

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Methods for High-throughput, Cellular-Resolution Estimation of Synaptic Properties

Permalink

<https://escholarship.org/uc/item/6q92f176>

Author

Shababo, Ben

Publication Date

2022

Peer reviewed|Thesis/dissertation

Methods for High-throughput, Cellular-Resolution Estimation of Synaptic Properties

by

Benjamin M. Shababo

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Dr. Hillel Adesnik, Chair

Dr. Linda Wilbrecht

Dr. Frederic Theunissen

Dr. Ehud Isacoff

Summer 2022

Abstract

Methods for High-throughput, Cellular-Resolution Estimation of Synaptic Properties

by

Benjamin M. Shababo

Doctor of Philosophy in Neuroscience

University of California, Berkeley

Dr. Hillel Adesnik, Chair

Circuit-mapping experiments combining whole-cell electrophysiology with two-photon optical stimulation of potentially presynaptic neurons (“2p-mapping”) have produced rich data on monosynaptic connectivity of neural circuits. However, mapping densely-packed presynaptic populations at cellular resolution has proven challenging, making the precise localization and identity of connected neurons difficult. To interpret data resulting from these experiments, it is therefore critical to develop statistical methods which can infer the properties of neural circuits despite the limited spatial resolution of state-of-the-art 2p-mapping technologies. In this work, we present several novel statistical methods for mapping neural circuits at cellular resolution. In Chapter 1, we present background on mapping neural circuits and the methods used. In Chapter 2, we present a statistical method for inferring monosynaptic connections when the resolution of two-photon stimulation is not sufficient given the density of the opsin-expressing cells. In Chapter 3, we present a high-resolution approach for inferring common input to two simultaneously patched neurons during mapping. And in Chapter 4, we present a high accuracy method for inferring the timing of post-synaptic currents - the key response observation when mapping neural circuits.

To July

The brain is a curious thing.

Contents

Contents	ii
1 Introduction	1
1.1 Studying Neural Microcircuits	1
1.2 Physiological synaptic mapping	2
2 Single-cell Resolution Mapping With Two-Photon Holographic Optogenetics	6
2.1 Introduction	6
2.2 A high-throughput, temporally precise data acquisition system	7
2.3 Single-cell resolution optogenetics is not (currently) possible	10
2.4 Single-cell resolution analysis for mapping	14
2.5 Learning the neural response function	21
2.6 Validation	24
2.7 Discussion	28
2.8 Method Details	28
3 Mapping Complementary Networks of Cortical Somatostatin Interneurons	32
3.1 Mapping a cortical circuit with multiple modalities	32
3.2 Distinct subtypes of L5 SST cells receive complementary patterns of excitatory synaptic input	34
3.3 Common input mapping reveals subnetwork structure in L5 SST cell output	35
3.4 Paired recordings confirm dense, reciprocal, and selective intra- and translaminar connectivity	45
3.5 Discussion	47
3.6 Methods	49
4 Bayesian Methods For Event Analysis of Intracellular Currents	57
4.1 Introduction	57
4.2 Methods	60
4.3 Results	68
4.4 Discussion	75
4.5 Supplemental Material	77
Bibliography	79

5	Appendix	97
5.1	Technical Details for Gaussian Processes Learning	97
5.2	Implementation Details of Algorithm 1	99
5.3	Table of Notation	101

Chapter 1

Introduction

1.1 Studying Neural Microcircuits

At the level of microcircuits, the goal of experimental neuroscience is to validate theories which describe 1) how neural activity in a circuit relates to information that is useful for sensation, cognition, and behavior and 2) validate quantitative models of how that neural activity is generated from the properties of the cell types and synapses that comprise that circuit [186]. Examples of features that can be used to classify cell type include intrinsic electrophysiology, transcriptional and epigenetic profiles, morphology, and location [183, 131, 28]. The synapses between these cells can be defined by their pre- and postsynaptic properties like the probability of vesicle release, the magnitude and shape of the postsynaptic current, and its short-term dynamics [51]. Additionally, the cell types in a specific neural microcircuit form synapses with each other at different rates [29]. Together, these features lead to the flow of neural activity through the circuit and thus its computational properties.

The relationship between neural activity and functions like sensation and behavior is experimentally validated by temporally correlating measurements or estimates of the activity in a circuit with measurements of stimuli or behavior. Models of neural microcircuits are validated by demonstrating that measurements or estimates of cellular and synaptic properties lead to the appropriate neural activity when used to simulate the circuit. One hallmark examples of this process is the discovery of neurons that detect particular visual features by Hubel and Weisel. In their work, they demonstrated that neurons in striate cortex of the cat responded to particular visual features. They then combined knowledge (and some speculation) about the connections and cell types in the circuits which contained the feature detecting neurons to model how the feature selection computation arises from the neural activity in the circuit [62].

One of the core challenges in obtaining these types of measurements is the sheer scale of neural microcircuits. A single circuit is made up of thousands to millions of neurons with orders of magnitude more synapses, and it is not always straightforward to aggregate measurements across different circuits to obtain a sufficient description of the circuit in general (i.e. to estimate higher-order statistics). The mouse has proven to be a useful model for studying microcircuits, especially in cortex. Advances in tools for the readout of neural activity in the mouse have allowed for unprecedented investigation of local neural

circuits. For example, fluorescent imaging [94] has facilitated a dramatic scaling in the ability to estimate and perturb neural activity at the cellular scale while also knowing the location and transcriptional identity of individual neurons. In [81], Calcium imaging was combined with measurements of the synapses between the imaged neurons in mouse visual cortex to demonstrate, like Hubel and Wiesel, that the activity and functional identities of the neurons could be correlated with the structure of the local microcircuit. However, the number of connections probed per circuit was orders of magnitude lower than the number of possible connections between the functionally identified neurons.

While the readout of neural activity has scaled up due to advances in imaging, scaling up measurements of synaptic properties, sometimes called mapping, has been much more difficult. Electron microscopy (EM) has been used to reconstruct microcircuits, but the process of tracing neurons and identifying synapses from the data has been difficult to scale [93]. Additionally, EM reconstructions are a primarily morphological description of the circuit and do not provide much information about the physiological properties of synapses which are essential for modeling the circuit and understanding its computational properties. Measuring the physiological properties of synapses at scale is extremely difficult because it depends on precise experimental access to two neurons simultaneously. Specifically, it requires evoking spikes in a reliable and precise manner from a pre-synaptic cell and, for nearly all properties of interest, requires measuring the subthreshold response of post-synaptic neurons.

1.2 Physiological synaptic mapping

Paired recording, in which both the pre- and postsynaptic cells are recorded intracellularly using whole-cell patch clamp, provides an unambiguous determination of the presence or absence of a synaptic connection as well as unmatched measurements of physiological properties [89, 167, 39]. Dozens of studies across many labs have together provided a quantitative outline of the connectivity between various types of neurons in microcircuits. However, its main drawback is its extremely low throughput. In any of these studies, only a tiny number of presynaptic neurons for any given postsynaptic cell could be probed, and so the total map of any cell type had to be inferred by collating data across different cells, in different circuits, in different animals. Thus even heroic studies, obtaining thousands of such pairs, have only provided small snapshots of the cortical wiring diagram [89, 67, 102].

Optically mapping connectivity by patching one cell and using light stimulate others opened the possibility of mapping many connections in a single circuit. One of the earliest, if not the first, such experiments was conducted in invertebrate ganglion [42]. The authors created molecular probes which could depolarize membranes by creating transient channels when activated with light. They then recorded from one neuron using an electrode while searching for presynaptic partners by stimulating many nearby neurons. The authors note the dramatic increase in both throughput as well as access since recording from small cells with electrodes is more difficult and temperamental.

There are two main factors to consider when designing an optical stimulation method: 1) the technology for targeting light to a subvolume of the sample and 2) the actuator which transduces the light into a transmembrane current.

1.2.1 Targeting light

Initial synaptic mapping experiments used visible wavelengths of light to drive currents in neurons. Visible light is generally safe and affordable to work with, and many compounds and proteins in the natural world interact with these wavelengths. Perhaps the most common method of targeting visible light is laser-scanning microscopy where a focused beam of light is directed to positions in the focal plane by placing galvos in the beam path [159]. Alternatively, one can send a collimated beam out of the objective which fills the field and then select portions of that field using a digital micromirror device in the beam path [138].

Photo-stimulation with visible light can be extremely fast and reliable, and it has been used extensively to map the coarse-scale connectivity of circuits [158, 60, 26, 135]. However, it cannot achieve the requisite spatial precision to stimulate neurons at near-single cell resolution, and therefore it can only make statements about synapses between layers and cell types but not necessarily about individual connections.

Alternatively, two-photon excitation [38] can be used to drive currents [146, 117, 124]. In two-photon excitation, the light-matter interaction is dependent on the coincident arrival of two lower energy photons. The probability of such an event occurring falls off non-linearly away from the focal plane where the convergence of the beam creates the highest density of photons. As a result, the axial resolution for two-photon excitation is considerably better than traditional one-photon excitation. The trade-off for the increased resolution is that considerably more photons are necessary and thus two-photon excitation requires higher power pulse lasers which can damage tissue more easily and are more temperamental and expensive.

Like the targeting of visible light, two-photon excitation can be structured using laser-scanning approaches [146, 117, 124] or a shape-based approach called Computer Generated Holography (CGH) [128, 156, 130]. Holography is a method for encoding an image in the interference pattern between a reference beam and that same reference beam after scattering off of the objects to be imaged. If the reference beam is diffracted off of the recorded interference pattern, the resulting wavefront will generate a three dimensional image of the original objects. Because of this ability to generate three dimensional distributions of light in the image plane, CGH was developed for microscopy to compute interference patterns from desired distributions of light at the sample. Due to its ability to target larger volumes and therefore larger quantities of the actuating molecule, two-photon excitation with CGH has demonstrated the ability to drive extremely fast and reliable spiking in neurons [100, 156, 31].

Importantly, CGH suffers in its axial resolution due to the fact that extending the shape of the beam in the radially also extends the volume of excitation axially. Therefore temporal focusing is commonly used to improve axial resolution [120, 128]. Temporal focusing diffracts the beam in the image plane such that the short laser pulses are stretched in time in a manner that the duration of the pulse is shortest at the focal plane. Analogous to the spatial principle that results in high excitation at the focal plane using two-photon excitation in general, the difference in pulse widths created by temporal focusing increases the level of excitation at the focal plane relative to planes away from the focus by focusing the light not only in space but in time as well. One drawback of temporal focusing is that it eliminates the ability of CGH to distribute light away from the focal plane, but recently three-dimensional temporal

focusing has been demonstrated for two-photon optogenetic excitation [130, 58].

1.2.2 Optical transducers of current

The first optical actuators to become widely used to map synapses were caged neurotransmitters [73, 26]. These neurotransmitters become "uncaged" in response to light and then are free to interact with the endogenous receptors on the nearby neurons to drive optically evoked currents. Uncaging with visible light has proven useful in mapping connections across layers of cortex and to examine correlations in connectivity between simultaneously patched cells and different cell types [157, 158, 182, 6, 103]. Two-photon uncaging of glutamate [117, 10, 47] has also been used to map synapses. In these experiments, individual connections could be characterized resulting in the estimate of many fine-scale synaptic parameters in a single, local cortical circuit. By using genetically targeted fluorescent labels, two-photon uncaging can be used to map synapses between particular subtypes of cells [46, 123].

Because caged compounds can be applied via perfusion at the time of the experiment, they can be in some ways easy to use. However, this is challenging *in vivo* because commonly used forms of caged glutamate act as antagonists at GABA receptors [126, 47] which could lead to epileptiform activity. Additionally, in many experiments neurons are targeted by their soma but most glutamate receptors are on neuronal dendrites. This can lead to ambiguity in the spiking response and presents major challenges for preventing off-target stimulation via the density of dendrites in the nearby neuropil. Off-target currents can also affect the patched cell and contaminate recordings of postsynaptic responses.

Mapping with opsins [23, 37] as the optical transducer (optogenetics) offers key advantages, in particular that the actuator only affects the target cell. Furthermore, opsins can be targeted to genetically defined cell types which facilitates experimental precision and, in the case where the patched cell is non-expressing, eliminates contamination of the postsynaptic recording. Like uncaging, optogenetics with visible light has been widely used to map connectivity in local cortical circuits at the scale of layers [135, 74, 119, 138]. While these methods solve some of the drawbacks of uncaging, they still suffer from the low resolution provided by one-photon approaches.

In contrast, two-photon stimulation can achieve much better spatial precision [124, 146, 100]. However, it became clear early on that there were challenges to achieving cellular resolution excitation with optogenetics [146]. In particular, two-photon optogenetics suffers from the fact that excitation occurs near saturation of the opsin of the molecules. That is, the amount of excitation produced is limited by the amount of available opsin molecules. In a very thorough theoretical and experimental characterization, [146] showed that under these conditions the axial resolution gains of two-photon excitation begin to deteriorate and thus the optical sectioning of two-photon optogenetics will not be as dramatic as the optical sectioning of two-photon imaging.

Nonetheless, several groups have employed two-photon optogenetics to map physiological connections [14, 124, 156, 64, 53]. There are several key considerations when mapping synapses with two-photon excitation - that is, the goal is to identify the properties of individual connections. The sensitivity and kinetics of the opsin heavily influence the success of mapping. In an ideal world, one could spike a presynaptic neuron with arbitrary timing and precision. Put another way, the ideal actuator is infinitely strong and infinitely fast.

Therefore two-photon optogenetics has been most successful when using specially designed strong and fast opsins [100, 31] which allow for low input powers, lower expression levels, and faster trial rates. Nonetheless, all examples of two-photon optogenetic mapping to date use slower opsins [124, 64, 53] which limits the throughput of these methods. Another important factor when optimizing resolution is trafficking the opsin to the soma [14, 178] which is essential for mapping in densely expressing populations. Two-photon optogenetics has been demonstrated for locally mapping both inhibitory (Chapter 3) and excitatory connections [64]. Most recently, Hage et al. [53], used two-photon optogenetics to map putative single connections between many different cell types and multiple layers of mouse cortex in what is by far the most impressive mapping effort yet.

For a full review of optical approaches to mapping synapses in cortex, see [7].

1.2.3 The core challenge

As mentioned above, the key goal of two-photon optogenetic mapping is to characterize as many monosynaptic connections as possible in a single circuit. The key question for any mapping method then becomes, "How can one know if the evoked postsynaptic responses come from the targeted presynaptic cell and not some off-target stimulation?" The answer to this question comes down to the spatial resolution of the stimulation methodology. In all papers on two-photon excitation, the authors provide some characterization of their spatial resolution. In general, these characterizations are then used to estimate error rates given the density of excitable neurons [124, 53]. Often, these error rates are supported by supplementing with paired-patch data.

An alternative approach is to use the characterization of the neural response to optical stimulation to statistically infer if responses are likely to have come from particular excitable neurons [155], and thus far, only one paper has taken this approach with real data [64]. In this thesis, we build on this approach and present novel statistical methods which facilitate two-photon optogenetic mapping and demonstrate their efficacy with real data.

Chapter 2

Single-cell Resolution Mapping With Two-Photon Holographic Optogenetics

This chapter contains unpublished work that was a product of a collaboration with Dr. Shizhe Chen, a postdoc under the supervision of Dr. Liam Paninski, Karl Kilborn and his team at Intelligent Imaging Innovations (3i), as well as Xinyi Deng (Paninski Group), Johannes Friedrich (Paninski Group, Flatiron Institute), and Savitha Sridharan (Adesnik Lab).

2.1 Introduction

In the previous section, we noted that the key challenge of two-photon optogenetic mapping experiments is determining which presynaptic neuron produced the evoked postsynaptic events, and that understanding the resolution of the optogenetic system is key to making that determination. In general, the spiking probability as a function of distance and averaged across cells is used to summarize the excitation resolution. By comparing this to the density of the opsin expressing neurons, one can estimate the error rate of a mapping method.

For example, in Hage et al. [53], the resolution of excitation was estimated for each presynaptic population, and given the densities of each of these populations, they were able to infer that their error rate was low enough to be more or less ignored. Importantly, they also validated their optically estimated connectivity rates between cell types against ground truth paired-patch experiments to support their argument. In general, their method was quite successful. Only one cell type had a large disparity between the optically estimated connectivity rate and the paired-patch estimated rate, somatostatin (SST) expressing neurons. For SST neurons, the optical rate was estimated to be much higher than the rate estimated with patching alone; and not surprisingly, the resolution of optical stimulation for SST cells was the worst of all the cell types they tested (also see Chapter 3). This implies that the resolution for SST cells is not sufficient given the expression density, and it demonstrates that there are conditions under which two-photon optogenetic mapping breaks down when off target excitation is ignored.

Considering this and the fact that the most densely expressing cell type considered in [53], the excitatory Tlx3-Cre line, is still orders of magnitude less dense than the excitatory neuronal population in general, there would be a benefit to increasing the resolution of two-photon optogenetic mapping. While improvements in two-photon optogenetics *per se* have pushed the resolution quite far [100, 58], the current state of the art is likely not sufficient to map dense populations of opsin expressing cells. Additionally, the variance in response can vary significantly across cells due to differences in opsin expression and intrinsic properties which can lead to other difficulties when mapping dense populations of cells, as we will demonstrate later.

Similar to super-resolution microscopy techniques, one can also use computational or statistical methods to improve the resolution of optical mapping [155]. For example, in [64], we used a model of the axial resolution of their two-photon optogenetic system to "deconvolve" off target responses and identify which opsin expressing neuron was in fact connected. In this work, we build on this approach by developing a more detailed optogenetic neural response model and demonstrating how it can be used to perform inference for two-photon optogenetic mapping experiments.

2.2 A high-throughput, temporally precise data acquisition system

The general experimental design of a two-photon optogenetic mapping experiment involves obtaining a whole-cell patch of a designated postsynaptic cell, estimating the locations of the opsin expressing presynaptic neurons, rapidly directing light to those locations, and then analyzing the data to determine which cells are connected and what the properties of those synapses are (Figure 2.1B). Our particular mapping system is designed to maximize throughput both in terms of the trial rate and the number of connections probed.

2.2.1 Voltage-clamp recording

The vast majority of synaptic parameters of interest are estimated using subthreshold responses in postsynaptic cells, i.e. postsynaptic currents (PSCs) and postsynaptic potentials (PSPs). Both types of responses can be recorded using whole-cell patch clamp, with the former using a voltage-clamp configuration and the latter using a current-clamp configuration. For high-throughput mapping, voltage-clamp is the better choice for several reasons. First, there is a more stable baseline and no spiking which reduces non-signal variance the recording and facilitates detecting evoked events for analysis. Second, PSCs have faster kinetics than PSPs since the membrane does not low-pass filter the signal so there is less chance that events will overlap at high trial rates. And third, estimates of synaptic properties such as strength (e.g. amplitude, charge transfer) and paired-pulse ratio have less variance since the driving force remains constant from trial to trial reducing.

The downside to whole-cell patch clamp is that it is relatively low-throughput and a patch has a relatively short time to record before the cell dies. While other methods, such as voltage-imaging [80], may provide a less invasive method to reliably estimate subthreshold responses in neurons in the future, for now those methods are too new and temperamental to

employ in a high-throughput mapping pipeline. Additionally, voltage-clamp recordings suffer from what is known as space-clamp - the fact the membrane potential is only held constant near the patch electrode and not in distal processes where many synapses are located as well as voltage-gated conductances. As a result, many inputs to the cell will appear distorted or may become undetectable in a voltage-clamp recording.

2.2.2 Two-photon holography

In previous two-photon mapping experiments spiral scanning was used to stimulate neurons, but this method of stimulation is relatively slow [123, 64]. Instead, we use holography [129] to shape light into volumes which can activate many more opsin molecules simultaneously and therefore drive enough current to spike the cell more quickly [156, 130].

Our system is based on a custom Phasor Two-Photon Computer-Generated Holography System (3i) with temporal focusing (Figure 2.1A) which is controlled by the SlideBook software. For two-photon excitation, we use a 5W 1040nm femtoTrain laser (Spectra-Physics). Holographic patterns consisted of 10um diameter disks targeted to arbitrary positions in the XY-plane (i.e. the radial plane) at 1um resolution. Because our system uses traditional temporal focusing and not three-dimensional temporal focusing, we target neurons in the Z-dimension (i.e. the optical axis) by moving the objective with a fast and reliable piezo stage (Physik Instrumente). Each pattern is illuminated for 3 milliseconds at desired intensities using a combination of a Pockels cell (Conoptics) and a fast shutter (Uniblitz).

The holographic field of excitation was registered to the two-photon scanning field of view by burning holes in a thin layer of fluorescent substrate at twelve locations defined in the holographic coordinate system. The location of each hole in the imaging coordinate system was determined by user input, and both sets of locations were used to solve for the affine transform between the two coordinate systems. Output intensities of the holographic system were estimated using a calibration procedure based on measuring light intensity at the specimen as a function of Pockels cell input and holographic targeting location.

2.2.3 Two-photon optogenetics

Our goal is to quickly stimulate a dense set of opsin expressing neurons in our cortical slice preparation. The speed at which we can stimulate neurons is a function of both the kinetics of the opsin response as well its strength. Therefore, we use ChroME [99], an opsin developed specifically for two-photon excitation, which has very fast on kinetics and large currents. Spikes produced by combining ChroME and two-photon excitation are extremely precise having low latency and low jitter (Figure 2.2C). Additionally, we the Kv2.1 tag to traffic opsin to the soma and away of from the distal dendrites [14, 178]. Soma-targeting is essential because it increases the spatial resolution of the system by effectively shrinking each cell as well as increases the response relative to the expression level of the opsin. To obtain broad and dense opsin expression in cortical neurons by postnatal day 14, we delivered the construct described above using adeno-associated virus (AAV) injection at P0-P4. We used the AAV9 serotype as it has been shown to produce reliable and broad expression in the central nervous system [36].

Another challenge for two-photon optogenetics is to identify the locations of the opsin expressing cells. While this is not strictly necessary, knowing these locations can drastically reduce the number of stimuli needed to estimate a map. Additionally, being able to identify the neurons in the tissue allows the map to be registered against other modalities such as Calcium imaging or seqFISH [40]. When combined, these methods could align estimates of the functional and transcriptional properties of hundreds of neurons with a synaptic map of their circuit, a goal that neuroscience is making slow but steady progress on [81, 83]. To facilitate the detection of each cell's location, we co-express a fluorescent protein, mRuby3, alongside ChroME. Importantly, a combination of the P2A motif with a nuclear-localization sequence is used which results in the population of opsin expressing cells being visualized as fairly well separated sphere-like objects (Figure 2.1D). Under these conditions, we were able to quickly and reliably detect each neuron's position from a 3D volume imaged with two-photon scanning microscopy using the extremely fast cell detection method from [168] (Figure 2.1E).

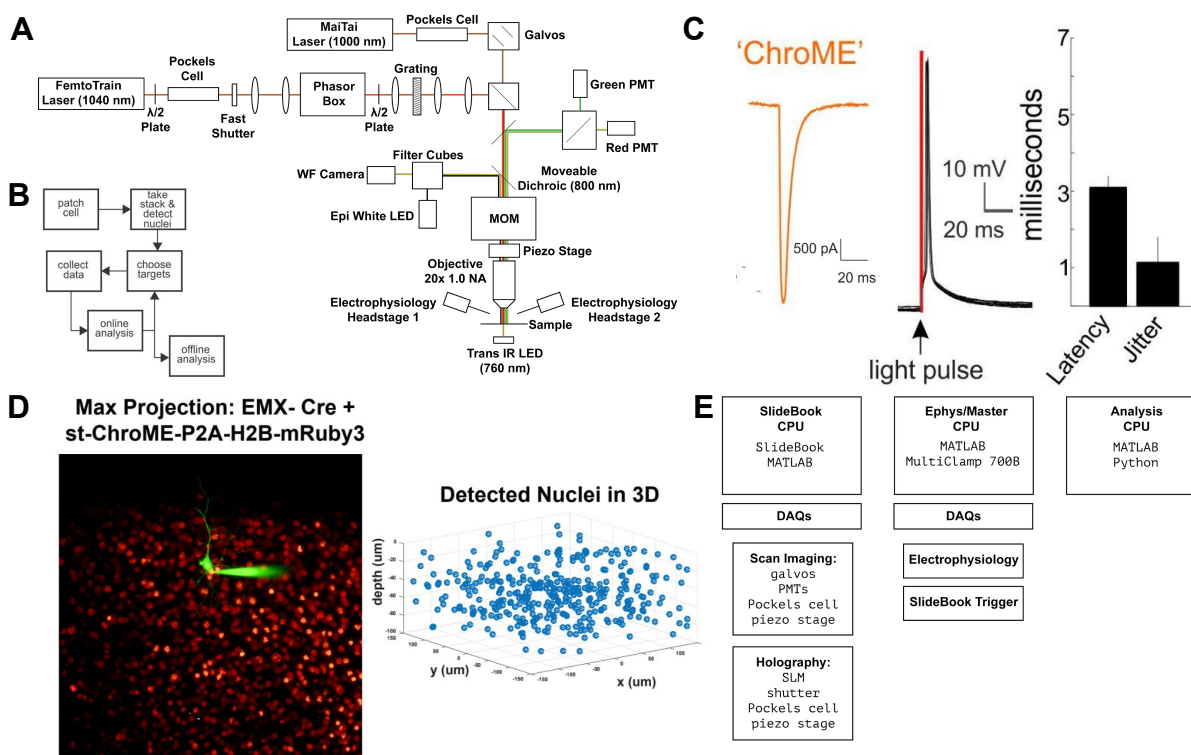


Figure 2.1: (A) A schematic of the microscope showing the holography, scanning, and wide field light paths as well as the headstages for electrophysiology. (B) A flowchart of the experimental workflow. (C) ChroME is a strong and fast opsin. Left, ChroME evokes strong and fast currents. Middle, ChroME induces reliable short latency spikes. Right, the average latency and jitter of a ChroME spike in a cell. (D) The construct delivery system produces broad, dense expression which facilitates 3D cell detection. Left, cells under the EMX-cre promoter in L2/3 of cortex with nuclear mRuby3. A patched neuron filled with Alexa488 dye. Right, the locations of cells in 3D from the cell detection algorithm. (E) A schematic of the computational and I/O resources of the system.

2.2.4 Optimized and automated mapping

The two-photon holographic optogenetic (TPHO) system described above can evoke temporally precise spikes in a dense population of neurons given a holographic pattern. However, to run this system at a high trial rate requires solving two bottlenecks. First, for each desired holographic pattern, one must compute the phase mask for that pattern. Second, one must coordinate the loading of phase masks into RAM with the execution of trial batches.

Computing arbitrary phase masks is achieved using the Gerchberg-Saxton [50] algorithm which can take hundreds of milliseconds or more to run. To speed up phase mask computation, we developed an algorithm which takes a precomputed phase mask for a single 10um disk at the center of the field and shifts it to an arbitrary location in the field of excitation at 1um resolution. The algorithm also uses precomputed phase masks for diffraction limited spots spaced at 10um intervals across the field and second set of phase masks for spots spaced at 1 um interval across the center 10um x 10um square of the field. By performing two convolutions (i.e. parallelizable elementwise multiplication the Fourier domain), we are able to quickly move the centered disk to the closest point in the 10um spacing grid and then shift it up to +/- 5um along either axis using the 1um spacing grid. Using this method allowed phase masks to be computed at rates faster than single trials.

To optimize acquisition time, we also coordinated the computation of phase masks on a separate analysis computer. After the system begins recording the first batch, it immediately triggers the analysis computer to begin computing the next batch of phase masks. Once these masks are ready, they are sent back to the computer which communicates with the SLM and loaded into RAM in preparation for the next batch. Additionally, a third computer acted as a master coordinator and handled the main DAQ I/O (Figure 2.1E). Under these conditions, we were able to perform mapping experiments with very little down time once mapping began.

The process of finding a cell to patch, patching the cell, imaging the 3D volume, detecting cells, defining stimuli, computing phase masks, and then executing those stimuli on the system can be quite cumbersome if every step requires user input. To speed up the entire process - which is important because the first step is obtaining the whole-cell patch - we designed the master control software to only require one click once the patch was acquired and the patched cell was centered in the field.

2.3 Single-cell resolution optogenetics is not (currently) possible

The fundamental challenge in analyzing the data resulting from the mapping experiment described above is that we do not know which presynaptic neuron fired the spike that led to an observed post-synaptic response. This is because the optogenetic stimulation may elicit spikes in nearby neurons even if it is intended to stimulate only one particular neuron. These off-target activations are a consequence of multiple factors.

2.3.1 Experimental parameters that affect resolution

Perhaps the most important parameter that affects the resolution of TPHO is the stimulation intensity. We illustrate this issue in Fig. 2.2A. Here we record spikes from opsin-expressing neurons while stimulating at locations near this neuron at various powers. Each color represents the data from a different neuron - aligned so they are all centered at the origin. Small, transparent dots show stimulation locations and solid, larger dots show locations which evoked spikes. The top row shows the projection onto the XY plane and bottom row shows the projection onto the the XZ plane.

