and deciphered into human comprehensible sound or text. To predict the behavior of weather systems, we first had to understand the phase transitions of water, and then how these interact with wind currents and atmospheric pressure to produce clouds and precipitation.

Likewise, to achieve a full understanding of the workings of the brain and the signals we measure in experimental neuroscience, we need to understand how the mechanisms and signals on the larger scale of structures and brain regions, or the whole brain, arise from interactions on the smaller scale of single neurons, and ultimately from ions gushing in and out of neurons and the extracellular space [20].

Neural activity can be measured from whole brain regions in a multitude of ways, from the hemodynamic responses in functional magnetic resonance imaging (fMRI), to the electrical or magnetic fluctuations in electroencephalograms (EEG) and magnetic encephalograms (MEG) [10], to name a few. This thesis is about the electric signals of the brain, so the measurement modality most relevant to exemplify the application of the signals studied here is the EEG. EEG measures the electric potential arising from electric currents in the brain tissue near the scalp via electrode contacts placed on the surface of the head [34]. Some advantages

afforded by EEG are that it can provide information about neural activity with high temporal resolution and that it records this activity non-invasively, which makes it suitable for recording brain activity in humans. However, due to the large number of neurons contributing to the potential at each electrode contact and the spatial averaging that happens as the electrical potential passes through the skull, the spatial resolution of the signal and the information about neural mechanisms that can be inferred from it is limited [34].

Instead of recording the electric potential with electrodes on the outside of the brain, we can insert a probe with electrodes into the brain tissue and measure the extracellular potential inside the brain [11, 46]. This produces a signal with higher spatial resolution that can both be more informative and may enable us to gain more insight into the neural and electric origins of EEG. As this is an invasive recording, it is typically done in animals, but it can also be performed in humans [74]. The high- and low-frequency components of the recorded extracellular potential are considered to reflect complementary aspects of neural activity (illustrated in Fig. 1.1). The high-frequency (above a few hundred Hz) component is thought to primarily contain information about the spiking activity of neurons, and is referred to as the multi-unit activity (MUA). The low-frequency (below a few hundred Hz) component is thought to mainly stem from transmembrane currents caused by synaptic input to the neurons and is referred to as the local field potential (LFP) [12, 19]. Thus, since the extracellular potential carries information both about the spiking activity of single cells and the ionic currents entering and leaving whole populations of cells, and it also underlies the EEG recorded across whole brain areas¹, it is a signal well suited for studying brain activity across multiple scales.

In single-neuron recordings, the spikes have a single source: a neuron that generates discrete action potentials over time [75]. Assuming an ideal recording with no pollution from other nearby neurons and no mistakes in the spike sorting procedure, spiking activity in the MUA can also be tied to single neurons [97]. (In reality, however, the sorted spikes will often contain spikes or at least have been influenced by activity recorded from other neurons in the vicinity of the electrode contact.) The LFP recorded *in vivo*, on the other hand, will reflect transmembrane currents from many neurons and a multitude of cell types and biophysical processes. In other words, it does not have a single source, it has many. Therefore, interpreting the LFP and uncovering the different origins of the signal can often be complicated [12, 19, 75]. The interpretation was made easier by the introduction of current source density (CSD) analysis of LFP [19]. The CSD is derived from the LFP, but affords a more localized measure of activity that is more readily interpreted in terms of the underlying neurophysiology [67, 72, 78]. A CSD plot provides a map of the different transmembrane currents that underlie the LFP, as it shows where ions enter or leave the extracellular medium.

Another approach to study the LFP that can facilitate its interpretation, is to construct mechanistic models of neurons and networks of neurons that simulate neural activity and the extracellular potential it generates [19, 40]. Mechanistic models are distinguished from statistical or descriptive models in that they are

¹LFP is in fact also known as the intracranial EEG [12].

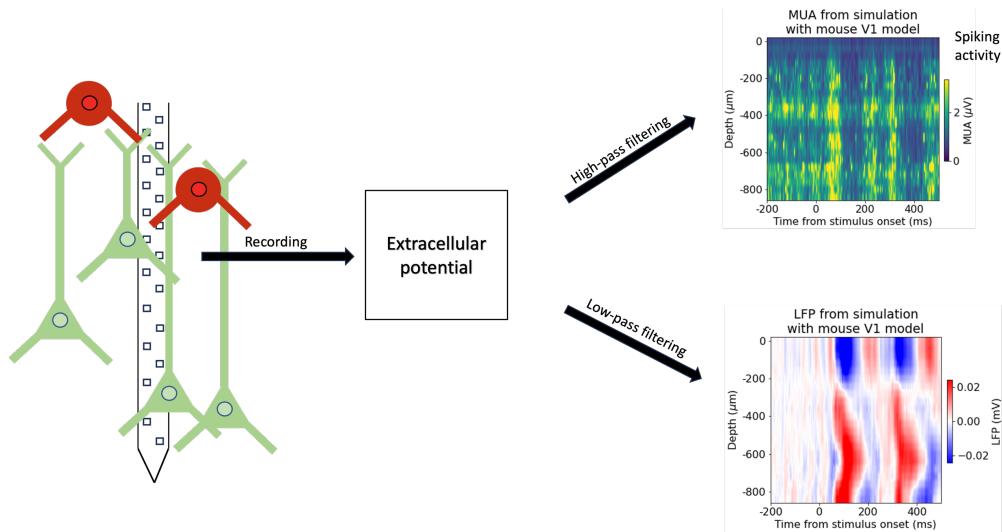


Figure 1.1: Recording extracellular potential to obtain MUA and LFP. Left: Illustration of the recording of extracellular potential arising from neural activity with multielectrode probes. Right: The recorded extracellular potential can be high-pass filtered to obtain the high-frequency part of the signal referred to as the multi-unit activity (MUA) or low-pass filtered to obtain the low-frequency part referred to as the local field potential (LFP). The signals in these example plots have been obtained from a simulation of the response to a full-field flash stimulus in a large-scale, biophysically detailed mouse V1 model.

built bottom-up from first principles, and mechanisms are modeled explicitly rather than given some abstract representation. The aim is to identify how a phenomenon is explained by physical processes [35]. In the context of modeling extracellular potentials, this means that the purpose is to explain the extracellular potentials by the ionic currents in the brain tissue and the measurement physics. This is achieved with biophysically detailed models of neurons - either simulated individually or in a population or network of neurons. This approach enables neuroscientists to establish a connection between the measured signal and the neural activity, biophysical processes, and circuit mechanisms underlying them.

Mechanistic modeling began with the pioneering work by Alan Hodgkin and Andrew Huxley (1952) [44], where they developed a mathematical model of how action potentials are instigated and propagated in neurons through the movement of ionic currents across the membrane. Rall (1959) [79] then initiated the quantitative framework for characterizing current flow in dendritic trees using cable theory. His work formed the foundation for volume conductor theory, the forward-modeling scheme for calculating extracellular potentials from neuron models with detailed morphologies [40]. Since then, we have achieved a fairly good understanding of how single neurons work [49] and how their activity generates extracellular potentials [45, 55, 77]. However, despite the significant progress on the single cell level, our understanding of how large populations of neurons behave in networks and how their activity produces the electric signals we record *in vivo* remains limited [19].

Utilizing the approach of simulating extracellular potentials from neural activity in mechanistic network models, the aim of this thesis is to increase our insight into how the extracellular electric potentials reflect neural activity. First by reproducing the circuit mechanisms generating extracellular potentials observed in *in vivo* experimental recordings with a large-scale, biophysically detailed model. Then by using data simulated with the validated model to test and develop a tool that allows for automatic discovery of the different presynaptic populations generating the recorded extracellular potentials. Lastly, by validating a method that can significantly reduce the computational resources required to simulate the LFP by using presynaptic firing rates and network properties to estimate the LFP.

1.1 The local field potential and the current source density

The LFP has been used to investigate sensory processing [6, 7, 47, 65, 66, 69, 73, 81, 90], motor planning [86, 88] navigation [22, 23, 59, 102, 106], and higher cognitive processing [51, 54, 57, 76, 107]. Furthermore, its relative stability in chronic recordings makes it a promising candidate signal for steering neuroprosthetic devices [1, 60, 64, 84, 94]. However, despite its broad application, the biophysical origins of LFPs are still not fully understood. In the following, we provide a brief overview of current knowledge of the LFP and related signals.

All transmembrane currents contribute to the generation of extracellular potentials [12]. How much transmembrane currents from different biophysical origins contribute depends on what part of the extracellular potential is studied and the specific conditions during which the membrane potential is recorded. At higher frequencies (above a few hundred Hz), the extracellular potential is predominantly made up of the transmembrane currents stemming from action potentials. At the lower frequencies of the LFP, the following biophysical processes are the primary sources of transmembrane currents underlying the potential [12, 36]:

- Synaptic currents and their associated return currents.
- Calcium spikes. Voltage-gated calcium channels are opened during membrane depolarization, leading to an influx of Ca^{2+} ions.
- Intrinsic currents and resonances can affect the membrane potential, and in certain conditions where it occurs synchronously in many nearby neurons, it can have a significant effect on the LFP.
- Membrane currents in glial cells.
- Diffusion potentials. If there are concentration differences of ions in the extracellular space, that will lead to extracellular diffusion currents, which contribute to the potential.
- Action potential currents. Even though the currents through active sodium and potassium channels during an action potential are thought to primarily

be visible in the high-frequency part of the extracellular potential, they can also shape the LFP [82, 85].

Though the other sources can contribute to the shape and magnitude of the LFP, the currents caused by synaptic input are still thought to exert the greatest influence on the LFP in most physiological conditions [19, 36].

The relationship between the electric potential and the ionic currents that generate it is governed by Maxwell's equations of electromagnetism, and the relationship can be expressed as:

$$\nabla \cdot (\sigma \nabla \phi) = -C \quad (1.1)$$

where ϕ is the electric potential, σ is the extracellular conductivity tensor, and C is the current source density (CSD) [78]. σ reflects how easily ions can move through the extracellular medium, and if the extracellular medium is sufficiently isotropic and homogeneous, σ can be treated as a constant. In these conditions, the solution for the potential generated by a single point source becomes:

$$\phi(r) = \frac{I_e}{4\pi\sigma r} \quad (1.2)$$

where I_e is the current from the point source and r is the radial distance from the point source to the point at which the potential is measured. Currents from multiple sources sum linearly to form the total electrical potential at the point r .

It is often easier to make inferences about the biophysical processes producing the transmembrane currents underlying the LFP if one examines the CSD instead of the LFP. Expanding eq. 1.1, we have that:

$$\sigma \left(\frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} + \frac{\partial^2 \phi}{\partial z^2} \right) = -C(x, y, z) \quad (1.3)$$

In most scenarios where the CSD is calculated, the CSD is assumed to be constant laterally to the probe, which means that only the variation in the depth direction (here denoted by z) contributes to the CSD. Eq. 1.3 then simplifies to:

$$\sigma \frac{\partial^2 \phi}{\partial z^2} = -C(z) \quad (1.4)$$

This is the traditional and simplest way of estimating the CSD, and it has often been referred to as the "standard" method. However, in many cases, the planar region of constant CSD is in fact too small for the errors from assuming the CSD to be laterally constant in an infinite plane to be negligible. Thus, calculating the CSD this way can lead to substantial deviations in magnitude from the true CSD as well as spurious sinks and sources near the top and bottom electrodes [78].

Attempting to resolve these issues, Pettersen et al. (2006) [78] developed a method for calculating the CSD from the LFP, referred to as the iCSD method, where the "i" stands for "inverse". The method is based on explicit inversion of the electrostatic forward solution characterizing how the extracellular potential arises from currents. Contrary to the standard method, the CSD is only assumed to be constant within cylindrical discs of radius R . In the limit where $R \rightarrow \infty$, the CSD estimated with the iCSD method and the standard method are equal.

The iCSD method comes in three variants. The δ -source iCSD method, which assumes that infinitely thin current source discs at the electrode contacts are the origin of the CSD. The step iCSD method, which assumes that the CSD is step-wise constant between electrode contacts. The spline iCSD method, which assumes that the CSD is smoothly varying from electrode to electrode. In this thesis, the δ -method was utilized, but the step and spline iCSD methods are known to produce similar estimates [78].

The relationship between LFP and CSD is illustrated in Fig. 1.2A. In the top panel, the evoked LFP response recorded in the primary visual cortex (V1) of a mouse from the Visual Coding dataset released by the Allen Institute is displayed [92]. In the panel below, the CSD calculated from this recorded LFP is plotted with the LFP traces in each channel superimposed on top. The blue blobs in this plot represent current sinks, which show where either *positive ions enter* or where *negative ions leave* the cells in this region. The red blobs represent current sources, which show where either *negative ions enter* or *positive ions leave* the cells.

Since CSD plots show where positive ions enter and leave cells, they can provide information about the position of synaptic input. In Fig. 1.2B-C, the CSD calculated from the LFP of a population of L4 excitatory cells in a mouse V1 model [8] is displayed in two scenarios. In Fig. 1.2B, all excitatory synapses impinging upon this population have been placed only on the basal dendrites, while in Fig. 1.2C, all excitatory synapses have been placed only on the apical dendrites. When all excitatory synapses are placed on the basal dendrites, a current sink appears at the bottom of this population with a current source above it. The current sink at the bottom reflects the input from excitatory synapses, which consists of positive ions rushing into the basal dendrites of this population. According to Kirchhoff's current law, the total amount of current entering or leaving neurons has to sum to zero, so if there is a stream of positive ions entering the cells, there has to be a return current of positive ions leaving the cells at the same time. This return current is the origin of the current source observed above the current sink in this simulation. When the excitatory synapses are placed on the apical dendrites, we get a current source at the bottom where the basal dendrites and somata are instead (as well as at the apical tufts at the very top of the apical dendrites), with the current sink appearing in the middle at the apical dendrites, where the excitatory synapses are predominantly placed.

Even though the CSD is easier to interpret in terms of neural activity than the LFP, it is still not straightforward. A current source can stem from return currents associated with excitatory input, as described above, but it can also arise from inhibitory input currents [12, 42]. The effect of inhibitory currents is opposite to the effects of excitatory currents, so an inhibitory current will generate a current source, not a sink (see Fig. 1.3). In a model, we can use knowledge about the synaptic placement and activity to determine the biophysical origin of a current source (or sink). This information is, however, typically not available in experiments, and acquiring it, or other information that allows for disambiguation, is not easy [10, 12]. Because of this, biophysical modeling of LFPs and CSDs has been utilized to facilitate the interpretation of these signals and help uncover their biophysical origins.

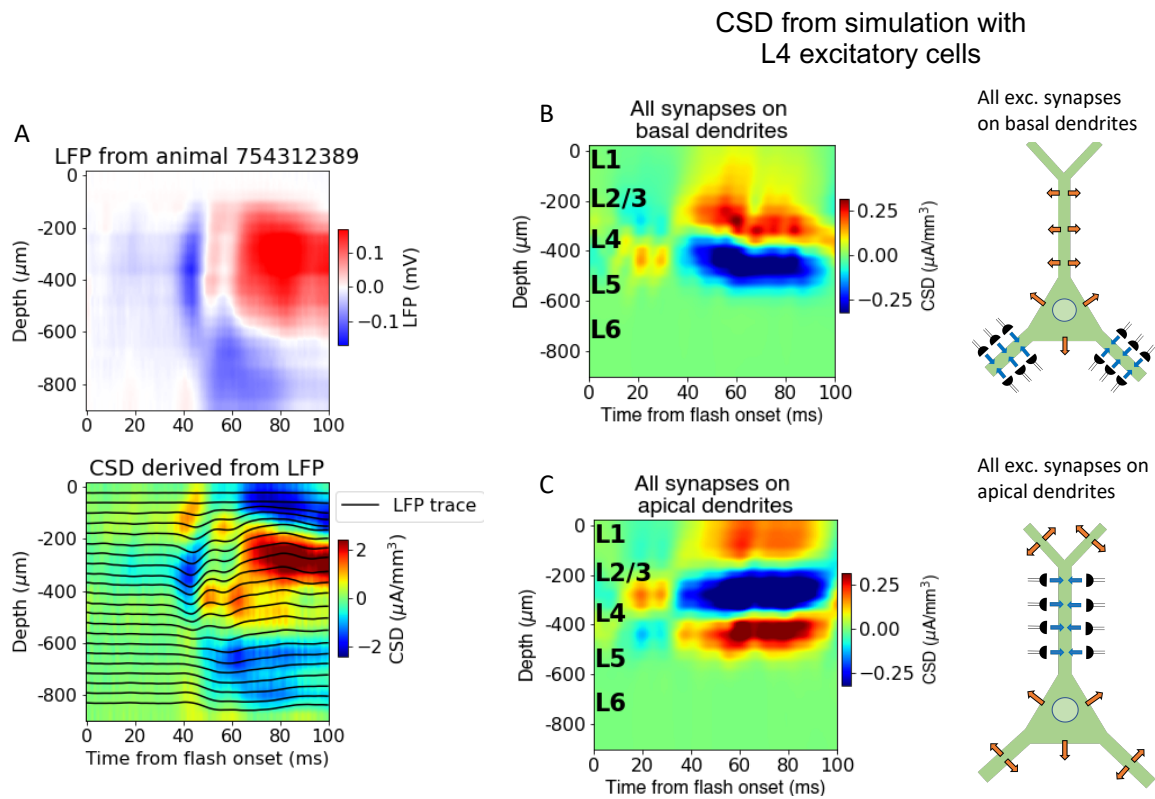


Figure 1.2: CSD illustrations. (A) Top: LFP recorded in V1 of a mouse from the Visual Coding dataset [92] during presentation of a full-field white screen flash, averaged over 75 trials. Bottom: CSD calculated from the recorded LFP using the δ -source iCSD method with CSD assumed to be constant within a radius of $800\ \mu\text{m}$ - roughly equivalent to the size of V1. (B-C) Left: CSD calculated from LFP of layer 4 (L4) excitatory cells from simulations of full-field flash stimuli presented to a modified version of the mouse V1 model developed by the Allen Institute [85] (radius of constant CSD assumed to be $400\ \mu\text{m}$, equal to the spatial extent of the cylinder of biophysically detailed neurons in the model), averaged over 10 trials. Right: illustrations of placement of excitatory synapses onto the L4 excitatory cells and with synaptic input currents shown as blue arrows and return currents as orange arrows. (B) All excitatory synapses placed on basal dendrites. (C) All excitatory synapses placed on apical dendrites. Figure adapted from [85].

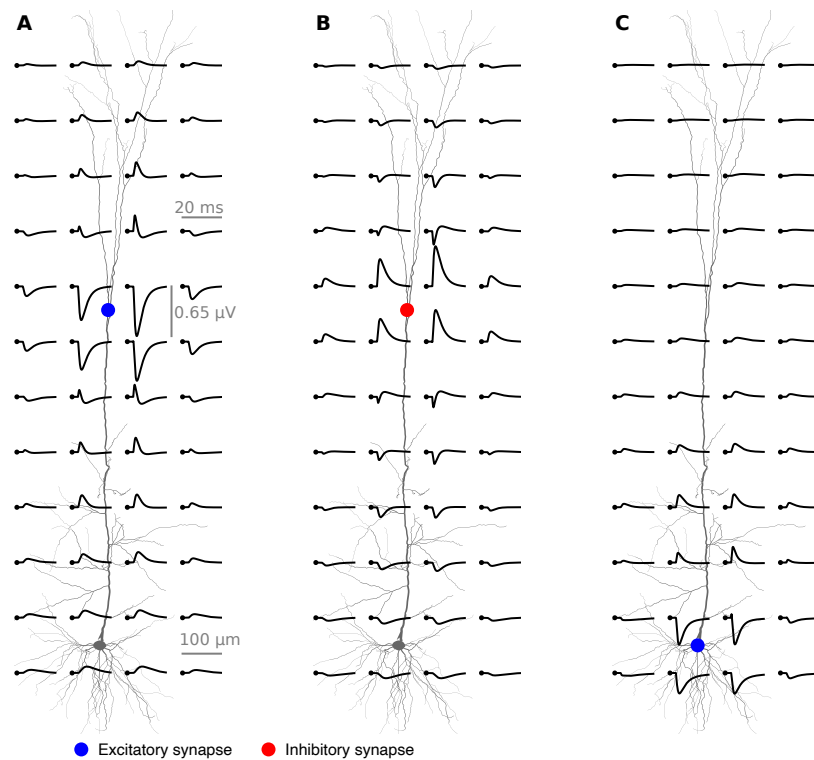


Figure 1.3: LFP from excitatory and inhibitory synapses at different positions on the Hay model of a L5 neuron[41]. (A) Excitatory synapse (blue dot) placed on apical dendrites of the neuron. The positive ions from an excitatory synapse will generate a current sink, which here is visible as a negative deflection in the LFP in the vicinity of the synapse. (B) Inhibitory synapse placed on apical dendrites of the neuron. Inhibitory currents generate a current source, which is visible as a positive deflection in the LFP near the synapse. (C) Excitatory synapse placed at the soma of the neuron. In this case, the negative deflection is observed near the soma. Adapted from figure made by Torbjørn V. Ness & Espen Hagen (2023) in chapter 8: Local Field Potentials in *Electric Brain Signals* [36].

1.2 Biophysical modeling

Neural modeling can roughly be divided into two major approaches [40]. One typically employs more abstract representations of neurons and seeks to identify the general principles and fundamental mechanisms that characterize information processing in the brain. The models in this regime usually put less emphasis on biological and physical realism, prioritizing rather to make models that are more amenable to analytical treatment. Neuron models of this type are often referred to as point-neurons because of their lack of spatial extent; i.e. each neuron only occupies a point in space, and the effects of neural processes in the spatially extended domains of the neurons are modeled implicitly rather than explicitly.

The other approach incorporates more biological detail, and these neuron models are therefore referred to as biophysically detailed models. They are, contrary to point-neurons, spatially extended and thus come with explicit (though simplified to different extents) representations of the neuron morphology, specifically the soma, dendrites, and axons. An important feature of these model types is that they can be used to simulate the extracellular potentials using volume conductor theory (see Chapter 2: Methods). Extracellular potentials arise when there is a spatial separation of currents, and since all currents are gathered in a single point for point-neurons, they cannot generate potentials. Point-neuron models can be used to approximate potentials, but only under certain assumptions [31, 63]. The higher level of detail in biophysical models does result in more degrees of freedom, so biophysically detailed models have to be constrained by the physical features of the neural system as well as the measurement physics of the signals they simulate.

This delineation is only approximate, and modeling studies can include techniques and features of both approaches, but it is a practical classification that can serve to clarify the focus and scope of a modeling study. Both have their advantages and limitations, and which approach to use ultimately depends on the research question [2, 8].

Biophysically detailed modeling has had a wide range of applications. Rall (1962) [80] used his theoretical foundation for how the movement of currents across and within neurons underlie electric signals to show that the dendritic tree of a neuron in many cases can be approximated by an equivalent cylinder. Then, later, Rall and Shepherd (1968) [81] analyzed the LFP recorded in the olfactory bulb of rabbits to uncover the circuit properties of the bulb. Holt & Koch (1999) [45] used volume conductor theory to show that ephaptic effects are minor except in special conditions. Up until the early 2000s, most biophysical models were small-scale models of a single or a few neurons. This was partly due to a lack of data and partly due to the computational requirements of larger models with this level of detail. Larger-scale models were made, but they typically consisted of the simplified point-neurons, which can be orders of magnitude cheaper and faster to run [8]. An early attempt at larger-scale modeling with biophysically detailed neurons was the thalamocortical loop model developed by Traub et al. (2005) [103]. It consists of 3560 neurons, which belong to four different major classes: the superficial layer (2/3), layer 4, or deep layer (5/6) in the cortex, and the

thalamus. The model exhibited oscillations in network states typically observed experimentally, such as persistent gamma oscillations and sleep spindles, and made predictions about the physiological mechanisms underpinning them. Reimann et al. (2013) [82] made a model of a neocortical column consisting of more than 12000 layer 4 (L4) and layer 5 (L5) neurons, and used it to investigate which types of transmembrane currents contribute the most to the LFP. They found that active currents crucially alter the pattern and magnitude of LFPs, and thus demonstrated that the assumption of dominance by passive synaptic currents in the LFP may not always be valid.

A general limitation of models of brain nuclei is that they have limited information on important network components like cell types, cell anatomy and physiology, and connectivity [98]. The Blue Brain project tried to resolve some of these limitations, which stemmed from limited experimental data, by developing an algorithm that estimates the missing information about parameters from sparse data [61, 83]. They used this approach to make a comprehensive model of rat somatosensory cortex, which consists of about 31,000 neurons of different morphological, electrical, and synaptic classes [61]. The Allen Institute, meanwhile, began systematically collecting large data sets on cell types, connectivity, and cell physiology, and used that data to fill some of the gaps in our knowledge of parameter values. They then constructed a large-scale, biophysically detailed model of mouse primary visual cortex which integrated much of this data with information from a literature review [8]. This model consists of more than 50,000 neurons positioned in five layers: layers 1, 2/3, 4, 5, and 6 (L1, L2/3, L4, L5, L6), where layers 2 and 3 are merged into a single layer. Each layer except L1 has a single class of excitatory neurons and three classes of inhibitory neurons (Pvalb, Sst, and Htr3a). L1 only has one inhibitory population (Htr3a) and no excitatory cells. The model reproduced experimentally observed levels of orientation and direction selectivity, and also serves as a database for the experimental data that were employed in the making of the neuron models and their network connectivity.

Though these models each reproduced important and interesting features observed in experimental data, none of them had quantitatively reproduced features of both spiking activity and LFP typically recorded in vivo. We therefore used the mouse V1 model presented in [8] as a starting point with the aim of developing a large-scale, biophysically detailed model that could reproduce both these signals simultaneously. The obtained model was then used to develop and validate a tool that automatically decomposes LFP/CSD to uncover LFP/CSD contributions generated by spiking in different presynaptic populations. Lastly, we also used the model as a testbed for a method that estimates the postsynaptic LFP from presynaptic firing rates and network properties, which could be used to reduce the large computational resources required to simulate LFP in a large-scale model with numerous parameters.

Chapter 2

Methods

This chapter provides more detail on the theoretical framework that enables mathematical modeling of neural activity and electrical signals. If the reader is already familiar with these topics, they may skip to the next chapter.

2.1 Modeling at different levels of detail

Neuronal networks can be modeled at different levels of detail depending on their purpose. At the coarsest level, individual neurons are not distinguished; their activity is merged with the activity of other neurons belonging to the same class or population. The spiking activity in these population models is only represented as population firing rates, and the timing of action potentials of individual neurons is not considered. These models are therefore referred to as *rate models* [37]. The second level is the so-called point-neuron models, which were briefly described in the introduction above. In these models, individual neurons are distinguished, but biological details such as cell morphology and distributions of various ionic currents are not modeled explicitly [37, 99]. The mathematical functions used to model point-neurons only characterize the behavior of the neuron as a whole rather than aiming for biological realism on the level of ion channels [99]. The level with the most comprehensive inclusion of biological knowledge is referred to as biophysically detailed models (they are also called multicompartment models, this designation will be explained later). Here, the neuronal geometries, ion channel distribution, and other spatial details of real neurons are modeled explicitly, though the degree of detail may vary considerably [19, 99].

One advantage of the less detailed neuron models - the population firing rate and the point-neuron models - is that they are computationally less demanding [99]. This broadens their potential user base as they are less likely to require access to high-performance computing centers. The lower computational demands also make them more amenable to large-scale network modeling. Population firing rate models can be used to simulate whole brain activity [37], and point-neuron networks have been simulated on networks consisting of up to 10^9 neurons and 10^{13} synapses, roughly on the scale of the cat brain [31]. Another advantage of less detailed models is that they may be easier to understand. The more complex models usually come with more parameters, which can make the model unwieldy

and obscure the relationship between the model’s components and its behavior. The many parameters may also be challenging to constrain properly if the data necessary to determine the parameters is not yet available [99].

An advantage of the more detailed models, however, is that they can be used to simulate the extracellular signals observed in experimental recordings [35]. This enables investigation into the relationship between the observed signals and the neural activity underpinning them, which can augment the insight into neural mechanisms that can be gained from experimental data [40, 41, 45, 82, 85, 93]. These models can also be used to develop tools to analyze experimental data, as the tools can be tested and validated on the signals simulated with the model where the ground truth is known [18, 25, 26, 59, 77]. Additionally, biophysically detailed modeling may allow for greater utilization of the data collected in experimental recordings; we are not limited to only or primarily studying the spikes or firing rates that can be modeled with point-neurons or population firing rate models. Finally, as outlined in the introduction, they can be used to bridge the scales of our understanding between events on the microscopic level of ion channels to signals and events on the macroscopic level of whole brain regions [20].

Population firing rate models have not been studied in this thesis (though population firing rates have been calculated or estimated). Therefore, they will not be described in more detail here. The focus will be on the framework for biophysically detailed modeling, but some of the essential features of point-neurons will also be characterized.