It is immediately clear that the resolution of TPHO is a function of stimulus intensity. That is, with more power, neurons spike in response to stimuli that are further away. Additionally, neurons do not show a stereotyped response as a function of either stimulus location or intensity. For example, stimuli as far as $30\ \mu\text{m}$ away can cause the red neuron to spike at 15 mW while the green neuron produced did not fire in response to stimuli at any locations at the same power. At the population level, we see that the fraction of cells which spike as a function of stimulus intensity slowly increases from about 30% at 5 mW until nearly all neurons spike at around 35 mW (Figure 2.2B). This means that the resolution of TPHO is also a function of the properties of each neuron, in particular the rheobase and opsin expression level which together determine each cell's sensitivity to light.

Fig. 2.2C shows the histogram of pairwise distances between neurons taken from one representative field of opsin expressing neurons in a cortical slice. In this particular sample, about 44.2% of neurons have a neighbour within $20\ \mu\text{m}$. Given the variance in light sensitivity across the field and tight packing of opsin expressing cells, it is essentially guaranteed that the minimum stimulus intensity to spike some cells will evoke spikes in nearby cells with either more opsin expression or lower rheobases. (Note that in this work we restrict out presynaptic population to cortical L23 excitatory cells, i.e. cells in L23 which express under the EMX1-Cre line, to reduce the variance due to intrinsic cell type properties.) Thus, simple approaches that estimate synaptic connections based solely on the stimulation location and the existence of postsynaptic events will be prone to false-positive errors. For example, the responses seen in a patched cell when stimulating the cells in the local area shown in Figure 2.2D reveal potential connections at 4 out of 5 of the cell locations. If we believed our mode of stimulation was truly cellular resolution, it would seem reasonable to predict that all four cells are presynaptic to the patched cell. If we want to create high-throughput system with a low false positive rate, we must have a way to deal with the uncertainty about which particular cells are connected when mapping.

2.3.2 A spike-timing solution

In previous work [64], we demonstrated that when everything else is held constant, evoked spike times are at a minimum when stimulation is directly on the cell. This lead to a relatively simple approach which was able to localize the presynaptic neuron by in the Z-axis by finding the plane which minimized the evoked PCS timing. Then due to the sparsity of expression it was straight-forward to determine which neuron in that plane produced the responses. We build upon this approach by developing a more detailed model which describes how two-photon holographic stimulation evokes spikes at varying times in a stimulus intensity and

distance dependant manner. We can then use this model to map circuits by extending it to describe how PCS timings arise in the postsynaptic cell given the evoked presynaptic spikes.

We can break down the model into three parts (Figure 2.2F). The first part of the model represents the transformation from TPHO to spike times or spike rates, the neural response model. The second part represents the transform from presynaptic spike rates to a postsynaptic PSC event rate, the connectivity model. And finally, the third part of the model represents the transform from the PSC event rate to a noisy voltage-clamp recording. In this chapter, we focus on the first two parts of the model and use the PSC event times as the response observation. For details on the voltage-clamp modeling and estimating the PSC event times see Chapter 4.

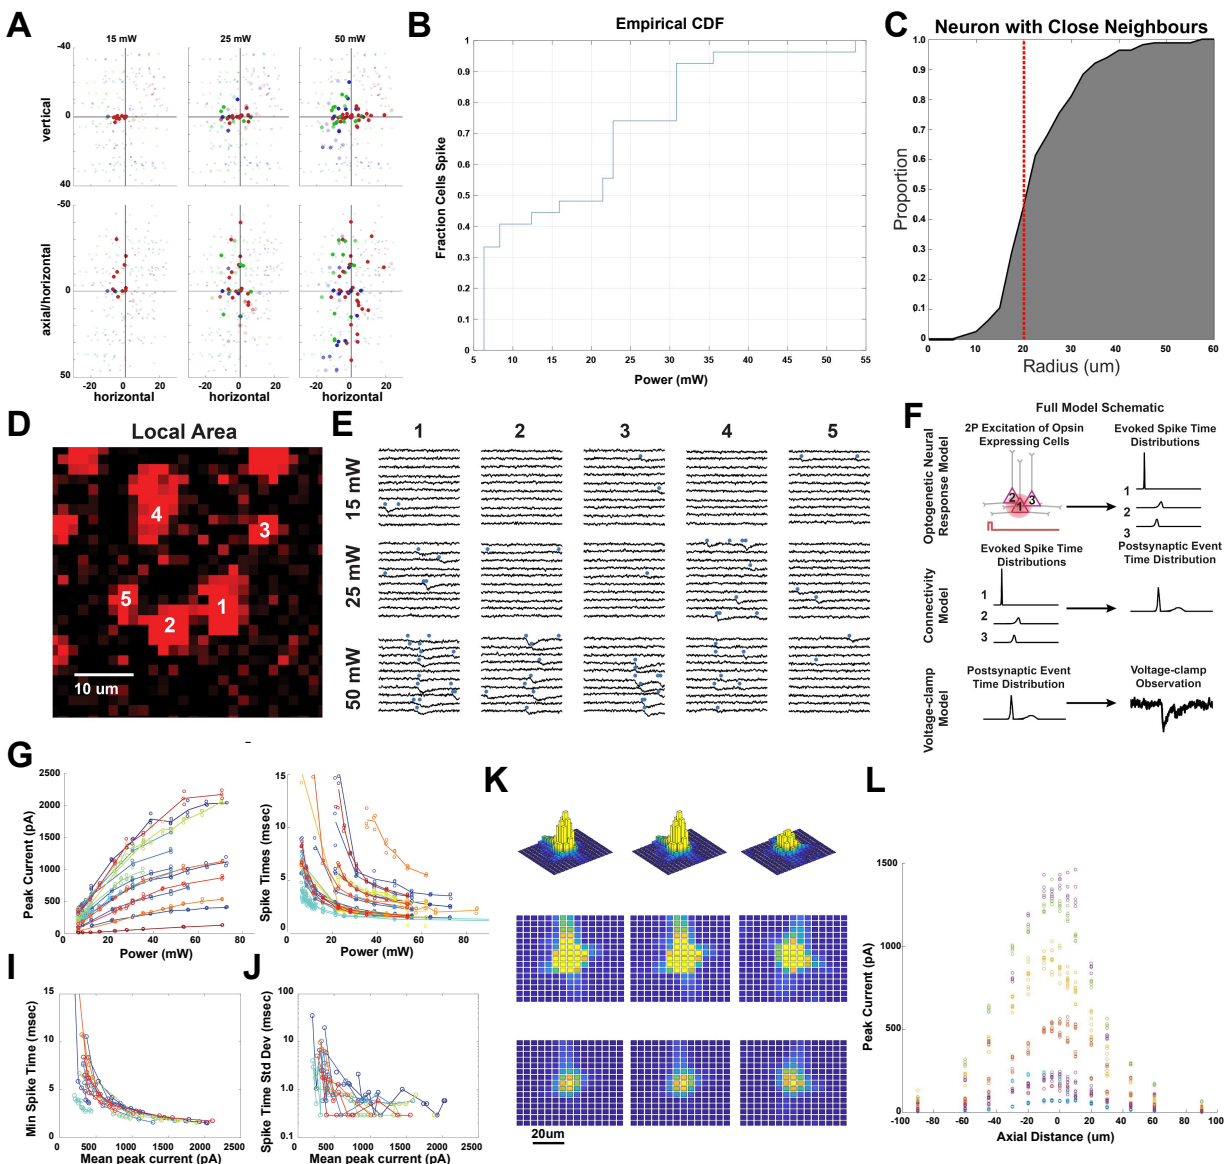


Figure 2.2: (A) Locations that evoke spikes from neurons at various powers. Each color represents a neuron. Smaller transparent dots are stimulus locations that did not evoke spikes and larger opaque dots are locations which did evoke spikes. The top row is the XY projection of the data and the bottom the XZ projection. (B) The empirical CDF of the distribution of minimum powers to spike a neuron with TPFO. (C) The empirical CDF of the distance to the closest neighbor neurons for a representative population. The vertical dashed red line helps indicate that almost half of all neurons have neighbor that is very close relative to the resolution of TPFO. (D) 5 opsin expressing neurons in a local area and (E) the postsynaptic response as each location when stimulated at varying powers. Each trace is one trial, blue dots represent the output of PSC event detection. (F) A schematic of our generative model for the TPFO mapping experiment. *Continued on next page.*

2.3.3 Building an intuition for the neural response model

Before we provide the mathematical details of the model, we will provide a basic intuition for the phenomenon we are trying to represent by examining some preliminary data of the neural response to our TPHO system. In the following section we will present a statistical analysis of this data as part of our presentation of our mapping model and inference procedure.

We start by considering the case where we only modulate the stimulus power and hold stimulus location directly on the soma, i.e. the location returned by the cell detection algorithm. We first stimulated while recording spikes in cell-attached mode, and then switched to voltage-clamp mode to repeat the same set of stimuli while recording evoked currents. As expected, the evoked current increases monotonically as the stimulus power increases and then eventually saturates (Figure 2.2G). Similarly, we see spike times decrease as a function of power and eventually hit a minimum time and exhibit extremely low jitter (Figure 2.2H,J). While the shapes of the response curves are similar across neurons, we also once again see a large variation of response sensitivity across the population. Based on the current responses, it seems possible that a scalar gain parameter per presynaptic neuron could be used to predict the evoked current given the stimulus power for a particular location. To investigate the transform from evoked current to spike time, we leveraged the fact that the data in Figures 2.2G and 2.2H were taken on the same cells and plotted the current-spike time pairs for each power and for each neuron (Figure 2.2I). The points closely followed a one-dimensional curve indicating that neurons expressing under the EMX-Cre line (mostly L23 excitatory cells) all have very similar rheobases, and that we can likely model the transform from stimulus to spike times as a scaling of the intensity which is then transformed by a function fit to the scattered data in Figure 2.2I.

Next we investigate the spatial response of neurons by keeping the stimulation power constant and raster a stimulus over the neuron while recording evoked currents (Figure 2.2K). The heatmaps generated by this procedure represent the spatial receptive field of the neuron which is the convolution of the stimulus shape with the spatial distribution of opsin in the neuron. The extensions of expression that reach out towards the top of the map, toward pia, are a result of the Kv2.1 tag also trafficking to the proximal apical dendrite which is consistent with [178, 14]. The non-Gaussian, non-convex, and relatively variable shapes of the opsin distribution in each neuron means that we cannot use a simple model to represent the shape of neurons in the radial plane. Measurements of evoked current along the Z-axis reveal a much simpler shape, but they also demonstrate how poor the resolution of TPHO can be along the optical axis (Figure 2.2L).

2.4 Single-cell resolution analysis for mapping

2.4.1 Generative statistical model

To use the intuitions we built from the optogenetic response data, we developed a Bayesian statistical model which describes the transforms from stimulus to spike times to postsynaptic responses using probability distributions. By modeling the gain and spatial receptive field for each opsin expressing neuron, fitting the model will (softly) assign postsynaptic responses back to single presynaptic neurons even if our stimulation is not single-cell resolution.

To formally introduce the generative model we need to begin by defining some notation for the three basic elements of the experiment. First, there are K opsin expressing neurons in the field of excitation. We denote the three dimensional location of the k -th neuron as \mathbf{w}_k for $k = 1, \dots, K$. Second, each stimulus on trial i has a location and intensity denoted as \mathbf{x}_i and I_i , respectively, and each experiment consists of M trials. (We note that while our model can handle multiple stimulus locations per trial, in this work we only consider one stimulus per location.) And finally, the output of each trial is a voltage clamp current recording, $J_i(t)$.

Optically evoked currents

The data presented above suggest that to develop an optogenetic neural response model, it is useful to first model the peak magnitude $m_{i,k}$ of the current elicited by the optogenetic stimulus on trial i in neuron k , and then in turn model the spiking of the cell at time t given $m_{i,k}$.

As demonstrated, $m_{i,k}$ is dependant on three critical factors: the stimulus location \mathbf{x}_i relative to the spatial profile of neuron k , the neuron's opsin expression level, denoted ϕ_k , and the stimulus intensity I_i . In the stimulus regime used here, it is reasonable to use a linear model to connect these ingredients; thus we model the mean evoked current $m_{i,k}$ as

$$m_{i,k} = s_k(\mathbf{v}_{i,k})\phi_k I_i, \quad (2.1)$$

where $\mathbf{v}_{i,k} = \mathbf{x}_i - \mathbf{w}_k$ and $s_k(\cdot)$ is a random function that reflects the spatial profile of neuron k , centered around its nucleus location \mathbf{w}_k . We assume that $s_k(\mathbf{0}) = 1$ so that it is identifiable with the total opsin expression level ϕ_k . Because the spatial profile is a three dimensional function that differs across neurons, it requires a large amount of trials at various locations in order to accurately learn $s_k(\cdot)$ for each neuron which is infeasible during a mapping experiment. Therefore, we propose to factorize the shape function of each neuron as the product between two a radial shape function $s_{k,xy}(\cdot)$ and an axial shape function $s_{k,z}(\cdot)$

$$s_k(\mathbf{v}_{i,k}) = s_{k,xy}^{a_{i,k,xy}}(\mathbf{v}_{i,k,xy})s_{k,z}^{a_{i,k,z}}(\mathbf{v}_{i,k,z}),$$

where $a_{i,k,xy} \equiv \mathbf{v}_{i,k,xy}^2 / \|\mathbf{v}_{i,k}\|_2^2$, and $a_{i,k,z} \equiv \mathbf{v}_{i,k,z}^2 / \|\mathbf{v}_{i,k}\|_2^2$.

Figure 2.2: *Continued from previous page.* (G) The maximum current evoked by stimulation as a function of power. Each color is an individual cell and each dot represents a trial. The line goes through mean for each condition. (H) Like G the evoked spike times. (I) The mean peak current plotted against the mean spike time for each cell and power. (J) The standard deviation of the spike times in H. (K) Peak current as a function of XY location, each columns is a cell. The first row shows the values as a 3D bar graph to help understand the scale within and across cells. The second and third row are just the flat heatmaps. The first and second row scale the color axis such that the color is clipped at roughly the rheobase of the cells to illustrate which locations are likely to evoke spikes for each cell. The bottom row is normalized within each cell. (L) Evoked currents as a function of axial distance. Each color is a unique cell and each circle represents a single trial.

Probabilistic spike timing

Next let $\tilde{t}_{i,k}$ be the evoked spike time of neuron k in the i -th trial. In the experimental conditions analyzed here, the stimulation intensities are chosen such that, usually, one spike is elicited per trial per presynaptic neuron. To further ensure that at most one evoked per neuron is observed per trial, we choose the shortest reasonable trial duration (in this work 7.0 ms). We approximate the distribution of $\tilde{t}_{i,k}$ given $m_{i,k}$ using a truncated normal distribution with mean $g(m_{i,k})$ and variance $h(m_{i,k})$, where $g(\cdot)$ and $h(\cdot)$ are both monotonically decreasing functions (i.e., the average spike latency and variance both decrease as the current magnitude increases) representing the trends in Figures 2.2I and 2.2J. Note that if we had more than one cell type in our opsin expressing population, we would likely need to model these functions for each type of cell.

Thus, the density of $\tilde{t}_{i,k}$ takes the form

$$\tilde{\lambda}(\tilde{t} | m_{i,k}) = \psi \left\{ \frac{\tilde{t} - g(m_{i,k})}{h(m_{i,k})} \right\}, \quad 0 < \tilde{t} < T. \quad (2.2)$$

Here ψ denotes the standard normal density, and note that $\tilde{\lambda}$ does not integrate to one on the given interval $[0, T]$; the complement $1 - \int_0^T \tilde{\lambda}(\tilde{t} | m_{i,k}) dt$ is simply the probability that neuron k does not spike on the i -th trial given $m_{i,k}$.

Transforming spikes into postsynaptic responses

Synaptic transmission is probabilistic; not every presynaptic event is transmitted to the postsynaptic neuron, and this synaptic release probability may vary from cell to cell. Therefore, we assume that a presynaptic spike at time $\tilde{t}_{i,k}$ induces a postsynaptic event at time t_{i,j_k} with probability γ_k , and that

$$t_{i,j_k} = \tilde{t}_{i,k} + d_k + e_{i,k}, \quad (2.3)$$

where d_k is a fixed synaptic delay between neuron k and the patched neuron, and $e_{i,k} \sim \mathcal{N}(0, \delta_k^2)$ is a random jitter. Note that we use j_k to indicate the j th postsynaptic event on that trial and that its source is neuron k . We refer to $\gamma_k \in [0, 1]$ as the *probability of transmission*; if $\gamma_k = 0$ then no connection exists between neuron k and the patched neuron. If $\gamma_k \in (0, 1]$, then it can be thought of as the probability that at least one vesicle is released from the presynapse given a spike. In addition to the evoked postsynaptic events, the patch recording may also contain spontaneous events that follow a homogeneous Poisson process with a small constant intensity λ_0 . In general, we treat λ_0 as known because it can be estimated from portions of the voltage-clamp recording between stimulation batches.

Finally, we observe the postsynaptic current $J_i(t)$ and estimate the postsynaptic event times from it using the method presented in Chapter 4. While event detection may introduce errors, we note that by biasing detection to false positives we simply increase the Poisson baseline intensity λ_0 which is preferred over false negatives which lead to decreases in the estimated γ_k . We denote the detected events on the i th trial as $\mathcal{T}_i = \{t_{i,j} : j = 1, \dots, n_i\}$ if $n_i > 0$, and $\mathcal{T}_i = \emptyset$ if $n_i = 0$.

The data on each trial i are summarized as $\mathcal{D}_i = \{\mathbf{x}_i, I_i, \mathcal{T}_i\}$, that is the stimulus locations, stimulus intensities, and the detected postsynaptic event times.

2.4.2 Inference

We take a Bayesian approach for inference in the proposed model. The two key ingredients in the Bayesian approach are the likelihood and prior distributions.

Likelihood

The likelihood of the model proposed in Section 2.4.1 takes the form

$$\mathcal{L}(\mathcal{D}; \Theta) \equiv \prod_{i=1}^M \left\{ \sum_{A \in \mathcal{A}(K+1, n_i)} \prod_{k \in A} \lambda_{i,k}(t_{i,k_A}) \prod_{k \notin A} \left[1 - \int_0^T \lambda_{i,k}(t) dt \right] \right\}, \quad (2.4)$$

where $\mathcal{A}(K+1, n_i)$ denotes all possible n_i permutations of the $K+1$ possible sources of the postsynaptic events, and k_A is a short-hand notation for the index of k in A . To be specific, each permutation $A \in \mathcal{A}(K+1, n_i)$ is an ordered set consisting of n_i entries, where each entry takes a value in $\{0, 1, 2, \dots, K\}$. For instance, if $n_i = 2$, then $A = \{7, 0\}$ means that the first postsynaptic event is induced by a spike from neuron 7 and the second event is a spontaneous event. $\lambda_{i,k}(t)$ represents the contribution of neuron k to the overall postsynaptic event rate in the i th trial

$$\lambda_{i,k}(t) = \gamma_k \int_0^\infty \tilde{\lambda}_{i,k}(t - \Delta; m_{i,k}) \psi([\Delta - d_k]/\delta_k) d\Delta, \quad (2.5)$$

where $\lambda_{i,k}(\cdot; m_{i,k})$ is the firing rate of neuron k in the i th trial given evoked current $m_{i,k}$. Recall that $\tilde{\lambda}_{i,k}$ follows a truncated normal distribution with mean $g(m_{i,k})$ and variance $h^2(m_{i,k})$. As a result, we can simply convolve to approximate (for positive t)

$$\lambda_{i,k}(t) \approx \gamma_k \psi \left(\frac{t - d_k - g(m_{i,k})}{[\delta_k^2 + h^2(m_{i,k})]^{1/2}} \right).$$

For the background events where $k = 0$, we set $\lambda_{i,0}(t) = \lambda_0$.

The unknown parameters in the likelihood are $\{\gamma_k, d_k, s_{k,xy}(\cdot), s_{k,z}(\cdot), \phi_k, \delta_k : k = 1, \dots, K\}$. Although $s_{k,xy}(\cdot)$ is a function that takes values on a two-dimensional space, we only need to learn its value at the unique locations stimulated in the experiment. Specifically, if there are $L_{k,xy}$ unique locations (projected to the $z = 0$ plane) that have been stimulated, we only need to learn the values of $s_{k,xy}(\cdot)$ at these $L_{k,xy}$ locations $\{s_{k,xy}(\mathbf{v}_{l,k,xy}) : l = 1, \dots, L_{k,xy}\}$. Similarly for the function $s_{k,z}(\cdot)$, we only need to learn its values at the unique stimulation locations, i.e., $\{s_{k,z}(v_{l,k,z}) : l = 1, \dots, L_{k,z}\}$.

Prior

The other ingredient is the prior distribution for the unknown parameters. For simplicity in notation, we introduce Θ_k to represent the unknown parameters associated with neuron k , i.e., $\Theta_k \equiv \{\gamma_k, \phi_k, s_k(\cdot), d_k, \delta_k : k = 1, \dots, p\}$, and we use $\theta \in \Theta_k$ to represent one of the unknown parameters in Θ_k . We introduce a prior distribution for each of the unknown

parameters in the likelihood, and we write the prior distribution for all unknown parameters as

$$\prod_{k=1}^K \mathbb{P}(\Theta_k) = \prod_{j=1}^K \prod_{\theta \in \Theta_k} p_0(\theta; \alpha_\theta^0, \beta_\theta^0, \mathbf{b}_\theta). \quad (2.6)$$

Briefly, we use the logit-normal distribution as the prior distribution for $\{\gamma_k, d_k, \phi_k, \delta_k : k = 1, \dots, K\}$, where the bounds \mathbf{b}_θ for the logit-normal distributions are set based on the response data. In other words, we assume that

$$p_0(\theta; \alpha_\theta^0, \beta_\theta^0, \mathbf{b}_\theta) = \psi \left(\frac{f(\theta; \mathbf{b}_\theta) - \alpha_\theta^0}{\exp(\beta_\theta^0)} \right), \quad (2.7)$$

where f represents the logit transformation with bounds $\mathbf{b}_\theta = (b_{\theta,1}, b_{\theta,2})^\top$, i.e.,

$$f(\theta; \mathbf{b}_\theta) = \log \left(\frac{\theta - b_{\theta,1}}{b_{\theta,2} - \theta} \right).$$

We use a Gaussian process to model the spatial profiles $s_{k,xy}(\cdot)$ and $s_{k,z}(\cdot)$. Recalling that we only need to learn values of the functions at the unique locations that have been stimulated, the Gaussian processes thus reduces to multivariate normal distributions. In particular, we have

$$(s_{k,xy}(\mathbf{v}_{1,k,xy}), \dots, s_{k,xy}(\mathbf{v}_{L_{k,xy,k,xy}}))^\top \sim \mathcal{N}(\boldsymbol{\mu}_{xy}(\mathbf{V}_{k,xy}), \boldsymbol{\Sigma}_{xy}(\mathbf{V}_{k,xy}, \mathbf{V}_{k,xy})) \quad (2.8)$$

$$\text{and } (s_{k,z}(v_{1,k,z}), \dots, s_{k,z}(v_{L_{k,z,k,z}}))^\top \sim \mathcal{N}(\boldsymbol{\mu}_z(\mathbf{V}_{k,z}), \boldsymbol{\Sigma}_z(\mathbf{V}_{k,z}, \mathbf{V}_{k,z})), \quad (2.9)$$

where $\mathbf{V}_{k,xy} = (\mathbf{v}_{1,k,xy}, \dots, \mathbf{v}_{L_{k,xy,k,xy}})$ is a matrix that contains all unique stimulation positions relative to the nucleus of neuron k projected on the plane $\{\mathbf{x} \in \mathbb{R}^3 : x_3 = \omega_3\}$, and $\mathbf{V}_{k,z} = (v_{1,k,z}, \dots, v_{L_{k,z,k,z}})$ is a vector containing all the unique z-coordinates of stimulation locations. We discuss the how we estimate the mean and variances functions for the two Gaussian processes from pilot studies in Section 2.5.2.

Approximating the posterior with stochastic variational inference

Combining the likelihood and the prior, the posterior distribution of the unknown parameter $\{\Theta_k\}_{k=1}^K$ satisfies that

$$\mathbb{P}(\Theta \mid \mathcal{D}) \propto \mathcal{L}(\mathcal{D}; \Theta) \prod_{k=1}^K \mathbb{P}(\Theta_k). \quad (2.10)$$

We approximate the true posterior distribution with mean-field variational inference. In variational inference, we aim to maximize the evidence lower bound

$$\text{ELBO}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \equiv \mathbb{E}_q \left[\log \mathcal{L}(\mathcal{D}; \Theta) + \sum_{k=1}^K \log \mathbb{P}(\Theta_k) - \sum_{k=1}^K \sum_{\theta \in \Theta_k} \log q(\theta; \alpha_\theta, \beta_\theta, \mathbf{b}_\theta) \right], \quad (2.11)$$

where $q(\theta; \alpha_\theta, \beta_\theta, \mathbf{b}_\theta)$ denotes the variational distribution for an unknown parameter $\theta \in \Theta_k$. We let $q(\theta; \alpha_\theta, \beta_\theta, \mathbf{b}_\theta)$ be a logit-normal distribution with bounds \mathbf{b}_θ for $\theta \in \{\gamma_k, d_k, \phi_k, \delta_k, \sigma_{k,e}^2 : k = 1, \dots, K\}$.

$k = 1, \dots, K\}$, where the bounds are the same as in the prior distributions. We approximate the posterior of $(s_{k,xy}(v_{1,k,xy}), \dots, s_{k,xy}(v_{L_k,xy}))^T \in \mathbb{R}^{L_{k,xy}}$ using a multivariate normal distribution $\mathcal{N}(\tilde{\boldsymbol{\mu}}_{k,xy}, \tilde{\boldsymbol{\Sigma}}_{k,xy})$, where $\tilde{\boldsymbol{\mu}}_k$ is an $L_{k,xy}$ -vector, and $\tilde{\boldsymbol{\Sigma}}_{k,xy}^{-1} = \boldsymbol{\Sigma}_{xy}^{-1}(\mathbf{V}_{k,xy}, \mathbf{V}_{k,xy}) + \mathbf{B}_{k,xy}^{-1}$ with $\mathbf{B}_{k,xy}$ being an $L_{k,xy} \times L_{k,xy}$ diagonal matrix. In this variational distribution, there are $2L_{k,xy}$ unknown parameters in $\tilde{\boldsymbol{\mu}}_k$ and $\mathbf{B}_{k,xy}$. We note that the assumed form of $\mathcal{N}(\tilde{\boldsymbol{\mu}}_k, \tilde{\boldsymbol{\Sigma}}_k)$ is commonly used in Gaussian approximation. Intuitively, the mean vector $\tilde{\boldsymbol{\mu}}_K$ allows the variational distribution to adjust its mean based on the data, and the additive diagonal matrix \mathbf{B} describes how reliable the information from the data is compared to the prior. When the entries in \mathbf{B} are sufficiently small, we know that $\tilde{\boldsymbol{\Sigma}}_k \approx \mathbf{B}$, which implies that the posterior distribution for each value $s_k(v_{l,k})$ is determined by the observed data rather than the prior. As a final note, although we introduce separate variational distributions for each unknown parameter, the variational parameters, i.e., the maximizers for the ELBO (2.11), are connected via the likelihood. Similarly for $(s_{k,z}(v_{1,k,z}), \dots, s_{k,z}(v_{L_k,z}))^T \in \mathbb{R}^{L_{k,z}}$, we approximate its posterior distribution using a multivariate normal distribution $\mathcal{N}(\tilde{\boldsymbol{\mu}}_{k,z}, \tilde{\boldsymbol{\Sigma}}_{k,z})$ with $\tilde{\boldsymbol{\Sigma}}_{k,z}^{-1} = \boldsymbol{\Sigma}_z^{-1}(\mathbf{V}_{k,z}, \mathbf{V}_{k,z}) + \mathbf{B}_{k,z}^{-1}$. With a slight abuse of notation, we denote the variational distribution for both vectors as $q(\theta; \alpha_\theta, \beta_\theta, \mathbf{b}_\theta)$ when $\theta = (s_{k,z}(v_{1,k,z}), \dots, s_{k,z}(v_{L_k,z}))^T$ and $\theta = (s_{k,xy}(v_{1,k,xy}), \dots, s_{k,xy}(v_{L_k,xy}))^T$.

We use gradient descent to find the maximizer of the evidence lower bound in (2.11). Briefly, we take the black-box variational inference approach proposed in [143]. This proposal uses a Monte-Carlo method to estimate the gradient of $\text{ELBO}(\boldsymbol{\alpha}, \boldsymbol{\beta})$ with respect to $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ as follows

$$\begin{aligned} & \hat{\nabla}_{(\boldsymbol{\alpha}, \boldsymbol{\beta})} \text{ELBO}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \\ &= \frac{1}{S} \sum_{s=1}^S \nabla_{(\boldsymbol{\alpha}, \boldsymbol{\beta})} \log q(\boldsymbol{\Theta}[s]) \left[\log \mathcal{L}(\mathcal{D}; \boldsymbol{\Theta}[s]) + \sum_{k=1}^K \log \mathbb{P}(\boldsymbol{\Theta}_k[s]) - \sum_{k=1}^K \sum_{\theta \in \boldsymbol{\Theta}_k[s]} \log q(\theta) \right], \end{aligned} \quad (2.12)$$

where we write $q(\boldsymbol{\Theta}[s]; \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{b})$ as $q(\boldsymbol{\Theta}[s])$ for simplicity, S is a user-specified parameter, and $\boldsymbol{\Theta}[s]$ is a sample from $q(\boldsymbol{\Theta}[s]; \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{b})$. The Monte-Carlo gradient is an approximation to the true gradient whose accuracy depends on the parameter S . To control the variability from Monte-Carlo, we employ the control variate as proposed in [143]. The algorithm is given in Algorithm 1.