2.1.1 Biophysically detailed modeling

Modeling a membrane patch as an RC circuit

The biophysically detailed modeling of neuronal activity begins not with a whole neuron, but with a patch of the neuronal membrane. A membrane patch with ion channels and a stimulating electrode can be modeled with an equivalent resistor-capacitor (RC) circuit (illustrated in Fig. 2.4) [100]. The membrane forms an insulating layer between ions settled on its extracellular and intracellular surface and can therefore be modeled as a capacitor. A stimulating electrode that provides current to the circuit is modeled as the current source. The ions in the extra- and intracellular media can cross the membrane through ion channels, which are modeled as resistors in parallel with the current source and the capacitor. For now, the ion channels are all passive in this model.

This circuit can be used to develop a mathematical formulation of the current flow through the membrane patch. Kirchhoff’s current law states that the total amount of current flowing in and out of a circuit has to sum to 0. For the passive membrane patch, this can be written as:

$$I_e/a - I_c - I_i = 0, \quad (2.1)$$

where a is the area of the membrane, I_e is the injected current, I_c is the capacitive current, and I_i is the ionic current through the ion channels.

The ionic current through the resistor follows Ohm’s law, i.e., it is given by the




Figure 2.4: Membrane patch as RC circuit. A patch of membrane (in blue) with ion channels (in green), and an electrode (in red) inserting ionic current can be represented with an equivalent RC circuit. The ion channels are modeled as resistors, while the membrane is modeled as a capacitor, and the electrode is modeled as a current source.

potential divided by the resistance. The resistance here refers to the membrane resistance - the resistance against ions crossing the membrane - and is denoted R_m . The potential is the difference between the potential over the membrane V and the equilibrium potential for the membrane E_m . Thus, the ionic current can be expressed as:

$$I_i = \frac{V - E_m}{R_m} \quad (2.2)$$

The capacitive current is found by starting with the fundamental equation for the relationship between the capacitance C set up by the charge q on a capacitor and the potential V over the capacitor:

$$C = \frac{q}{V} \quad (2.3)$$

By rearranging eq. 2.3 and taking the time-derivative on both sides we get:

$$\frac{dq}{dt} = C \frac{dV}{dt} \quad (2.4)$$

Since $\frac{dq}{dt}$ is the definition of current and the capacitance here is the membrane capacitance C_m , eq. 2.4 can be written as:

$$I_c = C_m \frac{dV}{dt} \quad (2.5)$$

By inserting eqs. 2.2 and 2.5 into eq. 2.1, we get the following differential equation characterizing the membrane potential over the patch over time following a current injection:

$$C_m \frac{dV}{dt} = \frac{E_m - V}{R_m} + I_e/a \quad (2.6)$$

We can find an analytical solution for the membrane potential over time by integrating eq. 2.6:

$$V(t) = E_m + \frac{R_m I_e}{a} (1 - e^{-\frac{t}{R_m C_m}}) \quad (2.7)$$




Figure 2.5: Multicompartment neurite as RC circuits in series. (A) A neurite divided up into cylindrical compartments of diameter d and length L connected by axial resistors with resistivity R_a . (B) Each compartment is isopotential and can be modeled as an RC circuit. The neurite is then modeled as RC circuits connected in series via axial resistors.

The cable equation and multicompartment modeling

The RC circuit modeling of the membrane patch assumes that membrane potential is isopotential - that is, that the potential is effectively constant over the whole membrane. That assumption can be valid for a small area of a cell but is usually not true over the whole membrane of a neuron. Fortunately, to model a whole neuron, we can simply divide it up into parts, or compartments, that each can be assumed to be effectively isopotential [100]. As illustrated in Fig. 2.5A, the compartments are modeled as cylinders and connected to the neighboring compartments through an axial resistor, here denoted by its resistivity R_a . The axial resistor represents the resistance to ionic movement in the cytoplasm of the neurite. Each compartment is then a membrane patch that can be modeled as an RC circuit, and the set of compartments that make up a neurite is a series of connected RC circuits (Fig. 2.5B).

This approach is the reason why biophysically detailed models are also referred to as multicompartment models. The number of compartments depends on how much detail is needed for the research question and on the size of the area for which the membrane can be assumed to be isopotential. Rall (1962) [80] showed that the dendrites of a whole motoneuron can, in fact, be modeled as a single equivalent cylinder. But there are also many situations where greater partitioning is necessary [50].

To formulate a mathematical description of a multicompartment model, we can begin with the result from the single patch of membrane (eq. 2.6), and add the

terms for the current flow between compartments. If we take the perspective of the middle compartment in Fig. 2.5A, labeling it j , then the current can both flow to or from the rightward compartment $j+1$ and to or from the leftward compartment $j-1$. The terms to be added, then, are the current to or from the left compartment $j-1$ and the current to or from the right compartment $j+1$. These currents are here denoted $I_{j-1,j}$ and $I_{j,j+1}$, respectively. They follow Ohm's law, so they are given by the potential difference between the compartments divided by the axial resistivity R_a , that is, $I_{j-1,j} = \frac{V_j - V_{j-1}}{R_a}$ and $I_{j,j+1} = \frac{V_{j+1} - V_j}{R_a}$, assuming that the axial resistivity is constant. The resistivity is given in Ωm , so to get the axial resistance, the resistivity must be multiplied by the length of the compartment per area, which for a cylinder of diameter d and length L is given by $\frac{4L}{\pi d^2}$. Adding these terms to eq. 2.6 (with the area of the cylinder $a = \pi dL$ placed on the left-hand side at first) we get:

$$\pi dL C_m \frac{dV}{dt} = \frac{E_m - V}{R_m / \pi dL} + I_e + \frac{V_{j+1} - V_j}{4R_a L / \pi d^2} + \frac{V_j - V_{j-1}}{4R_a L / \pi d^2} \quad (2.8)$$

Which can be rearranged to:

$$C_m \frac{dV}{dt} = \frac{E_m - V}{R_m} + \frac{I_e}{\pi dL} + \frac{d}{4R_a} \left(\frac{V_{j+1} - V_j}{L^2} + \frac{V_j - V_{j-1}}{L^2} \right) \quad (2.9)$$

This is the fundamental equation for simulating neural activity in the multicompartment modeling scheme, and it can be used to calculate the membrane potential over the whole neuron for neurons of arbitrary geometry.

If the aim is to study the electrical properties of neurons analytically, we can start with eq. 2.9 and let the length of each compartment become infinitesimally short, i.e. let $L \rightarrow \delta x \rightarrow 0$ [100]. In this case, eq. 2.9 becomes:

$$C_m \frac{\partial V(x, t)}{\partial t} = \frac{E_m - V(x, t)}{R_m} + \frac{I_e(x, t)}{\pi d} + \frac{d}{4R_a} \left[\frac{1}{\delta x} \frac{V(x + \delta x, t) - V(x, t)}{\delta x} - \frac{V(x, t) - V(x - \delta x, t)}{\delta x} \right] \quad (2.10)$$

Which can then be written as a partial differential equation:

$$C_m \frac{\partial V(x, t)}{\partial t} = \frac{E_m - V(x, t)}{R_m} + \frac{I_e(x, t)}{\pi d} + \frac{d}{4R_a} \frac{\partial^2 V(x, t)}{\partial x^2} \quad (2.11)$$

This expression is referred to as the cable equation. By moving the membrane resistance R_m to the left-hand side, we get the following formulation:

$$\tau_m \frac{\partial V}{\partial t} = E_m - V + \frac{I_e R_m}{\pi d} + \frac{d R_m}{4R_a} \frac{\partial^2 V}{\partial x^2} \quad (2.12)$$

The fraction in front of the double spatial derivative of the potential defines the electrotonic length constant λ :

$$\lambda = \sqrt{\frac{d R_m}{4R_a}} \quad (2.13)$$

Since the electrotonic length constant depends on the relationship between the specific membrane resistance (the transverse resistance of the membrane) and the axial resistivity (the longitudinal resistivity of the membrane), it describes the shape of the voltage decay over the length of the neuron. When the membrane resistance is large relative to the axial resistivity, λ will be large, which indicates that the potential will spread further along the length direction of the membrane.

Active channels

Thus far, we have described membranes with only passive channels. Real neurons, however, have active channels too. Active channels, contrary to passive channels, are not always open; they change their state from closed to open through a transfer of energy via either chemical reactions, mechanical work, or voltage changes [101]. An action potential is driven by the opening and closing of active channels, so to have a full description of the development of the membrane potential over time, a mathematical characterization of ionic currents through the active channels must also be included in the framework outlined above.

This mathematical characterization was developed by Hodgkin and Huxley (1952) in their pioneering work with the squid giant axon on how action potentials arise from ionic currents [44]. The total ionic current I_i can be broken down into the individual currents that make it up: the sodium, potassium, and what Hodgkin and Huxley called the leak current, which primarily consists of chloride ions [101]:

$$I_i = I_{Na} + I_K + I_L \quad (2.14)$$

The current from an ion type X will be the product of the membrane conductance g_X for that ion type and the difference $V - E_X$ between membrane potential and the equilibrium potential. Thus, for the three currents in eq. 2.14 we get:

$$I_{Na} = g_{Na}(V - E_{Na}) \quad (2.15)$$

$$I_K = g_K(V - E_K) \quad (2.16)$$

$$I_L = \bar{g}_L(V - E_L) \quad (2.17)$$

The bar above the membrane conductance for the leak current \bar{g}_L indicates that it is constant. The conductances g_{Na} and g_K are not constant since the Na^+ and K^+ channels are active, and their behavior over time is described by equations incorporating so-called gating particles which determine the state of the channel. The details of the equations governing the ion channels and the gating particles will not be described here, but can be found in the original articles by Hodgkin and Huxley or any textbook describing the computational aspects of neurophysiology. The leak conductance g_L is determined with experimental measurements. (Note that even though the role of the leak conductance g_L and the membrane resistance R_m in the all passive channels circuit (eq. 2.6) is similar, they are not simply the inverse of each other. In the Hodgkin-Huxley model, the resting membrane

resistance is determined by the sodium, potassium, and leak resting conductances in combination [96, 101].)

Volume conductor theory

In the preceding sections, we focused on the membrane potential, i.e. the potential across the membrane arising from a difference in ion concentrations between the interior and the exterior of a cell. This can be recorded by inserting one sharp electrode into the cell and measuring the potential difference to a reference electrode placed outside the cell [97]. However, the electrical activity of neurons is also often measured with recording electrodes positioned only in the extracellular space. Since there is some distance from the neurons to the electrode, the propagation of electric potentials through the extracellular medium needs to be included in the mathematical characterization of the recorded potential generated by neural activity. This is done with volume conductor theory [19, 38, 97].

For simplicity, we will first consider a neuron with only two transmembrane currents (this can also be seen as a multicompartiment neuron with only two compartments): a current sink ($-I_m$) from positive ions entering the cells and a current source ($+I_m$) from a return current of ions leaving the cell. This scenario is illustrated in Fig. 2.6. A recording electrode (we will only consider a single electrode here rather than the multiple electrodes of a multisite probe, again for simplicity) is positioned at \vec{r}_{el} . The sink and the source are positioned at \vec{r}_{sink} and \vec{r}_{source} , respectively. When the conductivity σ is constant¹, the potential recorded at \vec{r}_{el} from the two transmembrane currents is given by the following equation [19, 97]:

$$\phi(\vec{r}_{el}, t) = \frac{+I_m(t)}{4\pi\sigma|\vec{r}_{el} - \vec{r}_{source}|} + \frac{-I_m(t)}{4\pi\sigma|\vec{r}_{el} - \vec{r}_{sink}|} \quad (2.18)$$

To extend this to multicompartiment neurons with transmembrane currents in more than two compartments, one simply sums the current $I_{m,i}$ going through each compartment i up to all N compartments:

$$\phi(\vec{r}_{el}, t) = \frac{1}{4\pi\sigma} \sum_{i=1}^N \frac{I_{m,i}(t)}{|\vec{r}_{el} - \vec{r}_i|} \quad (2.19)$$

This equation enables calculation of the extracellular potential resulting from the transmembrane currents of a neuron at any position in the extracellular space. Furthermore, the total extracellular potential from multiple neurons is the superposition of the potential from each neuron, so eq. 2.19 will also give the potential from a population of neurons, as one would record in vivo.

¹This framework can be extended to the scenario where the conductivity is anisotropic, inhomogeneous, or frequency-dependent [33, 38]. In most situations, however, σ can reasonably be treated as constant, and it was treated as constant in all studies presented in this thesis.




Figure 2.6: **Potential measured by extracellular electrode arising from transmembrane currents can be calculated with volume conductor theory.** The current sink (blue arrow) and the current source (orange arrow) from ions crossing the membrane of the neuron each contribute to the potential that propagates through the extracellular medium with conductivity σ . The magnitude of the potential recorded at the electrode position depends on the distances r_{sink} and r_{source} from the transmembrane currents to the electrode position.

2.1.2 Point-neurons

Integrate-and-fire neurons

The multicompartiment modeling approach can be simplified by using fewer compartments. In some cases, it may even suffice with only two compartments [87]. However, if the dynamics arising from ion channel distribution or neuron geometry are not of interest or relevant to the problem to be investigated, the multicompartiment modeling approach can be dropped altogether in favor of so-called integrate-and-fire neuron models. With integrate-and-fire models, the time course of action potentials is not simulated. When the membrane potential of the neuron model reaches the threshold for spiking, the potential is simply reset to the resting potential automatically, and a spike is added to the total tally of spikes [24, 99]. The focus of this type of modeling is to reproduce the behavior of the neuron as a whole rather than aiming for similarity to the physics and biology on the level of channel dynamics.

The RC circuit used to model a membrane patch is also utilized in the modeling of integrate-and-fire neurons, but the circuit here models the behavior of the whole neuron, not just a patch of it. Thus, the resistance and capacitance here represent the resistance of all channels on the neuron and the capacitance of the whole membrane. Furthermore, a reset mechanism is included in the circuit to model spike occurrences. This reset mechanism is a switch that short-circuits the membrane resistance, which allows the membrane potential to return to rest and

makes the circuit ready to charge until a new spike is produced [24, 99].

To describe the behavior of an integrate-and-fire neuron mathematically, we start with an equation similar to eq. 2.6:

$$C_m \frac{dV}{dt} = \frac{E_m - V}{R_m} + I \quad (2.20)$$

The difference between eqs. 2.6 and 2.20 is that here I is the total current into the cell from either a stimulating electrode or a synapse and that R_m and C_m represent the resistance and capacitance of the whole neuron. This equation characterizes the membrane potential when it is below threshold. Moving R_m on the left-hand side, it can be written as:

$$\tau_m \frac{dV}{dt} = E_m - V + R_m I \quad (2.21)$$

where $\tau_m = R_m C_m$ is the time constant of the membrane potential. By integrating eq. 2.21, we get the following analytical expression for the membrane potential V over time:

$$V(t) = E_m + R_m I \left(1 - e^{-\frac{t}{\tau_m}}\right) \quad (2.22)$$

It is when the membrane potential described by this equation reaches a threshold, often denoted with θ , that the neuron fires and the potential is reset to the resting membrane potential [99].

2.2 Software

Though it is possible to study neural processes analytically, as outlined above, the biological complexity often necessitates numerical analysis. Numerous simulators and wrappers for those simulators have been developed to simplify and reduce the labor required of the modeler. A short overview of the software utilized in the projects in this thesis will be given here.

NEURON, developed by Hines and Carnevale (1997) [43], is one of the most widely used environments for running simulations of biophysically detailed neurons. It allows for specification of cell models - their morphological architecture, ion channel distribution, synaptic placement, and more - as well as network connectivity via the HOC scripting language. All biophysical simulations in this thesis were carried out using the NEURON environment.

NEST (the Neural Simulation Tool) [17] is a simulator for spiking neural networks. NEST provides a framework for developing models focusing on the dynamics and structure of neural systems, rather than the morphology of individual cells. NEST is therefore ideal for point-neuron simulations.

BMTK (Brain Modeling Toolkit) is a Python-based software package that is not in itself a simulator but rather serves as an interface for the NEURON and NEST simulators [14]. It provides capabilities for building and simulating neural models both across scales - from single cells to large-scale networks - and across levels of resolution - from filter-based and population firing rate-models to

Chapter 2. Methods

biophysically detailed neurons. BMTK is designed to facilitate easy incorporation of experimental data in models such that both cell models and networks can be constructed in a data-driven manner.

Part II

Results and Discussion

Chapter 3

Summary of papers

3.1 Paper I: Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex

Biophysical modeling has reproduced and shed light on the mechanistic origins of several important phenomena in systems neuroscience: from gamma oscillations in network activity [103], the contribution of active conductances to current sinks and sources [82], to the network architecture underlying orientation and direction selectivity in the visual cortex [8]. However, to our knowledge, no study has quantitatively reproduced in vivo-recorded spiking and LFPs at the same time. Past studies have either focused on one of these signals and/or only compared qualitatively to experimental data. This was largely because the data available was insufficient to do proper quantitative validation across scales.

In recent years, there has been an impressive increase in the data collected and made publicly available in neuroscience [9, 13, 92, 104, 105]. This has enabled data-driven modeling and quantitative comparison to experimental data on both single-cell and population levels on a scale not previously possible. Scientists at the Allen Institute leveraged extensive data on cell types to make a comprehensive library of biophysically detailed cell models [28], and then utilized this library together with both literature reviews and locally gathered data on network connectivity to construct a large-scale, biophysically detailed model of mouse primary visual cortex (V1) [8]. This model reproduced population firing rates and levels of orientation and direction selectivity observed experimentally during presentations of drifting grating stimuli. It was, however, not investigated whether the LFP simulated with the model matched the experimentally observed LFP.

In our first study, we used this published V1 model as a starting point for developing a model that could reproduce experimentally observed CSD (derived from LFP) and spiking activity in response to a full-field flash stimulus. To validate the model against experimental data, we used the Visual Coding data set containing in vivo recordings of extracellular potentials from 58 mice [92]. We discovered that there were substantial discrepancies between the original model CSD and the experimental CSD, and set out to investigate the cause of this

discrepancy and to find out what could be amended in the model to make it reproduce spikes and LFP from experiments simultaneously.

It was not only the simulated CSD that deviated from the experimental data, the spiking response of the model also deviated for some populations in response to full-field flashes. By adjusting the recurrent synaptic weights, however, we were able to get the spiking activity of the model within experimental variability. Still, even though the model now reproduced the spiking response, the discrepancy between the model CSD and the experimental CSD remained. The model CSD was also largely unaltered by the modification of the recurrent synaptic weights, even though the spiking response changed significantly. This aligns with findings in previous studies that have suggested that spiking and LFP represent complementary aspects of neural activity [53, 89].

The original version of the model received external input from LGN and background input representing input from the rest of the brain. It did not include feedback input from higher visual areas. In our attempts to reproduce the response to full-field flashes, we focused on the first 100 ms after stimulus onset. This choice was made partly because we surmised that the response in this time period would be dominated by the feedforward input and the activity generated by recurrent connectivity within V1, which were included in the original model. Our original intention was to introduce feedback from higher visual areas to V1 to the model in a separate, later project and present the results in a second paper.

However, we observed that the prominent, sustained current sinks and sources inside the 100 ms time window that were missing in the simulated CSD coincided in time with elevated activity in higher visual areas in the experiments. Therefore, we hypothesized that these sinks and sources were generated by feedback from these areas, and that it would be necessary to add feedback from a higher visual area to the model to reproduce the experimental response in the first 100 ms. We thus incorporated the feedback input already in this project. We did this in a data-driven manner, feeding spikes recorded experimentally in the lateromedial area (LM) of higher visual cortex to the model. This produced an immediate improvement in the correspondence between the simulated and experimental CSD. The sustained sinks and sources in upper layers that had been missing in the model were now present in the simulations too. Importantly, we were able to get the model to match the experimental CSD at the same time as it matched the experimental firing rates.

We also introduced the application of the Wasserstein distance (WD) as a metric to quantitatively compare CSD patterns. The WD has traditionally been used to compare probability distributions. Here, we used the WD to quantify the similarity between the distributions of current sinks and sources of two CSD plots. We first calculated the WD between individual CSD plots from animals and what we dubbed the *canonical* CSD pattern, which was the first principal component calculated from the CSD of all animals together. The distribution of WDs from individual animals to the canonical CSD pattern established the experimental variability. Then we calculated the WD between the simulated CSD and this canonical CSD pattern and compared it to the experimental variability to assess model performance. When the model was within the experimental

variability, we could claim that the model was indistinguishable from an individual animal in the experimental data set. We applied the same methodology with already established metrics for spike statistics to evaluate model performance and experimental variability for spiking.

3.2 Paper II: Uncovering population contributions to the extracellular potential in the mouse visual system using Laminar Population Analysis

Developing a large-scale biophysically detailed model, as we did in paper I, is one way to uncover the circuit mechanisms underpinning LFPs. However, constructing models to simulate extracellular potentials can be a resource-intensive endeavor, both with respect to human and computational resources. Thus, it would be advantageous to have tools that could disentangle LFP contributions from different biophysical processes without requiring the construction of a model.

Statistical tools such as independent component analysis (ICA) [21, 22, 52, 58, 59] and principal component analysis (PCA) [3–5, 16] have been used to decompose LFP with the aim of uncovering contributions from different pathways to the recorded LFP. The idea was that the different statistical features in the obtained components would correspond to the contributions from different populations or pathways, thereby separating and identifying these contributions. This approach may work if the assumptions of statistical independence, in the case of ICA, or orthogonality, in the case of PCA, are valid for the structure in question and experimental conditions in which the data was collected. However, these assumptions are unlikely to be generally valid [26, 29].

Seeking to avoid the issues with the validity of statistical assumptions, Einevoll et al. (2007) developed a method called laminar population analysis (LPA) [18, 25], which decomposes LFP based on physiological rather than statistical assumptions about how different populations generate LFP. LPA utilizes the multi-unit activity (MUA) and the LFP jointly in a spatiotemporal decomposition of the LFP. It first identifies the positions of layers in a laminar structure and temporal profiles of the firing rates of these layers from the MUA. Then, in the next step, the temporal profiles of the firing rates of these laminar populations are convolved with kernels to estimate the temporal profile of the postsynaptic LFP they generate. The firing rates and contributions of external populations can be included in the decomposition simply by appending them to the temporal profiles of the laminar populations. The convolution with kernels approximates the process by which action potentials in each presynaptic population produce postsynaptic LFP. The obtained postsynaptic temporal profiles of the population LFP are then used in the decomposition of the recorded or simulated LFP. The end result is the population contribution to the total LFP from each presynaptic population.

In this paper, we sought to validate LPA on data simulated with the mouse V1 model version we developed in paper I. We first found that LPA applied to MUA estimated both the positions of layers in the V1 model and the temporal profile of their population firing rates reasonably well. We next applied LPA to the CSD

derived from the simulated LFP, and found that we needed to add a regularization term to the cost function that penalized deviations from zero in the CSD summed across channels to improve the population estimates. After introducing this regularization term, LPA could estimate the salient current sinks and sources produced by feedforward inputs from LGN, feedback inputs from LM, as well as those stemming from recurrent activity within V1. LPA also partly recapitulated the contributions from some laminar populations. However, on the whole, the estimation of laminar contributions to the CSD was not satisfactory. It is likely that the poorer estimation of contributions from layers within V1 was caused by excessive synchrony in the population firing rates across layers. This synchrony could be diminished and the estimation of laminar contributions improved by the use of other stimulus protocols in the simulations and the experiments.

There were still spurious sinks and sources in the CSD estimates not observed in the ground truth CSD contribution from each population. There were also deviations in the magnitude for the sinks and sources that did have the correct position and timing. Nonetheless, LPA produced more accurate renderings of the CSD contributions from each structure than the alternative statistical decomposition tools, PCA and ICA, and had a clearer interpretation in terms of the underlying circuit mechanisms.

Lastly, we demonstrated the use of LPA on experimentally recorded extracellular potentials. We observed that some of the sinks and sources ascribed to the feedforward input from LGN, the recurrent activity within V1, and the feedback from higher visual areas corresponded to the CSD contributions expected from these structures based on anatomical data and our modeling results in paper I.

3.3 Paper III: Estimating simulated local field potentials from presynaptic firing rates and network properties

One bottleneck in studying extracellular potentials with large-scale, biophysically detailed models is the computational power required to simulate them. The biophysical mouse V1 model can only be simulated with the resources of a high-performance computing center, which makes it unfeasible to work with for those without access to such centers¹. The kernel methodology for estimating LFP from presynaptic firing rates, of which the LPA method in paper II is an example, can help substantially reduce the computational resources required to simulate extracellular potentials. The kernel approach can enable approximations of the LFP generated by models, obviating the need to run full-scale simulations. This would provide all labs with the opportunity to study the LFP of large-scale, biophysically detailed models.

In our third project, we employed a previously established framework [30] to construct kernels from detailed information about membrane potentials, synaptic

¹There is, however, a point-neuron version of the V1 model, which can be run on laptops.

parameters, and connectivity patterns in the model, and we used these kernels to estimate the LFP generated in L2/3 of the mouse primary visual cortex model version from paper I. We convolved the kernels with the firing rates of all presynaptic external inputs (LGN, LM, and the background) and the internal populations in L2/3 to estimate the LFP contribution from each population. We found that this approach produced estimates of LFP contributions that recapitulated the simulated LFP contributions with a high degree of accuracy. The kernel method estimated both the LFP patterns and magnitude generated by firing in each presynaptic population, and could be used to assess the relative strength of each contribution. The external inputs contributed the most to the LFP overall, while the parvalbumin-positive (PV) inhibitory population constituted the largest contributor among the internal sources.

The accuracy in the estimates with this approach - $R^2 > 0.9$ for all channels except those at the center point of the somatic distribution - is high enough that, at least for some research questions, this method could be utilized as a substitute for full simulations. The information required to construct the kernels is not concurrently available with LFP recordings in electrophysiological recordings, which may limit its application to current experimental data, but further investigation is required to conclusively determine this. We plan to extend the analysis done for L2/3 to all layers in the V1 model.

Chapter 3. Summary of papers

Chapter 4

Discussion

The thalamocortical loop model developed by Traub et al. (2005) [103], consisting of more than 3500 biophysically detailed cells, was an early effort to construct a large-scale, biophysically detailed network model. The article in which the model was presented was introduced with the following observation:

"The greatest scientific challenge, perhaps, in all of brain research is how to understand the cooperative behavior of large numbers of neurons. Such cooperative behavior is necessary for sensory processing and motor control, planning, and in the case of humans, at least, for thought and language. Yet it is a truism to observe that single neurons are complicated little machines, as well as to observe that not all neurons are alike — far from it; and finally to observe that the connective anatomy and synaptology of complex networks, in the cortex for example, have been studied long and hard, and yet are far from worked out. Any model, even of a small bit of cortex, is subject to difficulties and hazards: limited data, large numbers of parameters, criticisms that models with complexity comparable to the modeled system cannot be scientifically useful, the expense and slowness of the necessary computations, and serious uncertainties as to how a complex model can be compared with experiment and shown to be predictive. The above difficulties and hazards are too real to be dismissed readily. In our opinion, the only way to proceed is through a state of denial that any of the difficulties need be fatal. The reader must then judge whether the results, preliminary as they must be, help our understanding."

Though the amount and detail of data available today has increased prodigiously since this paragraph was written, these difficulties are still too real to be readily dismissed. They will presumably and hopefully abate as our methods for data collection become more precise and efficient, and more high-quality data are made available to the neuroscientific community. Until that time arrives, however, we remain firmly in a regime of modeling without knowing all the numbers.

The projects in this thesis all strive to address and improve the current state of computational neuroscience with respect to some of the challenges listed by Traub and colleagues. The first paper deals with how we can enhance constraints on model parameters and architectures by utilizing all the data we have available across multiple scales, and how a complex model can be compared to experiments

and shown to be predictive. The second paper deals with how to infer circuit mechanisms from recorded extracellular potentials, without having to construct a model, by taking advantage of physiological knowledge of how extracellular potentials are generated. The third paper aims to validate a method that circumvents the time- and computationally consuming simulations of extracellular fields by rather estimating them from presynaptic firing rates and known network parameters. As Traub et al. pointed out, the readers must be the judges as to whether the results are useful and augment our insight into neural mechanisms and experimental signals. In the following, the results from these projects and their limitations will be discussed, and potential paths forward will be outlined.

4.1 Validating large-scale models and making testable predictions

Models in computational neuroscience, like any theoretical endeavor in science, should produce testable predictions. To be able to test predictions, however, we must first determine how to compare model predictions with the experimental data.

Previous large-scale biophysical modeling efforts have often limited their comparison to signals on a single scale or compared simulated and experimental data qualitatively. In our first paper, we sought to establish a systematic framework to quantify experimental variability and assess model performance both with respect to spike data and CSD derived from LFP. We used the Kolmogorov-Smirnov similarity between distributions of firing rates across cells in different time periods as well as the moments of distributions of peak firing rates to quantify population spiking magnitudes. To evaluate temporal profiles of population spiking, we calculated the correlation between population firing rates and the moments of distributions of time to the initial peak. Lastly, to compare CSD patterns, we calculated the Wasserstein distance between distributions of sinks and sources, which tells us how much the sinks and sources of one pattern have to be moved around to match another pattern. The more sinks and sources have to be moved, the bigger the Wasserstein distance.