Computational details of Algorithm 1

The black-box variational inference algorithm (Algorithm 1) is computationally demanding due to the use of Monte-Carlo sampling for approximating the gradients. In this section, we discuss the computational details of Algorithm 1 and introduce our strategies for reducing the computational cost.

The main factors that affect the computational complexity of Algorithm 1 are the number of iterations to convergence C , the number of Monte Carlo samples S , the total number of trials M , and the total number of presynaptic neurons K . Roughly speaking, the computational complexity of Algorithm 1 is proportional to the product of these factors $\mathcal{O}(CSMK)$. One can see that the computational cost will quickly grow out of control when the size of the field of excitation increases: a larger field will include more presynaptic neurons, K . More

Algorithm 1: Variational inference for postsynaptic neural response model.

```

Input: data  $\mathcal{D}$ , likelihood function  $\mathcal{L}(\mathcal{D}; \boldsymbol{\theta})$ , prior distribution  $\mathbb{P}(\boldsymbol{\Theta})$  (2.6), mean field
variational family  $\mathcal{Q}_\theta = \{q(\theta; \alpha_\theta, \beta_\theta, \mathbf{b}_\theta) : \alpha_\theta, \beta_\theta\}$  for each  $\theta \in \boldsymbol{\Theta}$ ;
Initialize  $\boldsymbol{\alpha}, \boldsymbol{\beta}$  given prior knowledge or rough estimates from the data;
Set  $\epsilon = \kappa + 1$ ;
while  $\epsilon < \kappa$  do
  for  $s = 1, \dots, S$  do
     $\boldsymbol{\Theta}[s] \sim q_{\boldsymbol{\Theta}}(\boldsymbol{\Theta}; \boldsymbol{\alpha}, \boldsymbol{\beta})$ ;
    Calculate
     $\mathcal{E}[s] = \log \mathcal{L}(\mathcal{D}; \boldsymbol{\Theta}) + \sum_{k=1}^K \log \mathbb{P}(\boldsymbol{\Theta}_k) - \sum_{k=1}^K \sum_{\theta \in \boldsymbol{\Theta}_k} \log q(\theta; \alpha_\theta, \beta_\theta, \mathbf{b}_\theta)$ 
  end
  for  $k = 1, \dots, K$  do
    for  $\theta \in \boldsymbol{\Theta}_k$  do
      for  $s = 1, \dots, S$  do
         $\mathbf{a}_j[s] = \nabla_{(\alpha_\theta, \beta_\theta)} \log q(\theta; \alpha_\theta, \beta_\theta, \mathbf{b}_\theta) \mathcal{E}[s]$ ;
         $\mathbf{b}_j[s] = \nabla_{(\alpha_\theta, \beta_\theta)} \log q(\theta; \alpha_\theta, \beta_\theta, \mathbf{b}_\theta)$ ;
      end
    end
    for  $\theta \in \boldsymbol{\Theta}_k$  do
       $\mathbf{c}_\theta = [\text{var}(\mathbf{b}_\theta)]^{-1} [\text{cov}(\mathbf{a}_\theta, \mathbf{b}_\theta)]$ ;
       $\hat{\nabla}_{(\alpha_\theta, \beta_\theta)} \text{ELBO} = S^{-1} \sum_{s=1}^S (\mathbf{a}_\theta[s] - \mathbf{c}_\theta \mathbf{b}_\theta[s])$ ;
    end
  end
   $\rho = t^{-1}$ ;
   $(\alpha_\theta, \beta_\theta)^\top = (\alpha_\theta, \beta_\theta)^\top + \rho \hat{\nabla}_{(\alpha_\theta, \beta_\theta)} \text{ELBO}$ ;
   $t = t + 1$ ;
  Calculate the normalized change  $\epsilon = \|\rho \hat{\nabla}_{(\alpha_\theta, \beta_\theta)} \text{ELBO}\|_2 / \|(\alpha_\theta, \beta_\theta)\|_2$ ;
end

```

neurons naturally require more trials in order to map the circuit, and more neurons also introduces more unknown parameters to learn which may slow down convergence. We take the following strategies to address concerns about computational cost.

First, the number of iterations for Algorithm 1 to converge directly relates to the quality of initialization. Good initial estimates of the variational parameters $\boldsymbol{\alpha}, \boldsymbol{\beta}$ may significantly reduce the computing time for Algorithm 1. In addition, initialization helps find the global maximal, given that the evidence lower bound (2.11) is not necessarily a convex function of the variational parameters. Additionally, we develop a fast initialization strategy for the variational parameters based on prior knowledge and crude—but fast—analysis of data.

Second, we follow [143] to use a control variate to reduce the variability in the estimated gradients. Black-box inference approximates the gradients of the evidence lower bound using Monte Carlo sampling. The number of Monte Carlo samples S controls the accuracy for the estimated gradients. Increasing S yields more accurate estimators of the gradients, which

may reduce the number of iterations till convergence. However, a larger S also means higher computational cost. Thus, we implement the control variate to reduce the variance of the estimated gradients without increasing S .

Third, we propose to evaluate the likelihood of each trial by considering only the neurons that could plausibly respond to the stimulation. To be specific, when evaluating the likelihood for trial i , we can only consider the neurons that are within certain radius of the stimulation(s), instead of considering all K neurons, as potential sources of events. This simple but effective strategy will bring down the computational complexity for evaluating the likelihood of one trial from $\mathcal{O}(|\mathcal{A}(K + 1, n_i)|)$ to $\mathcal{O}(|\mathcal{A}(K + 1, n_i)|)$, where $|\mathcal{C}|$ represents the number of elements in a set \mathcal{C} and P is the number of neurons that may respond to stimulation in trial i . Here the number P is determined by the spatial resolution of stimulation and the density of cells in the field of excitation, which is a local parameter that is independent of K .

Lastly, we can further speed up the inference using distributed learning. The key observation is that, when learning the posteriors of neuron k , we only need to consider the neurons that live near neuron k , because they are the neurons that may overlap in their spatial profile and thus can potentially produce responses given the same stimuli. Therefore, we can divide the field into small, overlapping regions, where we analyze the trials and update the posteriors for each region in parallel. Moreover, there is little or no need for communication between each region during the experiment.

2.5 Learning the neural response function

2.5.1 Spike Time Transform

As we say in Figure 2.2I and 2.2J, given an evoked current m , all opsin expressing neurons produce similar mean spike times and spike time standard deviations. We can therefore use the optogenetic response data to learn the mean function $g(\cdot)$ and the standard deviation function $h(\cdot)$ which describe how evoked opsin currents relate to evoked spike times. We assume that the two functions take the form

$$g(m) = a + bm^{-1} \quad \text{and} \quad h(m) = \min\{\text{std.dev.}\} + cm^{-1}.$$

It is worthwhile to point out that the estimated standard deviation function h tends to underestimate the variability in a real mapping experiment, especially when the mean current is large. Therefore, we transform the standard deviation using the following function

$$\tilde{h}(x) = \min h + \frac{(\max h - \min h)}{1 + \exp(\beta_1(x - \beta_2))},$$

where β_1, β_2 are parameters, and x is the peak current.

2.5.2 Shape

In Section 2.4.1, we introduce two shape functions $s_{k,xy}(\cdot)$ and $s_{k,z}(\cdot)$ that capture the spatial distribution in the opsin on the membrane of a neuron. In this section, we discuss the

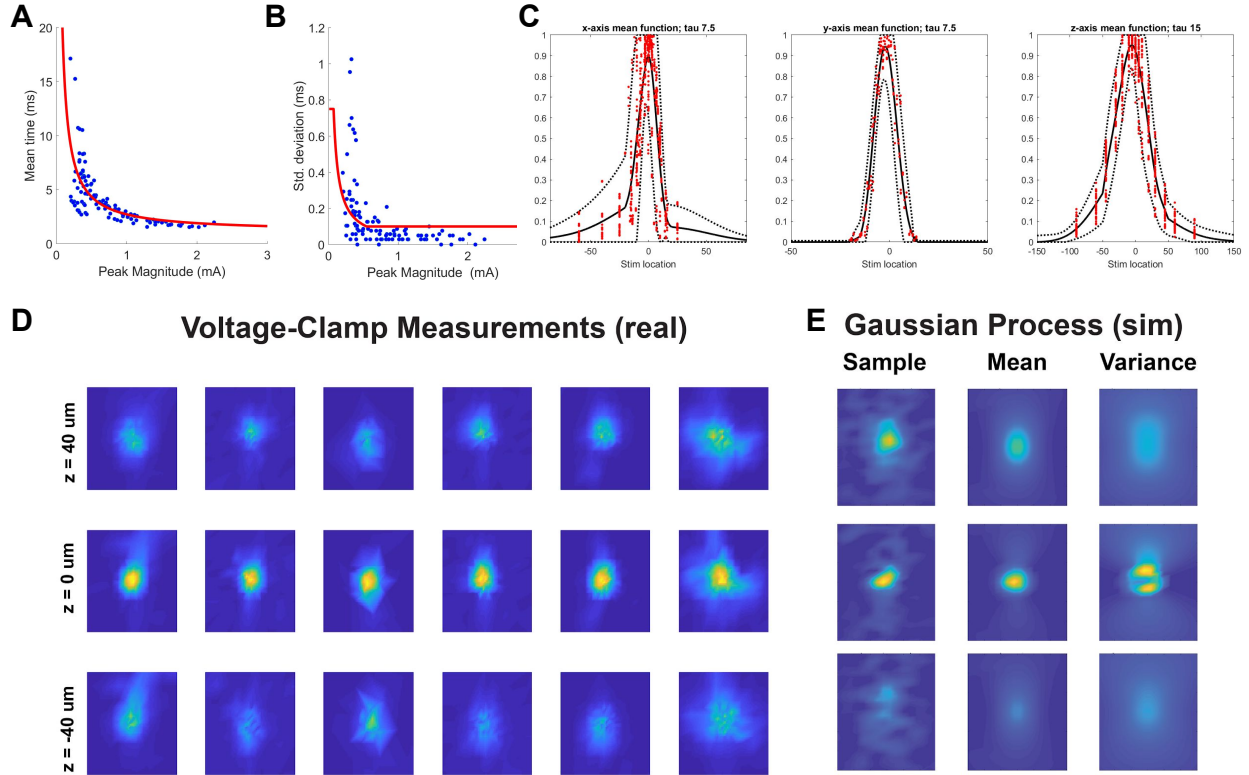


Figure 2.3: Data-driven priors. (A,B) Red lines indicate the learned functions, $g(\cdot)$ and $h(\cdot)$ and the blue dots are the data from Figures 2.2I and 2.2J. (C) Mean evoked currents at each location for several cells (red dots) and the mean (solid black line) and variance (dashed black line) fit for the Gaussian process shape prior. Right panel, along vertical axis of cortex where negative is closer to pia and positive is closer to L4. The extended shape into the negative direction is a result of the Kv2.1 tag trafficking to the apical dendrite. Middle, the horizontal axis of cortex in the radial plane. Left, the horizontal axis of cortex parallel to the optical axis. (D) Heatmaps of the normalized current response for six cells (each column). The rows correspond to different axial distances from the cell, and maps have been interpolated. (E) A sample from the fit Gaussian process prior at the same spatial locations as the data in D as well the mean and variance fit for each point.

Gaussian processes that we use as the prior distributions for these functions. We drop the subscript k which indexes a particular opsin expressing neuron in the discussion to follow.

From the response data (Figure 2.2K, 2.2L), it is clear that we can learn both functions $s_{xy}(\cdot)$ and $s_z(\cdot)$ if we record the evoked currents when a neuron responds to optogenetic stimulation targeted on a dense grid around the nucleus. Such an exhaustive approach is infeasible during a mapping experiment, but it can be conducted on several neurons to build a prior distribution for the shape functions.

Consider first the shape function $s_{xy}(\cdot)$ used to model the data in Figure 2.2K. In the prior distribution (2.8), we see that we need to learn the mean function $\mu_{xy}(\cdot)$ and the variance function $\Sigma_{xy}(\cdot, \cdot)$. Following common practice with Gaussian processes, we further assume

that the covariance function $\Sigma_{xy}(\cdot, \cdot)$ takes the form

$$\Sigma_{xy}(\mathbf{v}_{1,xy}, \mathbf{v}_{2,xy}) = \sigma_{xy}(\mathbf{v}_{1,xy})K_{xy}(\mathbf{v}_{1,xy}, \mathbf{v}_{2,xy})\sigma_{xy}(\mathbf{v}_{2,xy}), \quad (2.13)$$

where $K_{xy}(\mathbf{v}_{1,xy}, \mathbf{v}_{2,xy}) = K_{\tau_x}(v_{1,x}, v_{2,x})K_{\tau_y}(v_{1,y}, v_{2,y})$ is the kernel function and $\sigma_{xy}(\cdot)$ is the standard error function of the Gaussian process. We set the kernel $K_\tau(a, b)$ as

$$K_\tau(a, b) = \exp\{- (a - b)^2 / \tau^2\}, \quad (2.14)$$

and we choose $\tau_x = 7 \mu\text{m}$ and $\tau_y = 7.5 \mu\text{m}$.

Allowing for measurement errors, we assume that

$$m_i = s_{xy}(\mathbf{v}_{i,xy})\phi I_i + \epsilon_i, \quad (2.15)$$

where ϕ is the overall optical gain, I_i is the intensity of the stimulation, and ϵ_i is the measurement error in the i th trial. Recall that we assume the maximum of the shape function is unity for identifiability with ϕ . Therefore, we can estimate ϕ using the maximum ratio between the elicited current and the power level, i.e., $\hat{\phi} = \max_i(m_i/I_i)$. Note that the maximum may not be achieved at the nucleus $\mathbf{0}$ (see Remark 1 below). Plugging in the $\hat{\phi}$, we can rewrite (2.15) as

$$s_{xy}(\mathbf{v}_{i,xy}) = (\hat{\phi}I_i)^{-1}m_i - (\hat{\phi}I_i)^{-1}\epsilon_i. \quad (2.16)$$

We can then learn the mean and variance on the radial plane by using a smoothing method on the response data (Figure 2.2C, 2.2E). We take a similar strategy to build the prior for the shape function $s_z(\cdot)$ (Figure 2.2C). In this case, we estimate one-dimensional functions $\mu_z(\cdot)$ and $\sigma_z(\cdot)$ which represents the response to stimulation along the axial dimensions of the microscope, and we set $\tau_z = 15$. We defer the details of the smoothing to Appendix 5.1.

Inspecting the mean functions reveals that the Gaussian Process prior is able to capture important properties of the data such as the localization of the opsin to the apical dendrite (Figure 2.3C, first panel). We also note that the radial variance function (Figure 2.3E, right column) captures the variance in neuron size as well as the variance in apical dendrite expression indicated by the high variance around the edges of the soma and

Finally, it is important to notice that the shape functions should both take values in $[0, 1]$, but the range of values of the Gaussian process will clearly expand beyond the bounds. We emphasize here that the Gaussian process is probably not an ideal data generative distribution for simulating shape functions, but we merely use it for facilitating the learning of posteriors during the mapping experiments.

Remark 1. (*Potential shift in the nucleus detection and optogenetic stimulation.*) In the experiment, one recurring phenomenon is that the maximum elicited current may not be achieved when stimulating directly at the detected nucleus. This phenomenon can be the joint result of the errors in the nucleus detection and in the tuning of optogenetic stimulation even in the state-of-art experimental conditions. To accommodate the “shift” of the neuron from the detected and targeted location, we propose to include a new random variable $\mathbf{r}_k \in \mathbb{R}^3$ such that the true location $\boldsymbol{\omega}_k^* = \boldsymbol{\omega}_k + \mathbf{r}_k$. The shift variable \mathbf{r}_k is a neuron-specific parameter, and we can build the prior distribution for each entry of \mathbf{r}_k . Given the priors, we can easily add \mathbf{r}_k in the proposed framework to learn the posteriors.

Remark 2. (*Relaxing the priors.*) We introduce the prior distributions to help the inference of posterior distributions. However, the prior distribution built from a limited number of neurons may not present the distributions of the shapes of neurons in a mapping experiment. To allow for more flexibility during the mapping experiment, we can relax the prior distributions for $s_{xy}(\cdot)$ and $s_z(\cdot)$ by simply adding constants to the variance functions $\sigma_{xy}^2(\cdot)$ and $\sigma_z^2(\cdot)$ of the Gaussian processes. From Gaussianity, the addition of a constant to the variance is effectively adding a noise term to the prior distribution. The larger the constant is, the more flexibility we allow from the priors.

2.6 Validation

2.6.1 Simulation

Characterizing the resolution of inference

To better understand the resolution of our model and inference method, we simulate experiments with only two opsin expressing neurons where the two evaluating are separated by some amount along one of the three spatial axes. By evaluating the performance of inference as a function of this distance, we can characterize the resolution of our method.

Specifically, we generate two neurons that live at $\mathbf{w}_1 = (0, 0, 0)$ and $\mathbf{w}_2 = (x, y, z)$, respectively. The coordinates of the second neuron will be specified later. All parameters of the two neurons, except for the probability of release, are randomly generated from the prior distributions built from pilot studies. The probability of release for Neuron 1 at the origin is set to be from a uniform distribution, where $\gamma_1 \sim \text{Uniform}(0.6, 1)$. For Neuron 2, we draw the probability of release from a spiked uniform distribution where $\text{prob}(\gamma_2 > 0) = 0.5$ and $\gamma_2 \sim \text{Uniform}(0.6, 1)$ given γ_2 is non-zero. We consider three scenarios for \mathbf{w}_2 , the location of Neuron 2, where each one of the three coordinates is non-zero in each scenario.

In this simple two-neuron system, we consider an experimental design where stimulation is targeted directly at the nucleus of each neuron. Specifically, in each trial, the selected nucleus is stimulated at a power level from $\{30, 40, 50, 60\}$ mW. Each power level is repeated in 3 trials, which results in 12 trials at each nucleus in total. We compare our proposed inference with a naive benchmark approach. The benchmark approach calculates the probability of release as the proportion of trials with postsynaptic events when targeted at this neuron. We evaluate the performance of each method using the misclassification rates of postsynaptic events.

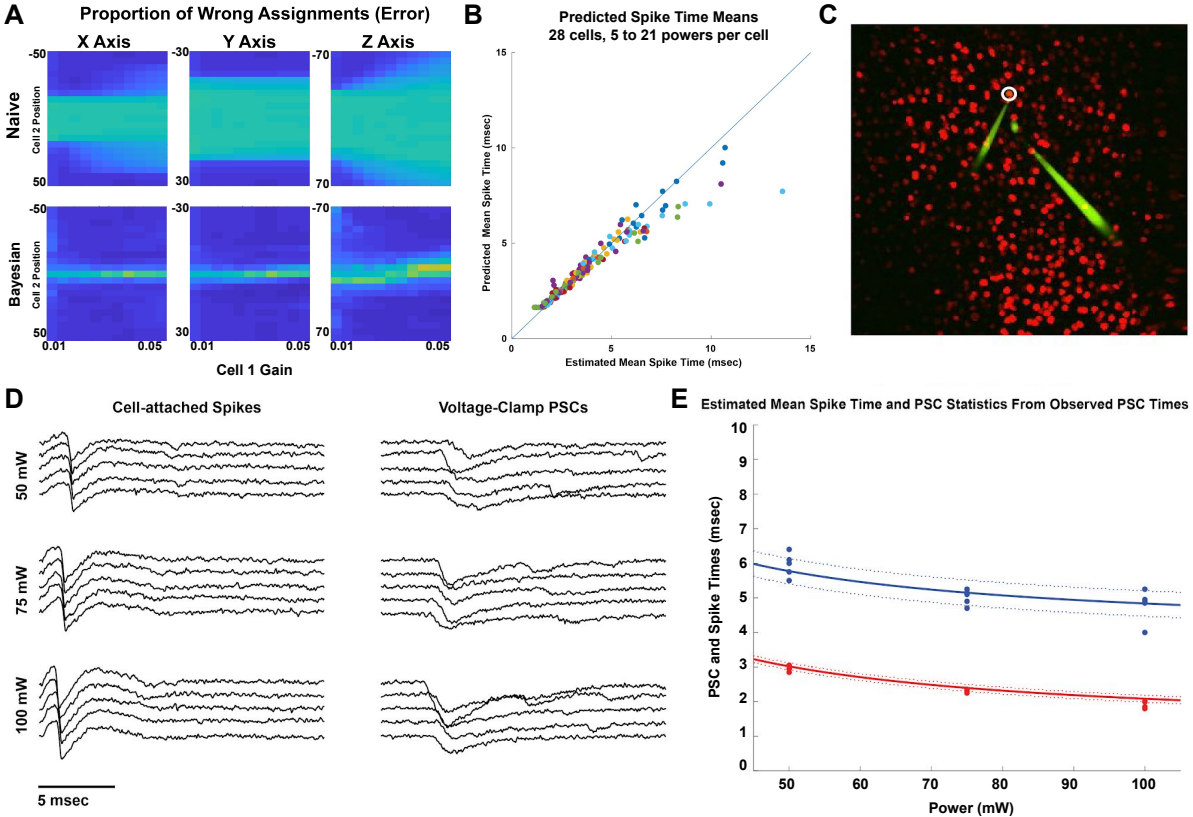


Figure 2.4: (A) The misclassification rate as a function of the distance between two presynaptic neurons for our proposed method and a naive benchmark approach for a set of simulated experiments. (B) Predicted spike times against true spike times. Each dot represents a particular cell at a particular power. (C) Simultaneous loose-patch recording of an opsin expressing cell with a voltage-clamp recording of non-expressing cell. (D) Simultaneously recorded raw traces from each cell during different stimulation intensities. (E) The spike times, red, and PCS times, blue. Dots are the data and the solid/dashed lines represent the mean/standard deviation of the posterior distribution.

Results of the simulation are shown in Figure 2.4A. The Gaussian process inference achieves smaller misclassification rates than the benchmark procedure. Misclassification rates for both approaches decrease sharply along the x- and y- axes as expected, since response data indicated that our optogenetic stimulation has relatively high spatial resolution in the radial plane.

2.6.2 Ground Truth Data

We next compare inference in our model with ground truth data. First we test the first part of the model which predicts spike times from TPHO stimulation given a static location. In this experiment we fit the gain of each cell given the spike times for some stimulus conditions and then tested the fit by comparing the expected spike times with data for held-out conditions.

We see that our model does an excellent job of predicting spike times up until around 7 milliseconds or so after stimulation which is sufficient for time windows we consider.

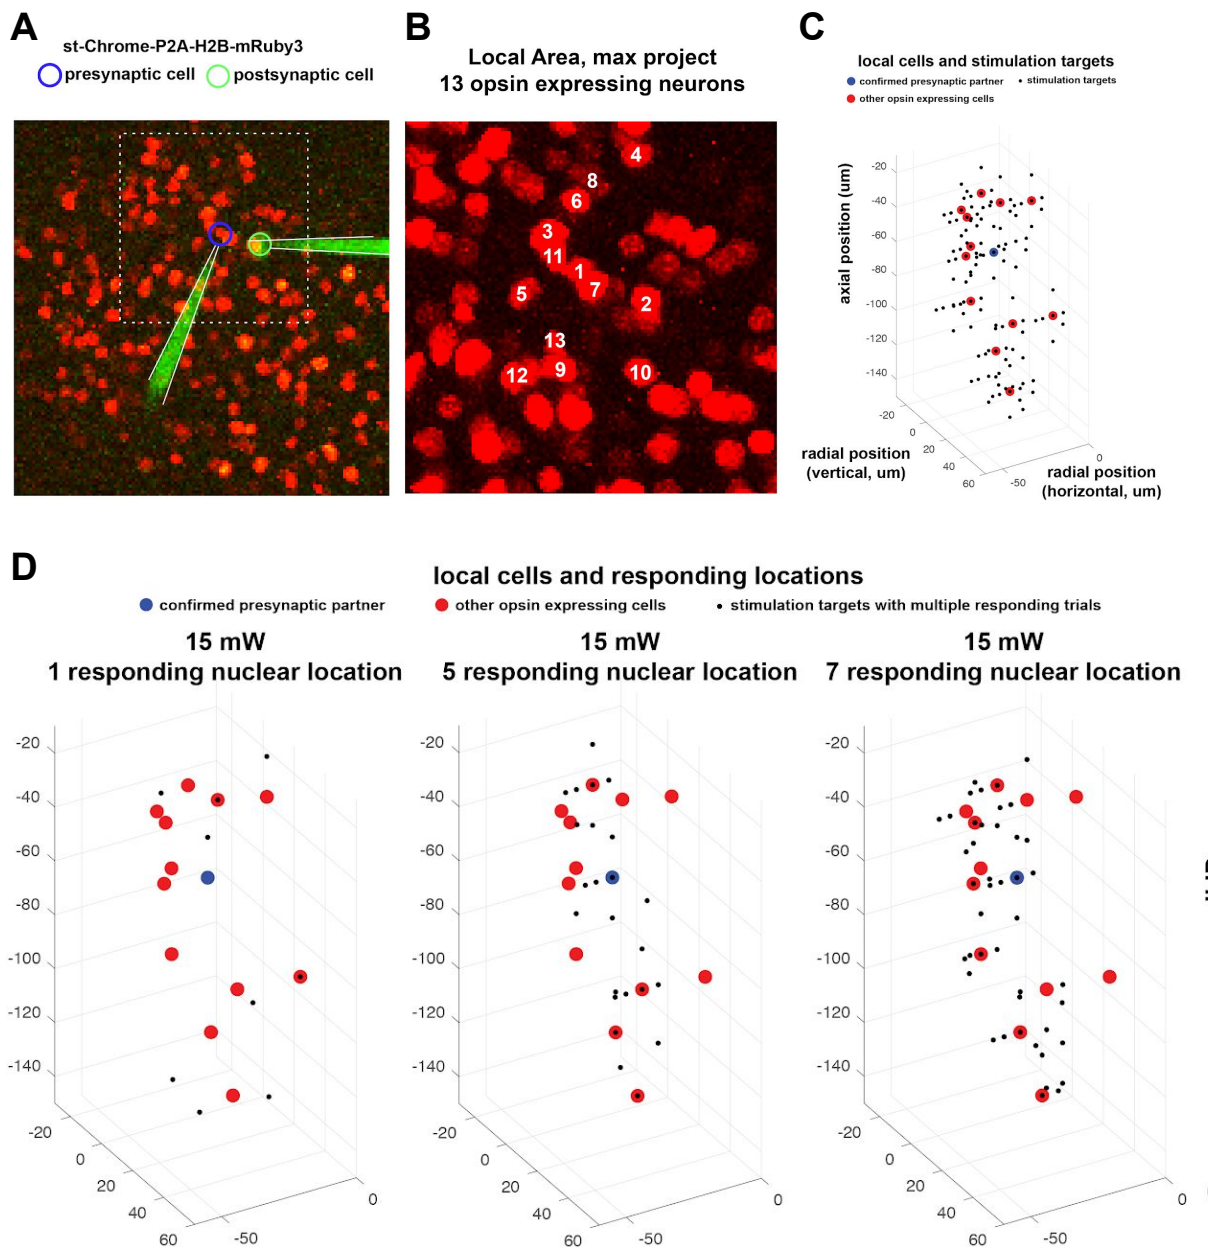


Figure 2.5: (A) Simultaneously recorded opsin-expressing cell in loose patch (blue circle) and a non-expressing cell in voltage-clamp (green circle). (B) Higher resolution image of the local area around the loose-patched cell (cell 1). (C) Scatter plot of the local opsin expressing cells (red, blue for loose-patched cell) and all stimulation locations (block). (D) The opsin expressing cells as in C but here each volume corresponds to a stimulation power, and the black dots show targets where PSCs were detected on multiple trials for that power.

Next we demonstrate that we are able to transform evoked spike time distributions into PCS time distributions. In this experiment we observe the PSC times in a post synaptic cell and then infer the latent spike time distributions. By loose-patching the presynaptic cell as well, we are able to observe the true spike times. We see that inferred spike time distribution matches the ground truth data quite well (Figure 2.4E).

Finally, we test our full inference procedure. As in the last experiment, we loose-patch a potential presynaptic cell along with the postsynaptic cell in voltage-clamp (Figure 2.5A). However in this case we target at and around a cluster of 13 neurons near the potential presynaptic cell (Figure 2.5B,C). As expected, we see the number of stimuli that evoked reliable PSCs in the postsynaptic cell increases with power, and the structure of those locations does not make it make obvious which cells are connected (Figure 2.5D).