For all these metrics, we first established the experimental variability. One way we did this was to calculate the relevant similarity or distance metric between each individual animal and the target - that is, the average population firing rate or the first principal component of trial-averaged CSD from all animals. Alternatively, we calculated the similarity or distance pairwise between animals. When we then computed the same metrics between the model and the experimental data, we could compare the obtained values to the experimental variability. If the model was within the experimental variability, we could argue that the model was indistinguishable from an individual mouse in the experimental data set, and thus that the experimental feature in question had been reproduced by the model.

The final model was within the experimental variability for all but a few measures. However, this does not preclude that other metrics would have revealed greater discrepancies between the model and the experiments. Different

metrics have different advantages and drawbacks, or different trade-offs, in other words. That is not to say that some metrics are not more appropriate than others: some metrics can be more informative or have fewer serious limitations. Thus, a systematic investigation into what could be the optimal metrics for comparing large-scale, biophysically detailed models to data would be a valuable and important avenue for future research. Additionally, an exhaustive evaluation of a model that includes biological detail on the subcellular level should ideally also be compared to in vivo experimental data on this scale as well. This was not feasible with the data currently available. Nonetheless, the metrics and approach applied in this study laid the groundwork for cross-scale quantitative comparison of simulated to recorded signals that can be used by other studies in the future.

The original V1 model received feedforward input from a structure representing LGN and background input from a Poisson source representing the influence of the rest of the brain. With this model configuration, there were major discrepancies between the simulated and experimental CSD. We explored whether the model with only these inputs and recurrent connections within V1 could be modified to match the experimental CSD, but did not succeed. Past studies on the effect of feedback input on LFP have suggested that it primarily exerts its influence only a few hundred ms after stimulus onsets [65, 66], indicating that the feedback should not be involved in the generation of major sinks and sources in the first 100 ms after stimulus onset that we focused on. However, a recent study on the effects of cortico-cortical feedback on LFP in the V1 area of monkeys suggested that feedback can crucially alter LFP already 80 ms after stimulus presentation [39]. In order to test the idea that feedback from higher visual areas is involved in the generation of the sinks and sources missing in the original model simulations, we decided to introduce feedback to the model in a data-driven manner. We added another input to the model consisting of spikes recorded experimentally in the lateromedial area (LM) of higher visual cortex, and it engendered an immediate improvement in the correspondence between the simulated and experimental CSD. The salient sinks and sources that had been observed in experiments but were missing in the original simulations were now present.

The idea that cortical feedback sculpts the LFP and CSD already in the first 100 ms of evoked responses, and the specific ways it does so, constitutes a testable hypothesis formulated by our model simulations. We suggested an experimental protocol that could falsify the hypothesis, where the influence of feedback is removed by, e.g., optogenetic silencing of the relevant higher cortical areas 60 ms after stimulus onset (the time at which the effect feedback seemed to arise). We simulated this experimental protocol by removing the feedback at this time point in the model, and with the framework of metrics described above, the resulting CSD (and spiking activity) can be compared to the experimental results for verification or falsification.

Another question we briefly attempted to address with our simulations, was the role of active channels in LFP generation. LFP has been considered to primarily originate from currents caused by synaptic input, and the contribution from active channels in shaping the LFP has been seen as less important in most physiological conditions [19, 56, 67]. This stemmed in part from the observation that the action

potentials during which most active channels are open are short in duration - lasting only a few milliseconds, which would probably be too short for them to summate constructively except in highly synchronous conditions [12, 30]. Synaptic currents, on the other hand, are slower, which allows for greater overlap in time and therefore a greater potential for constructive summation [30]. Reimann et al. (2013) [82], however, constructed a model of L4 and L5 in rat somatosensory cortex (S1) consisting of about 12000 neurons, and found that in their simulations, the active currents actually contributed a substantial portion of the transmembrane currents underpinning the LFP, and they posited that it is in fact the active currents that dominate LFP generation, not synaptic currents.

In our simulations with the mouse V1 model, we also observed that the CSD magnitude and patterns could be significantly altered by the inclusion or exclusion of active channels, thus aligning with the findings of Reimann and colleagues. However, our simulations also demonstrated the crucial importance of synaptic inputs in shaping the CSD, and therefore the LFP. The model did not reproduce the experimental CSD until synaptic input from LM was added. Furthermore, we investigated the origins of the major sinks and sources in the canonical CSD pattern in the validated model by selectively removing certain inputs, connections, or ion channel types to eliminate their effects, and we found that several of the prominent sinks and sources stemmed from either external or recurrent synaptic input. Thus, even though our results did support the importance of active channels in generating the observed LFP patterns, they also highlight the critical role of currents through passive channels caused by synaptic input.

Both our V1 model and the rat S1 model were constructed with simplifications that may affect the conclusions that can be made regarding the role of active channels with either model. Neurons in our V1 model only have active channels at the soma, and none in the dendrites. Dendritic calcium spikes, for example, may constitute a significant share of the LFP due to their relatively large amplitude and longer duration [12], and this effect would not be captured by our model. In the rat S1 model, L1, L2/3, and L6 were not modeled, though anatomical studies of S1 have shown that there is substantial synaptic innervation in these layers [68]. Furthermore, both our simulations and anatomical and physiological data suggest that L1 and L2/3 are major target areas for cortico-cortical input in V1 [27, 39, 48, 62, 91]. The findings for V1 may not generalize to S1, but are at least an indication that the absence of these layers may have significant effects on the estimate of synaptic contributions to the LFP. In summary, neither model is yet sufficient to give a precise and final answer to the question of whether active currents or currents associated with synaptic input have the largest influence on LFP, but both models have made contributions that can motivate future experimental investigations and modeling studies where the effects of the abovementioned limitations are eliminated. Indeed, some such experimental studies have been conducted. Haider et al. (2016) [32] found a strong coupling between LFP and synaptic currents in individual cells, suggesting that at least for subthreshold activity, the LFP is better explained by synaptic than active currents.

Lastly, the findings in our study also demonstrated that by using the LFP in addition to spikes in the optimization of the V1 model, we obtain more constraints

to put on model parameters and architecture. The firing rates of the model reproduced experimental firing rates even before the feedback from higher visual areas was included in the model, which means that it was not apparent from the spikes alone that the effects of feedback were absent in the model. That only became apparent by studying the LFP/CSD. Furthermore, we found that the CSD pattern could be significantly altered by changing synaptic positions with minor effects on population firing rates, while the population firing rates could be significantly altered with minor effects on CSD by modifying synaptic weights. Together, these observations support the notion that LFP and spiking represent complementary aspects of neural activity, and that by using LFP in addition to spikes in the configuration and optimization of the model we reduce the degrees of freedom, and may get a model that is one step closer to real, biological networks.

4.2 Inferring population contributions and the validity of assuming linear LFP generation

Whether we can rigorously compare models to experimental data and show them to be predictive when there are more parameters than we can properly constrain is one challenge, another is whether we can infer neural activity from recorded signals when we have many degrees of freedom.

This too is a question of constraints: How to add constraints and to what extent they are sufficient to ensure accurate estimates. When decomposition tools such as ICA [21, 22, 52, 58, 59] and PCA [3–5, 16] have been applied to separate and identify contributions from different populations or pathways to recorded LFP, the intention was to take advantage of statistical constraints. The statistical constraints in ICA and PCA are that the sources to be uncovered generate LFP in a way that is, respectively, statistically independent or orthogonal. In certain structures and experimental conditions, this may indeed map onto the way that different pathways generate LFP [58].

However, these assumptions are unlikely to always be valid [26, 29]. Different populations may generate LFP in ways that are highly interdependent. In mouse V1, for example, the contributions arising from recurrent activity within V1 depend on prior input from LGN, so the contributions from LGN and V1 are not generated independently. Likewise, the contribution from higher visual areas depends on prior input from V1, and as such, is not independent from the V1 contribution. Thus, for a structure like V1, the assumption of statistical independence should be invalid.

A different approach to constraining the decomposition is to take advantage of physiological knowledge. When the process by which the LFP in a postsynaptic population is produced by presynaptic spikes can be assumed sufficiently linear, the LFP generation can be approximated by the convolution of firing rates with an appropriate kernel [18, 30]. This is the assumption underlying the methods in both the second and the third paper of this thesis. Both estimate LFP contributions from presynaptic populations by convolving their firing rate with kernels, but they differ in their construction of those kernels and, partially, in their purpose.

Laminar Population Analysis (LPA), the method expanded upon in the second paper, convolves the firing rates with exponential¹ functions that represent effects such as synaptic delays and filtering when presynaptic spikes generate postsynaptic LFP. The linearized framework in the third paper leverages information about membrane potentials, synaptic parameters, and connectivity patterns from the model to construct more detailed kernels for convolution with firing rates.

Since the linearized framework requires more detailed information about network properties and membrane potentials than is typically available in today's electrophysiological recordings, its application for the estimation of LFP contributions from presynaptic populations in experiments may be limited with current data. If the inter-animal variability in network properties is small enough, it may be possible to apply the linearized framework to experimental data too by utilizing the same kernel parameters, estimated from anatomical and physiological literature, for all animals and let the recorded presynaptic spike rates dictate differences in LFP estimates across animals. Adjudication of whether this approach is feasible awaits future investigation. LPA, on the other hand, can already be used on experimental data, and at least some of the idiosyncrasies of network properties of different animals can, in principle, be captured by the parameters of the exponential function kernels.

The greater detail in the kernels constructed in the linearized framework, however, appears to result in significantly higher precision in the estimated population LFP for the model - potentially high enough that it can reduce, and in some cases eliminate, the need for running simulations of the full biophysically detailed model. This would not only save substantial computational resources, it would also make the ability to predict extracellular potentials with large-scale models available to the whole neuroscientific community, not only those with access to high-performance computing centers.

Can LFP generation from presynaptic input in fact be assumed to be sufficiently linear? Many aspects of neuronal processing are not linear [49]. For example, synaptic currents have a non-linear dependence on synaptic conductance since it is influenced by membrane potential changes, and active channels in the dendrites can make the integration of input from multiple synapses non-linear [30]. However, if the influence of the non-linear processes is small enough or they can be well explained by linear approximations around typical values of e.g. membrane potentials or firing rates, the assumption of linearity can be acceptable [30].

As described in the preceding section, the extent to which currents through active channels shape the LFP generation is not fully settled. The mouse V1 model used in the validation of the kernel methods does not have active channels in the dendrites, so the validation on the V1 model is not affected by this particular aspect. The question is, however, still relevant when the LPA method is applied to experimental data, where neurons do have active channels in the dendrites. Furthermore, the neurons in the V1 model do have active channels at the soma, and our own simulations did demonstrate a role for active channels in shaping the LFP.

¹Other mathematical functions, such as the alpha function, can also be utilized, but exponential functions have produced the best results.

This suggests that either the kernel methods must be restricted to estimating contributions from synaptic currents and their associated return currents through passive channels, or it must be possible to approximate currents through active channels with linear processes. Some studies have indeed shown that the subthreshold dynamics of active channels can be approximated with linear models [70, 71], indicating that at least the subthreshold contribution of active channels can reasonably be treated linearly. Additionally, the study that used experimental data to demonstrate a strong coupling of synaptic currents and LFP through a linear regression indicates that the relationship can be approximated by linear processes also in experiments [32]. In the application of the linearized framework to L2/3 LFP in paper III, we found that the LFP was well estimated ($R^2 > 0.9$) for all channels except those nearest the center point of the somatic region. The poorer fit near the center point of the somatic distribution is most likely due to the fact that it coincides with the inversion point of the LFP. This means that the LFP amplitude is very small in this region and that noise makes up a significant part of the signal. Additionally, however, the contribution from active channels is probably the largest in this region, which means that the validity of the linear assumption may be limited here. However, the fact that the linearized framework produced precise estimates of LFP for most channels across different stimuli and for several different biophysically detailed models - both the network utilized in [30] and the layer 2/3 network of the Allen V1 model - supports to the contention that the assumption of largely linear LFP generation is often valid. Nevertheless, a systematic investigation into precisely when the assumption of a linear relationship between spikes and LFP generation is valid is an important line for future studies.

Motivated by the observation that current sinks and sources calculated from LFP are typically balanced across depth, we added a regularization term to the cost function in LPA that penalized deviations from zero in the CSD summed across channels. In other words, we amended the original LPA formulation to include another constraint based on physiology. After making this modification, we saw an improvement in the LPA-estimates of salient sinks and sources generated by external input from LGN and LM as well as those generated by recurrent activity within V1.

However, the decomposition with LPA was still not perfect. The estimates contained sinks and sources that were not present in the ground truth CSD, and the sinks and sources that did have the correct position and timing sometimes deviated in magnitude from the true CSD. The LPA-estimates may be ameliorated by other stimulus protocols [15], but the inaccuracies may also stem from simplifications or from restricted validity of linearity assumptions, as discussed above.

One such simplification is that when the temporal profiles of firing rates are estimated from the MUA, the firing rates of excitatory and inhibitory cells are merged into the firing rates of one population (except if they happen to be spatially non-overlapping). Excitatory currents generate a current sink while inhibitory currents generate a current source, i.e., they generate currents with opposite signs [12, 36]. But a population firing rate containing spikes from both excitatory and inhibitory currents will still only generate currents with one sign at a time, so it cannot estimate both effects simultaneously even though spikes from both

excitatory and inhibitory neurons are represented in the firing rate. In the V1 model, the excitatory neurons make up the majority of the neurons in the model ($\sim 85\%$), and we found that the excitatory populations explained 88.5% of the MUA variance. This suggests that the error from this simplification will be small for the V1 model. However, this may not be the case for other models or for experimental data.

Still, the CSD contributions from different structures were rendered more accurately than they were with the alternative statistical approaches, ICA and PCA. As such, it is a step forward, that can be improved upon, and it can already provide a first estimate of circuit architectures. The linearized framework may be another step forward in terms of accurate estimation of LFPs with the kernel methodology, at least for model data, and potentially also for experimental data in the future. The findings in these projects demonstrate the utility of exploiting physiological knowledge when developing analysis tools and attempting to infer neural activity from recorded signals. This shows that not only our models, but also our tools will improve as our physiological understanding improves and we have more constraints on our approximations and inferences of neural activity in the brain.

4.3 Outlook

In the introduction of this thesis, we highlighted EEG, which can be recorded non-invasively in humans, as a signal that could be better understood and interpreted by the study and modeling of LFP. None of the projects here investigates or attempts to model EEG, but an interesting and valuable continuation of this work would be to extend the model to generate simultaneous EEGs from the simulations, in addition to LFPs and spiking data. EEG data would add a macroscopic level to the existing microscopic level of single-cell activity and mesoscopic level of LFP. This would enlarge the pool of data that can be used in model development and open up the potential for further top-down constraints on model parameters and architecture.

This extension could be combined with simultaneous experimental recordings of LFP and EEG, such that the link between the experimental LFP and EEG can be investigated with simulations. The simulations could then further motivate adjustments to the experimental setup to answer questions that may arise, bolstering our understanding of the relationship between neural activity and the observed LFP and EEG signals. Adding another scale to an already complex model will come with its own set of challenges. To model the EEG, the effect of the propagation of electrical signals through the cerebrospinal fluid, the skull, and the scalp would have to be incorporated. Additionally, the signals recorded with EEG electrodes are expected to reflect a larger area than the 1 mm scale of the current V1 model [34], so the influence of neighboring areas would have to be approximated. Nonetheless, it would be an intriguing step toward linking all scales of neural activity and measurement signals.

The construction of the large-scale, biophysically detailed mouse V1 model was made possible by data collected in recent years. In our validation of the model

against experimental data, we used the Visual Coding data set obtained with Neuropixels 1.0 probes, which records from 384 channels over a shank 1 cm in length [46]. Since the Visual Coding data set was released, both Neuropixels 2.0 [95] and Neuropixels Ultra [108] probes have been developed. These probes have greater channel density (1333/mm for Neuropixels Ultra, 133/mm for Neuropixels 2.0, vs. 100/mm for Neuropixels 1.0) and come with software for automatic motion correction to increase stability in chronic recordings, resulting in both greater yield and quality of cell recordings. The continued evolution of recording devices will help to alleviate some of the difficulties listed in Traub et al. (2005). While these devices are being developed and put to use, the models and tools we develop in parallel can hopefully expand our insight into the neural circuits and both inspire and inform future studies that bridge our understanding of the microscopic and the macroscopic parts of the brain.

Chapter 4. Discussion

Paper I

**Uncovering circuit mechanisms of
current sinks and sources with
biophysical simulations of primary
visual cortex**

I

Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex

Atle E Rimehaug^{1*}, Alexander J Stasik², Espen Hagen^{2,3}, Yazan N Billeh⁴, Josh H Siegle⁴, Kael Dai⁴, Shawn R Olsen⁴, Christof Koch⁴, Gaute T Einevoll^{2,5}, Anton Arkhipov^{4*†}

¹Department of Informatics, University of Oslo, Oslo, Norway; ²Department of Physics, University of Oslo, Oslo, Norway; ³Department of Data Science, Norwegian University of Life Sciences, Ås, Norway; ⁴MindScope Program, Allen Institute, Seattle, United States; ⁵Department of Physics, Norwegian University of Life Sciences, Ås, Norway

Abstract Local field potential (LFP) recordings reflect the dynamics of the current source density (CSD) in brain tissue. The synaptic, cellular, and circuit contributions to current sinks and sources are ill-understood. We investigated these in mouse primary visual cortex using public Neuropixels recordings and a detailed circuit model based on simulating the Hodgkin–Huxley dynamics of >50,000 neurons belonging to 17 cell types. The model simultaneously captured spiking and CSD responses and demonstrated a two-way dissociation: firing rates are altered with minor effects on the CSD pattern by adjusting synaptic weights, and CSD is altered with minor effects on firing rates by adjusting synaptic placement on the dendrites. We describe how thalamocortical inputs and recurrent connections sculpt specific sinks and sources early in the visual response, whereas cortical feedback crucially alters them in later stages. These results establish quantitative links between macroscopic brain measurements (LFP/CSD) and microscopic biophysics-based understanding of neuron dynamics and show that CSD analysis provides powerful constraints for modeling beyond those from considering spikes.

*For correspondence: atleeri@ifi.uio.no (AER); antona@alleninstitute.org (AA)

†Lead contact

Competing interest: See page 24

Funding: See page 24

Preprinted: 25 February 2022

Received: 23 February 2023

Accepted: 10 July 2023

Published: 24 July 2023

Reviewing Editor: Tirin Moore, Howard Hughes Medical Institute, Stanford University, United States

© Copyright Rimehaug *et al.* This article is distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use and redistribution provided that the original author and source are credited.

Editor's evaluation

The study demonstrates that utilizing the LFP and/or the CSD in modeling can facilitate model configuration and implementation by revealing discrepancies between models and experiments. The analysis of the biophysical origin of the canonical CSD using the model is an interesting and worthy line of investigation. The dissection of CSD components is detailed and exhaustive. A key novelty of this article is the addition of CSD patterns as another constraint to more accurately infer the model parameters beyond its prior state.

Introduction

The local field potential (LFP) is the low-frequency component (below a few hundred Hertz) of the extracellular potential recorded in brain tissue that originates from transmembrane currents in the vicinity of the recording electrode (Lindén *et al.*, 2011; Buzsáki *et al.*, 2012; Einevoll *et al.*, 2013; Pesaran *et al.*, 2018; Sinha and Narayanan, 2022). While the high-frequency component of the extracellular potential, the single- or multi-unit activity (MUA), primarily reflects action potentials of one or more nearby neurons, the LFP predominantly stems from currents caused by synaptic inputs

(Mitzdorf, 1985; Einevoll et al., 2007) and their associated return currents through the membranes. Thus, cortical LFPs represent aspects of neural activity that are complementary to those reflected in spikes, and as such, they can provide additional information about the underlying circuit dynamics from extracellular recordings.

Applications of LFP are diverse and include investigations of sensory processing (Rall and Shepherd, 1968; Di et al., 1990; Victor et al., 1994; Kandel and Buzsáki, 1997; Mehta et al., 2000a; Mehta et al., 2000b; Henrie and Shapley, 2005; Einevoll et al., 2007; Belitski et al., 2008; Montemurro et al., 2008; Niell and Stryker, 2008; Nauhaus et al., 2009; Bastos et al., 2015; Senzai et al., 2019), motor planning (Scherberger et al., 2005; Roux et al., 2006), navigation (Tort et al., 2008; Makarova et al., 2011; Fernández-Ruiz et al., 2012; Watrous et al., 2013; Fernández-Ruiz et al., 2017), and higher cognitive processes (Pesaran et al., 2002; Womelsdorf et al., 2006; Liu and Newsome, 2006; Kreiman et al., 2006; Liebe et al., 2012). The LFP is also a promising signal for steering neuroprosthetic devices (Mehring et al., 2003; Andersen et al., 2004; Rickert et al., 2005; Markowitz et al., 2011; Stavisky et al., 2015) and for monitoring neural activity in human recordings (Mukamel and Fried, 2012) because the LFP is more easily and stably recorded in chronic settings than spikes. Due to the vast number of neurons and multiple neural processes contributing to the LFP, however, it can be challenging to interpret (Buzsáki et al., 2012; Einevoll et al., 2013; Hagen et al., 2016). While we have extensive phenomenological understanding of the LFP, less is known about how different cell and synapse types and connection patterns contribute to the LFP or how these contributions are sculpted by different information processing streams (e.g., feedforward vs. feedback) and brain states.

One way to improve its interpretability is to calculate the current source density (CSD) from the LFP, which is a more localized measure of activity, and easier to read in terms of the underlying neural processes. The current sinks and sources indicate where positive ions flow into and out of cells, respectively, and are constrained by Kirchoff's current law (i.e., currents sum to zero over the total membrane area of a neuron). However, the interpretation of current sinks and sources is inherently ambiguous as several processes can be the origin of a current sink or source (Buzsáki, 2006; Pettersen et al., 2006; Einevoll et al., 2007). For example, a current source may reflect an inhibitory synaptic current or an outflowing return current resulting from excitatory synaptic input elsewhere on the neuron. There is no simple way of knowing which it is from an extracellular recording alone (Buzsáki, 2006).

Another approach to uncovering the biophysical origins of current sinks and sources, and by extension the LFP, is to simulate them computationally (Pettersen et al., 2008; Einevoll et al., 2013). Following the classic work by Rall in the 1960s (Rall, 1962), a forward-modeling scheme in which extracellular potentials are calculated from neuron models with detailed morphologies using volume conduction theory under the line source approximation has been established (Holt and Koch, 1999). With this framework, we have achieved a good understanding of the biophysical origins of extracellular action potentials (Koch, 1998; Holt and Koch, 1999; Pettersen and Einevoll, 2008; Hay et al., 2011; Lindén et al., 2010). Expanding on this understanding, models composed of populations of unconnected neurons (e.g., Pettersen et al., 2008; Lindén et al., 2011; Schomburg et al., 2012; Łęski et al., 2013; Sinha and Narayanan, 2015; Hagen et al., 2017; Ness et al., 2018) and recurrent network models (e.g., Traub et al., 2005; Vierling-Claassen et al., 2010; Reimann et al., 2013; Głabaska et al., 2014; Tomsett et al., 2015; Hagen et al., 2016; Hagen et al., 2018; Chatzikalymniou and Skinner, 2018) have been used to study the neural processes underlying LFP.

While interesting insights about CSD and LFP were obtained from these computational approaches, establishing a direct relationship between the biological details of the circuit structure and the electrical signal like LFP remains a major unresolved challenge. One reason is that the amount and quality of data available for modeling the circuit architecture in detail have been limited. This situation improved substantially in recent years, and a broad range of data on the composition, connectivity, and physiology of cortical circuits have been integrated systematically (Billeh et al., 2020) in a biophysically detailed model of mouse primary visual cortex (area V1). In addition, significant improvements were achieved in experimental recordings of the LFP and the simultaneous spiking responses. In particular, the Neuropixels probes (Jun et al., 2017) record LFP and hundreds of units across the cortical depth in multiple areas, with 20 μm spacing between recording channels allowing for an unprecedented level of spatial detail. These developments provide unique opportunities to improve our understanding of circuit mechanisms that determine LFP patterns.

Here, we analyze spikes and LFP from the publicly available Allen Institute's *Visual Coding* survey recorded using Neuropixels probes (<https://www.brain-map.org>; [Siegle et al., 2021](#)) and reproduce these using the mouse V1 model developed by [Billeh et al., 2020](#). The model is comprised of >50,000 biophysically detailed neuron models surrounded by an annulus of almost 180,000 generalized leaky-integrate-and-fire units. The neuron models belong to 17 different cell type classes: one inhibitory class (Htr3a) in layer 1, and four classes in each of the other layers (2/3, 4, 5, and 6) where one is excitatory and three are inhibitory (Pvalb, Sst, Htr3a) in each layer. The visual coding dataset consists of simultaneous recordings from six Neuropixels 1.0 probes across a range of cortical and subcortical structures in 58 mice while they are exposed to a range of visual stimuli (about 100,000 units and 2 billion spikes over 2 hr of recording).

In our analysis of this dataset, we identified a canonical CSD pattern that captures the evoked response in mouse V1 to a full-field flash. We then modified the biophysically detailed model of mouse V1 to reproduce both the canonical CSD pattern and laminar population firing rates in V1 simultaneously. We reproduce, in a quantitative manner, the shape and timing of the pattern of current sources and sinks that have been described in considerable detail by experimentalists (e.g., [Mitzdorf, 1987](#); [Swadlow et al., 2002](#); [Senzai et al., 2019](#)). This shows that adjustments to synaptic parameters such as weights and placement in addition to a circuit architecture that included feedback are sufficient to reproduce experimental findings on both single-cell measures such as spikes and population-level measures such as CSD. We use this model to explain, in a highly mechanistic manner, the biophysical origins of the various ionic current sinks and sources and their location across the various layers of visual cortex.

In the process of obtaining a model that could reproduce both spikes and CSD, we discovered that the model can be modified by adjusting the synaptic weights to reproduce the experimental firing rates with only minor effects on the simulated CSD, and, conversely, that the simulated CSD can be altered with only minor effects on the firing rates by adjusting synaptic placement. Furthermore, we found that comparing the simulated CSD to the experimental CSD revealed discrepancies between model and data that were not apparent from only comparing the firing rates. Additionally, it was not until feedback from higher cortical visual areas (HVAs) was added to the model that simulations reproduced both the experimentally recorded CSD and firing rates, as opposed to only the firing rates. This bio-realistic modeling approach sheds light on specific components of the V1 circuit that contribute to the generation of the major sinks and sources of the CSD in response to abrupt visual stimulation. Our findings demonstrate that utilizing the LFP and/or the CSD in modeling can aid model configuration and implementation by revealing discrepancies between models and experiments and provide additional constraints on model parameters beyond those offered by the spiking activity. The new model obtained here is freely accessible (<https://doi.org/10.5061/dryad.k3j9kd5b8>) to the community to facilitate further applications of biologically detailed modeling.

Results

Spikes and LFP were recorded across multiple brain areas, with a focus on six cortical (V1, LM, AL, RL, AM, PM) and two thalamic (LGN, LP) visual areas, using Neuropixels probes in 58 mice ([Siegle et al., 2021](#)).