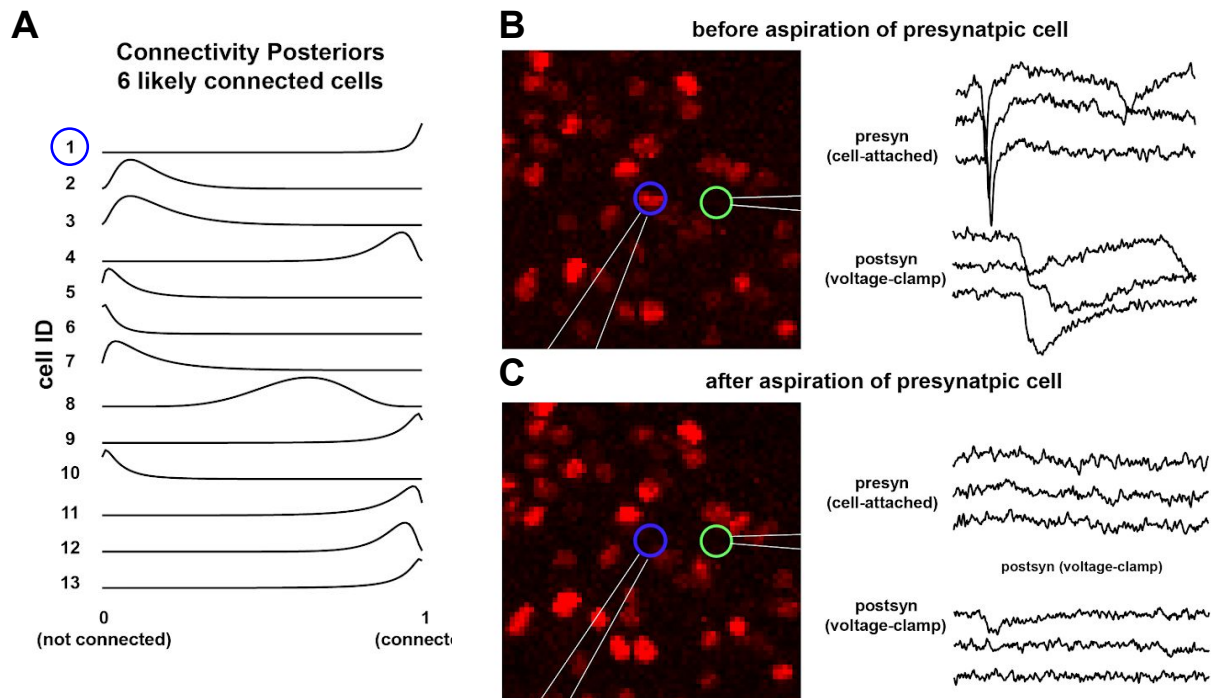


Figure 2.6: (A) Posterior distributions for the probability of transmission, γ_k , for each of the 13 neurons. (B) An image of the field before aspiration showing the nucleus of neuron 1, and raw traces of light evoked spikes in neuron one with the simultaneously recorded postsynaptic traces. (C) As in B but post aspiration. Note that now the nucleus of cell 1, the evoked spikes, and the PSCs have all been eliminated.

When we run inference on these data, we see that the algorithm predicts that the putative presynaptic cell is in fact connected (neuron 1 in Figure 2.6A). To verify this prediction we first stimulate the presynaptic cell with light to directly observe both the evoked spikes in that cell and reliable PSCs in the postsynaptic cell (Figure 2.6B). We then aspirate the presynaptic neuron through the loose-patch pipette and confirm that the nucleus is no longer

visible with fluorescent imaging (Figure 2.6C). We then target that location one last time to observe that both the spikes and the PSCs have been eliminated. Although we didn't have any other stimulated cells patched in this experiment, we did observe that targets on cell 7 (the cell closest to the confirmed presynaptic cell) at the highest power resulted in reliable responses in the postsynaptic cell. Despite this, inference did not predict that cell 7 was connected (Figure 2.6A) and we were able to verify that those PSCs in fact correlated with off-target spikes evoked in the patched presynaptic neuron.

2.7 Discussion

In this chapter, we have presented a statistical method for increasing the native resolution of two-photon holographic optogenetics when mapping synapses in dense populations of opsin expressing neurons. Our statistical model and inference method is able to analyze single trials and individual cells together to determine which opsin expressing neurons are responsible for evoked postsynaptic events. The key to our method is a thorough analysis and characterization of the neural response to our TPHO system. However, we note that our method is highly tuned to the properties of our particular system which are a function of the holographic microscope, the opsin and its expression system, and the stimulation protocol. To map a new cell types or use a new microscope or opsin, one would have to collect the response data and re-learn the priors and potentially even some of the functional forms. We also note that our validation data is limited to a single example for fitting the synaptic delay and, more importantly, for inferring synaptic transmission. It would be worthwhile to validate this method further before applying it at scale.

Another future direction for this work is to incorporate online experimental design to minimize the number of trials needed to infer synaptic connections [155]. One advantage of a Bayesian model is that the posterior uncertainty can be quantified and utilized during an experiment to guide the choice of stimuli. If synaptic connections can be detected quickly, then it could be possible, especially with faster opsins, to learn more complex models for each synapse such as the short-term dynamics.

2.8 Method Details

All experiments were performed in accordance with the guidelines and regulations of the Animal Care and Use Committee of the University of California, Berkeley.

2.8.1 Transgenic mice

The following mouse lines were used for this study: the *emx1-IRES-Cre* line (JAX stock 005628). Mice were housed in cohorts of five or fewer with a light:dark cycle of 12:12 hours, and were used for experimentation during their subjective night.

2.8.2 Viral Infection

Neonatal mice (p0-3) were deeply cryo-anesthetized and placed in a head mold. Viral aliquots were loaded into a Drummond Nanoject injector and injected into 4 sites in the barrel cortex of the left hemisphere. At each site, virus was injected at multiple depths in increments of 18.4 nL or 36.8 nL (for multiphoton experiments), for a total of 150-440 nL of virus injected per mouse. Following injections, mice were moved to an incubation chamber for recovery, and were returned to the dam once they regained color and began to move. Viruses used were pAAV-CAG-DIO-ChroME-ST-P2A-H2B-mRuby3.

2.8.3 Brain slice recording

Acute thalamocortical slices were prepared from mice (ages p14-21) as previously described. Slices were placed in a recording chamber and constantly perfused with oxygenated artificial cerebro-spinal fluid (NaCl 119 mM, KCl 2.5 mM, MgSO₄ 1.3 mM, NaH₂PO₄ 1.3 mM, glucose 20 mM, NaHCO₃ 26 mM, CaCl₂ 2.5 mM) maintained at 32° C (21° C for multiphoton mapping experiments). Slices were oriented with the caudal surface facing up in the recording chamber. To ensure minimal disruption of vertical connectivity, all slices used for recording were inspected under infrared illumination at 40x magnification and/or post-hoc confocal imaging to confirm that pyramidal cell apical dendrites stayed roughly parallel with the surface of the slice or receded slightly deeper as they progressed apically. Whole cell recordings were performed using glass micropipettes (2-5M Ω resistance) pulled on a Sutter P-1000 Micropipette Puller. Pipettes were filled with a Cs⁺ based internal (CsMeSO₄ 135 mM, NaCl 8 mM, HEPES 10 mM, Na₃GTP 0.3 mM, MgATP 4 mM, EGTA 0.3 mM, QX-314-Cl 5 mM, TEA-Cl 5mM) or a potassium gluconate based internal (K-gluconate 135 mM, NaCl 8 mM, HEPES 10 mM, Na₃GTP 0.3 mM, MgATP 4 mM, EGTA 0.3 mM). Data were analyzed from recordings in which series resistance remained stable and below 30M Ω . Data were acquired and filtered at 2.2 kHz using a Multiclamp 700B Amplifier (Axon Instruments) and digitized at 20 kHz (National Instruments). All data were acquired using custom written MATLAB (Mathworks) software.

Space clamp error will inevitably affect somatic measurements of currents from distally located SST- \bar{i} PC synapses; however, we recorded IPSCs using a cesium-based internal solution (which included the ion channel blockers tetraethylammonium and QX-314) and performed experiments at room temperature, which ameliorate this to some extent 115. We also used a holding potential of +40 mV to increase the IPSC driving force. In these experiments, internal solutions also contained 5 M Alexa 488 hydrazide (ThermoFisher Scientific) to aid visualization with multiphoton imaging, and 5 mM kynurenic acid Sodium salt (abcam) was added to the external ACSF to block glutamatergic activity.

2.8.4 Two-photon Holographic Excitation

Laser light was generated using a 5W 1040 nm femtoTrain laser (Septra-Physics) and power was modulated on short time scales using a Pockels cell (Conoptics) and a high speed shutter (UniBlitz). Light was delivered to the sample using a VIVO 2- Photon workstation (3i) based on a Sutter Moveable Objective Microscope (Sutter) and the hologram was created using a

Phasor 2-Photon computer-generated holography system (3i) with Slidebook software (3i). The holograms used for stimulation were 2D discs of diameter 15 μm centered at arbitrary points at 1 μm resolution within a 300 μm x 300 μm grid in the focal plane (Fig. 3.4b,c). Stimulation consisted of e ms square pulses to the Pockels cell with voltages calibrated to produce a range of average power on sample. Power for each hologram was calibrated empirically to account for power loss due to diffraction efficiency degradation away from the zero-order of the SLM.

2.8.5 Automated workflow

See Figure 2.1E for reference.

The Master CPU initializes the experiment, and then a patch is acquired - either voltage-clamp or loose patch depending on the experiment. Once the stimulation field is selected by moving the objective, a single click begins the experiment, orchestrated by the Master CPU.

First, the SlideBook CPU takes a 200 μm stack at 1 μm spacing was taken with 1000 nm light of the nuclear mRuby3 signal. The 3D volume is then sent to the analysis computer for cell detection [168]. While this happens, the Master CPU runs an electrophysiological protocol on the patched cell if it is in voltage-clamp. Once the cell locations are detected, the user can pause to look at and adjust the results as well as identify the patched cell in the image.

After locations are confirmed, the the experimental protocol (e.g. measure spike times, measure shape, map connections) generates one batch of locations and powers for stimulation and send them to the analysis computer to compute phase masks. When the first batch is done, they phase masks are loaded onto the SlideBook CPU as well as information about trial timing and hardware coordination. The first batch begins, and immediately the design of the second batch is sent to the analysis computer to prepare. This cycle allows for almost no down time in acquisition until the end of the experiment.

2.8.6 Experimental Protocols

Spiking, Spike Times, Current, and Shape

In these experiments, spikes from a cell in loose-patch are recorded as a function of stimulus intensity. The location for all of these experiments is the location returned from cell detection for the patched cell. The powers used varied from 5 mW to 70 mW and was chosen based on the expression level of the patched cell. For some experiments, an adaptive procedure was used to focus trials around the rheobase of the cell and also cover desired ranges.

For current measurements, everything is the same but the cell is in voltage-clamp. For many cells we recorded in both modes. In these cases we first obtained a gigaseal in attached mode, recorded spikes, broke in, and then recorded currents in voltage-clamp.

In the shape learning experiments, the experiment was similar to the previous current measurements but in this case the power was held constant to 25 mW and the location of stimulation was either rastered at 5 μm spacing on a 70 x 70 μm grid in the radial plane, or at non-uniform spacing along the axial axis of the microscope.

2.8.7 Synaptic Mapping

In these experiments, we have two patched cells: one in loose-patch and one in voltage-clamp. We center the excitation field on the loose-patch cell. Locations are chosen based on the detected cells' locations as well as randomly chosen locations near, but not directly on these cells.

Chapter 3

Mapping Complementary Networks of Cortical Somatostatin Interneurons

This chapter contains a work from a mapping collaboration which itself was part of a larger project. It is presented here as an excerpt from the larger work. References to supplementary figures (e.g. S1 or S4) can be found with the full publication:

Naka, A., Veit, J., Shababo, B., Chance, R., Risso, D., Stafford, D., Snyder, B., Egladyous, A., Chu, D., Sridharan, S., Mossing, D., Paninski, L., Ngai, J., and Adesnik, H. (2019). Complementary subnetworks of cortical somatostatin interneurons enforce layer specific control. *Elife*, e43696.

3.1 Mapping a cortical circuit with multiple modalities

Excitatory and inhibitory synaptic connectivity are fundamental to computation in neural circuits. In sensory regions, such as the primary somatosensory cortex (S1), excitatory circuitry follows a modular architecture centered on the 6-layered structure of neocortex. Excitatory principal cells (PCs) in each layer differentially encode and process sensory information, and their inter-connectivity helps define the basic input/output logic of the S1 microcircuit [43, 54]. ‘Bottom-up’ sensory inputs enter S1 via thalamocortical projections that most densely innervate L4, and cortical output emerges primarily from PCs in L2/3 and L5, feeding forward to other brain regions. In between, sensory-driven activity propagates through the different layers via highly specific translaminal pathways. Together, these excitatory pathways comprise a stereotyped microcircuit, which is thought underlie hierarchical transformation and processing of sensory information [12, 104].

In contrast, no comparable framework exists for inhibitory circuit organization. While it is well established that the connectivity of many inhibitory circuits is selective with respect to cell-type (especially interneuron-to-interneuron connectivity; [76, 136]), and that many interneurons target specific subcellular compartments of other neurons [101, 169], it remains unclear how inhibitory circuits integrate into the layer-based logic of the excitatory microcircuit. One possibility is that interneurons non-selectively interconnect with PCs [21,

151, 59, 46, 122], irrespective of layers; this would suggest that cortical inhibition primarily performs ‘global’ operations on excitatory neurons, such as non-specific normalization [141, 142, 30]. Alternatively, interneurons might form synaptic connections with specific layers within the microcircuit, which would suggest that different functional components of cortical circuits can be independently modulated by specific inhibitory pathways, enabling more sophisticated computations.

To resolve this, it is necessary to assess connectivity in inhibitory circuits with respect to both layer and cell-type, since different classes of interneurons likely obey different principles for translaminar wiring. For example, many interneurons, including parvalbumin (PV) cells, the most numerous class of cortical interneuron, are largely local, intralaminar inhibitors ([24, 82, 91, 122]; but see [22, 25]). However, some types of interneurons - most notably, dendrite-targeting somatostatin (SST) interneurons- appear specialized to mediate translaminar inhibition [74, 103, 66, 57, 173, 179, 68]. SST cells are critical players in many aspects of neural computation, perception, and cognition [3, 85, 77, 98, 87, 170]. Since they integrate excitatory inputs from both intralaminar and translaminar sources and since many SST cells, project extensive translaminar axons they are likely key players in translaminar inhibitory circuitry [9, 181, 161, 70, 121, 185, 72]. How do SST inhibitory networks integrate into the different layers of the cortical microcircuit? The intralaminar outputs of SST cells onto nearby PCs are thought to be extremely dense [46, 66], which has led to the hypothesis that SST cells might generate a ‘blanket of inhibition’ that overlays nearby excitatory circuits [71, 45]. One possibility is that the translaminar connectivity of SST cells is similarly non-selective. However, more sophisticated principles might govern the organization of laminar SST circuits. SST cells are highly heterogeneous [170, 184] and can be subdivided into electrophysiologically and genetically distinct subtypes [165]. Emerging evidence indicates that SST subtypes exhibit distinct sensory tuning and coupling to behavior [85, 77, 113, 96, 144, 115], indicating that different SST subtypes might serve different computational functions. Intriguingly, SST subtypes target their axons to different laminar domains [113, 97, 112]. While most SST cells are Martinotti cells (MCs) with axons that ascend to L1, there is also a substantial population of non-Martinotti cells with axons that primarily target L4 instead of L1. It has been hypothesized that these two distinct SST subtypes could uniquely influence cortical processing by targeting specific components of the cortical microcircuit. However, we lack basic knowledge of the circuit organization of different SST subtypes, which precludes mechanistic understanding of their computational contributions to cortical function. What are the patterns of synaptic input to different types of SST cells, and how might these give rise to differential activation of SST subtypes during sensation and behavior? Do all SST cells generate non-selective inhibitory outputs, or can SST cells provide targeted inhibition to specific components of the cortical microcircuit, such as certain cortical layers?

To address these issues, we used a combination of high-resolution one and two photon optogenetic mapping, paired intracellular recordings, and anatomical reconstructions to probe the logic of SST inhibitory circuits across the six cortical layers. We focused our investigation on L5, since this layer hosts a more diverse population of SST cells than the upper layers, which includes large fractions of both Martinotti and non-Martinotti cells [101, 148]. We found that two subtypes of SST cells exhibit exquisitely specific and strikingly complementary laminar patterns of connectivity. L5 Martinotti cells receive input from L2/3 and L5, whereas L5 non-Martinotti cells receive input from L4 and L6. In turn, L5 Martinotti cells

provide reciprocal inhibition to L5 PCs but not L4 PCs, while L5 non-Martinotti cells inhibit L4 PCs but not L5 PCs. Furthermore, in vivo optogenetic manipulation of Martinotti and non-Martinotti cells modulated cortical activity in a layer specific manner. These results demonstrate that subtypes of SST cells are wired into separate layer and cell-type specific subnetworks which independently shape different aspects of cortical computation.

3.2 Distinct subtypes of L5 SST cells receive complementary patterns of excitatory synaptic input

To probe the structure of L5 SST circuits, we first verified approaches to identify and target different subtypes of SST cells in the barrel cortex. The widely used SST-Cre line allows labeling of SST cells irrespective of subtype (by crossing to the LSL-tdTomato mouse line; hereafter referred to as ‘SST-TdT’). Additionally, there exist transgenic GFP reporter lines which have been characterized as labeling smaller subsets of the SST population. We chose two of these lines, the ‘X94’ line and the ‘GIN’ line in order to respectively target non-Martinotti and Martinotti cells for whole cell recording in acute barrel cortex slices [97, 118, 164]. GFP+ neurons were clearly visible within L5 in slices from both X94 and GIN animals, consistent with reports that L5 contains both Martinotti and non-Martinotti SST cells. Biocytin fills confirmed that the vast majority of L5 GIN cells were MCs (41/44; 93%) which exhibited an ascending translaminal axon that ramified in L1, L2/3, and L5a (Fig. 3.1a) and intrinsic properties classically associated with MCs (Fig. S1; [41, 75, 79, 173]). Similarly, the vast majority of L5 X94 cells were non-Martinotti cells (32/35; 91%) which formed a dense axonal plexus in L4 rather than L1 (Fig. 3.1a, Fig. S2) and exhibited quasi-fast-spiking electrophysiological phenotypes (Fig. 3.1b, Fig. S1). While the X94 line labels only 15% of SST cells in L5 [51], additional recordings from L5 SST-TdT cells suggested that a larger fraction (30 - 40%; Fig. S1c-g) of L5 SST cells are non-Martinotti, X94-like cells (hereafter termed ‘XCs’) with most of the remainder being MCs. Thus L5 contains at least two intermingled but qualitatively distinct SST cell subtypes, MCs and XCs, which can reliably be identified using the GIN and X94 lines, respectively [97].

Aligning reconstructions of L5 MCs and XCs revealed that these two populations had strikingly complementary vertical profiles of neurite density: MCs primarily innervated layers 1, 2/3 and 5, and XCs primarily innervated L4 and the L5/L6 border (Fig. 3.1c). This complementary pattern bears a strong similarity to the axonal projection profile of the two major thalamic input pathways to the barrel cortex: the profile of XCs aligns with the lemniscal pathway projection from the ventro-posterior medial nucleus (VPM), whereas the profile of MCs aligns with the paralemniscal pathway projection from the posterior medial nucleus (POm). These pathways, which carry different aspects of whisker related information, continue intracortically via specific, translaminal pathways; thus this distinction represents an important organizing principle in the excitatory connectivity and function of barrel cortex circuits.

We next asked whether MCs and XCs might differentially wire in to the laminar networks defined by these pathways. To do this, we first sought to determine whether the excitatory

inputs to MCs and XCs might arise from different laminar sources. We virally expressed channelrhodopsin-2 (ChR2) in all excitatory neurons and photo-stimulated PCs in different regions of the slice to map the excitatory inputs to L5 MCs and XCs (Fig. 3.1d,e,f, AAV-flexed-ChR2 in the emx1-Cre line; see methods and Fig. S4). Remarkably, we found that L5MCs and XCs receive inputs from highly specific but largely non-overlapping sources that aligned with the anatomical distributions of their neurites. L5 MCs, but not L5 XCs, frequently received excitatory inputs from either upper L5, L2/3, or from both L2/3 and L5, but received little input from L4 or L6. On average, MCs received stronger inputs from L2/3 than L5 ($49 \pm 9\%$ of total input from L2/3 versus $26 \pm 7\%$ from L5 in $n = 15$ MCs; $p = 0.008$, paired t-test). These results are consistent with previous studies [66, 70, 5, 138]. In contrast, L5 XCs received strong input from L4 and/or the L5B/L6 border (Fig. 3.1e,f; input from L4 and L6 was $62 \pm 7\%$ of total input for $n = 14$ XCs versus $25 \pm 3\%$ for $n = 15$ MCs; $p = 6.5 \cdot 10^{-10}$; two sample T-test; see also Fig. S3) but relatively little input from L2/3 and L5. This input architecture indicates that these two subtypes of L5 SST cells receive almost completely distinct patterns of excitatory innervation from their local circuit. Comparing the average laminar profiles of mapped excitatory input and anatomical reconstructions revealed a striking symmetry in the overall structure of XC and MC intracortical circuits: XCs received input from L4 and 6, and projected most densely to these two layers; MCs received input from L2/3 and 5, and projected preferentially to these same two layers (Fig. 3.1c,f).

The striking laminar differences in inputs to L5 XCs and MCs suggested that these two types of SOM cells should be differentially recruited by activity of different cortical layers. For instance, driving L4 should specifically recruit XCs, but not MCs. To test this possibility, we specifically photo-stimulated L4 excitatory neurons via Cre-dependent expression of ChR2 in *scnn1-Cre* mice (crossed to GIN or X94; Fig. 3.2a). L4-specific photostimulation (with two different stimulus protocols, across four different intensities; see methods) drove large EPSCs in XCs but evoked little to no input in MCs under identical conditions (Fig. 3.2b,c; Fig. S5). Current clamp recordings under the same conditions showed that L4 photo-stimulation reliably drove spiking in L5 XCs, but not in L5 MCs (Fig. 3.2b,d; Fig. S5), despite the fact that MCs are intrinsically more excitable than XCs (Fig. S1). The lack of evoked responses in MCs was not due to differences in the degree of L4 activation (see supplemental text and Fig. S5 for controls). Thus, these results indicate a stark difference between L5 XCs and MCs: L4 densely innervates and powerfully drives firing in L5 XCs, but not L5 MCs.

3.3 Common input mapping reveals subnetwork structure in L5 SST cell output

Having found that XCs and MCs in L5 receive highly complementary patterns of excitatory synaptic input, we next asked whether XCs and MCs also exhibit layer-specificity in their inhibitory outputs. Since SST cells have been implicated in generating feedback inhibition [3, 161, 70], we hypothesized that MCs and XCs might target their inhibitory outputs in order to reciprocally inhibit the same PC populations that excite them. For example, XCs but not MCs would inhibit L4 PCs, whereas MCs but not XCs inhibit L5 PCs. Alternatively, MCs,

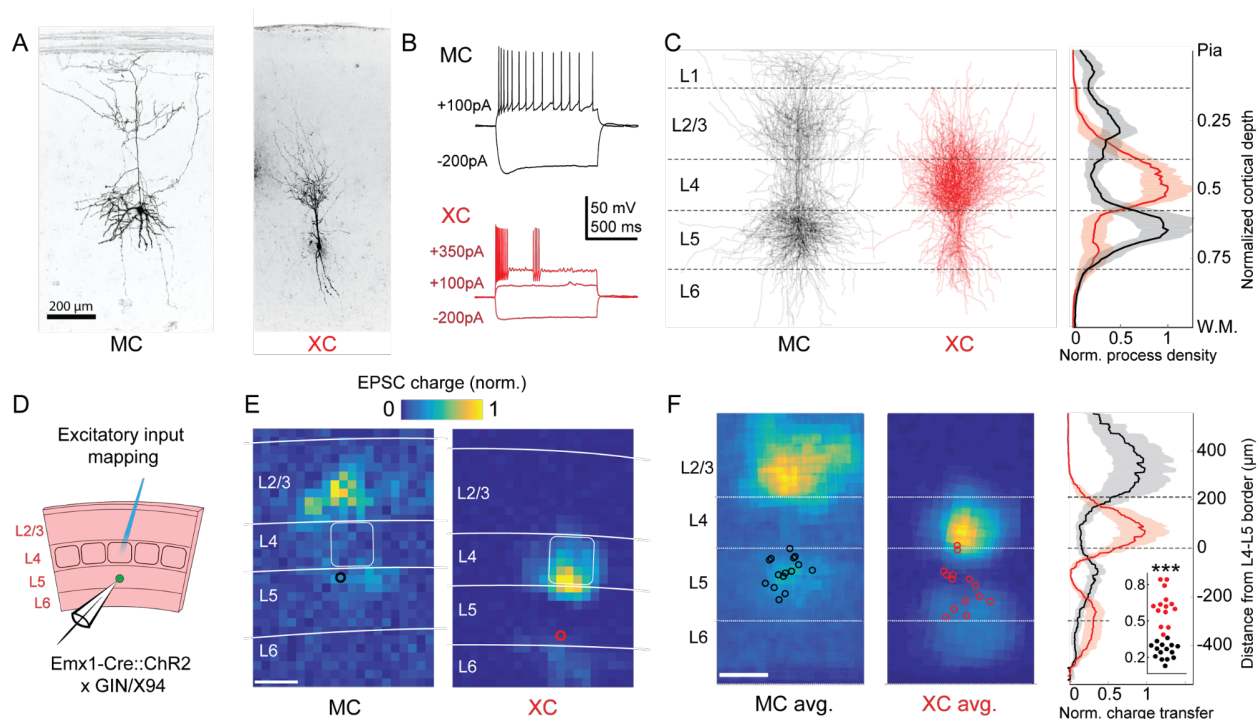


Figure 3.1: Optogenetic circuit mapping reveals complementary synaptic input patterns to two subtypes of L5 SOM cells. (A) Confocal images of dye filled neurons revealing two morphological phenotypes of L5 SST cells. Left: an L5 GIN cell. Right: an L5 X94 cell. Scale bar: 200 μm . (B) Example traces during current step injections from an L5 GIN cell (black) and an L5 X94 cell (red). (C) Left: Overlaid morphological reconstructions of L5 GIN/MC cells (black, $n=14$) and L5 X94/XC cells (red, $n=10$) showing differences in laminar distribution of neurites. Right: Normalized neurite density versus cortical depth for L5 GIN (black) and L5 X94 cells (red). Data are represented as mean \pm C.I. (D) Schematic of experimental configuration. A digital micromirror device was used to focally photo-stimulate excitatory cells in different regions of the slice in order to map the spatial profile of excitatory inputs to GFP+ L5 MCs (Emx1-Cre; GIN) or GFP+ L5 XCs (Emx1-Cre; X94). (E) Example heat maps of median EPSC charge transfer evoked at each stimulus site for example L5 SST cells. Left: An L5 MC that received inputs from L5 and L2/3. Right: An L5 XC that received inputs from L4 and the L5/6 border. Soma locations are indicated by red/black bordered white dot). Scale bar: 200 μm . (F) Left: Grand averages of input maps reveal cell-type specific patterns of laminar input. Soma locations are indicated as above. Right: Normalized charge transfer versus distance from L4-L5 border for MC (black) and XC (red) populations. Scale bar: 200 μm . Inset: Swarm plots showing the proportion of total evoked charge transfer in each map that originated from sites in L4+L6, i.e. $[L4+L6] / [L2/3+L4+L5+L6]$ for the MC (black; median, 27%; range, 13-36%) and XC (red; median, 62%; range, 38-84%) populations. Proportions were significantly different between L5 MCs and L5 XCs ($25 \pm 3\%$ in $n = 15$ MCs versus $62 \pm 7\%$ in $n = 14$ XCs, mean \pm C.I.; $p = 6.5 \cdot 10^{-10}$; two-sample t-test). See also Figures S1-4.