A schematic of the six probes used to perform the recordings in individual mice is shown in [Figure 1A](#), and the spikes and LFP recorded in V1 of an exemplar mouse during presentation of a full-field bright flash stimulus are displayed in [Figure 1B, C](#). The CSD can be estimated from the LFP (averaged over 75 trials) using the delta iCSD method to obtain a more localized measure of inflowing (sinks) and outflowing currents (sources) ([Pettersen et al., 2006](#); [Einevoll et al., 2013](#)). The biophysically detailed model of mouse V1 used to simulate the neural activity and the recorded potential in response to the full-field flash stimulus is illustrated in [Figure 1E](#). The model contains 230,924 neurons, of which 51,978 are biophysically detailed multicompartment neurons with somatic Hodgkin–Huxley conductances and passive dendrites, and 178,946 are leaky-integrate-and-fire (LIF) neurons. These neuron models are arranged in a cylinder with a radius 845 μm and a height 860 μm . The multicompartment neurons are placed in the 'core' with a radius of 400 μm , while the LIF neurons form an annulus surrounding this core. Cellular models belong to 17 different classes: one excitatory class and three inhibitory (Pvalb, Sst, Htr3a) in each of layers 2/3, 4, 5 and 6, and a single Htr3a inhibitory class in layer 1. The extracellular electric field in the model was recorded on an array of simulated

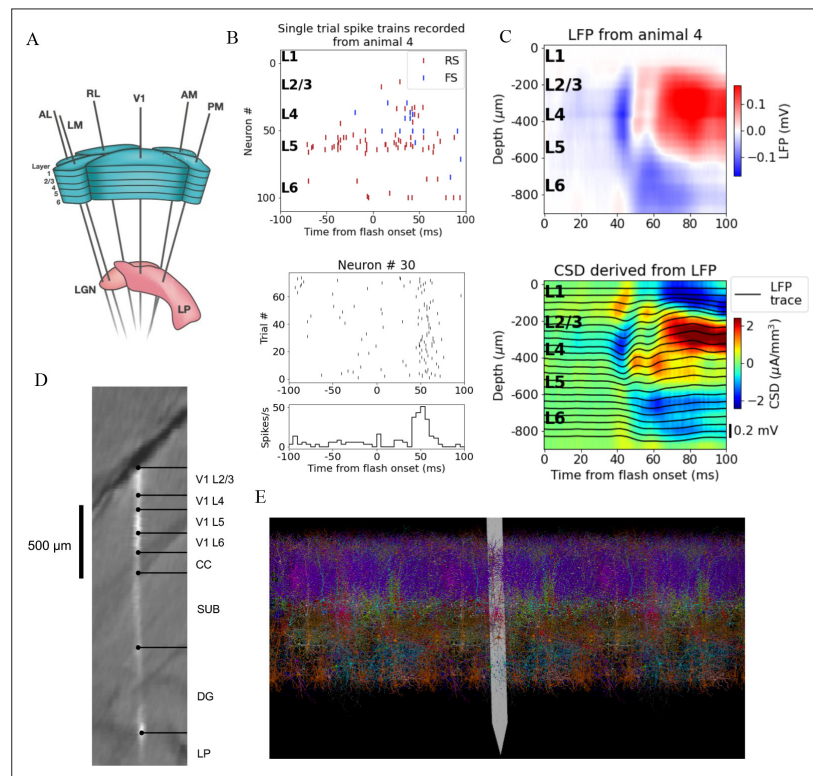


Figure 1. Illustration of experimental data and the biophysical model for mouse primary visual cortex (V1). **(A)** Schematic of the experimental setup, with six Neuropixels probes inserted into six cortical (V1, latero-medial [LM], rostro-lateral [RL], antero-lateral [AL], postero-medial [PM], AM) and two thalamic areas (LGN, LP). **(B)** Top: spikes from many simultaneously recorded neurons in V1 during a single trial. Bottom: spikes from a single neuron recorded across multiple trials. In both cases, the stimulus was a full-field bright flash (onset at time 0, offset at 250 ms). **(C)** Top: local field potential (LFP) across all layers of V1 in response to the full-field bright flash, averaged over 75 trials in a single animal. Bottom: current source density (CSD) computed from the LFP with the delta iCSD method. **(D)** Histology displaying trace of the Neuropixels probe across layers in V1, subiculum (SUB) and dentate gyrus (DG). **(E)** Visualization of the V1 model with the Neuropixels probe in situ. (Image made using VND.)

point electrodes (*Dai et al., 2020*) arranged in a straight line (*Figure 1D*) and separated by 20 μm , consistent with Neuropixels probes, shown in *Figure 1E* to scale with the model.

Uncovering a canonical visually evoked CSD response

We first established a ‘typical’ experimentally recorded CSD pattern to be reproduced with the model. Though there is substantial inter-trial and inter-animal variability in the evoked CSD response, we find that most trials and animals have several salient features in common. In *Figure 2A*, the trial-averaged evoked CSDs from five individual mice are displayed. In the first four animals (#1–4), we observe an early transient sink arising in layer 4 (L4) ~40 ms after flash onset, followed by a sustained source starting ~60 ms, which covers L4 and parts of layers 2/3 (L2/3) and layer 5 (L5). We also observe a sustained sink covering layers 5 and 6 (L6) emerging around 50 ms, as well as a sustained sink covering layers 1 and 2/3 around 60 ms. An animal that does not fully exhibit what we term the ‘canonical’ pattern is shown in the rightmost plot (#5 in *Figure 2A*); it has an early L4 sink arising at 40 ms, but this sink is not followed by the sustained sinks and sources from 50 to 60 ms and onward observed in the other animals. The timing and location of sinks and sources are, overall, similar to those described earlier by *Givre et al., 1994*; *Schroeder et al., 1998*, *Niell and Stryker, 2008*, and *Senzai et al., 2019*.

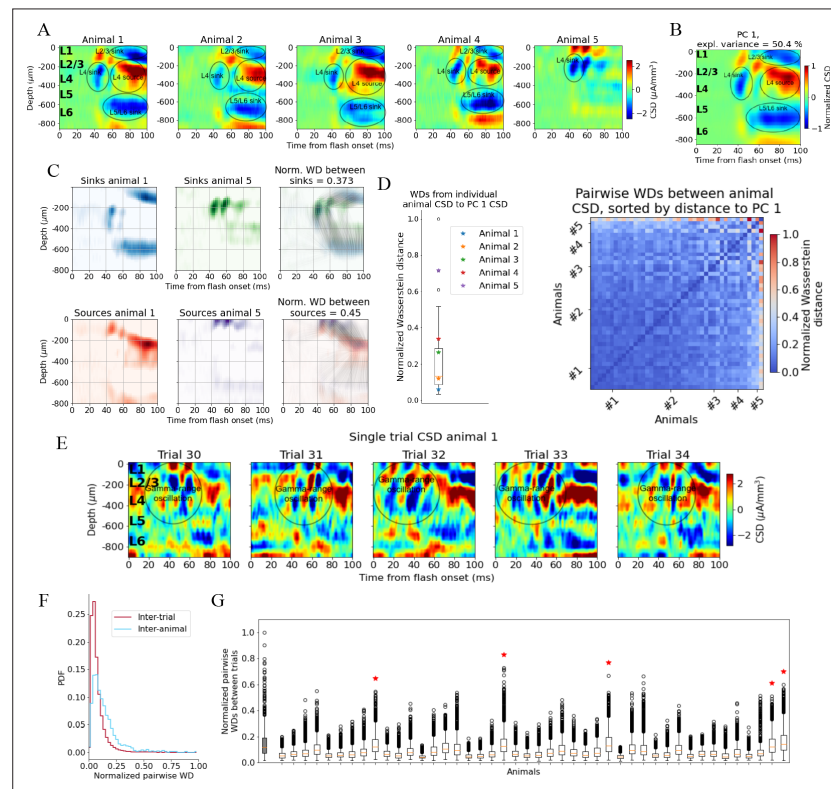


Figure 2. Variability in experimentally recorded current source density (CSD). (A) Evoked CSD response to a full-field flash averaged over 75 trials, from five animals in the dataset. (B) The first principal component (PC) computed from the CSD of all $n = 44$ animals, explaining 50.4% of the variance. (C) Illustration of movement of sinks and sources in the calculation of the Wasserstein distance (WD) between the CSD of two animals in the dataset. The gray lines in the rightmost panels display how the sinks or sources of one animal are moved to match the distribution of sinks or sources of the other animal. (D) Left: WDs from each animal to the PC 1 CSD. Right: pairwise WDs between all 44 animals sorted by their distance to the first PC. (E) CSD from five individual trials in example animal 1. (F) Distribution of pairwise distances between single-trial CSD (red) and pairwise distances between trial-averaged CSD of individual animals (blue). Both are normalized to the maximum pairwise distance between the trial-averaged CSD of individual animals. (G) Pairwise WDs between trials in each of 44 animals (white boxplots), normalized to maximal pairwise WDs between trial-averaged CSD of animals. Gray-colored boxplot shows the distribution of pairwise WDs between trial-averaged CSD of individual animals, and the red stars indicate the $n = 5$ animals for which the inter-trial variability was greater than the inter-animal variability (assessed with Kolmogorov–Smirnov [KS] tests, $p < 0.001$ in all cases, see **Figure 2—figure supplement 3**).

The online version of this article includes the following figure supplement(s) for figure 2:

Figure supplement 1. Trial-averaged current source density (CSD) during presentation of full-field flashes for all 44 animals in this study.

Figure supplement 2. Principal component analysis (PCA) on histology-aligned current source density (CSD).

Figure supplement 3. Comparing inter-trial and inter-animal pairwise Wasserstein distances (WDs).

To identify the robust features across animals in this dataset, we performed principal component analysis (PCA) on the trial-averaged evoked CSD from all animals. Out of the 58 animals in the dataset, 5 did not have readable recordings of LFP in V1 during the presentation of the full-field flash stimuli, and the exact probe locations in V1 could not be recovered for 9 other animals due to fading of fluorescent dye or artifacts in the optical projection tomography (OPT) volume (see ‘Materials and methods’). The remaining 44 (out of the 58) animals in the dataset were retained for the CSD analysis. The trial-averaged CSD plots of all these 44 animals are displayed in **Figure 2—figure supplement 1**. The first principal component (PC 1) (**Figure 2B**) constitutes a sum of weighted contributions of the

CSD patterns from all 44 animals and explains half (50.4%) of the variance. The salient features typically observed in individual animals are also prominent in the PC 1 CSD pattern (**Figure 2B**), that is, the canonical pattern. In **Figure 2—figure supplement 2**, the first 10 principal components cumulatively explaining 90% of the variance are plotted.

Quantifying CSD pattern similarity

We use the Wasserstein, or Earth Mover's, distance (WD) to quantify the differences in CSD patterns (see 'Materials and methods'), which can then be used to assess how well the simulated CSD matches the CSD typically observed in experiments. The WD reflects the cost of transforming one distribution into another by moving its 'distribution mass' around (**Rubner et al., 1998; Arjovsky et al., 2017**). An often-used analogy refers to the two distributions as two piles of dirt, where the WD tells us the minimal amount of work that must be done to move the mass of one pile around until its distribution matches the other pile (**Rubner et al., 1998**). In the context of CSD patterns, the WD reflects the cost of transforming the distribution of sinks and sources in one CSD pattern into the distribution of sinks and sources in another pattern, with larger WD indicating greater dissimilarity between CSD patterns. The WDs are computed between the sinks of two CSD patterns and between the sources of two CSD patterns independently, and then summed to form a total WD between the CSD patterns (**Figure 2C**). The sum of all sinks and the sum of all sources in each CSD pattern are normalized to -1 and $+1$, respectively, so the WD only reflects differences in patterns, and not differences in the overall amplitude. The WD scales linearly with shifts in space and time.

When computing the WDs between the evoked CSD patterns of individual animals and the canonical pattern, we find that the animals with CSD patterns that, by visual inspection, resemble the canonical pattern (**Figure 2A**, animals 1–4), are indeed among animals with smaller WD, while the animal with the more distinct CSD pattern (**Figure 2A**, animal 5) is an outlier (**Figure 2D**).

The onset of the evoked response is less conspicuous in the single-trial CSD due to pronounced, ongoing sinks and sources, but there is still a visible increase in magnitude from 40 to 50 ms onward (**Figure 2E**), compatible with the latency of spiking responses to full-field flashes in V1 (**Siegle et al., 2021**). An oscillation of sinks and sources with a periodicity of ~ 20 ms, that is, in the gamma range is apparent in the region stretching from L2/3 to the top of L5, which appears to be either partially interrupted or drowned out by more sustained sinks and sources emerging at about 60 ms. At least some of this gamma-range activity derives from the visual flash that covers the entire visual field and that drives retinal neurons and postsynaptic targets in the lateral geniculate nucleus (LGN) in an oscillatory manner (see the pronounced gamma-range oscillation in the LGN firing rate in **Figure 3D**).

The inter-trial variability is roughly comparable to the inter-animal variability of the trial-averaged responses. By computing the pairwise Wasserstein distances between single trial CSDs within each animal and comparing it to the pairwise WD between the trial-averaged CSD of each animal, we find that inter-trial variability in CSD is significantly lower than the inter-animal variability in trial-averaged CSD (Kolmogorov–Smirnov distance = 0.33; $p < 0.001$) (**Figure 2F**).

The majority of animals (39 out of 44) have a WD to the first principal component, PC 1, of the CSD that is less than half of the greatest WD between the CSD of individual animals and the PC 1 CSD (**Figure 2D**); the pairwise WDs between animals are also less than half of the maximum pairwise WD for most animals (921 out of the total 946 pairwise WDs; **Figure 2E**). This supports the view that most animals exhibit the canonical CSD pattern captured by the PC 1 CSD (**Figure 2B**). The total inter-trial variability is smaller than the inter-animal variability, both estimated by pairwise WDs (**Figure 2F and G**), though there are $n = 5$ animals for which the inter-trial WDs are larger than the inter-animal WDs (**Figure 2G**, marked by red stars; determined with KS tests on the distribution of pairwise WDs between animals and pairwise WDs between trials in each animal; see **Figure 2—figure supplement 3**).

Quantifying firing rate variability

For the spike analysis, we distinguish between fast-spiking (FS; putative Pvalb inhibitory) neurons and regular-spiking (RS; putative excitatory and non-Pvalb inhibitory) neurons (see 'Materials and methods' and **Figure 3—figure supplement 1**). All FS-neurons are grouped together into one population across all layers, while the RS-neurons are divided into separate populations for each layer (**Figure 3A**). The FS-neurons are merged across layers because we set a criterion of at least 10 recorded neurons in any

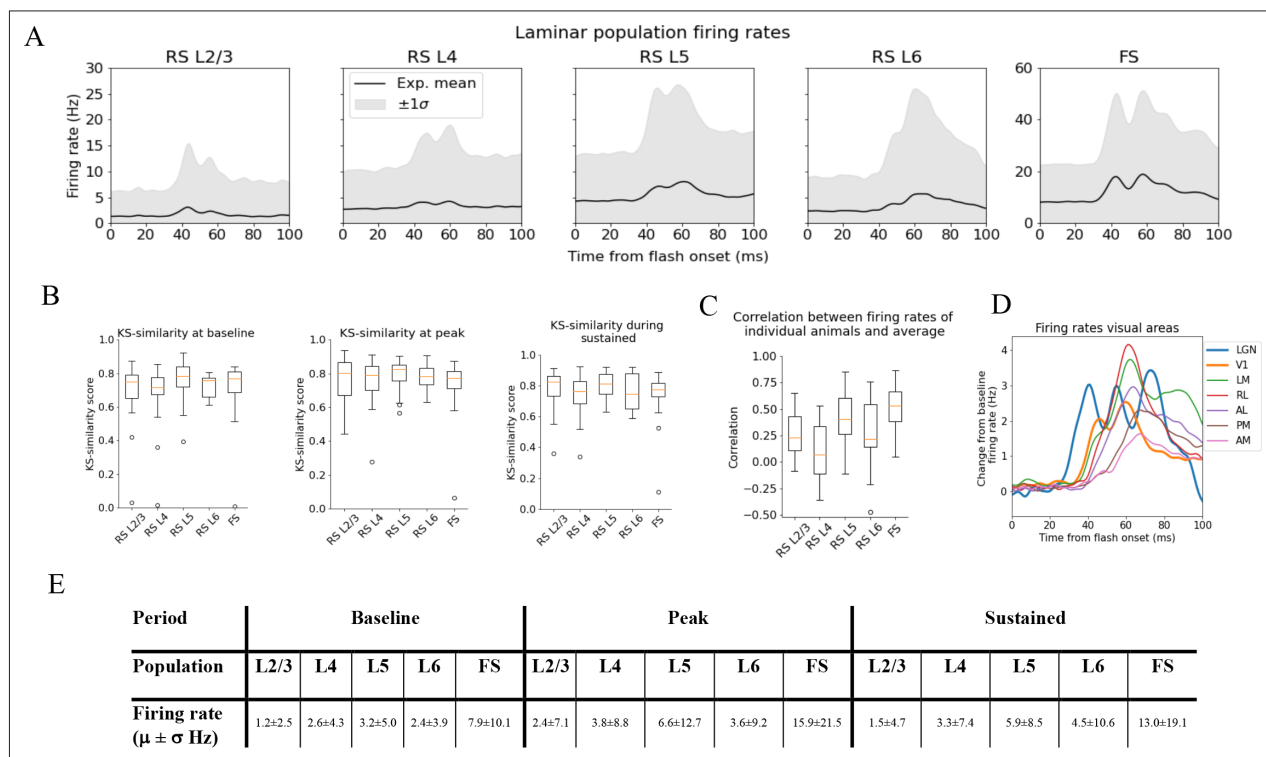


Figure 3. Variability in experimentally recorded spikes. **(A)** Trial-averaged laminar population firing rates of regular-spiking (RS) cells, differentiated by layer, and fast-spiking (FS) cells across all layers in response to full-field flash. Black line: average across all animals. Gray shaded area: ± 1 standard deviation. **(B)** Kolmogorov–Smirnov (KS) similarities (see ‘Materials and methods’) between the trial-averaged firing rates of each individual animal and the average firing rate over cells from all animals (black line in **A**) at baseline (the interval of 250 ms before flash onset), peak evoked response (from 35 to 60 ms after flash onset), and during the sustained period (from 60 to 100 ms). **(C)** Correlations between trial-averaged firing rates of individual mice and all mice (0–100 ms after flash onset). **(D)** Baseline-subtracted evoked firing rates for excitatory cells in seven visual areas (average over trials, neurons, and mice). Note the strong, stimulus-triggered gamma-range oscillations in the firing of lateral geniculate nucleus (LGN) neurons (blue). **(E)** Mean (μ) \pm standard deviation (σ) of population firing rates during baseline, peak evoked response, and the sustained period. Averaged across trials, neurons and time windows defined above.

The online version of this article includes the following figure supplement(s) for figure 3:

Figure supplement 1. Classifying cell types in experimental data.

Figure supplement 2. Number of cells in each population in experimental data.

one layer when comparing the population firing rate in individual animals to the average population firing rate in all animals, and only *one* animal had 10 FS-neurons or more in any layer (**Figure 3—figure supplement 2**). This criterion was set to have a more reliable estimate of the population firing rates in individual animals.

We use the KS similarity (defined as one minus the KS distance, see ‘Materials and methods’) and correlation to quantify the variability in firing rates. We use the experimental variability as a reference to assess whether the model reproduces firing rates typically observed in experiments. The KS similarity gives the similarity between the distributions of average firing rates across neurons in two populations in selected time windows, with KS similarity = 1 implying identity. As such, KS similarity provides a metric to compare the magnitudes of firing rates in certain time periods. We defined the ‘baseline’ window as the period over 250 ms before the flash onset, the ‘initial peak’ window as 35–60 ms after flash onset, and the ‘sustained’ window as 60–100 ms after flash onset. The KS similarity score during baseline is denoted ‘KSS_b’, during the ‘initial peak’ ‘KSS_p’, and ‘sustained’ ‘KSS_s’. The correlation, on the other hand, is computed between two population firing rates throughout the 100 ms window. The correlation thus gives us a measure of the similarity in the temporal profile of firing rates in this interval, independent of magnitudes. We establish the experimental variability in KS similarities

and correlation by computing these metrics between the population firing rates of each individual animal and the average population firing rates of all other animals (averaged over trials for both the individual animals and the average over all other animals) (Figure 3B and C).

The population firing rates for FS neurons are more than twice as high as RS cells during baseline, peak, and sustained. Among the RS populations, the firing rate in L5 is the highest in all periods, followed by L4 and L6, while L2/3 has the lowest firing rates (Figure 3E).

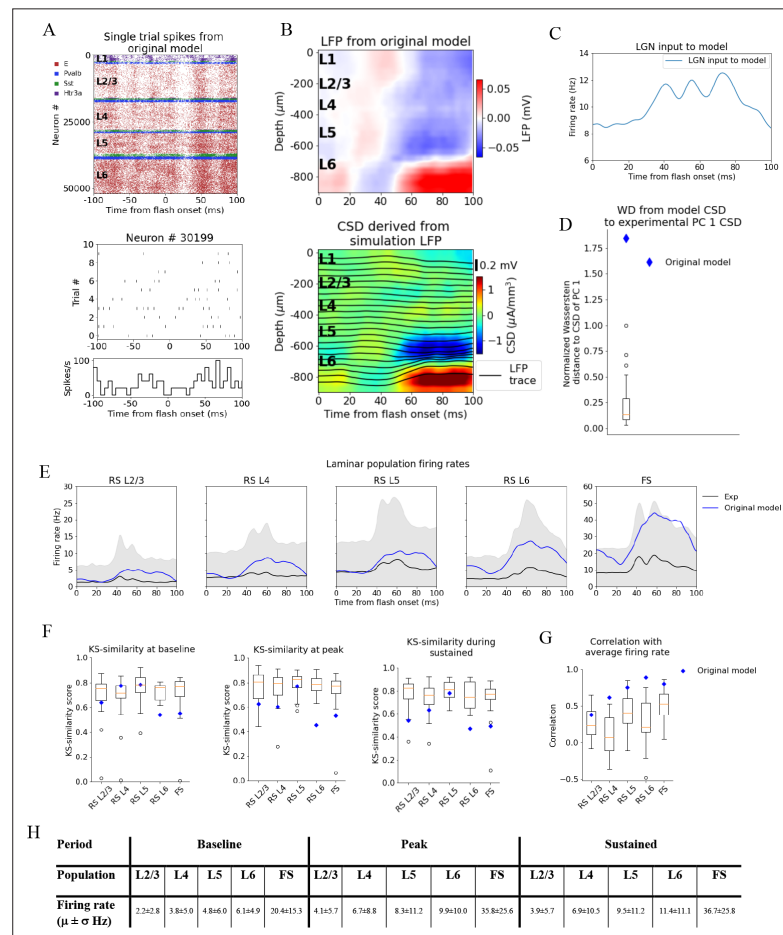


Figure 4. Local field potential (LFP), current source density (CSD), and spikes from simulations with the original model. (A) Top: raster plot of all ~50,000 cells in the model's 400 μm radius 'core' region spanning all layers, in a simulation of a single trial with the flash stimulus. Bottom: raster plot and histogram of spikes from 10 trials for an example cell. (B) Top: simulated LFP averaged over 10 trials of flash stimulus. Bottom: CSD calculated from the LFP via the delta iCSD method. (C) Firing rate of experimentally recorded lateral geniculate nucleus (LGN) spike trains used as input to the model. (D) Wasserstein distance between CSD from the original model (blue diamond) and PC 1 CSD from experiments together with the Wasserstein distances from experimental CSD in every animal to PC 1 CSD (boxplot), normalized to maximal distance for animals. (E) Experimentally recorded firing rates (black) and simulated firing rates (blue). (F) Kolmogorov–Smirnov (KS) similarity between firing rates in original model (blue diamond) or individual animals (boxplots) and firing rates in experiments at baseline, peak evoked response, and during the sustained period (defined in Figure 3). (G) Correlation between firing rates of model (blue diamond) or individual animals in experiments (boxplots) and average population firing rates in experiments (0–100 ms). (H) Mean (μ) \pm standard deviation (σ) of model firing rates during baseline, peak evoked response, and the sustained period. Averaged across trials, neurons and time windows defined above.

The online version of this article includes the following figure supplement(s) for figure 4:

Figure supplement 1. Effect of reducing recurrent inhibition.

Discrepancy between the original model and experimental observations

We simulated the response to a full-field flash stimulus with the biophysical network model of mouse primary visual cortex as presented in *Billeh et al., 2020*. As input to the model, we used experimentally recorded LGN spike trains (*Figure 4C*; see 'Materials and methods'). A Poisson source, firing at a constant rate of 1 kHz, provides additional synaptic input to all cells, representing the influence from the rest of the brain ('background' input). The thalamocortical input consists of spike trains from 17,400 LGN units (*Arhipov et al., 2018; Billeh et al., 2020*). The public Neuropixels data contain recordings from 1263 regular-spiking LGN neurons across 32 mice during 75 trials of full-field bright flash presentations, resulting in 94,725 spike trains. To construct the input for each of our 10 simulation trials, we randomly sampled 10 unique subsets of spike trains from this pool until all 17,400 units had been assigned a spike train in each trial.

Figure 4A, B displays the resulting spiking pattern across all layers with its associated LFP. The inferred CSD exhibits a strong sink in the L5 and L6 region, matched by a strong source below it, both starting at ~50 ms after flash onset (*Figure 4B*, bottom). However, the early L4 sink, the later sustained L4 source, and the sustained L2/3 sink typically observed in the experimental CSD (*Figure 2A, B*) are either absent or too weak compared to the sink and source in L5 and L6. The WD from the simulated CSD to the experimental PC 1 CSD is greater than the WD between the CSD of the farthest outlier animal and the PC 1 CSD (WD = 1.84, normalized to the largest WD between CSD of individual animals and PC 1 CSD). Thus, using experimental variability as a reference, the CSD from this simulation is an outlier (*Figure 4C*).

The population firing rates of the model, the KS similarities and correlation between the model and the data, are plotted together with the data in *Figure 4D–F*. The magnitudes of the model firing rates are higher than the experimental firing rates in all populations and time windows (*Figure 4H*). However, the KS similarities between the model firing rates and the experimental firing rates are still within the minimum to maximum range of the boxplots for the RS L2/3, RS L4, and RS L5 cells in all time windows (*Figure 4F*), and during baseline for the FS cells. For RS L6 neurons, the KS similarities were among the outliers of the experiments in all time windows, while for FS neurons they were among the outliers during the peak and sustained windows (RS L2/3: $KSS_b = 0.62$, $KSS_p = 0.63$, and $KSS_s = 0.54$; RS L4: $KSS_b = 0.77$, $KSS_p = 0.60$, and $KSS_s = 0.63$; RS L5: $KSS_b = 0.77$, $KSS_p = 0.77$, and $KSS_s = 0.78$; RS L6: $KSS_b = 0.54$, $KSS_p = 0.45$, and $KSS_s = 0.47$; FS: $KSS_b = 0.54$, $KSS_p = 0.53$, and $KSS_s = 0.49$). The temporal profile of the model firing rates is above the minimum of the boxplots for all populations (RS L2/3: $r = 0.38^{***}$, RS L4: $r = 0.62^{***}$, RS L5: $r = 0.75^{***}$, RS L6: $r = 0.90^{***}$, FS: $r = 0.80^{***}$, $^{***}p < 0.001$). For RS L4 and RS L6, the model is in fact outside the experimental distribution, but it is an outlier in a positive sense; the model firing rates for these populations are more similar to the experimental average than the corresponding population firing rates in individual animals.

The original model studied in *Figure 4* produced firing rates and orientation and direction tuning consistent with recordings in vivo (*Billeh et al., 2020*) with some shortcomings, such as relatively slow responses of V1 to the onset of visual stimuli (*Arhipov et al., 2018; Billeh et al., 2020*). Here, we see even more inconsistencies reflected clearly in the CSD pattern. This demonstrates the importance of multi-modal characterization of such biologically detailed models. To investigate the properties of the cortical circuit that sculpt the CSD, we manipulated the model and observed how both the CSD and firing rate responses were improved to match the experimental data.

Adjusting the model to fit experimental firing rates

Due to the discrepancy between the magnitudes of the model firing rates and the experimental firing rates, especially with respect to the outliers of the modeled RS L6 and FS neurons, we selectively adjusted the recurrent synaptic weights. We left the synaptic weights between LGN and the V1 model unchanged since they were well constrained by data (*Billeh et al., 2020*).

We first reduced the synaptic weights from all excitatory populations to the FS PV-neurons by 30% to bring their firing rates closer to the average firing rate in this population in the experiments. This resulted in increased firing rates in all other (RS) populations due to the reduced activity of the inhibitory Pvalb-neurons (*Figure 4—figure supplement 1*). Therefore, we further applied reductions in the synaptic weights from all excitatory neurons to RS neurons and increases in the synaptic weights from inhibitory neurons to the RS neurons to bring their firing rates closer to the experimental average

firing rates. We multiplied the recurrent synaptic weights with factors in the [0.2, 2.5] range until we arrived at a set of weights where none of the model firing rates were among the experimental outliers in any time window ($KSS_b = 0.73$, $KSS_p = 0.77$, and $KSS_s = 0.70$; average across RS populations and the FS population) and temporal profiles (RS L2/3: $r = 0.49^{***}$, RS L4: $r = 0.63^{***}$, RS L5: $r = 0.71^{***}$, RS L6: $r = 0.87^{***}$, FS: $r = 0.86^{***}$; $*** p < 0.001$) (Figure 5A–C).

The resulting pattern (but not the magnitude) of the CSD, however, was largely unchanged (Figure 5D) compared to the original CSD (Figure 4B). The overall magnitude was reduced, and there were some traces of a sink arising at 40 ms after flash onset, and a L2/3 (and L1) sink after 60 ms, but they were substantially weaker relative to the L5/L6 dipole than they were in the experiments. Furthermore, the large and sustained L4 source after 60 ms was still either absent or too weak to be visible. The WD between the CSD from this version of the model and the experimental PC 1 CSD remained among the outliers of the animals (Figure 5E) (normalized WD = 1.26).

Two-way dissociation between spikes and CSD

Simulations demonstrate that the LFP, and the associated CSD, can be significantly altered by changes to synaptic placement (Einevoll et al., 2007; Pettersen and Einevoll, 2008; Lindén et al., 2010; Lindén et al., 2011; Łęski et al., 2013; Hagen et al., 2017; Ness et al., 2018). As observed in Figure 5A–E and Figure 5—figure supplement 1, adjustments to synaptic weights can modify the population firing rates substantially, yet without substantially changing the pattern of the CSD, that is, the placement and timing of sinks and sources. The inverse can also occur; that is, the CSD pattern can be altered extensively with only minor effects on firing rates (Figure 5F and G, Figure 5—figure supplements 2–4).