XCs, or both cell types could globally target PCs within and across layers non-selectively. These hypotheses are difficult to test using conventional optogenetic approaches since, to our knowledge, Cre-driver lines do not exist for all SST subtypes. An alternative approach is to use high precision photo-stimulation to ask whether the outputs of individual SST cells (in the non-specific SST-Cre line) diverge onto PCs in multiple layers. If individual SST cells target either L4 or L5 PCs, but not both, then we should never observe common input to

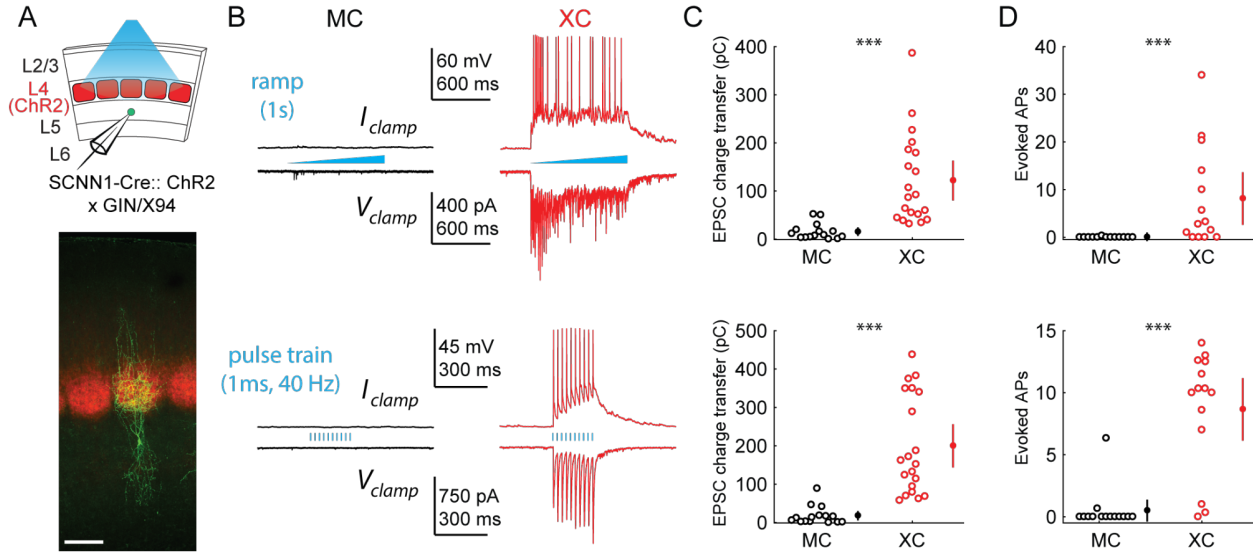


Figure 3.2: L4 photo-stimulation excites L5 XCs but not L5 MCs (A) Top: Schematic of the experimental configuration. L5 X94 or GIN cells were recorded during photo-stimulation of L4 excitatory neurons. Bottom: Confocal image of a filled L5 X94 neuron (green) with ChR2-TdTomato expression (red) visible in L4. Scale bar: 150 μm . (B) Top row: Example traces recorded in the current clamp (upper traces) or voltage clamp (lower traces) configurations during a 1-second ramp photo-stimulation. Bottom row: As above, but for photo-stimulation with a 40 Hz train of ten 1ms pulses. (C) Quantification of excitatory charge transfer during maximum intensity 1-second ramp stimulation trials. Mean 122 ± 41 pC in $n = 20$ XCs versus 15 ± 8 pC in $n = 15$ MCs; $p = 3.9 \cdot 10^{-6}$, Wilcoxon rank sum test. (D) Quantification of the mean number of evoked action potentials during maximum intensity 1-second ramp stimulation trials. Mean 8.1 ± 5.5 spikes per trial in $n = 15$ XCs versus 0.03 ± 0.05 spikes per trial in $n = 15$ MCs; $p = 6.6 \cdot 10^{-4}$, Wilcoxon rank sum test. (E) As in C, for maximum intensity 40Hz pulse train stimulation. Mean 200 ± 56 pA in $n = 20$ XCs versus 18 ± 12 pA in $n = 15$ MCs; $p = 2.8 \cdot 10^{-6}$, Wilcoxon rank sum test. (F) As in D, for maximum intensity 40Hz pulse train stimulation. Mean 8.7 ± 2.4 spikes per trial in $n = 15$ XCs versus 0.5 ± 0.9 spikes per trial in $n = 15$ MCs; $p = 1.5 \cdot 10^{-6}$, Wilcoxon rank sum test. Error bars denote mean \pm 95% confidence interval. Three asterisks denotes $p < 0.001$. See also Figure S5.

pairs of L4 and L5 PCs when photo-stimulating single SST neurons. This can be tested by mapping SST inhibitory connections onto multiple PCs simultaneously and analyzing the spatiotemporal coincidence of evoked IPSCs onto different pairs of PCs, thereby measuring the amount of common input shared between pairs of PCs in different layers [185, 111]. Although this approach does not discriminate between MCs and XCs directly, it performs a more stringent test by extending our hypothesis to apply to the structure of the outputs of the L5 SST population as a whole.

In order to conduct comprehensive mapping of individual SST outputs, we developed a novel approach to map neural circuits at high spatiotemporal resolution using two photon optogenetics and a statistical pipeline for detecting synaptic connections [108] (see Chapter 4) and evoked inhibitory postsynaptic current (IPSC) synchrony. In addition, we employed a soma-targeted opsin [178], which has the advantage of providing far superior spatial resolution during photo-stimulation than non-targeted opsins (Fig. 3.3a,b,c; [13]). We expressed a soma-targeted variant of the red-shifted opsin, ChrimsonR, in SST cells using the SST-

Cre line, and photo-stimulated SST cells using 2-photon excitation with computer generated holography (Fig. 3.3a,b; Fig. 3.4a). Since SST \rightarrow PC synapses are often located on the distal dendrites of PCs, we recorded IPSCs using a cesium-based internal solution and took additional steps to minimize false negatives (see Methods). Using this method, we simultaneously mapped SST inputs to pairs of L4-L5 PCs and L5-L5 PCs (Fig. 3.3d,e,f). We found that L5 PCs received SST inputs at more locations than L4 PCs on average (45.3 ± 5.5 inputs out of 441 locations tested per map in $n = 30$ L5 PC maps versus 12.8 ± 2.3 inputs in $n = 10$ L4 PC maps, Fig. 3.5e), but that both cell types received inputs from locations throughout L5 (Fig. 3.5f).

If individual L5 SOM cells target either L4 PCs or L5 PCs but not both, we would expect maps from L4-L5 PC pairs to contain spatially intermixed but non-overlapping input locations. Conversely, when recording a pair of L5 PCs, we would expect a much higher rate of overlap due to common input from MCs. Consistent with this hypothesis, we observe that the probability of detecting an input at the same location in L4-L5 pairs was very low and substantially smaller than for L5-L5 pairs ($2.4 \pm 1.3\%$ spatially coincident inputs out of all input locations in $n = 10$ L4-L5 pairs, versus $28 \pm 6.7\%$ in $n = 10$ L5-L5 pairs; $p = 1.2 \cdot 10^{-3}$, Wilcoxon rank sum test). Given that occasionally more than one SST cell might be photostimulated at any given location (Fig. 3.4d,g), we employed a statistical test for fine time scale synchrony of IPSCs between the patched cells at each candidate location (where both cells received input) to determine whether the IPSCs truly arose from a single SST cell diverging onto both recorded PCs ([4], see Methods, Fig. 3.3g, Fig. 3.5g,h,i,j). Using this far more stringent spatiotemporal test for the detection of common input, we detected no locations in which stimulation evoked common inputs for L4-L5 pairs, whereas we detected at least one common input in 7 of 10 L5-L5 pairs (Fig 3.3h; no locations in $n = 10$ L4-L5 pairs versus $13.7 \pm 5.1\%$ of all input locations in $n = 10$ L5-L5 pairs; $p = 1.1 \cdot 10^{-3}$, Wilcoxon rank sum test; see also Fig. 3.5k). These data argue that individual L5 SST cells connect to either L4 PCs or to L5 PCs, but never to both. In other words, L4 PCs and L5 PCs are inhibited by non-overlapping subnetworks of L5 SOM cells.

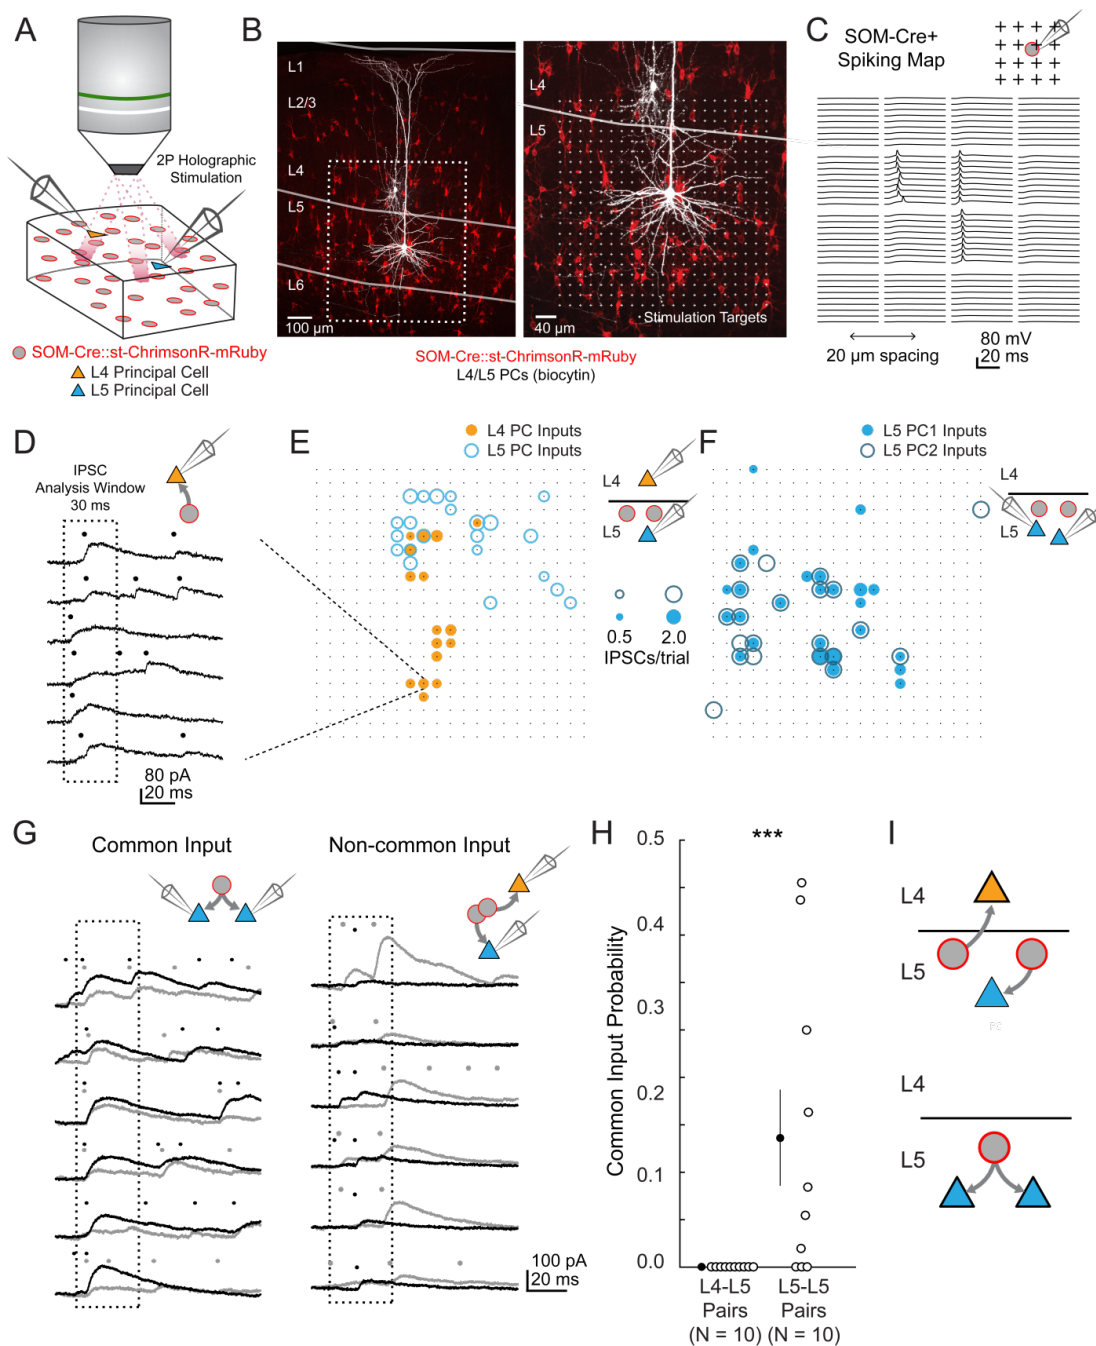


Figure 3.3: 2 photon optogenetic circuit mapping reveals that L4 and L5 PCs are inhibited by separate populations of L5 SST cells. (A) Schematic of the experimental configuration. IPSCs are recorded from a pair of PCs (either an L4/L5 pair or an L5/L5 pair) while SST cells expressing soma-targeted ChrimsonR-mRuby2 are focally activated using 2P photo-stimulation and computer generated holography. (B) Left: post-hoc confocal image showing SST cells expressing soma-targeted-ChrimsonR-mRuby2 (red) and biocytin fills of recorded PCs in L4 and L5 (white) at 10x magnification. Right: Confocal image at 20x magnification showing the grid of photo-stimulated target locations. Both images are max z-projections over 100 μm . *Continued on next page.*

Figure 3.3: *Continued from previous page.* (C) Spatial photo-excitation profile of a soma-targeted-ChrimsonR-mRuby2 expressing SST cell. Whole cell current-clamp recordings from this cell showing multiple trials of photo-stimulation at a 4 x 4 subsection of the photo-stimulation grid with 20 μm spacing between stimulation locations. The SST cell is recruited to spike only at a small number of stimulation sites, but does so reliably and with low jitter across trials at these sites. (D) Example traces showing IPSCs recorded from an L4 PC during SST photo-stimulation at a single site (corresponding to black boxed square in E) over multiple trials. Dots above each trace indicate the onset time of detected IPSCs ($p = .0003$, Poisson detection). (E) Example overlay of maps showing the mean number of IPSCs at detected input locations during photo-stimulation for a simultaneously recorded L4 PC-L5 PC pair. Bubble size indicates the mean number of IPSCs evoked (deviation from background rate) per trial. (F) As in E, but for an L5 PC-L5 PC pair. (G) Example traces illustrating method for detection of common SST-mediated inputs to pairs of simultaneously recorded PCs. Left: IPSC traces at a single site recorded simultaneously in two PCs (each PC is indicated by black or grey traces) and corresponding detected IPSCs. IPSCs with synchronous onset occur in many trials, despite the trial-to-trial jitter in IPSC onset, suggesting that a SST cell which diverges onto both recorded PCs is being stimulated at this site ($p = .0005$, synchrony jitter test). Right: IPSC traces from a different site. Evoked IPSCs are observed in both cells, but the lack of synchronicity suggests they arise from separate, neighboring SST cells ($p = .4$). Dots above each trace indicate the estimated onset time of detected IPSCs. (H) Probability of detecting common SST input per photo-stimulated site for pairs consisting of L4 PCs and L5 PCs versus pairs consisting of two L5 PCs. No common input locations were detected in $n = 10$ L4-L5 pairs versus $13.7 \pm 5.1\%$ of all input locations stimulated in $n = 10$ L5-L5 pairs; $p = 1.1 \times 10^{-3}$, Wilcoxon rank sum test. Data are summarized by mean \pm S.E.M. (I) Schematic of main result for SST outputting mapping. Individual L5 SST cells form inhibitory connections onto L4 PCs and or L5 PCs but not both.

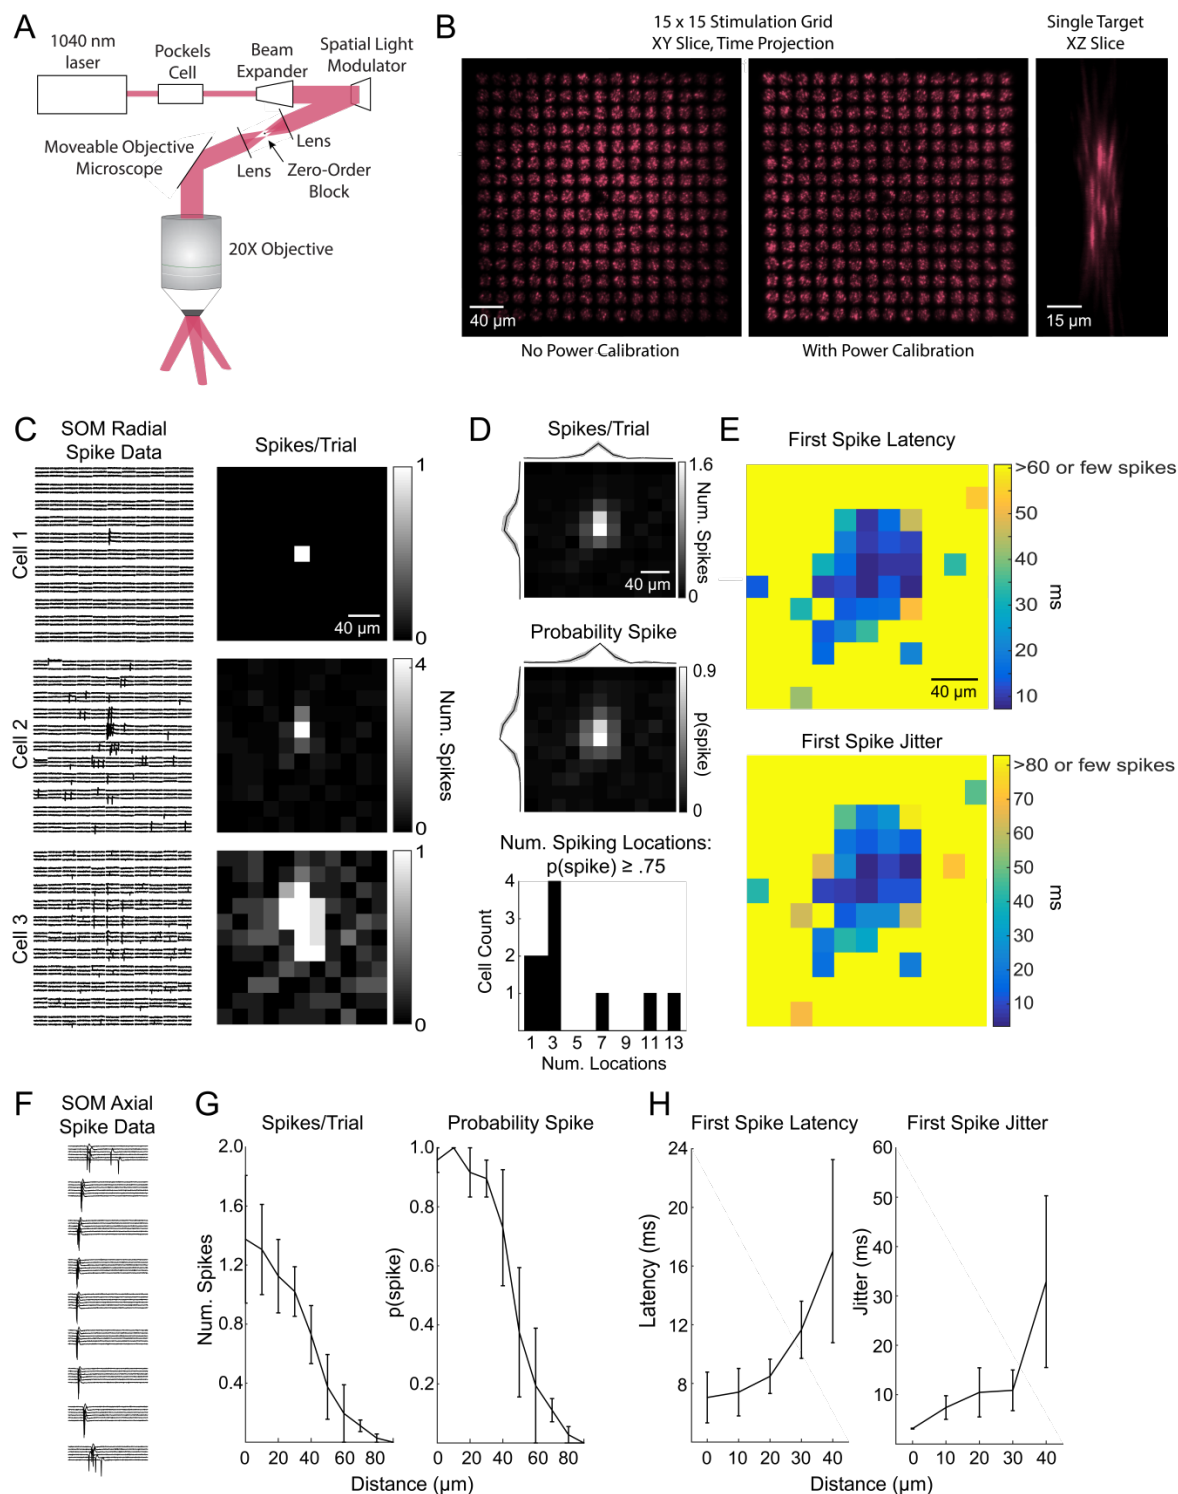


Figure 3.4: (A) Schematic of the CGH-based stimulation microscope. (B) Imaging stimulation holograms with two-photon induced fluorescence in a thin fluorescent slide. Left: Time projection of a stimulation sequence covering a 15 x 15 grid of targets at 20 μm spacing before power calibration. Middle: As in Left but with power calibration. Left: XZ slice of a single target. Note that excitation is confined to a small volume and that power calibration results in more uniform excitation radially. The decrease in fluorescence in the middle of the two stimulation grids is due to the zero-order block. *Continued on next page.*

Figure 3.4: *Continued from previous page.* (C) Example plots of the lateral resolution of photo-stimulation. Left column: Raw cell-attached data of light-evoked spiking for example SST cells. Right column: average number of spikes evoked per trial at each location. Locations are $20\ \mu\text{m}$ apart as in mapping experiments. (D) Radial spike count statistics. Top: Average map of the number of evoked spikes per trial in st-ChrimsonR+ cells ($n=11$). Maps are centered on the somata of recorded cells. Middle: As in top, but for probability of evoking at least 1 spike per trial. Bottom: Histogram of number of locations per cell which evoked a spike at least 75% of the trials. (E) Radial spike timing statistics. Top: Average first spike latency for st-ChrimsonR+ cells ($n=11$). Yellow indicates that the average was ≥ 60 ms or that too few spikes were observed across cells at those locations to obtain a good estimate. Bottom: Average first spike jitter, computed as the full-width half-maximum of the first spike times. Yellow indicates a jitter of greater than 80 msec or that too few spikes were observed across cells at those locations. (F) Data from an example cell as the hologram is moved axially. Distance between locations is $10\ \mu\text{m}$. (G) Axial spike count statistics. Left: Average number of evoked spikes as a function of axial distance for st-ChrimsonR+ cells ($n=4$). Right: As in left, but for the probability of at least one spike per trial. (H) Axial spike timing statistics. Left: Average first spike latency as a function of axial distance ($n=4$ cells). Right: First spike jitter as a function of axial distance for st-ChrimsonR+ cells ($n=3$).

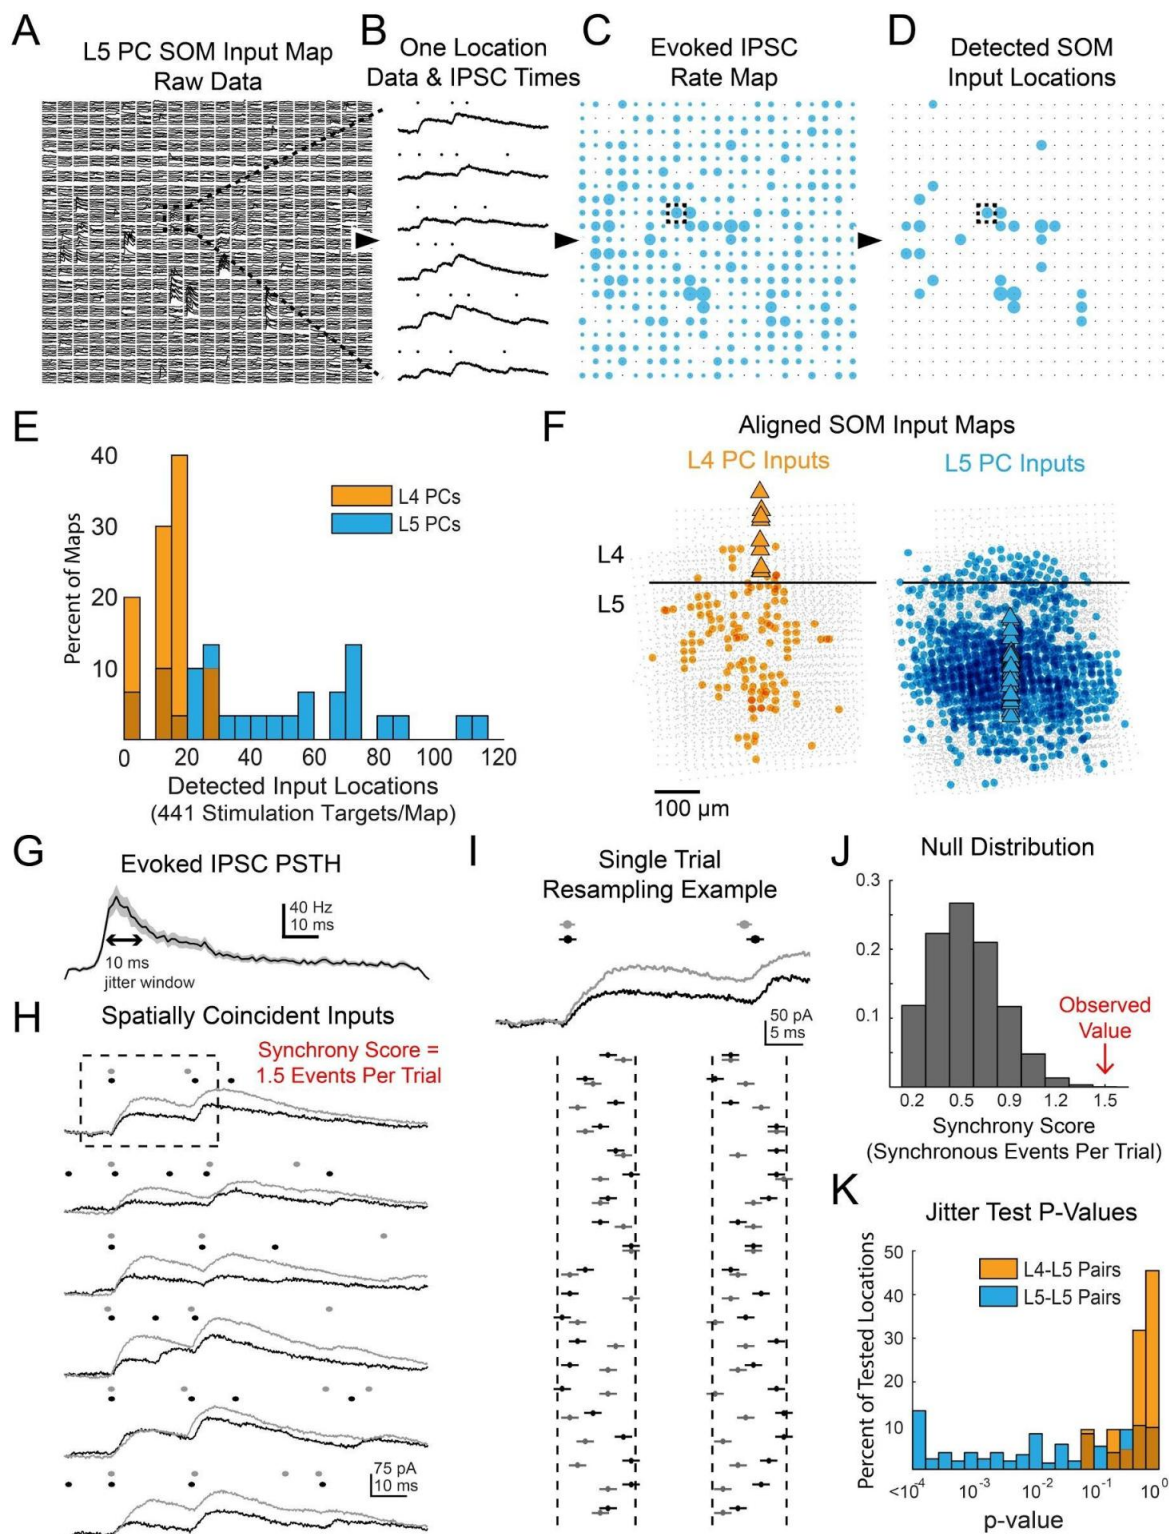


Figure 3.5: Data processing and additional results for multiphoton SST output mapping. *Continued on next page.*

Figure 3.5: *Continued from previous page.* (A) Example of raw data a single map onto a L5 PC. Dashed box shows location for which data is shown in B. (B) Example of Bayesian PSC detection on all 6 trials from a single location from the map in A. (C) A map showing the evoked IPSC rates at each location for the map from A. Dashed box shows location of data from B. (D) As in C except only locations which pass FDR detection are shown. (E) Histograms showing the number of locations with evoked IPSCs for both L4 and L5 PCs. (F) Overall spatial input distributions for SST cells in L5 to both L4 PCs (left, $n = 10$) and L5 PCs (right, $n = 28$). Maps are aligned vertically to the L4-L5 border and horizontally to the PC soma. For this representation all inputs are plotted with the same size circle. (G). PSTH of IPSC times aggregated from all locations with detected evoked IPSC rates. A 10 ms duration which matches the jitter duration for temporal synchrony is marked for comparison. (H) As in B except showing data and IPSC detection for two simultaneously recorded L5 PCs. The synchrony score for this location is 1.5 events/trial. Dashed box shows data used for I. (I) Example of 20 resamplings of the events during the analysis window for the first trial shown in H. Vertical dashed lines show discrete event jittering windows. Horizontal lines on each event span 2ms such that if the lines from two events overlap then they would be counted as synchronous. (J) Null distribution of synchrony score from event time series resampling of the data in H. The observed value is from the far right extreme of the distribution. (K) Histograms of p-values for all spatially overlapping locations for all pairs from jitter synchrony tests

3.4 Paired recordings confirm dense, reciprocal, and selective intra- and translaminar connectivity

We hypothesized that the L4-targeting and L5-targeting populations of L5 SST cells identified with two-photon optogenetic mapping corresponded to XCs and MCs. This would mean that SST cells in L5 are organized into at least two discrete subnetworks, each with its own unique input and output connectivity. If this were the case, one would expect L5 XCs to be reciprocally connected with L4 PCs, but not L5 PCs; conversely, L5 MCs should be reciprocally connected with L5 PCs but not L4 PCs. To address this, we conducted paired intracellular recordings to assess the fine-scale patterns of synaptic connectivity between L4/L5 PCs and L5 MCs/XCs. This strategy is low throughput, but provides unambiguous measurements of synaptic connectivity between identified cell types; furthermore, it complements the above experiments because it is unaffected by issues that potentially limit optogenetic approaches (for example, differences in intrinsic excitability between cell types or variability in opsin expression). In addition to targeting L5 MCs and XCs with the GIN and X94 lines as above, we also used the SST-TdT line to identify L5 SST cells, which we classified as putative MCs or XCs based on their electrophysiological properties (Fig. S1g; Table S1).