In the model's original network configuration, L4 excitatory neurons received geniculate input from synapses placed within 150 μm from the soma on both basal and apical dendrites, and excitatory, recurrent input from other V1 neurons within 200 μm from the soma on both basal and apical dendrites. We tested the effects of synaptic location by placing all synapses from both LGN and excitatory neurons onto the basal dendrites of L4 excitatory neurons (within the same ranges as in the original configuration). This increased the contribution from the L4 excitatory neurons to the total CSD (Figure 5F, middle row, leftmost plot) by a factor of ~ 2 and led to a dipole pattern with a single sink at the bottom and a single source at the top, as opposed to having two pairs of sinks and sources like in the case of the original synaptic placement (Figure 5F, top row; leftmost plot). The firing rate of the L4 excitatory cells, however, remained essentially unchanged by this modification (Figure 5G). On the other hand, placing all synapses from LGN and excitatory neurons onto the apical dendrites of L4 excitatory neurons resulted in even greater CSD magnitude from this population (Figure 5F, bottom row; leftmost plot), while the magnitude of its firing rates were reduced (Figure 5G). In this case, the pattern displayed a sink in the middle with a source above and below it. We also find that the somatic Hodgkin–Huxley channels significantly shapes the CSD pattern (Figure 5H), generating a source in the L4 region where the somata are localized.

We quantified the changes in CSD of the full model resulting from changes in recurrent synaptic weights in Figure 5—figure supplement 1 and the changes in firing rates resulting from the changes in synaptic placement onto L4 excitatory cells in Figure 5—figure supplement 2. Our analyses corroborate that the changes in CSD after changing only synaptic weights is minor, and that the changes in firing rate after changing only synaptic placement is minor. More examples of the disparate effects on CSD and spikes from adjusting placement of excitatory synapses on excitatory cells in L2/3 and L5 are displayed in Figure 5—figure supplement 3 and Figure 5—figure supplement 4, respectively.

These results indicate a two-way dissociation that can occur between CSD and firing rates of excitatory neurons. The firing rates can be changed without substantially changing the CSD by modifying the strength of synapses, while the CSD can be changed without substantially changing the firing rates by modifying synaptic location. This suggests that utilizing the CSD in the optimization of the model can provide constraints on the circuit architecture that could not be obtained from spikes alone.

Effects of feedback from higher visual areas to the model

Hartmann et al., 2019 found that feedback from higher visual areas (HVAs) can exert a powerful influence on the magnitude of the evoked LFP response recorded in V1 of macaque monkeys, particularly in the period 80–100 ms after stimulus onset. The sustained L2/3 sink and L4 source we observe in the

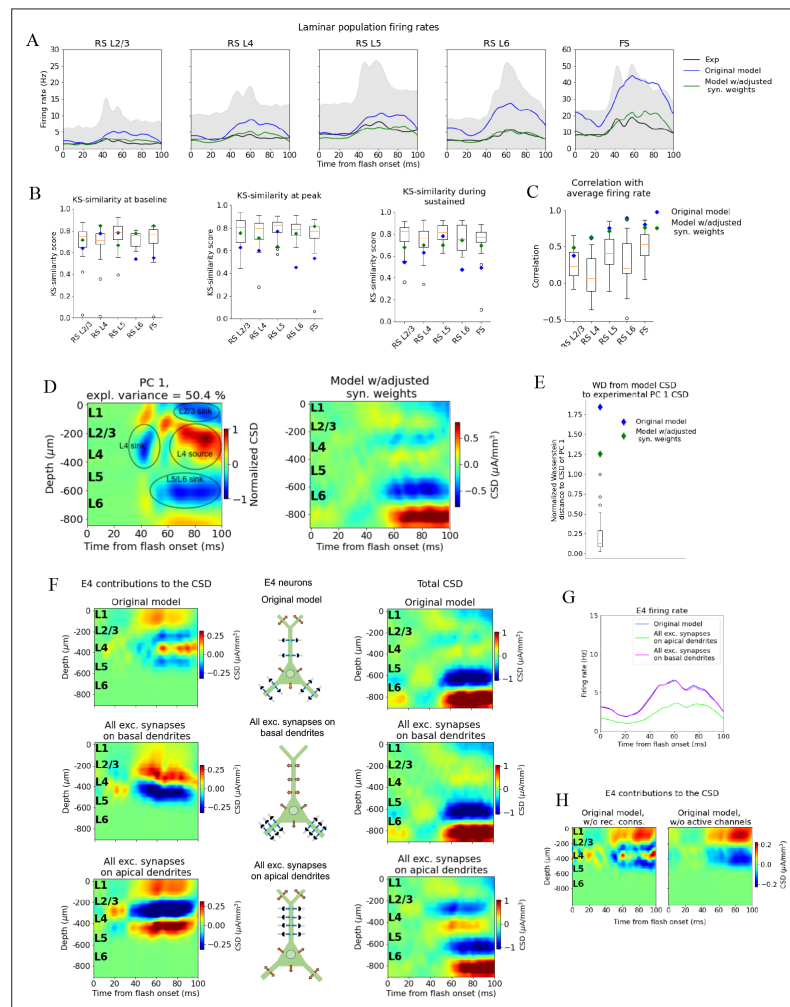


Figure 5. Adjusting the model to fit spikes or current source density (CSD). **(A)** Average experimentally (black) and simulated firing rates of experiments in the model with adjusted recurrent synaptic weights (green) and original model (blue). Synaptic adjustments included scaling the weights from all excitatory populations to the PV cells down by 30% to reduce the firing rates in these fast-spiking populations, reducing the synaptic weights from excitatory populations to all others and increasing synaptic weights from all PV cells to all other populations to compensate for the reduced inhibition. **(B)** Kolmogorov–Smirnov (KS) similarity between firing rates of model versions (markers) or individual animals in experiments (boxplots) and firing rates of experiments at baseline, peak evoked response, and during the sustained (defined in **Figure 3**). **(C)** Correlation between simulated firing rates or individual animals (boxplots) and measured firing rates (0–100 ms). **(D)** Left: PC 1 current source density (CSD) from experiments (see **Figure 2**). Right: CSD resulting from simulation on model with adjusted recurrent synaptic weights. **(E)** Wasserstein distance between CSD from model versions and PC 1 CSD from experiments together with Wasserstein distances from CSD in animals to PC 1 CSD (boxplot). **(F)** Effect of different patterns of placing excitatory synapses onto layer 4 excitatory cells on this population’s contribution to the simulated CSD (left) and to the total simulated CSD (right). These synaptic placement schemes with accompanying inflowing (blue arrows) and outflowing (orange arrows) currents are illustrated in the middle. **(G)** Effect of synaptic placement on the simulated population firing rate. **(H)** Contribution of L4 excitatory cells to the simulated CSD in the model where all recurrent connections have been cut (left) and when all active channels have been removed from all cells in the model (right).

The online version of this article includes the following figure supplement(s) for figure 5:

Figure supplement 1. Quantifying change in simulated current source density (CSD) with adjustments to synaptic weights.

Figure 5 continued on next page

Figure 5 continued

Figure supplement 2. Quantifying change in spiking of L4 excitatory cells after adjusting synaptic placement.

Figure supplement 3. Effects of manipulating synaptic placement onto L2/3 excitatory cells on population current source density (CSD) and spiking.

Figure supplement 4. Effects of manipulating synaptic placement onto L5 excitatory cells on population current source density (CSD) and spiking.

experimental CSD emerge at 60 ms (**Figure 2A and B**), which roughly coincides with the peak firing rates in the latero-medial (LM), rostro-lateral (RL), antero-lateral (AL), and postero-medial (PM) cortical areas (**Figure 3C**). Furthermore, anatomical data indicate that synapses from HVAs terminate on L1 and L2/3 apical dendrites of pyramidal cells (whose cell bodies reside in L2/3 or L5) (**Glickfeld and Olsen, 2017; Marques et al., 2018; Hartmann et al., 2019; Keller et al., 2020; Shen et al., 2020**). Together, these observations suggest that the sustained L2/3 sink and L4 source might, in part, be induced by feedback from HVAs, where the sink is generated from the input to the apical tufts in L1 and L2/3, and the source may be the return currents of this input.

Of these HVAs, the feedback from LM to V1 is best characterized (**Marques et al., 2018; Keller et al., 2020; Shen et al., 2020**) and has the highest connection density to V1 (**Harris et al., 2019**). Based on these considerations, we decided to test the hypothesis that the large sinks and sources in the upper layers were caused, at least in part, by feedback from LM. In addition to the earlier feed-forward LGN input and the background input representing the influence of the rest of the brain, we introduced a feedback input constructed from experimentally recorded spike trains in LM. In total, the public Neuropixels dataset has 2075 neurons recorded in LM (simultaneously with the recordings in LGN, V1, and other visual areas) from 42 animals during presentations of the full-field flash stimulus. 1823 of the 2075 neurons were classified as RS, and spike trains from these were used to generate the feedback input to the model (**Figure 6A**).

The synapses from this LM source were placed on the apical dendrites of L2/3 excitatory neurons (within 150 μm from the soma), on the apical tufts (>300 μm from the soma) and the basal dendrites (within 150 μm from the soma) of L5 excitatory neurons, and on the somata and basal dendrites of L2/3, and L5 inhibitory (Pvalb and Sst) neurons (at any distance from the soma). The input onto L2/3 excitatory neurons did generate a sink in L1 and L2/3 and a source below in L4 (**Figure 6B–E**).

The synaptic weights from LM to the populations targeted by the feedback were initialized at high values (see 'Materials and methods'), and then adjusted (decreased) by multiplying them with factors in the range [0.05, 0.5] (see 'Materials and methods'). The weights from the background to the feedback-targeted populations were also multiplied by factors in the range [0.2, 0.5], and the weights of connections from Pvalb neurons to L2/3 excitatory and L5 excitatory neurons were multiplied by factors in the range [0.8, 1.2]. This weight scaling was done until the population firing rates were within the experimental variability. Additionally, the synapses from excitatory populations onto L6 excitatory cells were restricted to be within 150 μm from the soma to reduce the magnitude of the L5/L6 dipole (**Figure 6—figure supplement 1**; see 'Materials and methods'). Greater separation between a sink and a source in a dipole moment increases its magnitude, so restricting the range within which synapses can be placed should diminish the L5/L6 dipole's dominance.

When the model received this feedback input together with the LGN input, the resulting CSD pattern reproduced the main features observed in the experiments (**Figure 6C**). The WD between the model CSD and the experimental PC 1 CSD was also no longer an outlier (Normalized WD = 0.41; **Figure 6D**), and the population firing rates remained within the minimum and maximum value of the experimental boxplots for the firing rates in all windows and all populations, both with respect to magnitudes ($KSS_b = 0.77$, $KSS_p = 0.70$, and $KSS_s = 0.68$; average across all populations) and temporal profiles (RS L2/3: $r = 0.36^{***}$, RS L4: $r = 0.64^{***}$, RS L5: $r = 0.69^{***}$, RS L6: $r = 0.87^{***}$, FS: $r = 0.77^{***}$, $^{***}p < 0.001$) (**Figure 6F and G**). Thus, when average responses to the full-field flash are considered, this final adjusted model exhibits both the CSD and firing rate patterns that are consistent with the experimental observations and are well within animal variability (**Figure 6F–H**).

Furthermore, we investigated whether the model could reproduce the stereotypical features of the single-unit firing response observed in experiments. To this end, we computed the moments of the distribution of peak firing rates as well as the distribution of latencies to the peak across cells both for individual animals and for the model versions (**Figure 6—figure supplements 2 and 3**). The original

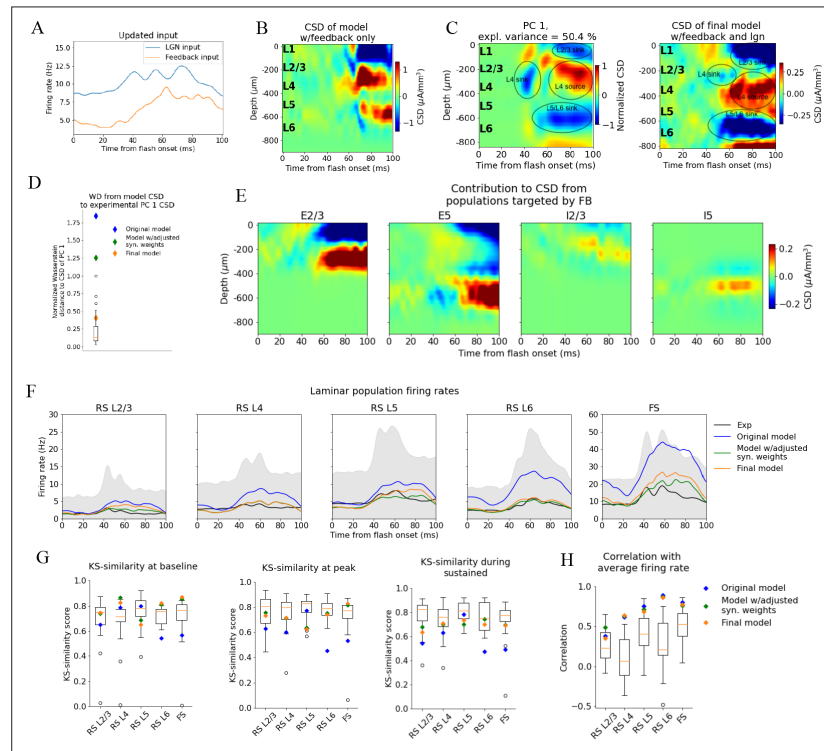


Figure 6. Introducing feedback from latero-medial (LM) to V1 in the model. **(A)** Firing rate of the experimentally recorded lateral geniculate nucleus (LGN) and LM units used as input to the model. **(B)** Total current source density (CSD) resulting from simulation with input only from the LM. **(C)** Left: PC 1 CSD from experiments (see **Figure 2**). Right: total CSD from simulation with both LGN input and LM input. **(D)** Wasserstein distance between CSD from model versions and PC 1 CSD from experiments together with Wasserstein distances from CSD in animals to PC 1 CSD (boxplot). **(E)** Population contributions from populations that receive input from LM. **(F)** Average population firing rates of experiments (black line) and model versions. **(G)** Kolmogorov–Smirnov (KS) similarity between simulated firing rates or individual animals (boxplots) and recorded firing rates at baseline, peak evoked response, and the sustained period (defined in **Figure 3**). **(H)** Correlation between simulated and experimentally recorded firing rates (0–100 ms).

The online version of this article includes the following figure supplement(s) for figure 6:

Figure supplement 1. Effect of adjusting synaptic placement onto L6 excitatory cells.

Figure supplement 2. Moments of distributions of peak firing rate in model versions and experiments for different populations.

Figure supplement 3. Moments of distributions of latency to peak of firing rates in model versions and experiments in different populations.

Figure supplement 4. Relative change in peak firing rates between neighboring populations.

Figure supplement 5. Moments of distributions of greatest curvature in firing rate across cells.

Figure supplement 6. Orientation and direction selectivity in final model.

Figure supplement 7. Current source density (CSD) analysis after aligning experimental CSD plots to landmarks rather than histology.

Figure supplement 8. Principal component analysis (PCA) on landmark aligned current source density (CSD).

Figure supplement 9. Effect of using plain average of trial-averaged current source density (CSD) from all animals instead of first principal component as target.

model was an outlier for the first moment of the peak firing rate distributions of RS L2/3, L4, L6, and FS populations, while the intermediate model was within the experimental variability for all populations, and the final model was within the experimental variability for all populations except RS L2/3, where it was just outside the maximum value of the experimental boxplot (**Figure 6—figure supplement 2A**). All three model versions were outside the experimental variability of the third moment (skewness) of the peak firing rate distributions for the RS L2/3 population, and the original model was just outside for the RS L5 and FS populations as well (**Figure 6—figure supplement 2C**). Otherwise, all model versions were within the experimental variability for all moments that we considered and for all populations.

With respect to the distribution of latencies to the peak firing rates, all model versions were within the experimental variability for all populations and all moments except the fourth moment (kurtosis) of the FS population, where all model versions were outside the experimental variability. We also computed the relative change in firing rates between neighboring populations (**Figure 6—figure supplement 4**) and the distributions of greatest curvature of the firing rate across cells (**Figure 6—figure supplement 5**), and found that all model versions were within the experimental variability on both of these metrics too. Thus, as with the analysis of population firing rates, the original model is an outlier when compared to the experiments on the features of unit firing, while the intermediate and final model versions reproduce most features observed in the experiments.

To check if the model continued to exhibit appropriate orientation and direction tuning after the adjustments made, we ran a simulation with the final model configuration and the same drifting grating stimulus that was utilized in *Billeh et al., 2020*. We found that the model still displayed firing rates at preferred directions and direction and orientation selectivity indices comparable to those observed experimentally (**Figure 6—figure supplement 6**).

Identifying the biophysical origins of the canonical CSD

With the canonical CSD (**Figure 2B**) reproduced, we can use the model to probe the biophysical origins of its sinks and sources. We began by removing all recurrent connections and only feeding the LGN input to the model to find the contribution from the thalamocortical synapses onto excitatory and inhibitory neurons (**Figure 7A**). The main thalamic contribution to the CSD is from synapses onto excitatory neurons, in line with the expectation that neurons with a spatial separation between synaptic input currents and the return currents dominate the cortical LFP generation (*Einevoll et al., 2013*). (Neurons without apical dendrites will have largely overlapping synaptic input currents and return currents, resulting in a cancellation of current sinks and sources.).

We further observed that the early L4 and the sustained L5/L6 sinks are present in the CSD contributions of excitatory neurons, though the magnitude of the L5/L6 sink is substantially reduced compared to its magnitude when the model is configured with recurrent synapses intact (**Figures 4B–6D**). The sustained L2/3 sink and L4 source, on the other hand, were not visible. This suggests that the early L4 sink and the L5/L6 sink are at least partly generated by thalamocortical synapses. However, the substantially diminished magnitude of the L5/L6 sink indicates that recurrent synapses also contribute significantly to the generation of this sink.

We then removed the LGN input and added the feedback (while keeping the recurrent connections cut), which resulted in a prominent upper layer dipole, with the sink residing in L1 and L2/3, and the source residing in L4 (**Figure 7B**). Together with their absence when input came from LGN only (**Figure 7A**), this suggests that the sustained L2/3 sink and the L4 source in the canonical pattern originate at least in part from the feedback synapses onto the apical dendrites of L2/3 and L5 pyramidal cells and the activity this input generates.

To assess the extent to which active channels at the somata contributed to the CSD pattern, we compared the CSD resulting from a simulation with both LGN and feedback input (where the recurrent connections were still cut) when we included or excluded the active channels (NaT, NaP, NaV, h, Kd, Kv2like, Kv3_1, K_T, Im_v2, SK, Ca_HVA, Ca_LVA; only at the soma [see supplementary information in *Gouwens et al., 2018* for definitions]) on all neurons in the model. The most prominent discrepancy between the CSD with and without active channels is the magnitude of the L4 source and the L2/3 sink (**Figure 7C**). In this all-passive setting, the L4 source is significantly attenuated, and the L2/3 sink is either absent or dominated by a source in the same region.

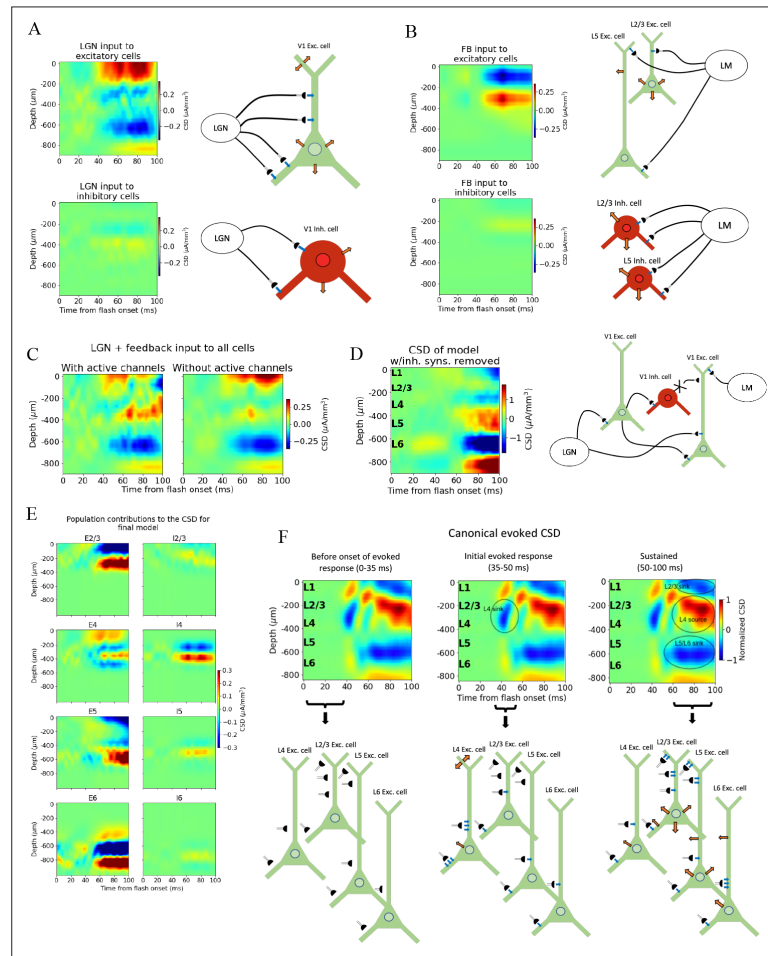


Figure 7. Biophysical origin of canonical current source density (CSD). **(A)** Sinks and sources generated from thalamocortical and **(B)** feedback synapses. The schematics illustrate which synapses cause the observed sinks and sources. Blue arrows indicate inflowing current (sinks), while orange arrows indicate outflowing current (sources). **(C)** Total CSD from thalamocortical and feedback synapses (without recurrent connections) with (left) and without (right) active channels in the V1 neurons. **(D)** Total CSD of model with both thalamocortical and feedback input when inhibitory synapses are removed (cross indicates removed connection). **(E)** Population contributions to the total CSD in final model with both lateral geniculate nucleus (LGN) and feedback input and recurrent connections. **(F)** Summary of biophysical origins of the main contributions to the sinks and sources in the canonical CSD in different periods of the first 100 ms after flash onset. More arrows mean more current. Left: before onset of evoked response (0–35 ms). The average inflowing and outflowing current in V1 neurons is zero in this time window. Middle: initial evoked response (35–50 ms). The L4 sink is primarily generated by inflowing current thalamocortical synapses onto L4 excitatory cells. Right: sustained evoked response (50–100 ms). The L5/L6 sink is primarily due to inflowing currents from thalamocortical synapses and recurrent excitatory synapses. Inflowing current at synapses from higher visual areas (HVAs) onto apical tufts of L2/3 and L5 excitatory cells generates, in part, the L2/3 sink, and the resulting return current generates, in part, the L4 source in this time window.

The online version of this article includes the following figure supplement(s) for figure 7:

Figure supplement 1. Effects of removing recurrent inhibition on population current source density (CSD) and firing rates.

Figure supplement 2. Effect of cell orientation on current source density (CSD) contributions of L4 inhibitory cells.

Figure supplement 3. Silencing feedback from latero-medial (LM) in model during evoked response.

We explored whether the contributions from currents in recurrent connections come primarily from excitatory or inhibitory synapses by removing all connections from inhibitory (Pvalb, Sst, Htr3a) neurons to all other neurons, so that all postsynaptic currents stem from excitatory thalamocortical synapses, excitatory synapses from HVAs, or recurrent excitatory synapses in V1 (**Figure 7D** and **Figure 7—figure supplement 1**). Note that inhibitory synaptic currents give rise to sources, while excitatory synaptic currents give rise to a sink. Of course, without inhibition, the network is unbalanced, which limits the conclusions that can be drawn from this simulation. However, the fact that the major sinks and sources are still present is an indication that the currents from excitatory input account for the majority of the sinks and sources observed in the experimental CSD.

The contributions from each population to the total CSD in the final model (**Figure 6D**) with both LGN and feedback input and intact recurrent connections are displayed in **Figure 7D**. From this, it is apparent that the L5/L6 dipole is mainly generated by L6 excitatory cells, the L2/3 sink stems from sinks at the apical tufts of the L2/3 and L5 excitatory cells, the L4 sink from both the L4 excitatory and inhibitory cells, while the L4 source is a mix of sources from mainly L2/3, L4, and L5 excitatory cells, as well as the L4 inhibitory cells. (The magnitude of the CSD contribution from L4 inhibitory cells is greater than anticipated. Given their lack of apical dendrites, we would expect their postsynaptic current sinks and sources to largely cancel [Einevoll et al., 2013]. Their contribution can be reduced by scrambling the 3-D orientation of these cells [**Figure 7—figure supplement 2**]. However, we cannot rule out that L4 inhibitory cells can have a contribution comparable in magnitude to the excitatory cells with the data we have available. We therefore let the L4 inhibitory cells keep their original orientation here.)

We investigated what would happen if we turned the feedback off again at 60 ms for the final model version (**Figure 7—figure supplement 3**). Most notably, we found that the sustained L4 source was replaced by a sink, and that the sustained L2/3 sink turned more transient as it was significantly reduced in magnitude after about 70 ms. This CSD pattern serves as a prediction that can be tested experimentally, for example, by optogenetic silencing of the HVAs in the sustained periods of the evoked response (e.g., see Keller et al., 2020).

We summarize the main contributions to the canonical CSD in **Figure 7F**. Before the onset of the evoked response (0–35 ms) there is, on average, no significant net inflow or outflow of current to any neurons. Around 40 ms, an inflow of current from excitatory thalamocortical synapses onto all excitatory neurons and all Pvalb inhibitory neurons appears, with the largest current coming from the synapses targeting basal and apical dendrites of L4 excitatory cells. This is the primary origin of the L4 sink. Following this initial L4 sink, there is a sustained sink in L5/L6 arising at ~50 ms, which originates partly from thalamocortical synapses onto L6 excitatory cells and partly from recurrent synapses from excitatory populations in V1 onto L6 excitatory cells. At ~60 ms, a sustained sink emerges in L1 and L2/3, which partly originates in synapses from HVAs targeting apical tufts of L2/3 and L5 excitatory cells. This feedback results in a stronger return current at the soma and basal dendrites of L2/3 excitatory cells and L5 excitatory cells.

Discussion

In the present study, we analyzed experimentally recorded spikes and LFP during presentation of full-field flashes from a large-scale visual coding dataset derived from mouse visual cortex (Siegle et al., 2021) and simulated the same experimental protocol using a biophysically detailed model of mouse V1 (Billeh et al., 2020). Our analysis of the experimental data focused on the responses in LGN and cortical visual areas V1 and HVAs. We found that the evoked CSD in V1, computed from the LFP, is captured by a canonical pattern of sinks and sources during the first 100 ms after stimulus onset (**Figure 2B**). This canonical CSD, in response to a flashed, bright field pattern, explains half (50.4 %) of the variance in the trial-averaged CSD responses across animals.

Both the early L4 sink with concurrent sources above and below and the L5/L6 sink with a source below were observed with a similar timing by Senzai et al., 2019. The L4 source and L2/3 sink were also observed in that study but emerge somewhat later than in our data – just after 100ms as opposed to ~60 ms in our canonical pattern. This discrepancy in onset might simply be due to differences in stimuli. In Senzai et al., 2019, the animals were exposed to 100 ms light pulses, while the animals in our data were presented with 250 ms whole-field flashes of a white screen. Nonetheless, the canonical CSD pattern exhibits good overall agreement with the pattern seen in Senzai et al., 2019. Studies in non-human primates where the animals were exposed to flashes also demonstrate good

spatio-temporal agreement with the CSD observed here, with sinks and sources occurring not only in the same layers and in the same order, but also at the same time after stimulus presentation (*Givre et al., 1994; Schroeder et al., 1998*).

We introduced the WD as a method to evaluate the difference between two CSD patterns and used it to quantify the variability in trial-averaged CSD between animals (*Figure 2D*), the trial-to-trial variability in CSD within animals (*Figure 2F and G*), and the difference between the model CSD, the trial-averaged CSD of individual animals, and the canonical CSD pattern. This application of the WD to compare CSD patterns comes with certain considerations that are important to note. First, although we compute WD for sinks and sources separately, sinks and sources do not arise independently. Current leaving the extracellular space in one place leads to current entering the extracellular space in another place, so current sinks and sources are inter-dependent. Second, the cost of shifting a sink or a source in space relative to shifting it in time is determined by the relative resolution in space vs. time. This relative cost does not necessarily correspond to the actual cost of changing the underlying physiology such that two distributions of sinks or sources match spatially vs. temporally. Determining the most appropriate relative cost of moving sinks and sources in space vs. time would require more detailed data than currently available and is beyond the scope of this study.

For the firing rate analysis, we utilized KS similarity and correlation to quantify experimental variability and model performance with regard to magnitude and temporal profile, respectively. We also investigated the statistics of unit spike firing by computing the moments of the distributions across cells in each population of peak firing and latencies to peak firing for both individual animals and model versions (*Figure 6—figure supplements 2 and 3*). Systematic use of quantitative metrics for biophysical modeling at this scale is still relatively uncommon, and our work establishes a set of measures for testing the model on LFP and spiking simultaneously, which can be useful for future studies in the field. Of course, there may well be other metrics that are equally or more suitable, and a systematic investigation into what would be the optimal metrics to apply is an important avenue for future work.