First, we address translaminar connections. Paired recordings between L5 XCs and L4 PCs revealed monosynaptic L5 XC→L4 PC connections, which we observed as short-latency IPSP/IPSCs in the L4 PC membrane potential following induced action potentials in the L5 XC. We observed frequent L5XC→L4PC connections (36/67 pairs tested; 54%; Fig. 3.6a,b), even across relatively long inter-somatic distances ($183 \pm 67 \mu\text{m}$, mean \pm S.D.; Fig. S6), suggesting that L5 XCs connect densely onto L4. We also frequently observed monosynaptic excitatory connections from L4 PCs onto L5 XCs, consistent with the optogenetic experiments above (39/72 pairs tested; 54%). These synapses exhibited profound facilitation during sustained high-frequency firing in the presynaptic cell (Fig. 3.6a; Fig. S8), the hallmark of excitatory connections onto SST cells [70, 17, 20]. In cases where we tested connectivity bidirectionally, we frequently observed reciprocal connections (23/56 pairs tested; 41%). Thus, L5 XCs and L4 PCs form a translaminar feedback inhibitory motif. We also observed frequent connections from L5 XCs onto L4 fast-spiking (FS) cells (12/23 pairs tested; 52%; Fig. S7), consistent with a known circuit in which L4 non-Martinotti SOM cells inhibit L4 FS cells [97, 180].

In contrast, we very rarely observed excitatory connections from L4 PCs to L5 MCs (1/95 pairs tested; 1%; Fig. 3.6a,b) or from L5 MCs onto L4 PCs (4/68 pairs tested; 6%), despite the fact that these pairs were separated by smaller inter-somatic distances than L4PC-L5XC pairs ($143 \pm 47 \mu\text{m}$, mean \pm S.D.; Fig. S6), likely due to the fact that MCs and XCs preferentially localize in the superficial and deeper portions of L5 respectively [51]. In a subset of these experiments, we recorded from L4 PCs in the voltage clamp configuration at +10mV (using a cesium-based internal solution), but did not observe connections any more frequently (0/38 pairs tested; 0%). These data suggest that L5 XCs are integrated into the densely interconnected network of L4 PCs and interneurons [16, 133], whereas L5 MCs are essentially isolated from it. Since our 2P mapping experiments indicated that distinct sets of SST cells inhibit L4 PCs and L5 PCs, we hypothesized that this pattern would be reversed

within L5, i.e. L5 MCs but not L5 XCs will interconnect with L5 PCs.

We next examined intralaminar connections between L5 PCs and L5 SST cells. We observed frequent inhibitory connections from L5 MCs onto L5 PCs (24/46 pairs tested; 52%; Fig. 3.6c,d), in agreement with a large body of work indicating that Martinotti cells diverge profusely onto nearby intralaminar targets [46, 66, 19]. We also observed excitatory connections from L5 PCs onto L5 MCs, albeit more rarely (4/29 pairs tested; 14%; Fig. 3.6c,d) but at a rate consistent with the literature [91, 66]. In contrast, we detected very few inhibitory outputs from L5 XCs onto L5 PCs (2/65 pairs tested; 3%; Fig. 3.6c,d) or excitatory connections from L5 PCs onto L5 XCs (1/60 pairs tested; 2%; Fig. 3.6c,d), despite the fact that L5 PCs were on average located much closer to L5 XCs than were L4 PCs. The surprising dearth of intralaminar connectivity between L5 PCs and L5 XCs stands in stark contrast to the dense intralaminar connectivity observed between L5 PCs and L5 MCs, as well as in other inhibitory circuits [46, 122, 91], lending further support to the notion that MC and XC circuits are uniquely and selectively wired (Fig. 3.6e).

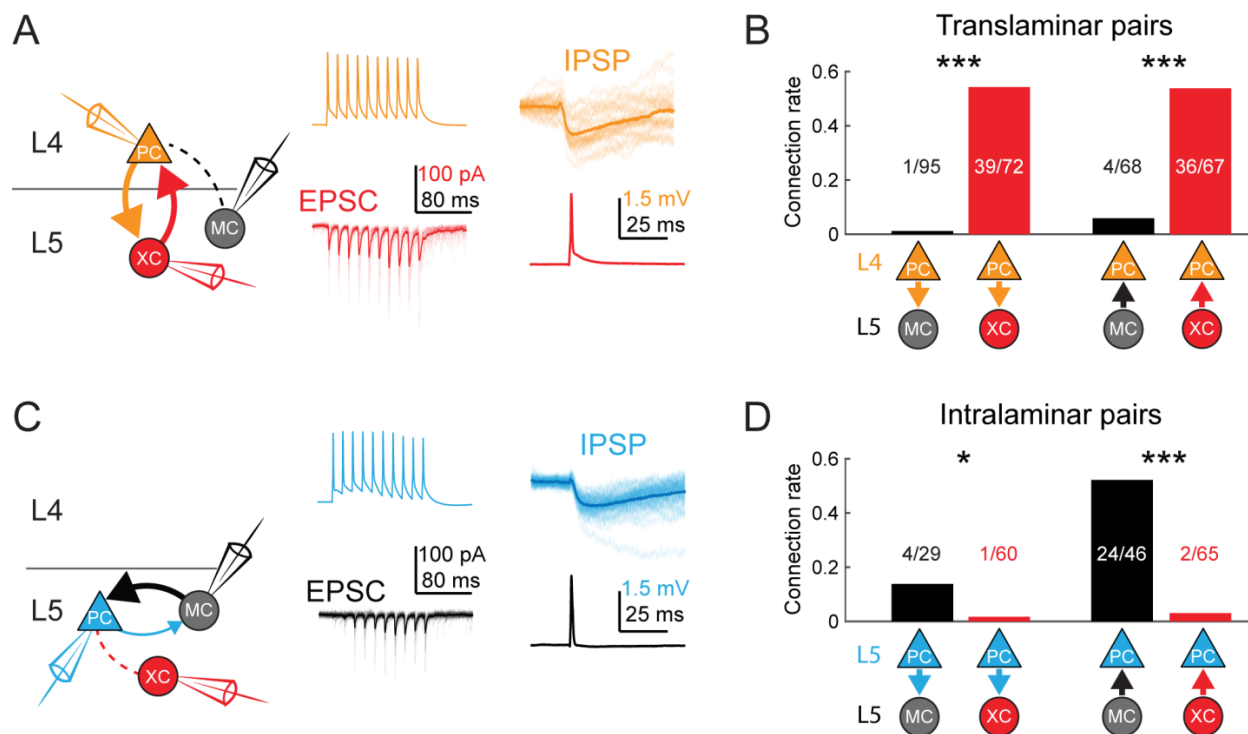


Figure 3.6: MCs and XCs exhibit different patterns of monosynaptic connectivity with L4 and L5 PCs. (A) Paired recordings of L4 PCs (orange) and L5 XCs/MCs (red/black). Left: schematic of the tested circuit. Middle: example traces of evoked spikes in a L4 PC (orange) and the excitatory synaptic current in a L5 XC (red). Right: example traces of evoked IPSPs in a L4 PC (orange) in response to a single action potential in a L5 XC (red). (B) Bar graph summarizing translaminar connection rates between L4 PCs and L5 MCs (black bars) and L4 PCs and L5 XCs (red bars). $p < 10^{-6}$ for L4PC→L5MC ($n = 95$ connections tested onto 39 MCs) versus L4PC→L5XC connection rate ($n = 72$ connections tested onto 51 XCs); $p = 2 \cdot 10^{-6}$ for L5MC→L4PC ($n = 68$ connections tested from 35 MCs) versus L5XC→L4PC connection rate ($n = 67$ connections tested from 51 XCs); Monte Carlo permutation test. (C) As in A, but intralaminar pairs between L5 MCs/XCs and L5 PCs (blue). (D) As in B, but for intralaminar connections with L5 PCs. $p = 0.020$ for L5PC→L5MC ($n = 29$ connections tested onto 20 MCs) versus L5PC→L5XC connection rate ($n = 60$ connections tested onto 35 XCs); $p < 10^{-6}$ for L5MC→L5PC ($n = 46$ connections tested from 30 MCs) versus L5XC→L5PC connection rate ($n = 65$ connections tested from 37 XCs); Monte Carlo permutation test.

3.5 Discussion

Understanding of the organization of cortical inhibitory circuits has surged in recent years, but the integration of these new architectures with the established, layer-centric framework of the excitatory cortical microcircuit has proved elusive. Our data establish previously unknown excitatory-to-inhibitory and inhibitory-to-excitatory pathways involving two sub-networks of L5 SST interneurons, which are both layer and cell-type specific. Optogenetic circuit mapping shows that L5 MCs receive excitatory inputs chiefly from PCs in L2/3 and L5, the primary cortical output layers, while L5 XCs receive inputs mainly from PCs in L4 and upper L6, the primary input zones for afferent input from the ventral posterome-

dial thalamus [177]. Paired recordings and 2-photon holographic optogenetic interrogation of common SST-mediated input indicate that in turn, these same SST subtypes selectively inhibit the same PC populations that excite them, at least within L4 and L5. Thus XCs and MCs, though spatially intermingled in L5, are functionally segregated into two disjoint networks with selective and complementary laminar connectivity (Fig. 3.6e).

3.5.1 Dense but selective inhibitory wiring

These results reveal a previously unknown, striking degree of specificity in the inhibitory cortical wiring diagram. In particular, the observation that L5 XCs exhibit nearly no intralaminar connectivity with L5 PCs, but do engage in dense, reciprocal connectivity with L4 PCs is inconsistent with the idea of a single, global blanket of SST-mediated inhibition. Instead, SST-PC circuits appear to more closely resemble a patchwork quilt, comprised of multiple networks of SST subtypes which independently service separate spatial domains. SST-PC connectivity can be extremely dense and non-selective within one of these domains (e.g. creating a blanket within a single laminar microcircuit), but highly selective on the scale of layers and columns. Though we did not investigate it in this study, it seems likely that selective connectivity on this scale could be achieved, at least in part, simply by SST subtypes having preferred laminar targets for axonal arborization [124].

3.5.2 Functional implications of separate, layer-specific SST feedback circuits

The circuit structure established here suggests that at least two different pools of SST-mediated inhibition separately modulate successive stages of cortical processing. XCs will chiefly be recruited as a function of the ‘input’ of the microcircuit, generating inhibition proportional to the amount of activity in the input layers and thalamus, which in turn varies widely depending on external sensory drive. Conversely, MCs will be recruited by activity in L2/3 and L5 PCs, whose projections represent the major feedforward pathways out of S1 [54]; consequently, MCs will generate feedback inhibition as a function of the ‘output’ of the microcircuit. This raises the possibility that MC and XC-mediated feedback inhibition could differentially contribute to various forms of population gain control: XCs could perform input-scaling on the gain, whereas MCs could regulate output-scaling. These different modes of gain control have been proposed to perform myriad functions in perception and cognition [30, 109, 149], and the circuit mechanisms of gain modulation are a topic of intense recent focus [87, 48, 176, 154, 142, 11, 95, 125]. Independent modulation of different SST circuits, either by the intracortical excitatory circuits described here, long range inputs from thalamus or higher cortical areas, or by other sources such as VIP-mediated inhibition [136, 113, 88, 137] or neuromodulation [179, 180, 140] could represent a mechanism by which the brain dynamically fine-tunes the cortical population input-output function.

Furthermore, these parallel subnetworks of SST inhibition might also fine-tune how the sensory cortex integrates ‘bottom-up’ and ‘top-down’ inputs, as bottom-up is conveyed via thalamocortical axons primarily (though not exclusively; [34]) to L4, and thus is most directly affected by XC-mediated inhibition, while top-down inputs are thought to be conveyed to the apical dendrites of L5 and L2/3, and thus are primarily modulated by MC-mediated

inhibition [114, 134]. Therefore, the parallel structure of MC and XC networks could allow S1 to independently alter its sensitivity to top-down and bottom-up input streams to optimize the behavioral responses to sensory stimuli. Supporting this proposal, recent studies have observed that subtypes of SST cells with either wide or narrow spike waveforms (which might correspond to MCs and XCs) are differentially activated during various behavioral contingencies [85, 77, 113, 144]. Pressing questions for future investigation will be to determine how excitatory drive from different laminar PC populations contributes to the unique activity patterns of SST subtypes, and conversely, to determine how distinct SST subtypes differentially shape the dynamics of the cortical microcircuit. Future experiments could address these questions by optogenetically manipulating specific SST subtypes using intersectional genetics [56] or with a Cre-DOG based approach [163]. Nearly all optogenetic manipulations of SST cells have employed the non-specific SST-Cre line, which could explain the somewhat heterogeneous results that have emerged, including both inhibition and disinhibition [87, 180, 176, 154, 35]. More precise manipulations of SST subtypes (e.g., MCs vs. XCs) could help reveal the functional impacts of distinct SST subnetworks.

The selective inhibitory ‘pathways’ described here also suggest specific predictions about how different SST subtypes shape the dynamics of cortical activity. For example, within layers, XCs and MCs will respectively mediate L4→L4 and L5→L5 feedback dendritic inhibition [161]. Across layers, SST cells should contribute to L2/3→L5 feedforward inhibition [142, 9, 2], but not to L4→L5 translaminar inhibition, implying that the latter operates solely through parvalbumin-expressing FS interneurons [138]. Furthermore, XCs (but not MCs) receive strong thalamocortical input [65, 61, 162] and are thought to generate feedforward inhibition in response to sustained thalamic activity; our data therefore imply that SST-mediated thalamocortical feedforward inhibition will impinge upon L4 neurons but not L5 PCs.

3.6 Methods

All experiments were performed in accordance with the guidelines and regulations of the Animal Care and Use Committee of the University of California, Berkeley.

3.6.1 Transgenic mice

The following mouse lines were used for this study: the *scnn1-tg3-Cre* line (JAX stock 009613), the *emx1-IRES-Cre* line (JAX stock 005628), the *PV-IRES-cre* line (B6;129P2-Pvalbtm1(cre)Arbr/J; JAX stock 008069), the *SST-IRES-cre* line (JAX stock 013044), the *GIN* line (FVB-Tg(GadGFP)45704Swn/J; JAX stock 003718), the *X94-GFP* line (Tg(Gad1-EGFP)94Agmo/J; JAX stock 006334), the *Ai9 Rosa-LSL-tdTomato* line (JAX stock 007909). Mice were housed in cohorts of five or fewer with a light:dark cycle of 12:12 hours, and were used for experimentation during their subjective night.

3.6.2 Viral Infection

Neonatal mice (p0-3) were deeply cryo-anesthetized and placed in a head mold. Viral aliquots were loaded into a Drummond Nanoject injector and injected into 4 sites in the barrel cortex of the left hemisphere. At each site, virus was injected at multiple depths (2 depths for *scnn1-tg3-cre* and *drd3-cre* mice, 3 depths for *emx1-IRES-Cre* and *SST-IRES-cre* mice) in increments of 18.4 nL or 36.8 nL (for multiphoton experiments), for a total of 150-440 nL of virus injected per mouse. Following injections, mice were moved to an incubation chamber for recovery, and were returned to the dam once they regained color and began to move. Viruses used were AAV9.CAGGS.Flex.ChR2-tdTomato.WPRE.SV40 (acquired from the University of Pennsylvania Vector Core; undiluted for *scnn1-tg3-cre* and *drd3-cre* mice, diluted 1:1 with PBS for *emx1-IRES-Cre* mice), and AAV9-2YF-hSyn-DIO-ChrimsonR-mRuby2-Kv2.1.

3.6.3 Brain slice recording

Acute thalamocortical slices were prepared from mice (ages p14-29, at least 14 days after viral injection; mice were selected randomly from litters and thus there should be no bias between sexes used) as previously described ([2]). Slices were placed in a recording chamber and constantly perfused with oxygenated artificial cerebro-spinal fluid (NaCl 119 mM, KCl 2.5 mM, MgSO₄ 1.3 mM, NaH₂PO₄ 1.3 mM, glucose 20 mM, NaHCO₃ 26 mM, CaCl₂ 2.5 mM) maintained at 32° C (21° C for multiphoton mapping experiments). Slices were oriented with the caudal surface facing up in the recording chamber. To ensure minimal disruption of vertical connectivity, all slices used for recording were inspected under infrared illumination at 40x magnification and/or post-hoc confocal imaging to confirm that pyramidal cell apical dendrites stayed roughly parallel with the surface of the slice or receded slightly deeper as they progressed apically. Whole cell recordings were performed using glass micropipettes (2-5M Ω resistance) pulled on a Sutter P-1000 Micropipette Puller. Pipettes were filled with a Cs⁺ based internal (CsMeSO₄ 135 mM, NaCl 8 mM, HEPES 10 mM, Na₃GTP 0.3 mM, MgATP 4 mM, EGTA 0.3 mM, QX-314-Cl 5 mM, TEA-Cl 5mM) or a potassium gluconate based internal (K-gluconate 135 mM, NaCl 8 mM, HEPES 10 mM, Na₃GTP 0.3 mM, MgATP 4 mM, EGTA 0.3 mM). In some experiments, biocytin (0.4-1%) was dissolved into the internal solution to enable morphological recovery. Voltage recordings were not corrected for the junction potential. Series resistance was monitored with negative voltage steps during each trial, and was compensated up to 60%. Data were analyzed from recordings in which series resistance remained stable and below 30M Ω . Data were acquired and filtered at 2.2 kHz using a Multiclamp 700B Amplifier (Axon Instruments) and digitized at 20 kHz (National Instruments). All data were acquired using custom written MATLAB (Mathworks) software.

3.6.4 Characterization of intrinsic properties

In all recordings using K-based internal solution, an F-I curve was measured at the start of the experiment using a series of 1-second current injections, at -200 pA, -100 pA, and then proceeding in 50 pA increments from +50 to +500 pA. In some experiments, additional current steps were manually designated and performed online to aid in estimation of

rheobase. Resting membrane potential was defined as the median membrane potential during a baseline period measured immediately after break-in. Input resistance was calculated with Ohm’s law using the steady state membrane potential during subthreshold current injections steps (current clamp) and/or the steady state current during 5 mV voltage steps (voltage clamp). Action potential onset was detected using code adapted from the Berg lab’s Spike_threshold_PS function, which defines onset as the point of maximum positive slope in the phase space of the membrane potential and its first derivative [153]. Spike width was measured as the full-width of each spike at the voltage halfway between the action potential threshold and the peak amplitude (half-max). Rheobase was estimated using the average of 1) a linear fit (with coefficients constrained to be nonnegative using the lsqnonneg function in MATLAB) of the F-I relation during the last subthreshold current injection step and the first few suprathreshold steps and 2) linear extrapolation of the current necessary to reach threshold based on measurements of the resting membrane potential, input resistance, and average threshold value of the first action potentials evoked during suprathreshold injections. These two measures were usually in good agreement. Adaptation index was calculated (following the Allen Brain Institute’s Cell Types Database protocol) for each current injection using the expression:

$$\frac{1}{N-1} \sum_{n=1}^{N-1} \frac{ISI_{n+1} - ISI_n}{ISI_{n+1} + ISI_n} \quad (3.1)$$

Where N is the number of spikes during that current step and ISI is the interspike interval.

3.6.5 Paired recording connectivity testing

We first targeted whole-cell recordings to a fluorescent (GFP+ or TdTomato+) SST cell, and then subsequently patched nearby neurons in the same slice. In some cases, we recorded serially from several neurons while maintaining the recording of the first neuron, in order to test multiple connections. Monosynaptic excitatory connectivity onto SST cells was tested by driving trains of 10 spikes in the presynaptic cell at 70 Hz via current injection, while monitoring for EPSCs in the postsynaptic cell. Stimulation was repeated at least 15 times in all pairs tested. Monosynaptic inhibitory connectivity from SST cells onto other neurons was tested by driving spikes in the presynaptic cell while monitoring postsynaptically for IPSCs (Cs-based internal, postsynaptic cell held at +10mV) or IPSPs (K-based internal, postsynaptic cell depolarized to approximately -52mV). Electrical connectivity between SST cells was tested by hyperpolarizing each cell with 1-second current injections (at least 15 trials) while monitoring for hyperpolarization in the other cell.

For L5SST-L5PC pairs, we recorded from both pyramidal tract and intratelencephalic type PCs, which could be distinguished by their laminar positions (preferentially L5B versus L5A), morphology visualized via infrared (large soma versus smaller soma) and post-hoc confocal imaging (thick-tufted apical dendrites versus slender-tufted), and/or their intrinsic properties (initial burst/doublet spiking followed by non-adapting spikes versus continuously adapting regular-spiking phenotype; [55, 78, 152]). We did not observe any significant differences in the connectivity of either L5PC type with L5MCs or XCs. For L5SST-L4PC pairs, we did not distinguish between spiny stellate and pyramidal/star-pyramidal excitatory cells.

For paired recordings between L5 SST cells and L4 FS/PV cells, we identified FS/PV cells using PV-Cre; LSL-TdTomato mice in some experiments (Fig. S7i). However, it was often difficult to visualize X94 cells using these animals due to the TdTomato fluorescence being much brighter than the GFP fluorescence. In other experiments (Fig. S7h), we targeted FS/PV cells in L4 by looking for L4 neurons with large cell bodies under IR, and then confirmed the identity of these cells electrophysiologically, with the primary criteria separating them from being narrow spike widths (slightly shorter than the average XC spike) and little or no spike frequency accommodation during high amplitude steps of current injection.

To classify SST-TdT cells as putative XCs or MCs, we fit a support vector machine (cross validated 10-fold) to perform binary classification of L5 GIN cells and L5 X94 cells using only their intrinsic electrophysiological properties. We found that a classifier based on only two measures (spike width and estimated rheobase) performed just as well as multivariate classification based on a large number of metrics (85% accuracy). We then used this classifier to predict the identity of a different dataset of L5 SST cells recorded in SST-TdT mice. This approach is likely to have resulted in a small number of SST-TdT cells being misclassified; however, the connectivity of putative XCs and MCs were highly similar to the connectivity of XCs and MCs identified using the X94 and GIN lines. Furthermore, our conclusions about the differences in connectivity rates of L5 MCs and XCs with L4 and L5 PCs are unchanged by the exclusion of the SST-TdT dataset, with the exception of L5PC→L5SST connections – a circuit which has been studied in some detail by others. This approach also effectively assumes a dichotomy in L5 SST cells, since we have only two labels (MC and XC) to provide as training data, which is an important caveat since it is likely that further subdivisions of SST cells exist in L5 [97]. In a handful of cases, we recorded from SST-TdT cells which appeared to be FS cells 114, with very narrow spikes, low input resistances, and a near complete lack of spike-frequency accommodation during high amplitude current injection steps; these neurons were excluded from further analysis.

Using paired recordings, we tested 544 total possible connections between 146 L5 SST cells (39 L5 GIN cells, 53 L5 X94 cells, 54 L5 SST-TdT cells) and PCs/FS cells in L4 and L5. Data from a subset of these neurons ($n = 17$ L5 GIN cells) were included in a previous study [138]. This dataset was unbalanced, and because in some cases we tested multiple connections onto the same L5 SST cell, included some non-independent observations. Because of this, we used Monte Carlo permutation tests to test for significant differences between the connectivity rates of MCs and XCs. We generated a permuted dataset with the same observation structure (same number of L5 SST cells and same number of connections tested per L5 SST cell) in place for MCs and XCs by randomly resampling with replacement at both levels. We then measured the difference in observed connectivity rate for the MC and XC groups, and repeated this procedure 100,000 times to generate a null distribution of rate differences. We used this distribution to perform a 1-tailed test for significant differences between MC and XC connectivity rates for each type of connection tested (Table S1).

3.6.6 Optogenetic connectivity mapping in vitro

Experiments were done in slices from Emx1-Cre; GIN or Emx1-Cre; X94 mice injected with an AAV driving Cre-dependent expression of ChR2 in all excitatory cells. Whole cell voltage

clamp recordings were performed in GFP+ L5 cells to target L5 MCs (Emx1-Cre; GIN) or L5 XCs (Emx1-Cre; X94). A digital micromirror device was used to focally photo-stimulate excitatory cells in different regions of the slice in order to map the spatial profile of excitatory inputs to recorded MCs and XCs.

Prior to experiments, slices were briefly visually inspected with epifluorescence under a 5x objective to confirm that a wide area containing dense, even expression of fluorescence (tagged to an opsin) was present in the barrel cortex. Recordings were targeted to within this region, which typically covered the entire lateral extent of barrel cortex in 4-5 slices. Slices in which expression appeared faint or uneven were discarded.

In some experiments, it was necessary to locate fluorophore-positive cells in slices also containing an excitatory opsin. To avoid excitotoxicity that can result from excessive illumination of opsin-containing neurons, we limited illumination to very brief intervals (1-2 seconds) while searching for fluorophore-positive cells. In some cases where the target cells were weakly fluorescent (young GIN and X94 animals), we searched for these cells while keeping the slice submerged in sucrose-substituted ACSF. Once target cells were located, this solution was washed out and replaced with normal recording ACSF prior to patching these cells and starting experiments.

3.6.7 DMD-based excitatory input mapping

Laser light was generated using a 1W 445nm diode laser (Ultralasers) and routed via a liquid light guide into a CEL5500 digital micromirror device (DMD) (Digital Light Innovations). The projection from the DMD was then collimated and integrated into the light path of the microscope, before being focused onto the slice chamber using a 5x (Olympus). For experiments using widefield illumination, the DMD passively reflected but not spatially modulate light. Prior to photo-stimulation, infrared and epifluorescence images were captured using an IR-1000 CCD camera (DAGE-MTI) and imported into MATLAB.

Excitatory mapping experiments were performed using a modified version of a previously described protocol [138]. Mapping was performed over an area extending from pia to the white matter, covering 2-4 barrel columns laterally (400 to 800 μm). For mapping excitatory inputs to GIN and X94 cells, the DMD was used to pattern light into a square region (75 μm x 75 μm). Each stimulation site was spaced 40 μm apart from adjacent ones, resulting in some overlap of adjacent stimuli. We chose to ‘ramp’ our photostimulation, starting each stimulus with the light off and linearly increasing the light intensity over time. Ramping in this manner minimizes activation of fibers of passage [2]. In each trial, a ‘sawtooth’ light stimulus composed of three successive 25ms ramps of light (1.25 mW/mm² final intensity) was applied to one stimulus site. This protocol was chosen in order to maximize the short-term facilitation of excitatory inputs to L5 SST cells, though in practice we found it was usually possible to observe responses during the first ramp alone. Ten regions were stimulated per second in a serial, pseudorandom order, with 4 second breaks after every 10 seconds of mapping. Control experiments were performed using identical stimulation conditions while recording from ChR2+ neurons in all layers. These experiments determined the spatial resolution of photostimulation and confirmed that spiking was elicited in ChR2+ neurons only when regions very close to the soma were stimulated. We also included $n = 2$ experiments mapping

inputs to L5 X94 cells which were performed using the exact mapping protocol described in [138], though our results and conclusions were not substantially altered by their exclusion.

All data were analyzed using custom written MATLAB software. Data preprocessing consisted of removing baseline offsets and slow fluctuations from recordings by subtracting a down-sampled and median-filtered versions. Charge was calculated as the integral of the preprocessed recordings during photo-stimulation and the subsequent 25 milliseconds. To aggregate maps across cells, we first rotated the average map collected in each experiment in order to horizontally orient the laminar boundaries of the mapped area. Maps were next translated vertically to align the L4-L5 laminar boundary, and translated horizontally to align either the home column or the soma position of the recorded cell, before being horizontally cropped to an area $\pm 300\mu\text{m}$ of their center and then averaged to yield a summary map.