Our aim was to simultaneously reproduce experimentally recorded spikes and the CSD in our simulations. The original model captured spiking responses to gratings well (reproducing, e.g., direction selectivity distributions for different neuronal populations) with variable success when applied to other visual stimuli (*Billeh et al., 2020*). It was not originally tested on LFP/CSD. We found that, for the full-field flash stimulus, this model did not reproduce the CSD pattern in the upper layers of V1, and the spiking responses for this stimulus also exhibited a number of discrepancies.

After making selective adjustments to the recurrent synaptic weights, the model could reproduce the experimental firing rates (*Figure 5A–C*), though the discrepancy between the model CSD and the canonical CSD remained (*Figure 5D and E*), with only minor differences relative to the CSD of the original model (*Figure 4B*). The fact that the model can capture the experimental firing rates without capturing the experimental CSD and that adjustments to the synaptic weights yielded significant alterations in firing rates with only small changes in the CSD supports the point that LFP/CSD reflects aspects of circuit dynamics that are complementary to those reflected in locally recorded spikes.

Previous simulation studies demonstrated the importance of synaptic placement in shaping the LFP and CSD signature (*Einevoll et al., 2007; Pettersen et al., 2008; Lindén et al., 2010; Lindén et al., 2011; Łęski et al., 2013; Hagen et al., 2017; Ness et al., 2018*). To uncover the model adjustments that capture firing rates and CSD simultaneously, we explored the effects of changes in the synaptic positioning. In one case, we placed all excitatory synapses onto only basal or apical dendrites of L4 excitatory cells, as opposed to their original placement on both apical and basal dendrites. Moving all excitatory synapses onto basal dendrites resulted in substantial changes in both the pattern and magnitude of the CSD contribution from these L4 excitatory cells, with only minor changes to their firing rates (*Figure 5F and G* and *Figure 5—figure supplement 1*). Placing all excitatory synapses on apical dendrites led to somewhat larger changes in firing rates, though still similar to the firing rate of the original model, and to even bigger changes in the CSD magnitude. Performing the same manipulations of synaptic placement on L2/3 or L5 excitatory cells resulted in a similar decoupling of CSD and firing rates (*Figure 5—figure supplements 3 and 4*). It should, however, be noted that the decoupling is neither perfect nor universal.

This demonstrates a two-way dissociation of the firing rates and the pattern of sinks and sources in the CSD: The firing rates can be substantially altered with small effects on the CSD by adjusting the

synaptic weights, and the CSD can be substantially altered with only small effects on the firing rates by adjusting synaptic placement. These results align with findings in, for example, *Schroeder et al., 1998* and *Leszczyński et al., 2020*, where a lack of correlation between MUA and CSD in the upper layers of V1 during flash exposure to non-human primates suggested that these signals can sometimes be decoupled. Our findings imply that the LFP can reveal deficiencies in the model architecture that would not be evident from the firing rates alone, and that, to a certain extent, models can be optimized for firing rates and CSD independently.

Past studies have suggested that LFP can be modulated by attention through feedback during evoked responses (*Mehta et al., 2000a; Mehta et al., 2000b*). However, their findings indicated that V1 was not significantly affected by this modulation until the period 250 ms or later after stimulus onset. A recent study showed that feedback from higher visual areas can in fact exert a strong influence on the magnitude of LFP already at around 80 ms after stimulus onset (*Hartmann et al., 2019*). To investigate whether such cortico-cortical influence can contribute to the sinks and sources in the later stages (>50 ms) of the canonical CSD pattern, we added feedback consisting of experimentally recorded spikes from the higher cortical visual area LM (*Siegle et al., 2021*) impinging on synapses placed onto V1 neurons in our model, using anatomical data (*Glickfeld and Olsen, 2017; Marques et al., 2018; Hartmann et al., 2019; Keller et al., 2020; Shen et al., 2020*).

We found that the feedback can play a significant role in shaping the sustained sinks and sources (*Figure 6B–E*). The resulting model CSD reproduced the major sinks and sources identified in the canonical CSD pattern and was no longer among the outliers compared to the experimental variability (*Figure 6D*). Interestingly, absence of the feedback was not apparent from analysis of the firing rates alone, as the firing rates were already within the experimental variability before adding the feedback, further underscoring the utility of the LFP in illuminating structure–function relations in the circuit. Contributions from other visual cortical areas were not included, even though they too impinge upon neurons in V1 (*Harris et al., 2019; Siegle et al., 2021*), due to the lack of data characterizing such connections. This awaits future work. Finally, turning off the feedback from LM at 60ms disrupted the CSD pattern in layers 2/3 and 4 during the sustained period (*Figure 7—figure supplement 3*), which serves as a prediction for what will be observed if HVAs, and particularly LM, are silenced in this period of an evoked flash response. Our findings accord with the view that basal dendrites are the main targets for feedforward input, while the tufts of apical dendrites are the main targets for feedback input, even though basal dendrites can also receive input from long-range feedback connections and apical dendrites receive input from feedforward connections (*Aru et al., 2020*).

In our analyses, we aligned the CSD plots to depths obtained from histology and provided in Allen Common Coordinate Framework (CCF) coordinates. An alternative approach is to align the CSD to landmarks in the data (*Senzai et al., 2019*). In *Figure 6—figure supplement 7*, we explored whether utilizing landmarks rather than histology for alignment would affect our results. We found that this approach did not change our conclusion as the final model CSD was still within the experimental variability while the original and intermediate model CSD were outliers both when using the first principal component and when using the plain average as the target (*Figure 6—figure supplements 7B and C and 8*). Lastly, using the plain average rather than the first principal component as the target did not significantly affect our results when we aligned to histology (*Figure 6—figure supplement 9*).

With the major sinks and sources of the canonical CSD pattern reproduced, we explored their biophysical origins. We found that the initial L4 sink originates in the thalamocortical input to L4 excitatory cells, which aligns with suggestions made in *Mitzdorf, 1987; Swadlow et al., 2002*, and *Senzai et al., 2019*. The sustained L5/L6 sink comes from postsynaptic currents in L6 excitatory cells triggered by a combination of thalamocortical and recurrent excitatory inputs. The sustained L2/3 sink stems, in part, from input from LM onto the apical tufts of L2/3 and L5 excitatory cells. The sustained L4 source has its origins in a mixture of return currents from L2/3 and L5 excitatory cells resulting from the abovementioned feedback onto the apical dendrites of these cells, as well as contributions from L4 excitatory and inhibitory cells (*Figure 7A, B, D and E*).

In line with observations made by *Reimann et al., 2013*, we found that the somatic voltage-dependent membrane currents significantly shape the CSD signature (*Figures 5H and 7C*). Even so, our findings still emphasize the importance of synaptic inputs in sculpting the CSD, as the addition of synaptic input (*Figure 6A–E*) and changes to synaptic placement (*Figure 5F*) substantially altered the CSD pattern.

Recent investigations into the unitary LFP (the LFP generated by a single neuron) from inhibitory and excitatory synapses have suggested that inhibitory inputs exert a greater influence on LFP than excitatory input (Bazelot et al., 2010; Teleńczuk et al., 2017; Telenczuk et al., 2020). While this may be true of unitary effects, the total effect of excitatory input can still be greater if there are significantly more excitatory than inhibitory cells, and, correspondingly, significantly more excitatory synapses. In this V1 model, inhibitory cells make up about 15% and excitatory cells about 85% of the total number of cells, reflecting the cellular composition in mouse V1. Whether the pronounced dominance of excitatory cells is enough to make up for the reduced unitary influence of excitatory cells is an interesting area for further research.

This investigation into the biophysical origins of sinks and sources is limited by the fact that the contributions from recurrent connections are difficult to estimate precisely due to the nonlinear effects of these connections. That is, their contribution cannot simply be found by subtracting the CSD from thalamocortical and feedback synapses with all recurrent connections removed (Figure 7A and B) from the total CSD with the same input and recurrent connections intact (Figure 6C, right). Still, this analysis provides an initial estimate into the biophysical origins of the sinks and sources observed experimentally and demonstrates the insights that can be obtained from modeling of extracellular signals.

There is ample evidence that firing rates and LFP are modulated by the behavioral state of the animal, including measures like the pupil size (considered to be a proxy for arousal level) or running speed (Niell and Stryker, 2010; McGinley et al., 2015; Vinck et al., 2015; Saleem et al., 2017). In this study, the responses averaged over all trials have been the target for the modeling, without regard to any state-dependence of the responses. Our understanding of the state-dependent responses could benefit from the potential to probe the biophysical origins of extracellular signals. Therefore, reproducing these state-dependent responses is an important avenue for future research.

Note that the set of synaptic weights and other parameters that can reproduce the experimental firing rates and CSD is unlikely to be unique. This is a consequence of the degeneracy inherent to biological neural networks, known from both simulation and experimental studies, as many different parameterizations of neuronal networks can perform the same functions (Prinz et al., 2004; Marder and Goaillard, 2006; Drion et al., 2015; O'Leary, 2018). Thus, our network should only be considered an example of a circuit model that can produce firing rates and CSD that match the experimental observations. Obtaining multiple solutions and characterizing their diversity using automatic searches of the parameter space will be an interesting direction for future work. We did not utilize such an approach here because the number of simulations required (typically, many thousands or more for automatic optimization approaches) would currently be prohibitively expensive on a model of this scale and level of complexity: running a 1 s simulation with this model takes ~90 min on 384 CPU-cores (Billeh et al., 2020); a single trial in this study simulates 0.75 s of activity.

The original model used as a starting point here produced firing rates and direction and orientation tuning consistent with recordings during presentations of drifting gratings (Arkhipov et al., 2018; Billeh et al., 2020). In this study, we focused on the analysis and modeling of the response to full-field flashes, but when the final model was tested with the drifting gratings stimulus utilized in Billeh et al., 2020, we found that the present, revised model continued to exhibit orientation and direction tuning even though the adjustments were made with the aim to reproduce the observed CSD and population firing rates during full-field flashes (Figure 6—figure supplement 6). Ideally, the model should reproduce both firing rates and LFP simultaneously not only for flashes or drifting gratings, but for any visual stimulus (out-class generalization). This is a long-term goal and can be called 'the holy grail' of visual system modeling.

In this study, we developed a systematic framework to quantify experimental variability in both LFP/CSD and spikes and to evaluate model performance. We identified a canonical CSD pattern observed during presentations of full-field flash stimuli and obtained a bio-realistic model that reproduced both the canonical CSD pattern and spikes simultaneously. This model thus reproduces, in a quantitative manner, the shape and timing of current sinks and sources observed experimentally. We utilized this validated model to explain, mechanistically, the biophysical origins of the various current sinks and sources and their location across the layers of visual cortex. Our models are freely shared and should be useful for future studies disentangling the mechanisms underlying spiking dynamics and electrogenesis in the cortex.

Materials and methods

Experiments

Quality control

Of the 58 mice in the visual coding dataset, 9 were excluded because the exact probe location could not be recovered due to fading of fluorescent dye or artifacts in the OPT volume (Siegle *et al.*, 2021). Another five animals were excluded because they were missing LFP recordings from V1 during presentation of the flash stimulus. Thus, data for 44 animals were retained for the CSD analysis.

For the spike analysis, the same nine animals for which the exact probe location could not be recovered were excluded, and two additional animals were excluded because they did not have any cells recorded in V1, leaving a total of 47 animals for this part of the data analysis.

Neuronal classification

We distinguished between RS and FS cells by the time from trough to peak of the spike waveforms (Barthó *et al.*, 2004). For cortical cells, the spike duration was bimodally distributed with a dip at ~0.4 ms, while for thalamic cells, it was bimodally distributed with a dip at ~0.3 ms (Figure 3—figure supplement 1). Thus, the cutoff in the classification of cells as RS or FS was set at 0.4 ms for LM and V1, and at 0.3 ms for cells in LGN. Note that several studies have demonstrated that some pyramidal neurons may have spike waveforms short enough to be classified as FS cells (Vigneswaran *et al.*, 2011; Lemon *et al.*, 2021). Thus, some caution is warranted when interpreting the population firing rates.

When comparing the model firing rates to the experimental firing rates, the excitatory and non-Pvalb populations were grouped together in each layer of the model to make up the RS cells in L2/3, L4, L5, and L6, while the Pvalb cells across all layers were grouped together to make up the FS cells of V1. The layer boundaries were taken from the Allen CCF (Oh *et al.*, 2014), allowing for the assignment of each neuron's position to a specific cortical layer (Siegle *et al.*, 2021).

Model

The model consists of both biophysically detailed multicompartment neurons and leaky-integrate-and-fire (LIF) point-neurons. In total, there are 51,978 multicompartment neurons with Hodgkin–Huxley conductances at the soma and only passive conductances at the dendrites. These are arranged in a cylinder of radius of 400 μm and height 860 μm (corresponding to the average cortical thickness of V1 taken from the Allen CCF; Billeh *et al.*, 2020; Oh *et al.*, 2014). This cylinder makes up the 'core' of the model and is surrounded by an annulus of 178,946 LIF neurons which has the same height and a thickness of 445 μm . This makes the total number of neurons in the model 230,924 and the radius of the whole cylinder with both biophysically detailed and LIF neurons 845 μm . There are 17 different classes of neuron models. In each layer from 2/3 to 6, there are one excitatory and three inhibitory classes (Pvalb, Sst, Htr3a), while in layer 1 there is a single Htr3a class. The LGN module providing thalamocortical input to the model consists of 17,400 units selectively connected to the excitatory neurons and Pvalb neurons in L2/3 to L6, as well as the non-Pvalb neurons in L1. The background input to all neurons in the model comes from a single Poisson source firing at 1 kHz and represents influence from the rest of the brain. The feedback input to L2/3 and L5 excitatory, Pvalb, and Sst neurons comes from a node representing LM.

Simulation configuration

Instructions on how to run simulations of the model are provided in Billeh *et al.*, 2020. The files and code necessary to run the model versions presented in Figures 4–6 are provided in the directories `old_model_fig4`, `intermediate_model_fig5`, and `final_model_fig6`, respectively, on Dryad (see 'Data availability').

Data processing

LFP and CSD

The LFP in simulations was obtained from the extracellular potential by first downsampling to every other electrode along the probe (resulting in a spatial separation of 40 μm between each recording electrode, equal to the spacing in the public Neuropixels data) and using a low-pass fifth-order

Butterworth filter with a cutoff frequency of 500 Hz (utilizing functions `scipy.signal.butter` and `scipy.signal.filtfilt`). The same filtering was applied to get the experimental LFP. The CSD was calculated from the experimental and model LFP using the delta iCSD method (*Pettersen et al., 2006*), where the radius of laterally (orthogonal to the probe axis) constant CSD was assumed to be 400 μm – the radius of the V1 model's 'core' region consisting of biophysically detailed multicompartment neurons. For the experimental CSD, this radius was set to 800 μm , roughly corresponding to the size of mouse V1.

Visual stimulus

The stimulus used to compare the model and the experiments was full-field flashes. In the experiments, the mice were presented with gray screens for 1 s, followed by 250 ms of white screen, and then 750 ms of gray screen over 75 trials. In the simulations, both the stimulus presentation and the pre- and the poststimulus gray screen periods lasted 250 ms, and the number of trials was 10.

Input from LGN)

Originally, the LGN spike trains used as input to the model were generated with the FilterNet module provided with the model, using 17,400 'LGN units' (*Billeh et al., 2020*). However, when this input was used for simulations, the onset of the evoked response in V1 was 20–30 ms delayed in comparison with experiments. Therefore, we used experimentally recorded LGN spike trains as input to the model instead. We assigned a recorded spike train to each of the 17,400 LGN units in all trials. In total, the public Neuropixels data contain recordings from 1263 regular-spiking LGN neurons across 32 animals during 75 trials of full-field flash presentations. We divided the total pool of spike trains into 10 subsets, and then randomly sampled spike trains from one subset in each trial until all 17,400 LGN units had been assigned a spike train in all trials.

Input from LM

The experimentally recorded spike trains in the LM were used to construct the feedback input to V1. In total, the public Neuropixels data contain recordings from 1823 RS LM neurons across 42 animals during presentations of the full-field flash stimulus. Spikes were randomly sampled from the pool of all spike trains to construct a spike train that was used as input to all the cells that were targeted by the feedback in the model. All neurons received the same spike train.

Background input

The input from the Poisson source firing at 1 kHz was not stimulus dependent. It is a coarse representation of the continuous influence of the rest of the brain on V1.

Dendritic targeting

The rules for placement of synapses were set in *Billeh et al., 2020* and were based on reviews of literature on anatomy.

LGN to V1

In the original model, the synapses from LGN onto excitatory V1 neurons were placed on apical and basal dendrites within 150 μm from the soma, while synapses onto inhibitory V1 neurons were placed on their soma and on their basal dendrites without distance limitations (*Billeh et al., 2020*). This placement was left unchanged in this study.

V1-V1

The synapses for recurrent connections were placed according to the following rules in the original model (*Billeh et al., 2020*):

Excitatory-to-excitatory connections

All synapses from excitatory V1 neurons onto other excitatory V1 neurons were placed along the dendrites and avoided the soma. In layers 2/3 and 4, the placement of synapses was restricted to be

within 200 μm from the somata, while in layers 5 and 6, they could be placed anywhere along the dendrites.

Excitatory-to-inhibitory connections

All synapses from excitatory V1 neurons onto inhibitory V1 neurons were placed on their somata or dendrites without any distance limitations.

Inhibitory-to-excitatory connections

Synapses from Pvalb neurons onto excitatory V1 neurons were placed on the soma and on the dendrites within 50 μm from the soma. Synapses from Sst neurons were placed only on dendrites and only more than 50 μm from the soma. Synapses from Htr3a neurons were placed on dendrites between 50 and 300 μm from the soma.

Inhibitory-to-inhibitory connections

Synapses from inhibitory neurons to other inhibitory neurons were placed according to the same rules as the inhibitory-to-excitatory connections described above.

These placement rules were kept in this study, except for the synapses from excitatory neurons to excitatory L6 neurons. Here, they were restricted to be within 150 μm of the soma. The purpose of this restriction was to reduce the spatial separation between the current sink and source, and thereby decrease the magnitude of the L6 sink-source dipole.

LM-V1

The synapses from the node representing LM to V1 were placed on the apical dendrites of L2/3 neurons (within 150 μm from the soma), on the apical tufts (>300 μm from the soma) and the basal dendrites (within 150 μm from the soma) of L5 excitatory cells, and on the somata and basal dendrites of L2/3 and L5 inhibitory cells (at any distance from the soma).

Adjusting synaptic weights

In the original model, the synaptic weights of thalamocortical connections were based on experimental recordings of synaptic current, while the synaptic weights for recurrent connections were initially set to estimates from literature, then optimized to a drifting gratings stimulus until the model reproduced experimental values of orientation and direction selectivity. In this study, the synaptic weights for thalamocortical connections were left unchanged from the original model. Before the addition of feedback from higher visual areas to the model, the synaptic weights for recurrent connections in V1 were multiplied by factors in the range [0.2, 2.5].

In the original model, the input from the background node represented the influence of the rest of the brain on V1, which included the influence from higher visual areas such as LM. This means that some of the feedback influence from LM on V1 should be present (though coarsely represented) in the input from the background node. When the influence of input from LM to feedback-targeted cells is modeled on its own, the influence of the background node must be updated accordingly. Thus, after the addition of feedback, the synaptic weights from the background node to the populations targeted by feedback (the L2/3 and L5 excitatory, Pvalb, and Sst cells) were multiplied by factors in the range [0.2, 0.5]. The synaptic weights from the node representing LM were initially set equal to the original weights between the background node and the populations targeted by the feedback, but this led to too high firing rates compared to the experimental firing rates in these populations, so they were multiplied by factors in the range [0.2, 0.5]. Finally, the connections from Pvalb neurons in V1 to L2/3 excitatory neurons and L5 excitatory cells were re-scaled in the range [0.8, 1.2] times the weights set prior to the addition of feedback. The ranges reported here were set after experimenting with different ranges to find what would allow the model to reproduce the experimental observations. Only a single value within each range was used in the final model.

Quantification and statistical analysis

Firing rates

The time-resolved population firing rates (bin size 1 ms, filtered using `scipy.ndimage.gaussian_filter` with `sigma = 2`) were computed by averaging the spike count over all cells in a population and over all

trials (10 trials in the simulations and 75 trials in the experiments). The distribution of firing rates across cells used in the calculation of the KS similarities was computed by averaging over the time windows baseline, initial peak, and sustained activity (defined in **Figure 3**) and over all trials.

Kolmogorov–Smirnov similarity

The KS similarity scores (**Billeh et al., 2020**) were computed by first calculating the KS distance (using the function `scipy.stats.ks_2samp`) between two distributions of firing rates across cells, and subtracting this number from 1, such that a KS similarity score of 1 implies identity and a score of 0 implies no overlap between the two distributions. In the comparison of the model to the experimental data, the KS similarity was computed between the distribution of firing rates across cells in each RS and the FS population of the model and the distribution of firing rates across cells from all animals in the corresponding populations. To assess the variability in the experiments, the KS similarity was calculated between the distribution of firing rates across cells in the same RS and FS populations in individual animals, provided there were more than 10 cells recorded in a given population in this animal, and the distribution of firing rates across cells from all other animals.

Correlation

We computed the similarity in the profile of time-resolved population firing rates with the Pearson correlation coefficient (using the function `scipy.stats.pearsonr`). The correlation between the model and the experimental firing rates was calculated between model population firing rates and the population firing rates averaged across cells from all animals. The level of experimental variability was assessed by calculating the correlation between population firing rates in each animal and the population firing rates averaged across cells from all other animals.

CSD analysis

Since the number of recording electrodes in V1 are not the same in all animals, we interpolated the CSD of each animal and the CSD from simulations onto dimensions of the same lengths ($M = 30$ points along the depth and $K = 100$ points along the time axis for 100 ms time windows) before we quantitatively analyzed the CSD.

PCA

The trial-averaged CSD of each animal was flattened into a vector of length $M \times K = 3000$, and the vectors of all $N = 44$ animals were stacked together into a matrix of size 44×3000 . Then, we performed PCA (using `sklearn.decomposition.PCA`) on this matrix to obtain the principal components that would constitute sums of weighted contributions of the trial-averaged CSD patterns.

Wasserstein distance (WD)

The first Wasserstein distance $W(P_1, P_2)$ between two distributions P_1 and P_2 is defined as

$$W(P_1, P_2) = \inf_{\gamma \in \Gamma(P_1, P_2)} \int c(x, y) \gamma(x, y) dx dy$$

where $c(x, y)$ is the cost of moving a unit ‘mass’ from position x to y following the optimal transport plan $\gamma(x, y)$ in all transport plans $\Gamma(P_1, P_2)$ (**Rubner et al., 1998; Arjovsky et al., 2017**).

In the utilization of WD to quantify the similarity between two CSD patterns, the distance between the distribution of sinks in the two patterns $W(P_{sinks,1}, P_{sinks,2})$ and the distance between distribution of sources of the two patterns $W(P_{sources,1}, P_{sources,2})$ are calculated separately and summed to form a total WD between the two CSD patterns:

$$W_{CSD}(P_1, P_2) = W(P_{sinks,1}, P_{sinks,2}) + W(P_{sources,1}, P_{sources,2})$$

where P_1 and P_2 refer to the two CSD patterns. The Python Optimal transport library (<https://pythonot.github.io/index.html>) was used to implement this calculation.

Acknowledgements

Research reported in this publication was supported the Simula School of Research and Innovation, CINPLA, the European Union Horizon 2020 Research and Innovation Program under Grant Agreement No. 785907 and No. 945539 (Human Brain Project [HBP] SGA2 and SGA3), the National Institute of Neurological Disorders and Stroke of the National Institutes of Health under Award Number R01NS122742, and by the National Institute of Biomedical Imaging and Bioengineering of the National Institutes of Health under Award Number R01EB029813, as well as by the Allen Institute. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. We acknowledge the use of Fenix Infrastructure resources, which are partially funded from the European Union's Horizon 2020 research and innovation program through the ICEI project under the grant agreement No. 800858. We thank the Allen Institute founder, Paul G Allen, for his vision, encouragement, and support.

Additional information

Competing interests

Christof Koch: holds an executive position, and has financial interest, in Intrinsic Powers, Inc, a company whose purpose is to develop a device that can be used in the clinic to assess the presence and absence of consciousness in patients. This does not pose any conflict of interest with regard to the work undertaken for this publication. The other authors declare that no competing interests exist.

Funding

Funder	Grant reference number	Author
Simula School of Research		Atle E Rimehaug
European Union Horizon 2020 Research and Innovation program	785907	Espen Hagen
European Union Horizon 2020 Research and Innovation program	945539	Espen Hagen
Research Council of Norway	COBRA - project number 250128	Alexander J Stasik
IKTPLUSS-IKT and Digital Innovation	300504	Alexander J Stasik
National Institute of Neurological Disorders and Stroke	R01NS122742	Kael Dai Josh H Siegle Shawn R Olsen Christof Koch Anton Arkhipov Yazan N Billeh
National Institute of Biomedical Imaging and Bioengineering	R01EB029813	Kael Dai Josh H Siegle Shawn R Olsen Christof Koch Anton Arkhipov Yazan N Billeh
Allen Institute		Kael Dai Josh H Siegle Shawn R Olsen Christof Koch Anton Arkhipov Yazan N Billeh

The funders had no role in study design, data collection and interpretation, or the decision to submit the work for publication.

Author contributions

Atle E Rimehaug, Conceptualization, Software, Formal analysis, Investigation, Visualization, Methodology, Writing – original draft, Writing – review and editing; Alexander J Stasik, Conceptualization, Software, Formal analysis, Investigation, Visualization, Methodology; Espen Hagen, Conceptualization, Resources, Software, Formal analysis, Investigation, Methodology, Writing – review and editing; Yazan N Billeh, Conceptualization, Software, Methodology; Josh H Siegle, Data curation, Investigation, Visualization, Writing – review and editing; Kael Dai, Resources, Software, Methodology; Shawn R Olsen, Resources, Supervision; Christof Koch, Gaute T Einevoll, Anton Arkhipov, Conceptualization, Resources, Supervision, Funding acquisition, Methodology, Project administration, Writing – review and editing

Author ORCIDs

Atle E Rimehaug  <https://orcid.org/0000-0002-8312-9875>
 Alexander J Stasik  <https://orcid.org/0000-0003-1646-2472>
 Josh H Siegle  <https://orcid.org/0000-0002-7736-4844>
 Shawn R Olsen  <https://orcid.org/0000-0002-9568-7057>
 Gaute T Einevoll  <https://orcid.org/0000-0002-5425-5012>
 Anton Arkhipov  <https://orcid.org/0000-0003-1106-8310>

Decision letter and Author response

Decision letter <https://doi.org/10.7554/eLife.87169.sa1>

Author response <https://doi.org/10.7554/eLife.87169.sa2>

Additional files**Supplementary files**

- MDAR checklist

Data availability

The files necessary to run simulations of the different model versions presented in the paper as well as data resulting from simulations of those model versions are publicly available in Dryad: <https://doi.org/10.5061/dryad.k3j9kd5b8>. The experimental data set utilized is publicly available at: <https://portal.brain-map.org/explore/circuits/visual-coding-neuropixels>. The code generated for data analysis and producing the figures in this manuscript is publicly available at: https://github.com/atleer/CINPLA_Allen_V1_analysis.git (copy archived at **Rimehaug, 2023**).