For L4 stimulation experiments, we used widefield photostimulation through a 5x objective. We used two stimulation protocols: prolonged, 1-second ramps of linearly increasing light intensity and trains of ten pulses (1ms duration) at 40 Hz. We stimulated at 4 different intensities for each protocol. Since we sometimes recorded multiple neurons in the same slice (see Fig. S5), we fit generalized linear mixed effects models to the dose-response function of light-intensity versus evoked response (EPSC charge transfer or number of spikes), with fixed effects coefficients for the slope of this function for each cell-type and random effects slope coefficients for each slice and neuron in the dataset as well as a constant intercept term. F-tests were used to test for differences in fixed effects coefficients. For paired analysis of L4 XCs and L5 MCs/XCs (Fig. S5), paired t-tests were used to test for differences in L4-evoked responses at maximum stimulus intensity.

3.6.8 Multiphoton CGH-based inhibitory output mapping

Laser light was generated using a 5W 1040 nm femtoTrain laser (Septra-Physics) and power was modulated on short time scales using a Pockels cell (Conoptics) and a high speed shutter (UniBlitz). Light was delivered to the sample using a VIVO 2- Photon workstation (3i) based on a Sutter Moveable Objective Microscope (Sutter) and the hologram was created using a Phasor 2-Photon computer-generated holography system (3i) with Slidebook software (3i) (Fig. 3.4a). The holograms used for stimulation were 2D discs of diameter 15 μm centered on points with 20 μm spacing, making a 400 μm x 400 μm grid in the focal plane (Fig. 3.4b,c). Stimulation consisted of 4 or 10 ms square pulses to the Pockels cell with voltages calibrated to produce 200 or 250 mW average power on sample, respectively. The choice of 4 ms at 200 mW or 10 ms at 250 mW stimulation was determined slice to slice based on opsin expression. Power for each hologram was calibrated empirically to account for power loss due to diffraction efficiency degradation away from the zero-order of the SLM. There was an inter-trial interval of 100 ms between the end of one stimulation and the start of the next stimulation. Under these conditions, SST cells spiked reliably and with high radial resolution (Fig. 3.3c, Fig. 3.4c,d) and moderate axial resolution (Fig. 3.4f,g). Given the sparsity of SST neurons (Fig. 3.3b), this level of spatial resolution provided a good tradeoff between sampling many cells with fewer targets and spiking cells with high spatial resolution. In addition, reliably evoked spikes were produced with low latency and jitter when stimulating randomly through the target grid at 10 Hz. Under these conditions, most evoked spikes occurred in the first 20 ms after the onset of stimulation (Fig. 3.4e,h).

Space clamp error will inevitably affect somatic measurements of currents from distally located SST- \rightarrow PC synapses; however, we recorded IPSCs using a cesium-based internal solution (which included the ion channel blockers tetraethylammonium and QX-314) and performed experiments at room temperature, which ameliorate this to some extent [175]. We also used a holding potential of +40 mV to increase the IPSC driving force. In these experiments, internal solutions also contained 5 M Alexa 488 hydrazide (ThermoFisher Scientific) to aid visualization with multiphoton imaging, and 5 mM kynurenic acid Sodium salt (abcam) was added to the external ACSF to block glutamatergic activity.

To determine which locations evoked responses in the voltage-clamp recordings, first we detected IPSCs using a Bayesian modeling approach via Gibbs sampling [108]. Event times were estimated by first binning all of the time samples at 1 ms resolution and then finding maxima of those timeseries (using findpeaks in MATLAB). Because the vast majority of evoked spikes recorded from opsin expressing SST cells occurred with short latency (Fig. 3.4c,f), we estimated the background rate of IPSCs for each patched cell from the last 25 ms of all inter-trial intervals and the evoked rate at each location from a time window of 30 ms starting 5ms after the onset of each stimulation. Taking a Poisson distribution of events with the estimated background rate as a null distribution for all locations for each cell, we could then calculate a p-value for the hypothesis that there is no evoked IPSCs each location. We then detected locations with evoked responses using the Benjamini-Hochberg False Detection Rate (FDR) procedure with $q = 0.1$ [18]. We chose this relatively liberal FDR rate because any false positives will likely be thrown out after the temporal statistics are taken into account.

To determine if a location with evoked rates in both simultaneously patched cells was in fact a common input from a single source, we employed a statistical test that compares a computed synchrony statistic against a null distribution computed from resampled event time series. Specifically, the test we use employs a null distribution where all synchrony is a result of some processes at timescales longer than some given duration [4]. The intuition is that the chosen duration should match the general timing of evoked IPSCs such that any synchrony under this null arises only because IPSCs are being generated by two separate presynaptic SST cells at roughly the same time. When we reject this null, we have evidence that the synchrony is coming from a process that operates at a finer timescale than the general evoked IPSC statistics: that is, a single presynaptic SST cell is generating highly time-locked IPSCs in two postsynaptic PCs. In our case, the duration of the timescale we want to test against can be estimated from both the timing statistics of evoked spiking of SST cells as well as the peristimulus time histogram (PSTH) of IPSCs for all trials at all detected input locations across all PC input maps (Fig. 3.4e,h, Fig. 3.5g). Using these statistics as guidance, we chose 10 ms as the timescale for our null distribution. In detail, we first summarize the synchrony of events between two simultaneously patched cells at each location where both cells receive input. The statistic we use is the sum of the center and two flanking bins of the cross correlation of the binary event time series for each cell. As mentioned, the event time series are binned at 1 ms. We then created a null distribution for this statistic at each of these locations using the event series resampling described in 66 which allowed us to estimate a p-value for each location as well (Fig. 3.5h,i,j,k). We then detected common spatiotemporal input using these p-values and the Benjamini-Hochberg FDR procedure with $q = .05$, aggregating all tests across all paired maps. The common

input probability for a simultaneously patched pair could then be computed as the total number of detected common input locations for that pair divided by the total number of unique detected input locations for the pair (i.e. the cardinality of the union of the sets of input locations for the two cells).

To align the input maps for each cell, we first aligned each input map to a two-photon image of the tissue taken at the time of recording based on previous calibrations between the SLM coordinate frame (e.g. the input map frame) and the two-photon imaging frame. Next, the tissue-aligned maps were then registered via an affine transform to a confocal image of the fixed slice which had been stained with DAPI and in which the opsin expressing cells could be visualized as well as the patched cells which had been filled with biocytin. This allowed each map to be registered to each other based on the laminar borders, in particular the L4-L5 boundary.

3.6.9 Biocytin staining and reconstruction

Following experiments, slices were transferred to 4% paraformaldehyde at 4° for several days. Slices were then repeatedly washed in TBS and subsequently incubated in block solution at room temperature for two hours. Next, 1:1000 streptavidin-Alexa647 conjugate was added to the solution and allowed to stain for 2 hours. Slices were then washed again and mounted/DAPI-stained on coverslips using VectaShield.

Stained neurons were imaged on a confocal microscope, along with the DAPI signal in order to identify laminar boundaries. These images allowed us to qualitatively assess whether recorded cells were L1-targeting MCs or L4-targeting XCs. We reconstructed a subset of filled neurons, with the goal of performing a bulk quantification of how MC and XC neurites are distributed with respect to the cortical layers (Fig. 3.1c, Fig S2e). Since detailed reconstructions of the morphologies of these neurons have already been carried out by others [173, 161, 97, 180, 56, 162, 105], we adopted a high-throughput, semi-automated approach to perform 2D reconstruct MCs and XCs (Fig. S2c). We imaged neurons using a 10x air objective and used the Imaris software package to automatically trace filled neurites. Subsequently, we manually edited these traces and annotated layer boundaries. These reconstructions did not distinguish between axon and dendrite and contained small scale errors (e.g. neurites passing near each other were sometimes spuriously connected). However, comparison of semi-automated reconstructions with detailed 3D reconstructions (performed manually in Imaris, after imaging with a 60x oil immersion objective and/or a 20x air objective) showed that the semi-automated approach yielded an accurate measurement of neurite density in each layer (Fig. S2a,b).

Chapter 4

Bayesian Methods For Event Analysis of Intracellular Currents

This chapter contains work that was a collaboration with Dr. Josh Merel, at the time a postdoc in Dr. Paninski's group. It was published as:

Merel, J.*, Shababo, B.*, Naka, A., Adesnik, H., Paninski, P. (2016). Bayesian methods for event analysis of intracellular currents. *Journal of Neuroscience Methods*, Volume 268, p 21-32.

4.1 Introduction

Subthreshold neuronal activity provides an unsurpassed richness of information about a single cell's physiological properties. Access to subthreshold activity allows for inference about intrinsic biophysical properties (e.g. membrane and ion channel parameters), circuit level properties (e.g. synaptic connectivity), neural coding (e.g. receptive fields), and synaptic properties (e.g. quantal properties & plasticity). At present, whole-cell patch-clamp stands alone in its ability to reliably access subthreshold activity owing to excellent signal-to-noise ratio (SNR) and very high temporal precision, as opposed to optical subthreshold measurements. At the same time, optical technologies have advanced to the point where we can observe the suprathreshold activity of hundreds of individual neurons simultaneously with calcium imaging and stimulate neurons by subtype or spatial location [146]. However, the limits on temporal resolution and the indirectness of the observations make inferring fine-scale network and cellular parameters difficult. Approaches which combine optical tools with electrophysiology offer unique advantages [150]. In this work we present new statistical techniques useful for analyzing whole-cell data as well as extensions demonstrating how our approach is particularly well-suited to settings where electrophysiology is combined with optical physiology.

4.1.1 Our setting and approach

Fundamentally, many of the subthreshold-based analyses mentioned above depend on the interpretation of the recorded time series as a sequence of *events*. In this setting, events are the successful transmission of neurotransmitter onto the recorded cell, and when this occurs, a transient current flows into or out of the cell, known as a postsynaptic current (PSC). The analyses of experiments designed to infer properties of evoked or spontaneous inputs to a cell (e.g. monosynaptic mapping or quantal/mini-PSC analyses) require determining when a postsynaptic event happened and describing that event. Estimating PSC properties is most straightforward when recordings are acquired using the voltage-clamp configuration which employs a feedback circuit to hold the membrane potential at a constant value thus mitigating variability in PSC properties due to the intrinsic biophysics of the cell (though see [15]).

In this work, we present a Bayesian approach for inferring the timing, strength, and kinetics of postsynaptic currents from voltage-clamp recordings, and we demonstrate on simulated and real data that this method performs better than standard methods for detecting PSCs. The improvement in single-trial accuracy with our method should allow for better estimation of physiological parameters with less data and under more variable conditions (e.g. when the exact timings of stimuli or its effects are unknown). In addition, the quantification of uncertainty over PSC features provided by Bayesian inference enables new experimental designs (e.g. [155]).

Bayesian approaches are naturally extensible, so the intuitive, generative model and straightforward inference procedure flexibly extend to include structure relevant to the analyses mentioned above. Specifically, we extend the core single-trial model to include types of data obtained in monosynaptic mapping experiments which may involve optical stimulation artifacts or combine voltage-clamp recordings and optical recordings. For this latter extension, we combine the single trace model presented in this work with related work on calcium imaging [139] to demonstrate a Bayesian approach to analyzing mapping experiments consisting of simultaneous population calcium imaging and single cell voltage-clamp recordings [1].

4.1.2 Review of other approaches

To our knowledge, all previous methods for inferring PSCs have relied on first inferring the timing of single events (i.e. event onsets), and then sometimes fitting per-event kinetics given that event time. These methods have tended to fall into two categories. The superficially simpler of the two approaches is find events by thresholding the trace or its first derivative (i.e. finite difference). In practice, such methods have extra parameters for smoothing, computing the appropriate offset, or post-processing. Implementations tend to over-detect candidate events and then evaluate candidates based on analysis of per event kinetics [69, 8, 63, 84]. For concreteness, consider a two-stage approach wherein a threshold is used to identify initial candidates, and then a model is fit to the transient dynamics in order to confirm or reject candidate events by comparison of the parameters of the dynamics against pre-determined criteria [8]. Even with post-processing, such methods can be non-selective

and tend not to exploit all of the available information (i.e. the transient dynamics aren't used to detect the events initially).

Threshold methods have been largely superseded by the second class of approaches, template-based methods [33, 132]. In these methods, templates are usually learned by averaging event-responses collected by a simpler method (e.g. thresholding and/or hand-curation). While template methods are straightforward, initial attempts to apply these methods failed when the amplitude of the events varied or where events overlapped - both common scenarios. The first commonly used algorithm for PSC detection that attempted to avoid issues related to amplitude variability introduced the idea of rescaling a fixed template at each time step [33]. Following this trend, template-matching approaches have gradually shifted towards deconvolution methods, which are a more well-founded way to use templates [132]. Deconvolution generally refers to methods that assume the observed trace is the result of convolving a template with unobserved events (of varying amplitude), and such methods invert this model to estimate the times from the template. Both of these template-based methods produce inferred events with different amplitudes and a threshold can then be used to screen out small events (see [52] for a Python implementation of [33] and [132], and see [145] for deconvolution of current clamp traces).

Methods that rely on fixed-shape templates can work very well when the shape of the event is consistent across events, but postsynaptic events can vary in shape and amplitude, especially for events from different pre-synaptic sources due to different dendritic filtering, issues with space-clamp, or different receptor subunit distributions. Indeed, a core rationale behind the initial preference for threshold based approaches was the recognition that events may vary too much for a single template. While it is possible to use approaches that employ multiple templates [92, 160], there are still potential issues related to the stage-wise separation between learning the template and subsequent detection causing a sub-optimal use of information.

We take a Bayesian approach, rooted in a probabilistic, generative model. Broadening the taxonomy, this approach is a type of deconvolution method. However, we do not consider a single template (or a handful of templates), but instead a distribution over templates through the use of prior distributions on the kinetics and amplitudes of individual PSCs. Importantly, we also model event timing in continuous time (i.e. without binning), and we incorporate an autocorrelated, $AR(p)$ noise process [32], which provides a more accurate description of the data. This leads to more precise detection of event times and inference that is more robust (i.e. less susceptible to noise). As such, our inference better leverages all available information (i.e. all events and full timecourse of each event). Given this probabilistic formulation of the noise process and the inclusion of priors on the PSC features, we can then perform posterior inference in this model using Markov chain Monte Carlo (MCMC, see methods).

A tradeoff is that the proposed approach is more computationally intensive than previous approaches. Nevertheless, we believe the flexibility and robustness that this approach affords makes up for this in many settings. Beyond handling overlapping events and variation in the shape of events, our method inherits advantages of probabilistic modelling. The method is extensible and amenable to serving as a modular component of hierarchical models, as we show. Moreover, while existing methods tend to produce all-or-none results and the precise timing of the event is a secondary consideration, using a probabilistic approach, it

is straightforward to consider posterior uncertainty. We essentially get a level of confidence for detection of each event and the level of uncertainty in the precise timing of the event. This posterior uncertainty in event times can translate into posterior uncertainty for other parameters of interest, such as synaptic weights.

4.1.3 Overview

In the following sections, we will first present the details of the model for single-trial voltage-clamp traces and extensions mentioned previously. We then provide the details of the inference scheme we use for sampling from the posterior distribution. In the results, we compare our approach with the standard template-based approach [33] as well as a Wiener Filtering approach [174], since these serve as competitive and robust baselines. Note that while we are not aware of the Wiener filter having been explicitly proposed for this application, the Wiener filter is a standard deconvolution approach (similar to [132]). Unlike the deconvolution approach in [132], the Wiener filter automatically and optimally determines deconvolution parameters from the power spectral density (PSD), and we found this performs better over a range of data than a deconvolution approach with hand-tuned parameters (not shown). We show inference results from our approach on simulated and real data for spontaneous EPSCs and IPSCs across several cell types, PSCs evoked via paired-patching as ground-truth validation, and PSCs evoked optically with one-photon and two-photon stimulation with stimulation artifacts. We also show results on simulated data for a mapping experiment which combines voltage-clamp recordings with calcium imaging, illustrating extensibility.

4.2 Methods

We draw on tools developed in statistics [110] and signal processing [162] to decompose a voltage-clamp recording into interpretable elements. In this application, our events are unitary synaptic currents and their features describe the strength and kinetics of each event. In previous work, we have found similar methods useful for inferring spiking events in calcium imaging data [139].

The framework involves (1) specifying a generative model for voltage-clamp recordings, the parameters of which describe event times, features, the noise model, etc., and (2) performing Bayesian inference on event times and features and the model parameters jointly. Theoretically, Bayesian estimators have nice guarantees (under a “true” model, see [90]). In practice, the generic Bayesian approach can fail if the model is inadequate (e.g. in our case, if the generative model does not capture the true statistics of voltage-clamp traces) or if the inference algorithm performs poorly and the true model posterior is not obtained. Motivated by these legitimate concerns, it is critical to validate that for our application, the model captures the statistics of real voltage-clamp data and that inference performs well on both simulated and real data (see Results).

4.2.1 Model of a single electrophysiological trace

In the simplest version of our model, the observed current trace, y_t , is a discrete time series composed of the sum of n unitary synaptic currents, a baseline (holding current), b , and observation noise ϵ_t (eq. 4.1).

$$y_t = \sum_{i=1}^n a_i f_i(t - t_i) + b + \epsilon_t. \quad (4.1)$$

$$f_i(t) = (e^{-t/\tau_i^d} - e^{-t/\tau_i^r}) \mathbb{1}(t \geq 0). \quad (4.2)$$

$$\epsilon_t = \sum_{j=1}^p \phi_j \epsilon_{t-j} + u_t, \quad u_t \sim \mathcal{N}(0, \sigma^2). \quad (4.3)$$

Each event, indexed by i , is characterized by an event time, $t_i \in \mathbb{R}^+$, which need not be aligned with the sampling time of y_t , its own kinetics determined by $f_i(\cdot)$, and a strength which we define as the amplitude, or peak current, of the event, a_i . For synaptic currents, we use a difference of exponentials for f_i which is parameterized by a rise time constant, τ_i^r , and a decay time constant, τ_i^d . For an example of the model, see Figure 4.1B-D. Observe in Figure 4.1E that recovered kinetics of individual events do vary significantly, suggesting that a model which captures this structure should perform better than methods which rely on a single (or handful of) template(s).

As opposed to an *i.i.d.* Gaussian noise process, the more general autoregressive, $\text{AR}(p)$, process better captures the noise in voltage-clamp recordings. We have found the noise model to be crucial for robust inference (see Results). In an $\text{AR}(p)$ noise model, the noise has temporal correlations due to direct dependencies between noise values for p timesteps. In this work, we use an AR noise model with $p = 2$ (eq. 4.3). Voltage-clamp recordings exhibit correlated noise whose source can be electrical hardware, changes in resistance between the electrode and the interior of the neuron, and other biological non-event contributions to the observation. In practice, it is these forms of temporally correlated noise which lead to many of the false positives since they are more likely to exhibit a similar shape to true events.

With eqs. 4.1 and 4.3 we can write down the likelihood of the observed, noisy data given the parameters ($\Theta \equiv \{\sigma, b, n, \{a_i, t_i, \tau_i^d, \tau_i^r\}_{i=1..n}, \{\phi_j\}_{j=1..p}\}$). We begin with the *i.i.d.*, i.e. $\text{AR}(0)$, case,

$$p(Y|\Theta) = \prod_{t=1}^T (2\pi\sigma^2)^{-1/2} \exp\left[-\frac{1}{2\sigma^2}(y_t - \hat{y}_t)^2\right], \quad (4.4)$$

where \hat{y}_t refer to the predicted noiseless trace:

$$\hat{y}_t = \sum_{i=1}^n a_i f_i(t - t_i) + b. \quad (4.5)$$

In equation 8 of [32], the likelihood is extended to the $\text{AR}(p)$ case

$$p(Y|\Theta) = \prod_{t=1}^T (2\pi\sigma^2)^{-1/2} \exp\left[-\frac{1}{2\sigma^2}(y_t - \hat{y}_{t|t-1})^2\right], \quad (4.6)$$

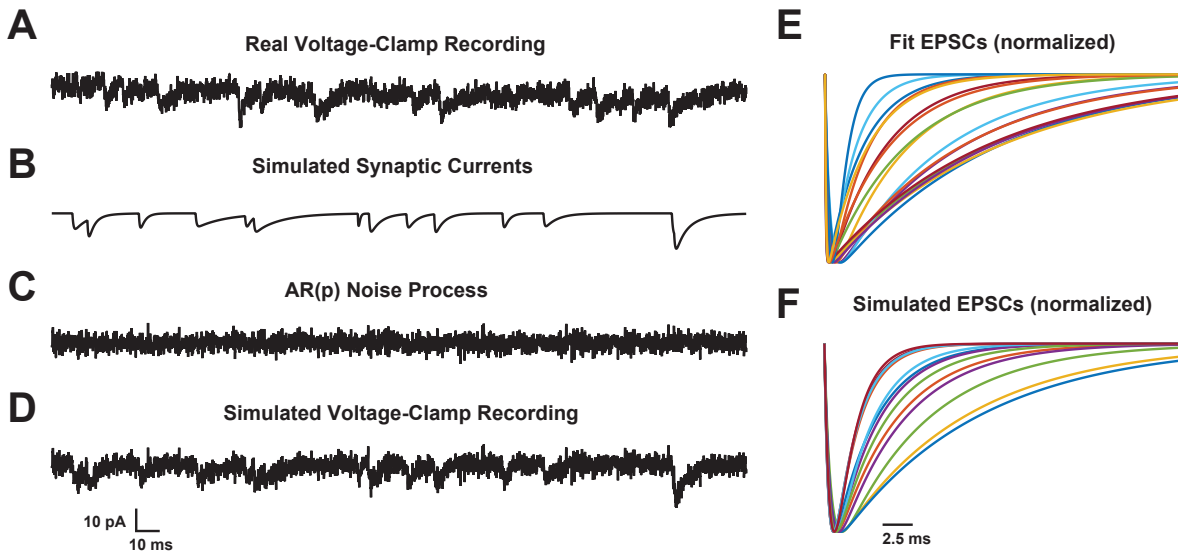


Figure 4.1: This figure depicts the correspondence between real data and the generative model. **A** Shows a real voltage-clamp recording. **B** Depicts simulated synaptic currents generated from a fit to the data in **A** (noiseless). **C** Depicts AR(p) noise process ($p = 2$) generated from a fit to the data in **A**. **D** Illustrates a model-based voltage-clamp recording simulation (the sum of the data in **B** and **C**), to illustrate that the simulation visually captures core features of the real data. **E** Shows individual events estimated from the data in **A** (with normalized amplitudes), and panel **F** shows the individual simulated events used in **B** (with normalized amplitudes).

where $\hat{y}_{t|t-1}$ is (adapted from equation 11 of [32], omitting boundary observations),

$$\hat{y}_{t|t-1} = \hat{y}_t + \sum_{j=1}^p \phi_j (y_{t-j} - \hat{y}_{t-j}). \quad (4.7)$$

The probabilistic model provides a natural objective function:

$$\mathcal{L}(\Theta|Y) \propto \ln p(Y|\Theta) + \ln p(\Theta). \quad (4.8)$$

It is possible to optimize this log-posterior directly, or inference can be performed to obtain an estimate of the posterior distribution. $p(\Theta)$ corresponds to the prior probability on the parameters. In a probabilistic formulation, it is worth explicitly keeping in mind that the posterior distribution for a given parameter can only have support where its prior distribution has support, so hard constraints (e.g. a parameter being positive or a minimum amplitude size) are naturally incorporated as prior information. For amplitudes, baselines, time-constants, and event times, we use non-informative, improper uniform priors over either real or bounded real numbers (we selected very broad ranges appropriate for each parameter). The prior on the event count is given by a Poisson distribution which has one free parameter corresponding to the prior expected number of events – this is a parameter that a user would tune depending on their general expectation about the number of events in their data. The

prior on the noise level σ^2 is a diffuse inverse gamma distribution (which is the conjugate prior), and we use a diffuse prior distribution over ϕ (i.e. a normal distribution subject to stability constraints, same as [32]) – see section 4.2.3 on inference for more details.

Additionally, we note that more sophisticated priors in a Bayesian model serve as natural ways to extend the core model. For example, distributions over event features could be modulated by clustering onto presynaptic sources or the rate of events in each trace could be time-varied based on stimulation of presynaptic cells in mapping experiments. Model extensions are explored in the next section.

4.2.2 Model extensions

Including optical currents

In this extension, we consider an active mapping experiment where we have some level of spatially structured optical stimulation (via optogenetics [44] or neurotransmitter uncaging [27]) of presynaptic cells while holding a postsynaptic cell in voltage-clamp [158, 74]. In this setting, detections of PSCs coincident with stimulations can be used to infer connectivity between neurons. However, in some protocols, the postsynaptic cell also responds to the stimulation¹. If this is the case, we will see a direct optically evoked current in our voltage-clamp recording when we attempt to stimulate cells near the patched cell. In order to remove artifacts of this sort, we can incorporate a parameterized, additive term in the generative model and then perform inference jointly with respect to these parameters. The choice of an additive term is more reasonable for optogenetic stimulation because the currents are carried through different channels whereas with neurotransmitter uncaging the direct stimulation current and the synaptic current may be competing for the same channels.

For example, it is straightforward to include a parameterized kernel for the optical response of the neuron, $h(\cdot)$, and then to convolve that response with the known optical input to the neuron (e.g. the laser or LED power),

$$y_t = \sum_{i=1}^n a_i f_i(t - t_i) + \sum_j^{n_s} a_j h(\theta_h) * d_j(t) + b + \epsilon_t, \quad (4.9)$$

where we have n_s optical inputs each with known timecourse $d_j(t)$ and the response kinetics are modelled with a convolution which is parameterized by θ_h (e.g. a set of time constants). Each stimulation will have its own gain, a_s , which depends on the density of the corresponding channels at the location of that stimulation. We have found this approach to be useful under certain conditions with optogenetics (Figure 4.5D). However, for some optogenetic currents, the multi-state kinetics of opsins can make it difficult to design a parameterized $h(\cdot)$ which can be sampled efficiently. In this case, the shape of the optical current could be measured empirically and then modelled as

$$y_t = \sum_{i=1}^n a_i f_i(t - t_i) + \sum_{j=1}^{n_s} a_j h(t - t_j^{(s)}) + b + \epsilon_t, \quad (4.10)$$

¹This is seen with glutamate uncaging at locations near the patched cell and could also occur with optogenetics when mapping connections between cells of the same transcriptional identity in the same location or when mapping many heterogeneous populations of cells (i.e. when a pan-neuronal promoter is used).

where $h(\cdot)$ is now the stereotyped shape of the current and we know the set of times $\{t_j^{(s)}\}_{j=1\dots n_s}$ at which we have stimulated. Since the currents are filtered in the dendrites, eq. 4.10 breaks down slightly when handling stimulations at locations at varying distances from the soma of the patched cell (Figure 4.5D). Nonetheless, it still provides a good trade off between accuracy and computational tractability. In addition, one could create several optical current templates based on the distance of the stimulation site from the soma [107].

Mapping with calcium imaging and voltage-clamp recordings

Next we extend the model to show how the inferred events could be identified with presynaptic sources when the local population is partially observed via calcium imaging [1] (alternatively, a similar approach could combine information from voltage imaging). In this setting our observations will consist of the fluorescence traces of the imaged cells and the voltage-clamp recording. Both the presynaptic fluorescence traces, c_t^k where k indexes the neurons observed via imaging, and the postsynaptic electrophysiological recording, y_t , can be interpreted as a sum of events [139],

$$c_t^k = \sum_{i=1}^{n_k} a_k^{(c)} g_k(t - t_{ki}) + b_k + \eta_t^k, \quad (4.11)$$

$$y_t = \sum_{k=1}^K \sum_{i=1}^{n_k} a_k^{(y)} f_k(t - t_{ki}) + \sum_{j=1}^{n_\emptyset} a_j^{(0)} f_j(t - t_j) + b + \epsilon_t, \quad (4.12)$$

with ϵ_t as before and $\eta_t^k \sim \mathcal{N}(0, \nu_k^2)$. Like $f_k(\cdot)$, $g_k(\cdot)$ is also a sum of exponentials but the timescale of the kinetics is much slower. $\{t_j\}_{j=1\dots n_\emptyset}$ are the times of PSCs with no corresponding imaged, presynaptic cell. We want to point out that because we model the process in continuous time, the observed calcium traces need not have the same sampling rate or observation times as each other or the voltage-clamp recording (which will be sampled many orders of magnitude faster).

In this demonstration, we have chosen to assume the calcium events for the same cell are all of equal amplitude – this would be appropriate if the spikes occur sparsely or if calcium transients sum roughly linearly. If the summation is known to be nonlinear with some biophysically plausible nonlinearity (e.g. see [172]), this could be straightforwardly accommodated by modifying the model and similar inference methods may still be applied. Alternatively, it is also straightforward for the calcium event amplitudes to vary across events.

We have similarly chosen for all postsynaptic events to be of the same amplitude and shape when they follow from a specific presynaptic cell, but in practice one can extend the clustering to impose cell-specific priors on these amplitudes and shapes. Some subset of events observed in the electrophysiological recording will arise from unobserved presynaptic inputs so we allow these to be explained by events with no observed presynaptic cause and with independent shape and amplitude per event (second summation term indexed $i = 1\dots n_\emptyset$). We could also incorporate a fixed delay between presynaptic events and postsynaptic events (this would be an additional parameter, pre-specified or inferred, in the postsynaptic transient response $f_k(\cdot)$).