The following dataset was generated:

Author(s)	Year	Dataset title	Dataset URL	Database and Identifier
Rimehaug AR, Stasik AJ, Hagen E, Billeh YN, Siegle JH, Dai K, Olsen SR, Koch C, Einevoll G, Arkhipov A	2022	Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex	https://dx.doi.org/10.5061/dryad.k3j9kd5b8	Dryad Digital Repository, 10.5061/dryad.k3j9kd5b8

The following previously published dataset was used:

Author(s)	Year	Dataset title	Dataset URL	Database and Identifier
Siegle JH, Wakeman W, Jia X, Heller G, Ramirez T, Graddis N, Mei N, Durand S	2020	20191003_AIBS_mouse_ecephys_brain_observatory_1_1	https://dandiarchive.org/dandiset/000021	DANDI Archive 000021, 000021

References

Andersen RA, Musallam S, Pesaran B. 2004. Selecting the signals for a brain-machine interface. *Current Opinion in Neurobiology* **14**:720–726. DOI: <https://doi.org/10.1016/j.conb.2004.10.005>, PMID: 15582374

- Arjovsky M**, Chintala S, Bottou L. 2017. Wasserstein generative adversarial networks. *Proceedings of the 34th International Conference on Machine Learning*. 214–223.
- Arkhipov A**, Gouwens NW, Billeh YN, Gratiy S, Iyer R, Wei Z, Xu Z, Abbasi-Asl R, Berg J, Buice M, Cain N, da Costa N, de Vries S, Denman D, Durand S, Feng D, Jarsky T, Lecoq J, Lee B, Li L, et al. 2018. Visual physiology of the layer 4 cortical circuit in silico. *PLOS Computational Biology* **14**:e1006535. DOI: <https://doi.org/10.1371/journal.pcbi.1006535>, PMID: 30419013
- Aru J**, Suzuki M, Larkum ME. 2020. Cellular mechanisms of conscious processing. *Trends in Cognitive Sciences* **24**:814–825. DOI: <https://doi.org/10.1016/j.tics.2020.07.006>, PMID: 32855048
- Barthó P**, Hirase H, Monconduit L, Zugaro M, Harris KD, Buzsáki G. 2004. Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *Journal of Neurophysiology* **92**:600–608. DOI: <https://doi.org/10.1152/jn.01170.2003>, PMID: 15056678
- Bastos AM**, Vezoli J, Bosman CA, Schoffelen JM, Oostenveld R, Dowdall JR, De Weerd P, Kennedy H, Fries P. 2015. Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron* **85**:390–401. DOI: <https://doi.org/10.1016/j.neuron.2014.12.018>, PMID: 25556836
- Bazelon M**, Dinocourt C, Cohen I, Miles R. 2010. Unitary inhibitory field potentials in the CA3 region of rat hippocampus. *The Journal of Physiology* **588**:2077–2090. DOI: <https://doi.org/10.1113/jphysiol.2009.185918>, PMID: 20403979
- Belitski A**, Gretton A, Magri C, Murayama Y, Montemurro MA, Logothetis NK, Panzeri S. 2008. Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *The Journal of Neuroscience* **28**:5696–5709. DOI: <https://doi.org/10.1523/JNEUROSCI.0009-08.2008>, PMID: 18509031
- Billeh YN**, Cai B, Gratiy SL, Dai K, Iyer R, Gouwens NW, Abbasi-Asl R, Jia X, Siegle JH, Olsen SR, Koch C, Mihalas S, Arkhipov A. 2020. Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex. *Neuron* **106**:388–403. DOI: <https://doi.org/10.1016/j.neuron.2020.01.040>, PMID: 32142648
- Buzsáki G**. 2006. *Rhythms of the Brain* Oxford university press. DOI: <https://doi.org/10.1093/acprof:oso/9780195301069.001.0001>
- Buzsáki G**, Anastassiou CA, Koch C. 2012. The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience* **13**:407–420. DOI: <https://doi.org/10.1038/nrn3241>, PMID: 22595786
- Chatzikalymniou AP**, Skinner FK. 2018. Deciphering the Contribution of Oriens-Lacunosum/Moleculare (OLM) Cells to Intrinsic θ Rhythms Using Biophysical Local Field Potential (LFP) Models. *eNeuro* **5**:ENEURO.0146-18.2018. DOI: <https://doi.org/10.1523/ENEURO.0146-18.2018>, PMID: 30225351
- Dai K**, Gratiy SL, Billeh YN, Xu R, Cai B, Cain N, Rimehaug AE, Stasik AJ, Einevoll GT, Mihalas S, Koch C, Arkhipov A, Marinazzo D. 2020. Brain Modeling Toolkit: An open source software suite for multiscale modeling of brain circuits. *PLOS Computational Biology* **16**:e1008386. DOI: <https://doi.org/10.1371/journal.pcbi.1008386>, PMID: 33253147
- Di S**, Baumgartner C, Barth DS. 1990. Laminar analysis of extracellular field potentials in rat vibrissa/barrel cortex. *Journal of Neurophysiology* **63**:832–840. DOI: <https://doi.org/10.1152/jn.1990.63.4.832>, PMID: 2341880
- Drion G**, O’Leary T, Marder E. 2015. Ion channel degeneracy enables robust and tunable neuronal firing rates. *PNAS* **112**:E5361–E5370. DOI: <https://doi.org/10.1073/pnas.1516400112>, PMID: 26354124
- Einevoll GT**, Pettersen KH, Devor A, Ulbert I, Halgren E, Dale AM. 2007. Laminar population analysis: estimating firing rates and evoked synaptic activity from multielectrode recordings in rat barrel cortex. *Journal of Neurophysiology* **97**:2174–2190. DOI: <https://doi.org/10.1152/jn.00845.2006>, PMID: 17182911
- Einevoll GT**, Kayser C, Logothetis NK, Panzeri S. 2013. Modelling and analysis of local field potentials for studying the function of cortical circuits. *Nature Reviews Neuroscience* **14**:770–785. DOI: <https://doi.org/10.1038/nrn3599>, PMID: 24135696
- Fernández-Ruiz A**, Makarov VA, Benito N, Herreras O. 2012. Schaffer-specific local field potentials reflect discrete excitatory events at gamma frequency that may fire postsynaptic hippocampal CA1 units. *The Journal of Neuroscience* **32**:5165–5176. DOI: <https://doi.org/10.1523/JNEUROSCI.4499-11.2012>, PMID: 22496562
- Fernández-Ruiz A**, Oliva A, Nagy GA, Maurer AP, Berényi A, Buzsáki G. 2017. Entorhinal-CA3 Dual-Input Control of Spike Timing in the Hippocampus by Theta-Gamma Coupling. *Neuron* **93**:1213–1226. DOI: <https://doi.org/10.1016/j.neuron.2017.02.017>, PMID: 28279355
- Givre SJ**, Schroeder CE, Arezzo JC. 1994. Contribution of extrastriate area V4 to the surface-recorded flash VEP in the awake macaque. *Vision Research* **34**:415–428. DOI: [https://doi.org/10.1016/0042-6989\(94\)90156-2](https://doi.org/10.1016/0042-6989(94)90156-2), PMID: 8303826
- Głabska H**, Chintaluri HC, Wójcik DK. 2014. Collection of Simulated Data for Validation of Methods of Analysis of Extracellular Potentials. *Neuroinformatics 2014*. Leiden, Netherlands. DOI: <https://doi.org/10.3389/conf.fninf.2014.18.00035>
- Glickfeld LL**, Olsen SR. 2017. Higher-Order Areas of the Mouse Visual Cortex. *Science* **3**:251–273. DOI: <https://doi.org/10.1146/annurev-vision-102016-061331>
- Gouwens NW**, Berg J, Feng D, Sorensen SA, Zeng H, Hawrylycz MJ, Koch C, Arkhipov A. 2018. Systematic generation of biophysically detailed models for diverse cortical neuron types. *Nature Communications* **9**:710. DOI: <https://doi.org/10.1038/s41467-017-02718-3>, PMID: 29459718
- Hagen E**, Dahmen D, Stavrinou ML, Lindén H, Tetzlaff T, van Albada SJ, Grün S, Diesmann M, Einevoll GT. 2016. Hybrid scheme for modeling local field potentials from point-neuron networks. *Cereb. Cortex* **26**:4461–4496. DOI: <https://doi.org/10.1093/cercor/bhw237>, PMID: 27797828

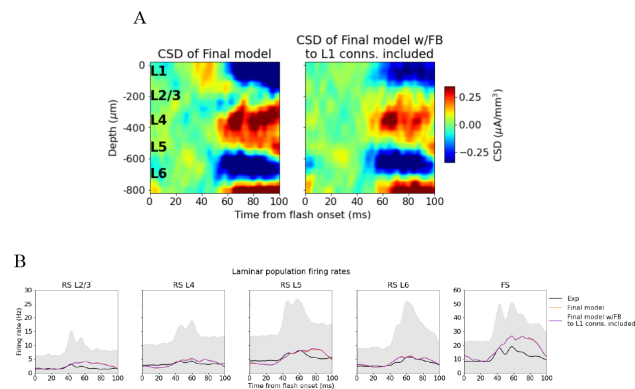
- Hagen E, Fossum JC, Pettersen KH, Alonso JM, Swadlow HA, Einevoll GT. 2017. Focal local field potential signature of the single-axon monosynaptic thalamocortical connection. *The Journal of Neuroscience* **37**:5123–5143. DOI: <https://doi.org/10.1523/JNEUROSCI.2715-16.2017>, PMID: 28432143
- Hagen E, Næss S, Ness TV, Einevoll GT. 2018. Multimodal Modeling of Neural Network Activity: Computing LFP, ECoG, EEG, and MEG Signals With LFPy 2.0. *Frontiers in Neuroinformatics* **12**:92. DOI: <https://doi.org/10.3389/fninf.2018.00092>, PMID: 30618697
- 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, et al. 2019. Hierarchical organization of cortical and thalamic connectivity. *Nature* **575**:195–202. DOI: <https://doi.org/10.1038/s41586-019-1716-z>, PMID: 31666704
- Hartmann TS, Raja S, Lomber SG, Born RT. 2019. Cortico-Cortical Feedback from V2 Exerts a Powerful Influence over the Visually Evoked Local Field Potential and Associated Spike Timing in V1. *bioRxiv*. DOI: <https://doi.org/10.1101/792010>
- Hay E, Hill S, Schürmann F, Markram H, Segev I. 2011. Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties. *PLOS Computational Biology* **7**:e1002107. DOI: <https://doi.org/10.1371/journal.pcbi.1002107>, PMID: 21829333
- Henrie JA, Shapley R. 2005. LFP power spectra in V1 cortex: the graded effect of stimulus contrast. *Journal of Neurophysiology* **94**:479–490. DOI: <https://doi.org/10.1152/jn.00919.2004>, PMID: 15703230
- Holt GR, Koch C. 1999. Electrical interactions via the extracellular potential near cell bodies. *Journal of Computational Neuroscience* **6**:169–184. DOI: <https://doi.org/10.1023/a:1008832702585>, PMID: 10333161
- Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, Lee AK, Anastassiou CA, Andrei A, Aydin C, Barbic M, Blanche TJ, Bonin V, Couto J, Dutta B, Gratiy SL, Gutnisky DA, Häusser M, Karsh B, Ledochowitsch P, et al. 2017. Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**:232–236. DOI: <https://doi.org/10.1038/nature24636>, PMID: 29120427
- Kandel A, Buzsáki G. 1997. Cellular-synaptic generation of sleep spindles, spike-and-wave discharges, and evoked thalamocortical responses in the neocortex of the rat. *The Journal of Neuroscience* **17**:6783–6797. DOI: <https://doi.org/10.1523/JNEUROSCI.17-17-06783.1997>, PMID: 9254689
- Keller AJ, Roth MM, Scanziani M. 2020. Feedback generates a second receptive field in neurons of the visual cortex. *Nature* **582**:545–549. DOI: <https://doi.org/10.1038/s41586-020-2319-4>, PMID: 32499655
- Koch C. 1998. *Biophysics of Computation: Information Processing in Single Neurons* Oxford university press. DOI: <https://doi.org/10.1093/oso/9780195104912.001.0001>
- Kreiman G, Hung CP, Kraskov A, Quiroga RQ, Poggio T, DiCarlo JJ. 2006. Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. *Neuron* **49**:433–445. DOI: <https://doi.org/10.1016/j.neuron.2005.12.019>, PMID: 16446146
- Lemon RN, Baker SN, Kraskov A. 2021. Classification of cortical neurons by spike shape and the identification of pyramidal neurons. *Cerebral Cortex* **31**:5131–5138. DOI: <https://doi.org/10.1093/cercor/bhab147>, PMID: 34117760
- Łęski S, Lindén H, Tetzlaff T, Pettersen KH, Einevoll GT. 2013. Frequency dependence of signal power and spatial reach of the local field potential. *PLOS Computational Biology* **9**:e1003137. DOI: <https://doi.org/10.1371/journal.pcbi.1003137>, PMID: 23874180
- Leszczyński M, Barczak A, Kajikawa Y, Ulbert I, Falchier AY, Tal I, Haegens S, Melloni L, Knight RT, Schroeder CE. 2020. Dissociation of broadband high-frequency activity and neuronal firing in the neocortex. *Science Advances* **6**:eabb0977. DOI: <https://doi.org/10.1126/sciadv.abb0977>, PMID: 32851172
- Liebe S, Hoerzer GM, Logothetis NK, Rainer G. 2012. Theta coupling between V4 and prefrontal cortex predicts visual short-term memory performance. *Nature Neuroscience* **15**:456–462. DOI: <https://doi.org/10.1038/nn.3038>, PMID: 22286175
- Lindén H, Pettersen KH, Einevoll GT. 2010. Intrinsic dendritic filtering gives low-pass power spectra of local field potentials. *Journal of Computational Neuroscience* **29**:423–444. DOI: <https://doi.org/10.1007/s10827-010-0245-4>, PMID: 20502952
- Lindén H, Tetzlaff T, Potjans TC, Pettersen KH, Grün S, Diesmann M, Einevoll GT. 2011. Modeling the spatial reach of the LFP. *Neuron* **72**:859–872. DOI: <https://doi.org/10.1016/j.neuron.2011.11.006>, PMID: 22153380
- Liu J, Newsome WT. 2006. Local field potential in cortical area MT: stimulus tuning and behavioral correlations. *The Journal of Neuroscience* **26**:7779–7790. DOI: <https://doi.org/10.1523/JNEUROSCI.5052-05.2006>, PMID: 16870724
- Makarova J, Ibarz JM, Makarov VA, Benito N, Herreras O. 2011. Parallel readout of pathway-specific inputs to laminated brain structures. *Frontiers in Systems Neuroscience* **5**:77. DOI: <https://doi.org/10.3389/fnsys.2011.00077>, PMID: 21949504
- Marder E, Goaillard JM. 2006. Variability, compensation and homeostasis in neuron and network function. *Nature Reviews. Neuroscience* **7**:563–574. DOI: <https://doi.org/10.1038/nrn1949>, PMID: 16791145
- Markowitz DA, Wong YT, Gray CM, Pesaran B. 2011. Optimizing the decoding of movement goals from local field potentials in macaque cortex. *The Journal of Neuroscience* **31**:18412–18422. DOI: <https://doi.org/10.1523/JNEUROSCI.4165-11.2011>, PMID: 22171043
- Marques T, Nguyen J, Fioreze G, Petreanu L. 2018. The functional organization of cortical feedback inputs to primary visual cortex. *Nature Neuroscience* **21**:757–764. DOI: <https://doi.org/10.1038/s41593-018-0135-z>, PMID: 29662217

- 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>, PMID: 26402600
- Mehring C**, Rickert J, Vaadia E, Cardoso de Oliveira S, Aertsen A, Rotter S. 2003. Inference of hand movements from local field potentials in monkey motor cortex. *Nature Neuroscience* **6**:1253–1254. DOI: <https://doi.org/10.1038/nn1158>, PMID: 14634657
- Mehta AD**, Ulbert I, Schroeder CE. 2000a. Intermodal selective attention in monkeys. I: distribution and timing of effects across visual areas. *Cerebral Cortex* **10**:343–358. DOI: <https://doi.org/10.1093/cercor/10.4.343>, PMID: 10769247
- Mehta AD**, Ulbert I, Schroeder CE. 2000b. Intermodal selective attention in monkeys. II: physiological mechanisms of modulation. *Cerebral Cortex* **10**:359–370. DOI: <https://doi.org/10.1093/cercor/10.4.359>, PMID: 10769248
- Mitzdorf U**. 1985. Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiological Reviews* **65**:37–100. DOI: <https://doi.org/10.1152/physrev.1985.65.1.37>, PMID: 3880898
- Mitzdorf U**. 1987. Properties of the evoked potential generators: current source-density analysis of visually evoked potentials in the cat cortex. *The International Journal of Neuroscience* **33**:33–59. DOI: <https://doi.org/10.3109/00207458708985928>, PMID: 3610492
- Montemurro MA**, Rasch MJ, Murayama Y, Logothetis NK, Panzeri S. 2008. Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Current Biology* **18**:375–380. DOI: <https://doi.org/10.1016/j.cub.2008.02.023>, PMID: 18328702
- Mukamel R**, Fried I. 2012. Human intracranial recordings and cognitive neuroscience. *Annual Review of Psychology* **63**:511–537. DOI: <https://doi.org/10.1146/annurev-psych-120709-145401>, PMID: 21943170
- Nauhaus I**, Busse L, Carandini M, Ringach DL. 2009. Stimulus contrast modulates functional connectivity in visual cortex. *Nature Neuroscience* **12**:70–76. DOI: <https://doi.org/10.1038/nn.2232>, PMID: 19029885
- Ness TV**, Remme MWH, Einevoll GT. 2018. h-type membrane current shapes the local field potential from populations of pyramidal neurons. *The Journal of Neuroscience* **38**:6011–6024. DOI: <https://doi.org/10.1523/JNEUROSCI.3278-17.2018>, PMID: 29875266
- Niell CM**, Stryker MP. 2008. Highly selective receptive fields in mouse visual cortex. *The Journal of Neuroscience* **28**:7520–7536. DOI: <https://doi.org/10.1523/JNEUROSCI.0623-08.2008>, PMID: 18650330
- 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>, PMID: 20188652
- 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, et al. 2014. A mesoscale connectome of the mouse brain. *Nature* **508**:207–214. DOI: <https://doi.org/10.1038/nature13186>, PMID: 24695228
- O’Leary T**. 2018. Homeostasis, failure of homeostasis and degenerate ion channel regulation. *Current Opinion in Physiology* **2**:129–138. DOI: <https://doi.org/10.1016/j.cophys.2018.01.006>
- Pesaran B**, Pezaris JS, Sahani M, Mitra PP, Andersen RA. 2002. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature Neuroscience* **5**:805–811. DOI: <https://doi.org/10.1038/nn890>, PMID: 12134152
- Pesaran B**, Vinck M, Einevoll GT, Sirota A, Fries P, Siegel M, Truccolo W, Schroeder CE, Srinivasan R. 2018. Investigating large-scale brain dynamics using field potential recordings: analysis and interpretation. *Nature Neuroscience* **21**:903–919. DOI: <https://doi.org/10.1038/s41593-018-0171-8>, PMID: 29942039
- Petttersen KH**, Devor A, Ulbert I, Dale AM, Einevoll GT. 2006. Current-source density estimation based on inversion of electrostatic forward solution: effects of finite extent of neuronal activity and conductivity discontinuities. *Journal of Neuroscience Methods* **154**:116–133. DOI: <https://doi.org/10.1016/j.jneumeth.2005.12.005>, PMID: 16436298
- Petttersen KH**, Einevoll GT. 2008. Amplitude variability and extracellular low-pass filtering of neuronal spikes. *Biophysical Journal* **94**:784–802. DOI: <https://doi.org/10.1529/biophysj.107.111179>, PMID: 17921225
- Petttersen KH**, Hagen E, Einevoll GT. 2008. Estimation of population firing rates and current source densities from laminar electrode recordings. *Journal of Computational Neuroscience* **24**:291–313. DOI: <https://doi.org/10.1007/s10827-007-0056-4>, PMID: 17926125
- Prinz AA**, Bucher D, Marder E. 2004. Similar network activity from disparate circuit parameters. *Nature Neuroscience* **7**:1345–1352. DOI: <https://doi.org/10.1038/nn1352>, PMID: 15558066
- Rall W**. 1962. Electrophysiology of a dendritic neuron model. *Biophysical Journal* **2**:145–167. DOI: [https://doi.org/10.1016/s0006-3495\(62\)86953-7](https://doi.org/10.1016/s0006-3495(62)86953-7), PMID: 14490040
- Rall W**, Shepherd GM. 1968. Theoretical reconstruction of field potentials and dendrodendritic synaptic interactions in olfactory bulb. *Journal of Neurophysiology* **31**:884–915. DOI: <https://doi.org/10.1152/jn.1968.31.6.884>, PMID: 5710539
- Reimann MW**, Anastassiou CA, Perin R, Hill SL, Markram H, Koch C. 2013. A biophysically detailed model of neocortical local field potentials predicts the critical role of active membrane currents. *Neuron* **79**:375–390. DOI: <https://doi.org/10.1016/j.neuron.2013.05.023>, PMID: 23889937
- Rickert J**, Oliveira S de, Vaadia E, Aertsen A, Rotter S, Mehring C. 2005. Encoding of movement direction in different frequency ranges of motor cortical local field potentials. *The Journal of Neuroscience* **25**:8815–8824. DOI: <https://doi.org/10.1523/JNEUROSCI.0816-05.2005>, PMID: 16192371

- Rimehaug AE.** 2023. CINPLA Allen V1 analysis. swh:1:rev:bc32ab40499ba9b7687a1388f461cdf158bba375. Software Heritage. https://archive.softwareheritage.org/swh:1:dir:2ace4a17f8fd874a86821913488320804ffcc8a;origin=https://github.com/atleer/CINPLA_Allen_V1_analysis;visit=swh:1:snp:f1b3d0f86648ca18ab6280411c9cc704e015f61f;anchor=swh:1:rev:bc32ab40499ba9b7687a1388f461cdf158bba375
- Roux S, Mackay WA, Riehle A.** 2006. The pre-movement component of motor cortical local field potentials reflects the level of expectancy. *Behavioural Brain Research* **169**:335–351. DOI: <https://doi.org/10.1016/j.bbr.2006.02.004>, PMID: 16530860
- Rubner Y, Tomasi C, Guibas LJ.** 1998. A metric for distributions with applications to image databases. IEEE 6th International Conference on Computer Vision. , 59–66. DOI: <https://doi.org/10.1109/ICCV.1998.710701>, PMID: 9482381
- Saleem AB, Lien AD, Krumin M, Haider B, Rosón MR, Ayaz A, Reinhold K, Busse L, Carandini M, Harris KD, Carandini M.** 2017. Subcortical source and modulation of the narrowband gamma oscillation in mouse visual cortex. *Neuron* **93**:315–322. DOI: <https://doi.org/10.1016/j.neuron.2016.12.028>, PMID: 28103479
- Scherberger H, Jarvis MR, Andersen RA.** 2005. Cortical local field potential encodes movement intentions in the posterior parietal cortex. *Neuron* **46**:347–354. DOI: <https://doi.org/10.1016/j.neuron.2005.03.004>, PMID: 15848811
- Schomburg EW, Anastassiou CA, Buzsáki G, Koch C.** 2012. The spiking component of oscillatory extracellular potentials in the rat hippocampus. *The Journal of Neuroscience* **32**:11798–11811. DOI: <https://doi.org/10.1523/JNEUROSCI.0656-12.2012>, PMID: 22915121
- Schroeder CE, Mehta AD, Givre SJ.** 1998. A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque. *Cerebral Cortex* **8**:575–592. DOI: <https://doi.org/10.1093/cercor/8.7.575>, PMID: 9823479
- Senzai Y, Fernandez-Ruiz A, Buzsáki G.** 2019. Layer-specific physiological features and interlaminar interactions in the primary visual cortex of the mouse. *Neuron* **101**:500–513. DOI: <https://doi.org/10.1016/j.neuron.2018.12.009>, PMID: 30635232
- Shen S, Jiang X, Scala F, Fu J, Fahey P, Kobak D, Tan Z, Reimer J, Sinz F, Tolia AS.** 2020. Distinct organization of two cortico-cortical feedback pathways. *bioRxiv*. DOI: <https://doi.org/10.1101/2020.02.27.968792>
- Siegle JH, Jia X, Durand S, Gale S, Bennett C, Graddis N, Heller G, Ramirez TK, Choi H, Luviano JA, Groblewski PA, Ahmed R, Arkhipov A, Bernard A, Billeh YN, Brown D, Buice MA, Cain N, Caldejon S, Casal L, et al.** 2021. Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**:86–92. DOI: <https://doi.org/10.1038/s41586-020-03171-x>, PMID: 33473216
- Sinha M, Narayanan R.** 2015. HCN channels enhance spike phase coherence and regulate the phase of spikes and LFPs in the theta-frequency range. *PNAS* **112**:E2207–E2216. DOI: <https://doi.org/10.1073/pnas.1419017112>
- Sinha M, Narayanan R.** 2022. Active dendrites and local field potentials: biophysical mechanisms and computational explorations. *Neuroscience* **489**:111–142. DOI: <https://doi.org/10.1016/j.neuroscience.2021.08.035>, PMID: 34506834
- Stavisky SD, Kao JC, Nuyujukian P, Ryu SI, Shenoy KV.** 2015. A high performing brain-machine interface driven by low-frequency local field potentials alone and together with spikes. *Journal of Neural Engineering* **12**:036009. DOI: <https://doi.org/10.1088/1741-2560/12/3/036009>, PMID: 25946198
- Swadlow HA, Gusev AG, Bezdudnaya T.** 2002. Activation of a cortical column by a thalamocortical impulse. *The Journal of Neuroscience* **22**:7766–7773. DOI: <https://doi.org/10.1523/JNEUROSCI.22-17-07766.2002>, PMID: 12196600
- Teleńczuk B, Dehghani N, Le Van Quyen M, Cash SS, Halgren E, Hatsopoulos NG, Destexhe A.** 2017. Local field potentials primarily reflect inhibitory neuron activity in human and monkey cortex. *Scientific Reports* **7**:40211. DOI: <https://doi.org/10.1038/srep40211>, PMID: 28074856
- Telenczuk B, Telenczuk M, Destexhe A.** 2020. A kernel-based method to calculate local field potentials from networks of spiking neurons. *Journal of Neuroscience Methods* **344**:108871. DOI: <https://doi.org/10.1016/j.jneumeth.2020.108871>, PMID: 32687850
- Tomsett RJ, Ainsworth M, Thiele A, Sanayei M, Chen X, Giesemann MA, Whittington MA, Cunningham MO, Kaiser M.** 2015. Virtual Electrode Recording Tool for EXtracellular potentials (VERTEX): comparing multi-electrode recordings from simulated and biological mammalian cortical tissue. *Brain Structure & Function* **220**:2333–2353. DOI: <https://doi.org/10.1007/s00429-014-0793-x>, PMID: 24863422
- Tort ABL, Kramer MA, Thorn C, Gibson DJ, Kubota Y, Graybiel AM, Kopell NJ.** 2008. Dynamic cross-frequency couplings of local field potential oscillations in rat striatum and hippocampus during performance of a T-maze task. *PNAS* **105**:20517–20522. DOI: <https://doi.org/10.1073/pnas.0810524105>, PMID: 19074268
- Traub RD, Contreras D, Cunningham MO, Murray H, LeBeau FEN, Roopun A, Bibbig A, Wilentz WB, Higley MJ, Whittington MA.** 2005. Single-column thalamocortical network model exhibiting gamma oscillations, sleep spindles, and epileptogenic bursts. *Journal of Neurophysiology* **93**:2194–2232. DOI: <https://doi.org/10.1152/jn.00983.2004>, PMID: 15525801
- Victor JD, Purpura K, Katz E, Mao B.** 1994. Population encoding of spatial frequency, orientation, and color in macaque V1. *Journal of Neurophysiology* **72**:2151–2166. DOI: <https://doi.org/10.1152/jn.1994.72.5.2151>, PMID: 7884450
- Vierling-Claassen D, Cardin JA, Moore CI, Jones SR.** 2010. Computational modeling of distinct neocortical oscillations driven by cell-type selective optogenetic drive: separable resonant circuits controlled by low-threshold spiking and fast-spiking interneurons. *Frontiers in Human Neuroscience* **4**:198. DOI: <https://doi.org/10.3389/fnhum.2010.00198>, PMID: 21152338

- Vigneswaran G**, Kraskov A, Lemon RN. 2011. Large identified pyramidal cells in macaque motor and premotor cortex exhibit “thin spikes”: implications for cell type classification. *The Journal of Neuroscience* **31**:14235–14242. DOI: <https://doi.org/10.1523/JNEUROSCI.3142-11.2011>, PMID: 21976508
- Vinck M**, Batista-Brito R, Knoblich U, Cardin JA. 2015. Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron* **86**:740–754. DOI: <https://doi.org/10.1016/j.neuron.2015.03.028>, PMID: 25892300
- Watrous AJ**, Lee DJ, Izadi A, Gurkoff GG, Shahlaie K, Ekstrom AD. 2013. A comparative study of human and rat hippocampal low-frequency oscillations during spatial navigation. *Hippocampus* **23**:656–661. DOI: <https://doi.org/10.1002/hipo.22124>, PMID: 23520039
- Womelsdorf T**, Fries P, Mitra PP, Desimone R. 2006. Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* **439**:733–736. DOI: <https://doi.org/10.1038/nature04258>, PMID: 16372022

Appendix 1



Appendix 1—figure 1. Effect of adding feedback connections from latero-medial (LM) to L1 inhibitory cells. **(A)** Left: current source density (CSD) from simulation of final model version without connections from LM to L1 inhibitory cells. Right: CSD from simulation of final model version with connections from LM to L1 inhibitory cells included. **(B)** Population firing rates of experiments (black line), final model without connections between LM and L1 inhibitory cells (orange line), final model with connections between LM and L1 inhibitory cells included (purple colored line).



Figures and figure supplements

Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex

Atle E Rimehaug et al.

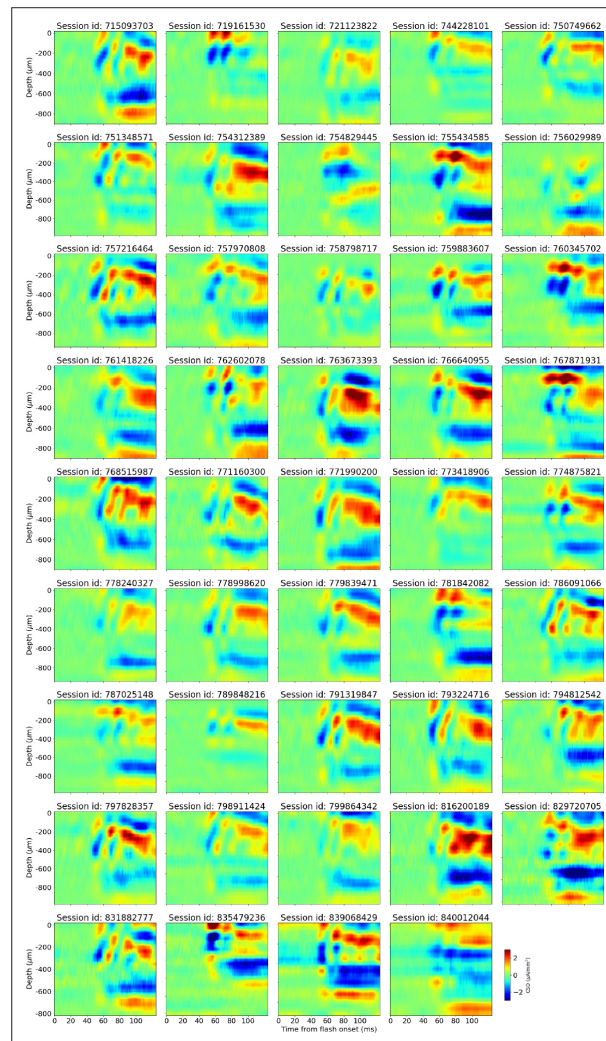


Figure 2—figure supplement 1. Trial-averaged current source density (CSD) during presentation of full-field flashes for all 44 animals in this study.

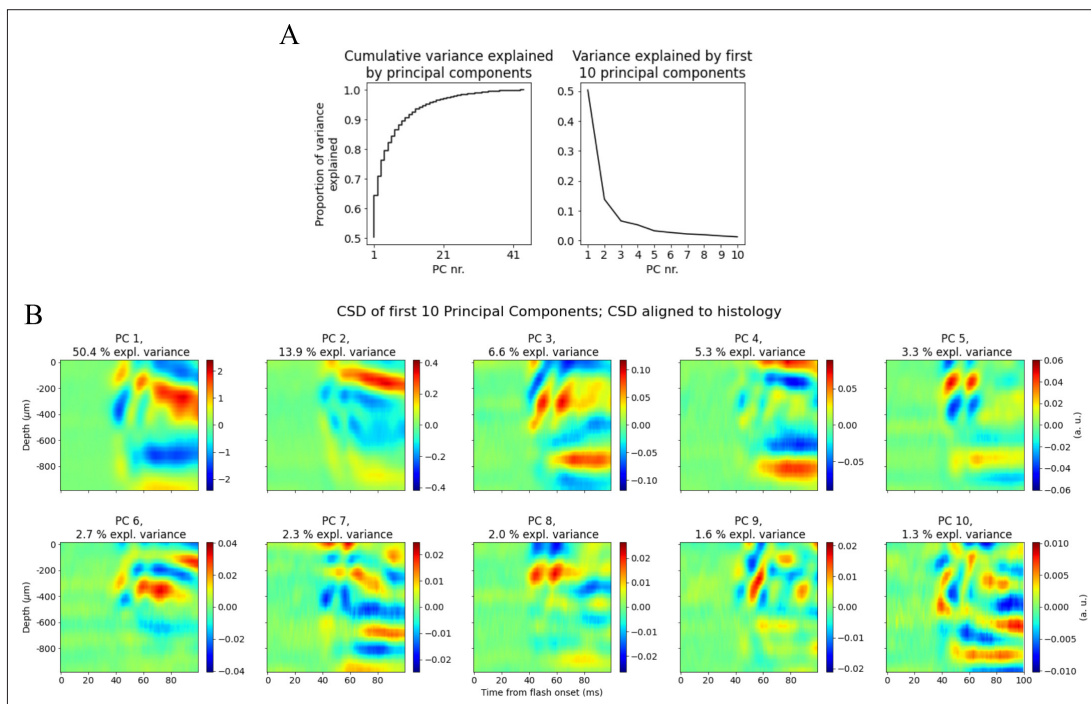


Figure 2—figure supplement 2. Principal component analysis (PCA) on histology-aligned current source density (CSD). **(A)** Left: cumulative variance explained by principal components. Right: variance explained by first 10 components. **(B)** CSD plots of the first 10 principal components explaining in total >90% of the variance in trial-averaged CSD across animals.




Figure 2—figure supplement 3. Comparing inter-trial and inter-animal pairwise Wasserstein distances (WDs). Cumulative distributions of pairwise WDs between trial-averaged current source density (CSD) of individual animals (blue line) and pairwise WDs between single trial CSD in each animal (red lines). The black line denotes the point of maximal distance between the two distributions.




Figure 3—figure supplement 1. Classifying cell types in experimental data. Distributions of waveform duration in cells from lateral geniculate nucleus (LGN), V1, and latero-medial (LM) and the threshold (red line) between classifying as regular-spiking (RS) or fast-spiking (FS).




Figure 3—figure supplement 2. Number of cells in each population in experimental data. Left: number of regular-spiking (RS) and fast-spiking (FS) in each layer in individual animals. Right: number of FS cells across all layers in V1 in individual animals.




Figure 4—figure supplement 1. Effect of reducing recurrent inhibition. Laminar population firing rates in the original model (blue line), in the model after reduction of recurrent excitatory synaptic weights to all Pvalb cells by 30% (brown line), and in experiments (black line).




Figure 5—figure supplement 1. Quantifying change in simulated current source density (CSD) with adjustments to synaptic weights. **(A)** CSD of original model. **(B)** CSD of intermediate model where synaptic weights between populations in V1 have been adjusted. **(C)** CSD of final model. **(D)** Pairwise Wasserstein distance between CSD different model versions normalized to the largest pairwise Wasserstein distance between trial-averaged CSD of individual animals. Green star: Wasserstein distance (WD) between CSD of original model **(A)** and intermediate model **(B)**. Blue star: WD between CSD of original model **(A)** and final model **(C)**. Orange star: WD between CSD of intermediate model **(B)** and final model **(C)**.

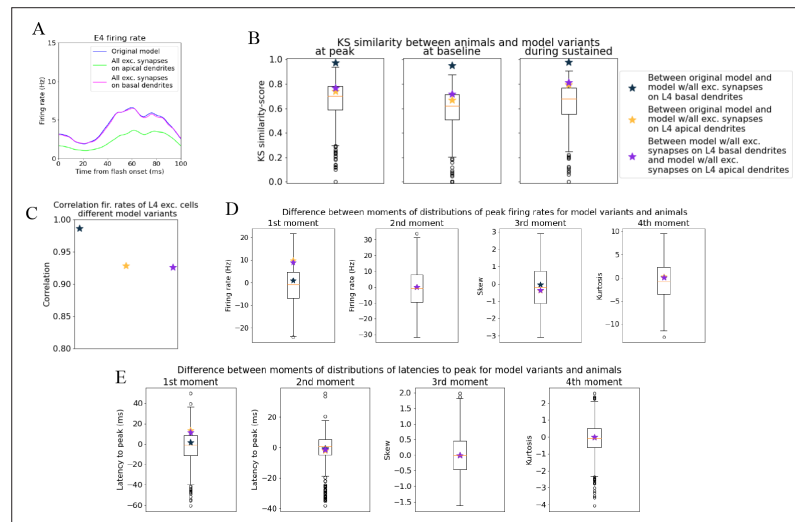


Figure 5—figure supplement 2. Quantifying change in spiking of L4 excitatory cells after adjusting synaptic placement. **(A)** Blue line: original model where excitatory synapses onto L4 excitatory cells were placed on both basal and apical dendrites. Green line: all excitatory synapses onto L4 excitatory cells were placed only onto apical dendrites. Pink line: all excitatory synapses onto L4 excitatory cells were placed only onto basal dendrites. **(B)** Kolmogorov–Smirnov (KS) similarity between the model variants in **(A)** for average firing rates across cells in different time periods (defined in **Figure 3**). Boxplots represent distribution of pairwise KS similarities between animals. **(C)** Correlation between firing rates of model variants in **(A)**. **(D)** Difference in moments of distributions of peak firing rates across cells between model variants in **(A)**. Boxplots represent pairwise differences between animals. **(E)** Difference in moments of distributions of latencies to peak firing rate across cells between model variants in **(A)** and pairwise differences between animals (boxplots).

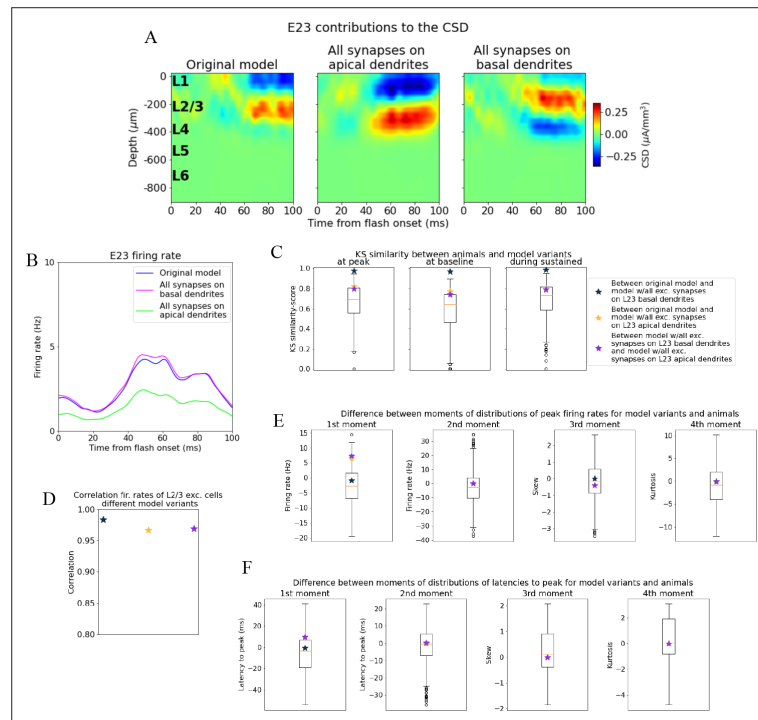


Figure 5—figure supplement 3. Effects of manipulating synaptic placement onto L2/3 excitatory cells on population current source density (CSD) and spiking. **(A)** CSD generated by L2/3 excitatory cells in (left) the original configuration with synapses on both apical and basal dendrites, (middle) all excitatory synapses placed on apical dendrites, and (right) all excitatory synapses placed on basal dendrites. **(B)** Population firing rates of L2/3 excitatory cells with different synaptic placement. Blue line: original model. Green line: all excitatory synapses onto L2/3 excitatory cells were placed only onto apical dendrites. Pink line: all excitatory synapses onto L2/3 excitatory cells were placed only onto basal dendrites. **(C)** Kolmogorov–Smirnov (KS) similarity between the model variants in **(B)** for average firing rates across cells in different time periods (defined in **Figure 3**). Boxplots represent distribution of pairwise KS similarities between animals. **(D)** Correlation between firing rates of model variants in **(B)**. **(E)** Difference in moments of distributions of peak firing rates across cells between model variants in **(B)**. Boxplots represent pairwise differences in moments between animals. **(F)** Difference in moments of distributions of latencies to peak firing rate across cells between model variants in **(B)** and between animals (boxplots).

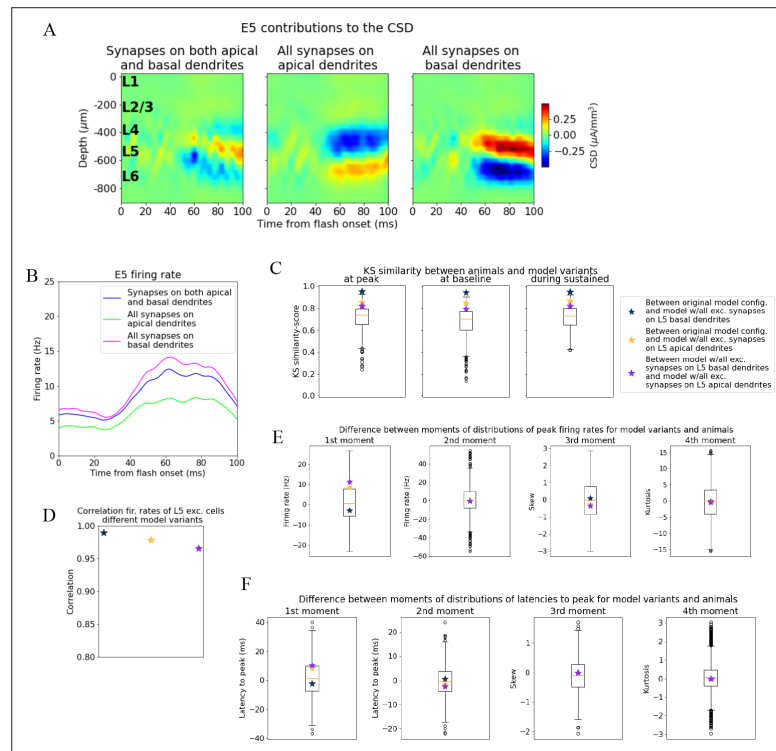


Figure 5—figure supplement 4. Effects of manipulating synaptic placement onto L5 excitatory cells on population current source density (CSD) and spiking. **(A)** CSD generated by L5 excitatory cells in (left) the original configuration with synapses on both apical and basal dendrites, (middle) all excitatory synapses placed on apical dendrites, and (right) all excitatory synapses on basal dendrites. In all cases, the synapses were placed within 200 μm from the soma to have the same ranges as was used for L4 and L2/3 cells in **Figure 5** and **Figure 5—figure supplement 3**, respectively. **(B)** Population firing rates of L5 excitatory cells with different synaptic placement. Blue line: original model configuration with synapses on both apical and basal dendrites, Green line: all excitatory synapses onto L5 excitatory cells were placed only onto apical dendrites. Pink line: all excitatory synapses onto L5 excitatory cells were placed only onto basal dendrites. **(C)** Kolmogorov–Smirnov (KS) similarity in average firing rates across cells in different time periods (defined in **Figure 3**) between the model variants in **(B)**. Boxplots represent distribution of pairwise KS-similarities in animals. **(D)** Correlation between firing rates of model variants in **(A)**. **(E)** Difference in moments of distributions of peak firing rates across cells between model variants in **(B)**. Boxplots represent pairwise differences in moments between animals. **(F)** Difference in moments of distributions of latencies to peak firing rate across cells between model variants in **(B)** and between animals (boxplots).




Figure 6—figure supplement 1. Effect of adjusting synaptic placement onto L6 excitatory cells. Contributions to the total current source density (CSD) from L6 excitatory cells with the original placement of recurrent excitatory synapses uniformly along the whole length of their dendrites (left) and after moving all recurrent excitatory synapses within 150 µm from the soma.

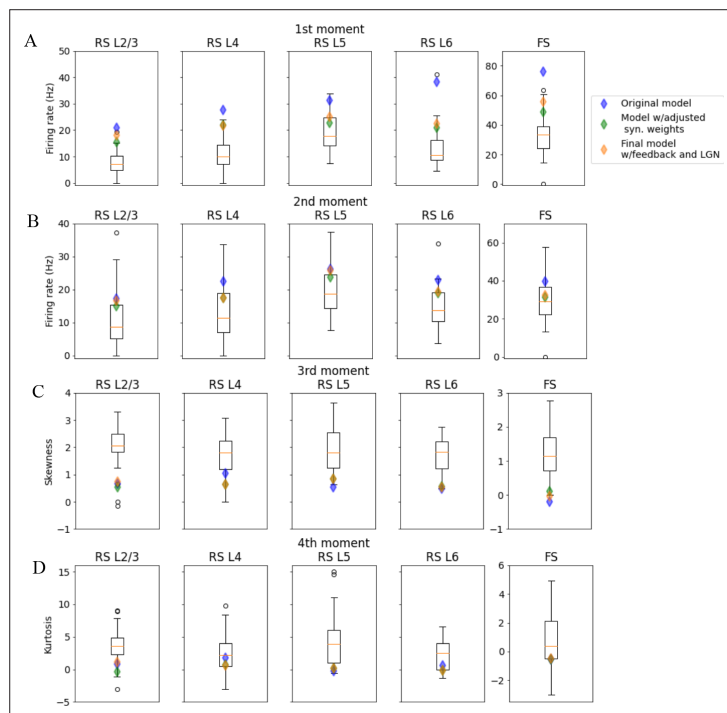


Figure 6—figure supplement 2. Moments of distributions of peak firing rate in model versions and experiments for different populations. Boxplots represent data for different animals, diamonds represent model versions. **(A)** First moment (mean). **(B)** Second moment (standard deviation). **(C)** Third moment (skewness). **(D)** Fourth moment (kurtosis).

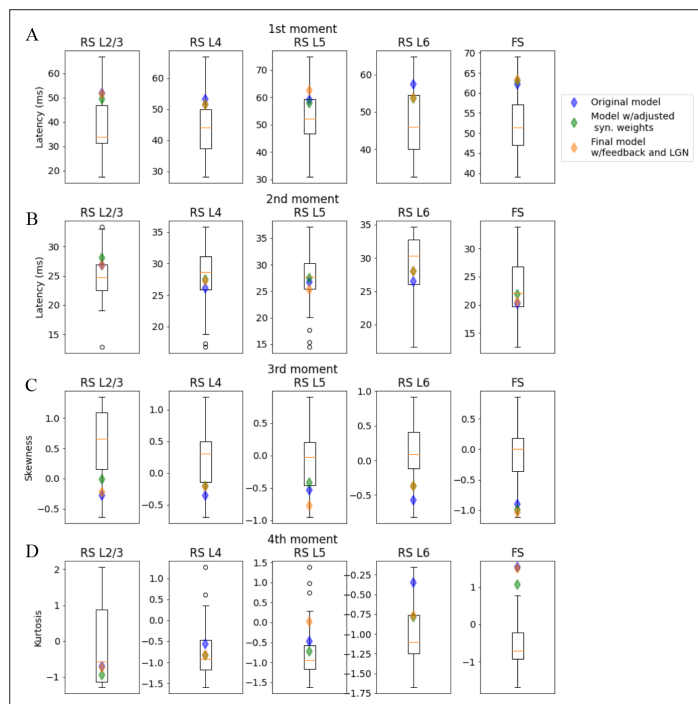


Figure 6—figure supplement 3. Moments of distributions of latency to peak of firing rates in model versions and experiments in different populations. Boxplots represent data for different animals, diamonds represent model versions. **(A)** First moment (mean). **(B)** Second moment (standard deviation). **(C)** Third moment (skewness). **(D)** Fourth moment (kurtosis).




Figure 6—figure supplement 4. Relative change in peak firing rates between neighboring populations. **(A)** From regular-spiking (RS) cells in L2/3 to RS cells in L4. **(B)** From RS cells in L4 to RS cells in L5. **(C)** From RS cells in L5 to RS cells in L6. **(D)** From all RS cells in V1 to all FS cells in V1. Boxplots represent data for different animals, diamonds represent model versions.

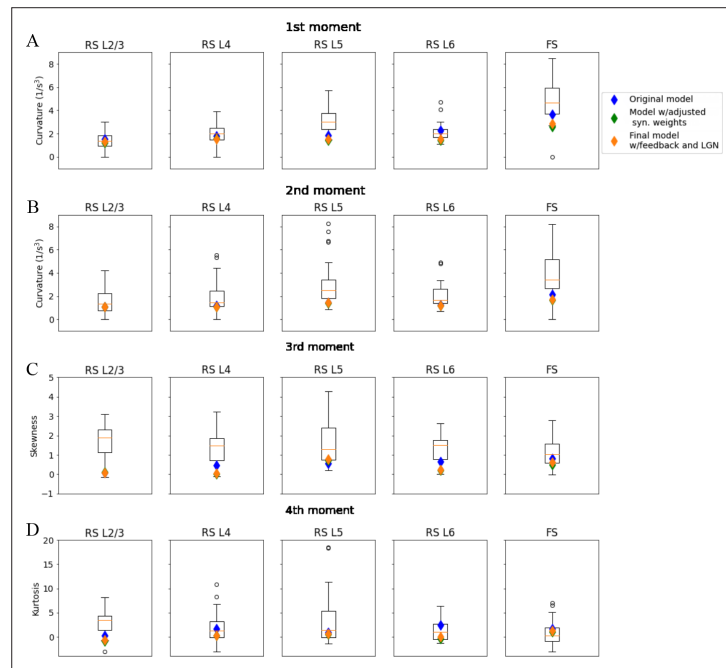


Figure 6—figure supplement 5. Moments of distributions of greatest curvature in firing rate across cells. Diamonds represent model versions and boxplots represent moments calculated for different animals. **(A)** First moment (mean). **(B)** Second moment (standard deviation). **(C)** Third moment (skew). **(D)** Fourth moment (kurtosis).




Figure 6—figure supplement 6. Orientation and direction selectivity in final model. Rate at preferred direction (A), direction selectivity index (B), and orientation selectivity index (C) in layer populations of model (blue) and in experiments (gray).

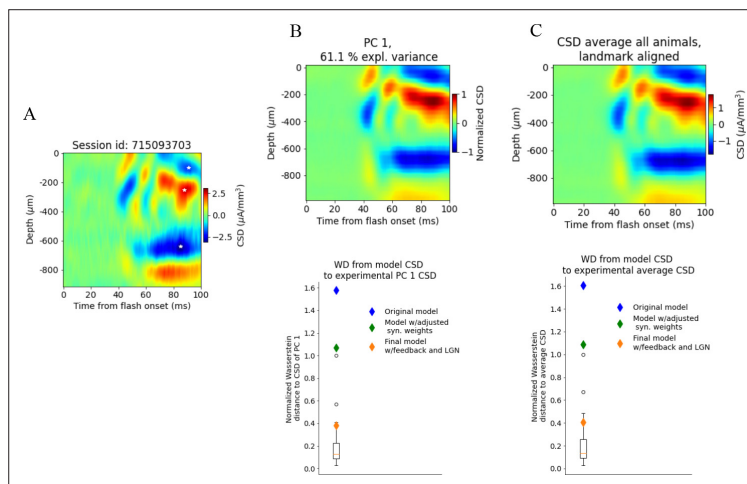


Figure 6—figure supplement 7. Current source density (CSD) analysis after aligning experimental CSD plots to landmarks rather than histology. **(A)** Example CSD plot with landmarks used for alignment marked (white stars). **(B)** Top: PC 1 CSD after application of principal component analysis (PCA) on trial-averaged CSD from all animals. Bottom: Wasserstein distances to PC 1 CSD after aligning CSD to landmarks for all animals. Diamonds denote WD from CSD of model versions and boxplots denote distances from individual animals. **(C)** Top: CSD average across all animals after aligning to landmarks. Bottom: Wasserstein distances to average CSD after alignment to landmarks for both model versions and animals.

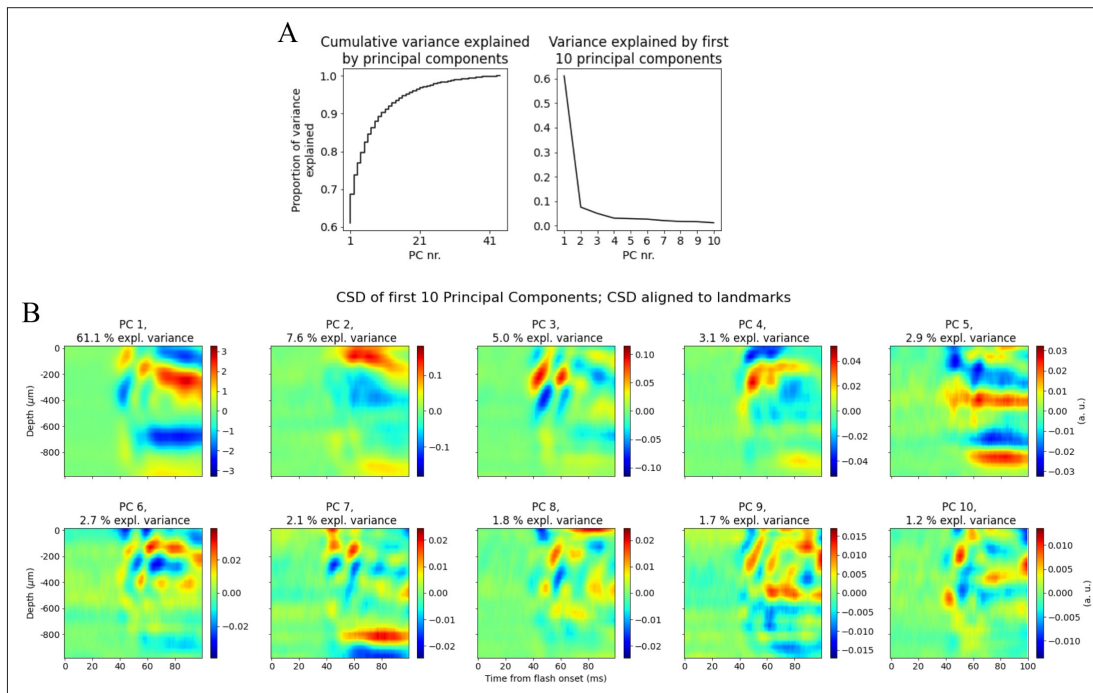


Figure 6—figure supplement 8. Principal component analysis (PCA) on landmark aligned current source density (CSD). **(A)** Left: cumulative variance explained by components. Right: variance explained by first 10 components. **(B)** CSD plots of the first 10 principal components explaining in total >90% of the variance in trial-averaged CSD across animals.




Figure 6—figure supplement 9. Effect of using plain average of trial-averaged current source density (CSD) from all animals instead of first principal component as target. Wasserstein distance from model versions and individual animal CSD to plain average of trial-averaged CSD from all animals. **(A)** CSD averaged over trial-averaged CSD from all animals. CSD aligned to histology. **(B)** Wasserstein distance from CSD of model versions (diamonds) and individual animals (boxplot) to CSD averaged over all animals (normalized to largest distance from the CSD of an individual animal to the average CSD).

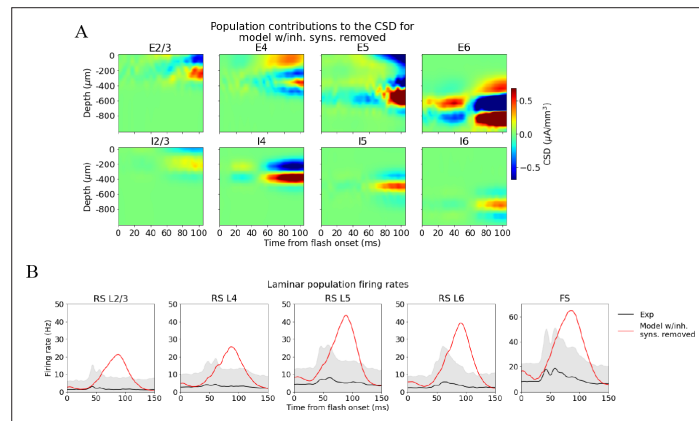


Figure 7—figure supplement 1. Effects of removing recurrent inhibition on population current source density (CSD) and firing rates. **(A)** Population contributions to the total CSD and **(B)** laminar population firing rates in a simulation where all inhibitory synapses have been removed.




Figure 7—figure supplement 2. Effect of cell orientation on current source density (CSD) contributions of L4 inhibitory cells. CSD contribution of L4 inhibitory cells with original orientation (left) and scrambled orientation (right).

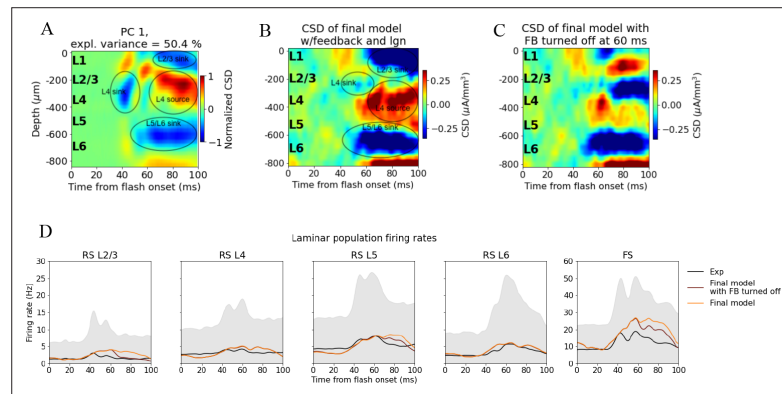


Figure 7—figure supplement 3. Silencing feedback from latero-medial (LM) in model during evoked response. Current source density (CSD) of (A) PC 1 computed from experiments, (B) final model, and (C) final model with feedback turned off at 60ms. (D) Trial-averaged population firing rates in experiments (black line), final model (orange line), and final model with feedback turned off at 60ms (brown line).

Paper II

**Uncovering population contributions
to the extracellular potential in the
mouse visual system using Laminar
Population Analysis**

II

II

Paper III

Estimating simulated local field potentials from presynaptic firing rates and network properties