Since we now observe multiple traces, our new objective function is simply the sum of the log probabilities of the various traces. The multiple traces are conditionally independent given event times $\{t_{ki}\}_{k=1\dots K, i=1\dots n_k}$ and $\{t_j\}_{j=1\dots n_\emptyset}$, so we have an overall objective corresponding to the log-posterior:

$$\mathcal{L}(\Theta|C^k, Y) \propto \ln p(Y|\Theta) + \sum_{k=1}^K \ln p(C^k|\Theta) + \ln p(\Theta), \quad (4.13)$$

again with $p(\Theta)$ corresponding to the prior probability on the full set of parameters (see Section 4.3.4 for results).

4.2.3 Inference

For this work, we perform inference using Markov chain Monte Carlo (MCMC) [116, 49]. MCMC techniques allow us to obtain samples from the posterior distribution over all unknown variables in the model and thereby approximate the posterior by a histogram of such samples. Specifically, we perform Gibbs sampling over all the parameters [49]. This means that for each parameter, we hold all other parameters fixed and conditionally update the focal parameter by sampling it from its conditional distribution. A ‘‘sweep’’ consists of an update of all parameters. While it is possible to compute conditional distributions analytically for sufficiently simple models, it is quite simple to use generic sampling methods to update parameters for any model for which one can compute the likelihood. For example, we use random-walk Metropolis (RWM) to update many parameters – this consists of updating parameters by proposing updates from a distribution centered on the current value and accepting or rejecting proposed updates such that the resulting set of samples are consistent with the conditional distribution [49]. Alternatively, more powerful samplers such as Hamiltonian Monte Carlo could be used [49] but simple RWM sufficed here.

In the single-trial, voltage-clamp case, we update $\{t_i, a_i, b, \tau_i^d, \tau_i^r\}_{i=1\dots n}$ by RWM. Inclusion of a direct optical current, in the simplest case, only contributes one additional amplitude parameter for each stimulation, and these can be inferred similarly. Through sampling, we also estimate the posterior distribution of the number of events in each trace, given a Poisson prior (note that this could be generalized to an inhomogenous Poisson process prior so that the event rate can vary based on inputs such as optical stimulation). To add and remove events (i.e. perform inference over n), we use birth-death moves which consist of the proposal of a new event time and the removal of an existing event time respectively [110].

The noise process parameters $\phi_{1..p}$ are sampled by rejection sampling from the constrained conditional distribution and σ is sampled from its conditional distribution. Specifically, we reproduce the updates for the $\phi_{1..p}$ and σ , which are provided in section 4.1 of [32]. $\phi_{1..p}$ is shown to be conditionally normal with a mean and posterior that depend on $\hat{e}_t = y_t - \hat{y}_t$ and σ^2 , but also with the constraint that the AR(p) process defined by $\phi_{1..p}$ is stable (note that stability here means that the AR process remains bounded on bounded intervals if run autonomously, which can be quickly checked by examining the magnitude of the roots of the

associated characteristic polynomial [86]). Omitting boundary observations:

$$E = [\hat{e}_{t-1} \dots \hat{e}_{t-p}], \quad (4.14)$$

$$\Phi_n = \Phi_0 + \sigma^{-2}(E'E), \quad (4.15)$$

$$\hat{\phi} = \Phi_n^{-1}(\Phi_0\phi_0 + \sigma^{-2}E'e), \quad (4.16)$$

$$\phi \sim \mathcal{N}(\hat{\phi}, \Phi_n^{-1}) \mathbb{1}_{S\phi}, \quad (4.17)$$

where $\mathbb{1}_{S\phi}$ is an indicator over the set of stable ϕ values and ϕ_0 and Φ_0 are set to weakly informative prior values. Following [32], we place an inverse-gamma prior on σ^2 , which is the conjugate prior given the likelihood (eq. 4.6) so that we can sample directly from the conditional posterior,

$$\delta_1 = \sum_{t=1}^T (y_t - \hat{y}_{t|t-1})^2, \quad (4.18)$$

$$\sigma^2 \sim \mathcal{IG}\left(\frac{1}{2}(T + \nu_0), 1/\left(\frac{1}{2}(\delta_1 + \delta_0)\right)\right), \quad (4.19)$$

where ν_0 and δ_0 are also set to weakly informative prior values.

For the full mapping case, inference is similar. However, additional moves need to be incorporated to improve mixing. That is, in addition to the single-variable updates, additional custom moves are performed each sweep. The additional moves correspond to proposing a swap for the presynaptic identity for a PSC event. That is, an MCMC step is proposed wherein an event associated with one presynaptic source is eliminated and an event at the same time is considered for another source. To implement such moves, we simply propose to drop an event at a given time for one presynaptic source and to add an event at the same time for another presynaptic source and evaluate the combined acceptance or rejection of these proposed moves using the Metropolis-Hastings ratio [49].

4.2.4 Specific implementation details

Code implementing our inference routine is available on github². When applying our inference method to data, we can run from a cold start, or we can initialize event times with those found via a simpler, faster method (e.g. [33] or [132]). In this work, to establish initial performance of our algorithm, we present results on simulated and real data using cold starts. MCMC methods ideally will be run long enough that a proper posterior distribution can be estimated. A standard convergence check for MCMC methods is to run multiple chains and assess that they provide similar estimates. We did not run multiple Markov chains in this case for all analyses. Rather, for ease of use, we followed a simple heuristic of initially tuning the total number of sweeps by running multiple Markov chains and then consistently performing a conservatively large number of sweeps to robustly achieve good performance (i.e. as measured relative to other methods).

For applications, it is pragmatic to adjust the total number of samples to trade off reliability (i.e. high probability of convergence) against computation time. For small timeseries

²https://github.com/jsmerel/joint_calcium_ephys_mapping

(roughly 1 second), inference consisting of 2000 sweeps of the sampler tends to provide an accurate posterior (see Results, simulated data) and requires anywhere from a few seconds to a few minutes on a contemporary computer depending on the number of inferred events. It is standard to discard the first fraction of the sweeps as *burn-in* (we discard between one-fifth to two-fifths of the total number of sweeps, depending on the event rate and its prior). For longer timeseries or many timeseries, we run smaller sections of the timeseries in parallel on a computing cluster (on a single CPU, the computation time would scale linearly with the number of requested samples and scale linearly in the number of inferred events, which is assumed to be proportional to the length of the trace). We also note that RWM requires a proposal width hyperparameter – this can be automatically tuned early in the sampling process by adjusting the proposal variance such that the accept rate is reasonable [147]. We perform this automatic tuning in an *ad hoc* adaptive fashion, incrementally increasing or decreasing the proposal variance when either too many or too few moves are being accepted.

We also note that in the implementation, we frequently rely on the log-likelihood, $\ln(p(Y|\Theta))$, to determine whether to accept or reject moves. In order to minimize redundant operations, it proves useful to store $\hat{e}_t = y_t - \hat{y}_t$ and we perform evaluations of $y_t - \hat{y}_{t-1}$ by observing that:

$$y_t - \hat{y}_{t-1} = y_t - (\hat{y}_t + \sum_{j=1}^p \phi_j(y_{t-j} - \hat{y}_{t-j})) \quad (4.20)$$

$$= y_t - \hat{y}_t - \sum_{j=1}^p \phi_j(y_{t-j} - \hat{y}_{t-j}) \quad (4.21)$$

$$= \hat{e}_t - \sum_{j=1}^p \phi_j \hat{e}_{t-j}. \quad (4.22)$$

Finally, note that key hyperparameter settings are presented in Supplementary Table 1.

4.2.5 Experimental methods

All experiments were performed in accordance with the guidelines and regulations of the Animal Care and Use Committee of the University of California, Berkeley. Mice used for experiments in this paper were either wild type (ICR white strain, Charles River), som-IRES-Cre (JAX stock #018973); Ai9 Rosa-LSL-tdTomato (JAX stock #007909), PV-Cre (JAX stock #008069); Ai9 Rosa-LSL-tdTomato, or emx1-IRES-Cre (JAX stock #005628).

Viral Infection: Neonatal emx1-cre mice were injected with AAV9-CAG-flexed-Chr2-tdTomato (for 1-photon experiments) or AAV9-syn-ChrimsonR-tdTomato (for 2-photon experiments) at P0-P4. Viruses were acquired from the University of Pennsylvania Vector Core. Undiluted viral aliquots were loaded into a Drummond Nanoject injector. Neonates were briefly cryo-anesthetized and placed in a head mold. With respect to the lambda suture coordinates for S1 were: 2.0 mm AP; 3.0 mm L; 0.3 mm DV.

2-Photon Optogenetic Stimulation: 1040 nm light (femtoTrain, Spectra-Physics) was delivered to the sample using a VIVO 2-Photon workstation (3i) based on a Sutter Moveable Objective Microscope (Sutter, Novato, CA) and the hologram was created using a Phasor 2-

Photon computer-generated holography system (3i). Light was delivered for 10 milliseconds at 100 mW power on sample. The hologram for these data was a disc of radius 15 μm .

For more detailed methods on brain slicing, *in vitro* and *in vivo* electrophysiology, and 1-photon optogenetic stimulation, see [138]. All electrophysiology was analyzed at 20 kHz, or equivalently, with timebins of 0.05 ms.

4.3 Results

4.3.1 AR noise model validation

We have found the choice of noise model to be critical when analyzing voltage-clamp data. It is common to use *i.i.d* Gaussian noise for neurophysiological time-series ([127, 145]), but very often the noise can exhibit temporal correlations. To obtain a recording of a voltage-clamp noise process which contains no events, we recorded from a neuron exposed to an excitatory synaptic blocker (Kynurenic Acid, 4mM) while holding the cell near the inhibitory reversal potential (-70 mV). Under these conditions, the recording should be relatively event free; nonetheless the recording shows clear temporal correlations (Figure 4.2A). In the context of deconvolving these data, it is primarily this correlated noise that drives false positives because it can have similar features to PSCs.

To better capture the structure of voltage-clamp noise, we used the more general AR process which we found was sufficiently flexible and expressive to represent the types of noise we encountered. Specifically, an AR(2) model balanced model expressiveness and computational cost. Figures 4.2A and 4.2B show that the extra structure in the AR(2) model does in fact provide a better description of the data.

For a systematic validation, we performed a comparison between the AR(0) and AR(2) inference for many simulated traces (10 x 1s traces per noise level). Specifically, we can simulate traces with random event times and with various levels of AR(2) noise added to the traces (i.e., varying SNR). In these simulations, events are naturalistic in that they have variability in their distribution of amplitudes and time constants, and events may overlap (4.2C). For inference with either noise model, performance is similar when there is low levels of noise (or for specific AR(2) noise process parameters that result in only weak noise autocorrelation, not shown), but the inference results diverge dramatically when the noise is larger in magnitude. To summarize inference, we examine the correlation between the posterior mean trace and the “true” simulated, noiseless event trace (4.2D). For biologically realistic AR(2) noise structure and magnitude (indicated by an asterisk in 4.2D), inference with the AR(2) model indeed performs considerably better than with the AR(0) model. Note that in the high-SNR limit, simple methods like template matching algorithms, greedy optimization, or other direct optimization of the model likelihood can perform well enough, so sample-based inference would be computationally excessive. However, this validation demonstrates that in the biologically realistic noise regime, noise is sufficiently large and structured for proper inference to be useful.

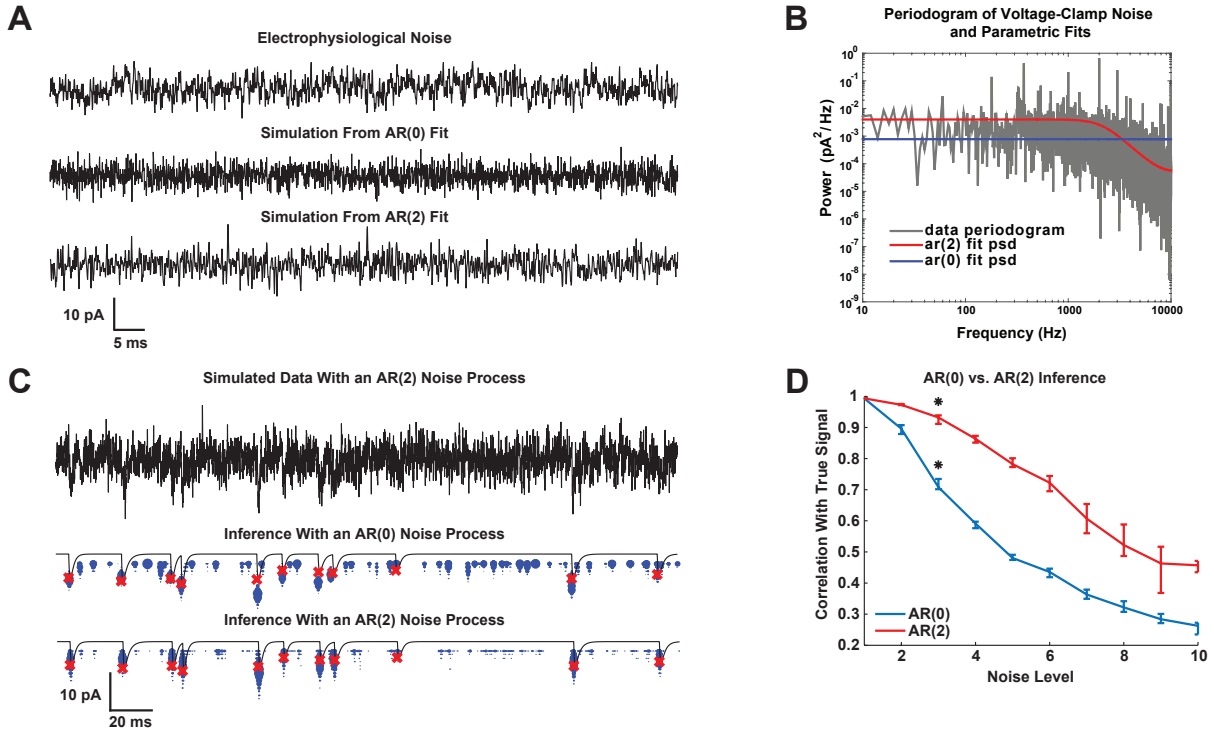


Figure 4.2: Validation of the $AR(p)$ noise model and inference. **A** Top: a real voltage-clamp recording under “event-free” conditions such that the recording is dominated by noise. Middle: simulated noise from an $AR(0)$ fit to the real example. Bottom: simulated noise from an $AR(2)$ fit to the real example. The $AR(2)$ example better captures the structure of real noise. **B** Grey: periodogram estimate of the power spectral density (PSD) of an “event-free” recording. Blue: a parametric fit to the PSD with an $AR(0)$ model. Red: a parametric fit to the PSD with an $AR(2)$ model. **C** Top: a simulated voltage-clamp recording with $AR(2)$ noise. Middle, black: The true current for the top trace. Red: The true time-amplitude coordinate for each PSC. Blue: A bivariate histogram reflecting the estimated time-amplitude posterior distribution using an $AR(0)$ noise model (uses a small, but non-zero minimum event size threshold). Bottom: same as in the middle trace but inference is performed with an $AR(2)$ noise model. **D** Curves depict accuracy of inference as a function of SNR level (ranging an order of magnitude, with 1 indicating low noise relative to size of events and 10 indicating noise that has marginal variance larger than the signal) for $AR(0)$ vs $AR(2)$ model-based inference. The measure of accuracy is correlation coefficient between true (simulated) trace and estimate of posterior mean. The asterisk indicates a biologically realistic SNR level equal to the example in **C**. For each point on the curve, we simulate timeseries with random event-times and amplitudes (traces are median and inter-quartile range over 10 repeats). Both algorithms do very well in the high SNR regime. As soon as noise level begins to make inference difficult, both algorithms begin to lose accuracy. However the $AR(2)$ model inference degrades much more slowly.

4.3.2 Comparison with other methods on spontaneous and evoked PSCs

We also compared our method to two common PSC detection algorithms: a standard template-based approach [33] and a deconvolution approach (a Wiener Filter [174], which we found tended to improve upon the slightly simpler inverse filter used in [132]). Although

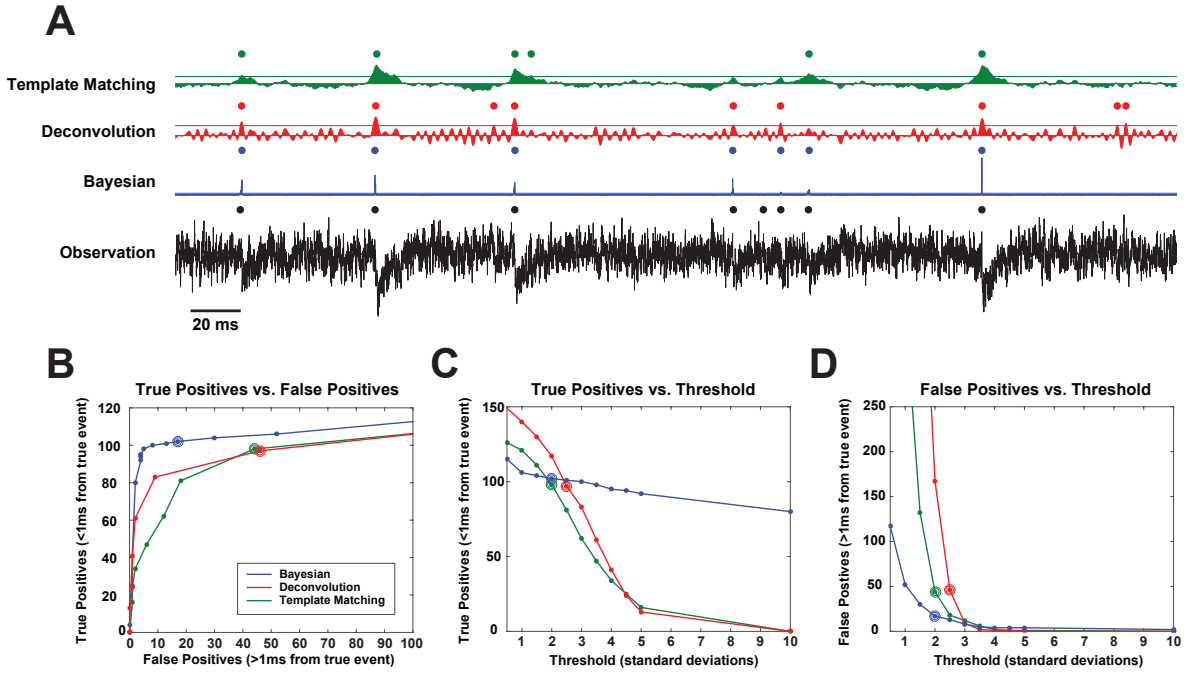


Figure 4.3: Comparison of methods on simulated data. **A** An example of inference/detection results for each method on a simulated voltage-clamp recording. For each method we show the output time-series for that method (i.e. the score template matching, deconvolved trace for deconvolution, and the posterior of event times for the Bayesian approach), estimated event times, and the threshold used to determine those event times (for the Bayesian trace the threshold is so low that it can't be seen above the baseline posterior). **B** Parametric curve showing the number of true positives and false positives as a function of peak detection threshold. The results are from detection across 10 simulated traces with a total of 180 events ranging from 0.5 to 10 pA. The highlighted point in each line corresponds the threshold used in **A**. **C** Same as in **B** except showing the true positive count as a function of the threshold. **D** same as in **C** except showing the false positive count.

[33] is nearly twenty years old and is often outperformed by deconvolution approaches, it is implemented in two commonly used software packages for the analysis of electrophysiological data: Axograph and pClamp. Therefore it is still regularly used when PSC detection is performed [99, 106]. When testing these algorithms, we attempted to give each its best chance to perform well. Specifically, that means that the template-based and deconvolution methods were given the average of the true underlying events as a template. We gave the Wiener Filter the true noise power spectral density for simulated data. For real data with high enough spontaneous rates, it was difficult to find a “quiet” section of the trace to estimate the noise power spectral density (PSD); therefore we provided an AR fit to the noise from Bayesian inference for the Wiener Filter’s noise PSD. For simulated data with Bayesian inference, we provided the true priors on τ^r and τ^d that were used to generate the data.

First, we simulated a test set of recordings with realistic levels of noise and with relatively low SNR PSCs. Each of the three algorithms produces a time series equal in length to the input recording that is something like a score that an event is happening at that point in

time. For the template-based method the output represents the goodness-of-fit to the event template at that point in time, and for the deconvolution the output is an estimate of the event amplitude at that point in time. For the Bayesian approach the output time series is the marginal posterior of an event at each sample. Importantly, in practice the output for the template-matching and deconvolution methods only have an empirically reasonable interpretation in the high SNR case. To obtain estimates of event times for each simulated trace for each method, we detected peaks above some threshold in these time series as a function of their standard deviations. Figure 4.3A shows an example of results for each algorithm on simulated data. The threshold for each method’s output is shown as a horizontal line, and inferred events are shown as dots. (The threshold for the Bayesian approach is difficult to see as it is very close to zero.)

By varying the threshold and counting the number of true positives and false positives, we can get a sense of how noisy the output is from each approach. The Bayesian approach is able to accurately detect more events while accumulating fewer false positives than the other methods (Figure 4.3B). Importantly, it also maintains similar levels of true and false positives as the threshold is varied (Figures 4.3C & 4.3D), indicating that our approach is more robust to the choice of threshold. This is achieved because of the more precise and less noisy representation of PSC timing produced by the Bayesian method which can be seen in the sharpness of the peaks in the Bayesian trace in Figure 4.3A as compared to the other methods. The average error for true positives was also smallest for the Bayesian method (see Supplemental Figure 1). While there exists a threshold for the template-matching and deconvolution methods that performs quite well, small deviations from this value lead to vastly more false positives or less true positives (Figures 4.3C & D). When applying these methods to real data, it may not be possible to finely tune the threshold parameter on a per dataset basis since ground truth information is not available. This can be especially troublesome when online or closed-loop analysis is desired.

We next compared methods on real voltage-clamp recordings in which we could modulate the SNR of an event physiologically. To achieve this, we made paired patch recordings of a PV+ cell and a layer V pyramidal cell until we found an inhibitory connection from the fast-spiking cell onto the pyramidal cell with a probability of a postsynaptic event extremely close to 1.0. We then moved the holding potential for the postsynaptic cell towards the reversal potential for the inhibitory current. In this way, we could obtain ground truth data in which we had direct control over the SNR (Figure 4.4A, B, & C, top traces).

We ran 50 trials each at three holding potentials representing relatively high, medium, and low SNR regimes and detected events using all three methods. Only the Bayesian approach was sensitive enough to detect events reliably in the low SNR case (Figure 4.4A, B, & C, rasters). Similar to the simulated data, the Bayesian approach was also the only method which was robust to the thresholding parameter, indicating that the posterior over event times is highly peaked. As expected, as the SNR decreases, the ability to accurately detect the timing of the event decreases.

Any approach will have hyperparameters that must be selected, and the Bayesian approach allows for tuning of hyperparameters corresponding to prior distributions on model parameters. We show that a single set of prior distribution hyperparameters can perform well across several physiological regimes by running inference on traces from different cell types and under different recording conditions while holding the hyperparameters constant.

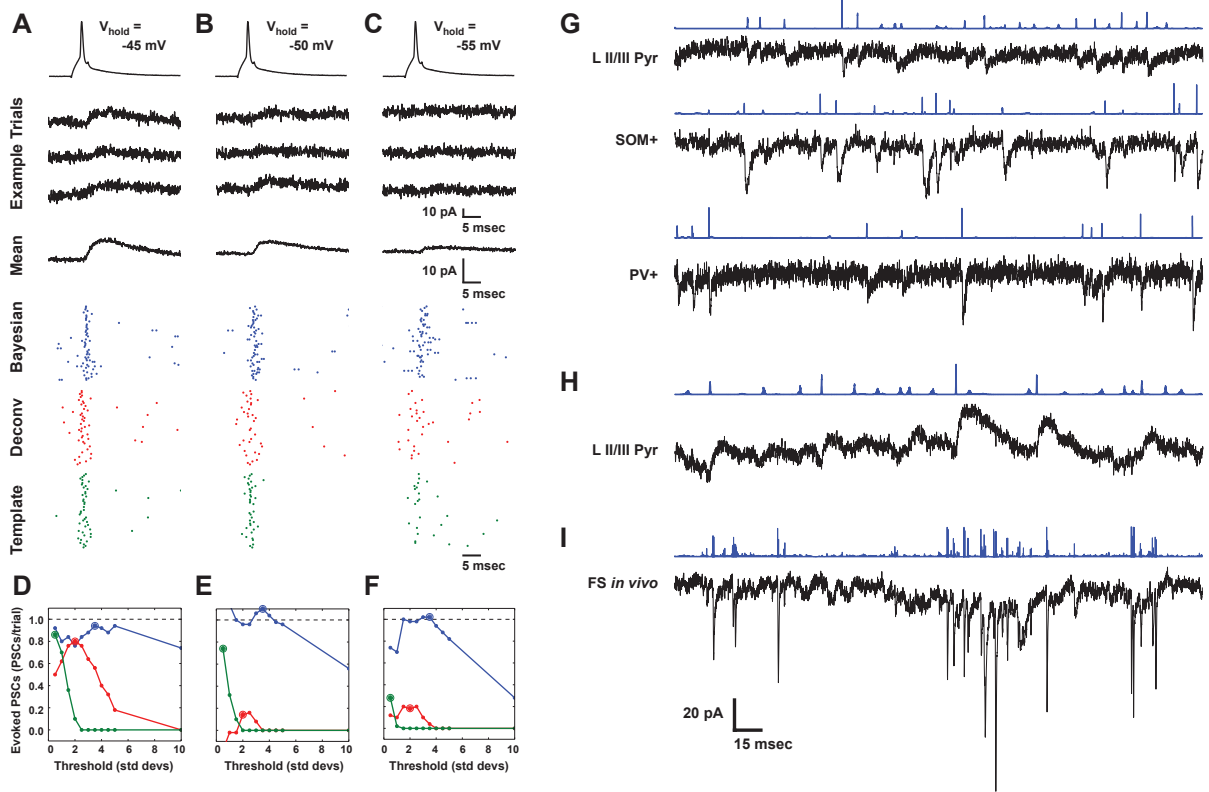


Figure 4.4: Results on real voltage-clamp recordings. **A-F** Comparison to other detection methods on real data while modulating the SNR. **A** Top: a single trial of the spiking presynaptic cell firing from direct current injection. Middle traces: three example trials of the connected postsynaptic cell responding with an IPSC. The postsynaptic cell is being clamped at -45 mV . Below the examples is the mean trace over all 50 trials. Bottom rasters: estimated events for 50 trials for each detection method. **B-C** Same as in **A** except with the postsynaptic cell held at -50 mV (**B**) and -55 mV (**C**). The scales for the example traces are all the same across **A-C** and likewise for the mean traces. **D-F** The average number of evoked events counted in a window around the spike time (1.0 msec before to 5.0 msec after the spike) above a baseline rate of detected events per trial as a function of threshold for each method. Colors and symbols as in **2C-B**. The dotted horizontal line represents perfect performance of one extra event detected per trial. **G-I** Examples of results on several types of real data. **G** Spontaneous EPSCs detected in a layer II/III pyramidal cell, top, a SOM+ cell, middle, and a PV+ cell, bottom. **H** Spontaneous IPSCs in a layer II/III pyramidal cell. **I** Spontaneous EPSCs and IPSCs detected in an *in vivo* recording from a FS cell.

In 4.4**G** we show that with a single set of prior parameters our method can detect EPSCs across three different cell types: a layer II/III pyramidal cell, a SOM+ interneuron, and PV+ interneuron. Despite the differing statistics in each of these traces (event features and noise), event detection performs well. In 4.4**H** we show results for spontaneous IPSCs in a layer II/III pyramidal cell. For these results, we used the same prior parameters as in 4.4**G** except that we increased the adjusted the bounds for the time constants to account for the slower kinetics of IPSCs. Finally, in 4.4**I** we show results on an FS cell recorded *in vivo*. The prior parameter settings here are similar to those in 4.4**G** except that we had to increase the rate prior on the number of events. We stress that there is no correct prior setting, but

that the priors allow the user to trade off between objectives like true versus false positives (e.g., through the rate parameter or minimum event amplitude). For reference, we include the hyperparameter settings for all inference results in Supplementary Table 1.

4.3.3 Extension 1: Direct optical stimulation artifact

For certain experiments, it may be productive to actively drive cells to fire in order to study circuit properties. In particular, we may want to combine voltage-clamp recordings with spatiotemporally structured optical stimulation of a putative presynaptic population of neurons. This can induce optically evoked artifacts from direct optical currents in the voltage-clamp recording. As an example, one could express an optogenetic channel pan-neuronally which would allow the simultaneous mapping EPSCs and IPSCs onto one cell. Under these conditions, the cell under voltage-clamp will also express opsin and respond to any stimulating light. Similarly, in neurotransmitter uncaging mapping experiments, there will be a direct stimulation of the postsynaptic cell at most stimulation sites close to the cell.

We show that our approach is able to decompose a simulated trace with a direct optical current in Figure 4.5. Specifically, we simulated a trace consisting of many EPSCs with an additive direct optical stimulation current, consistent with eqn. 4.10. On this simulated data, inference performs very well in terms of extracting events simultaneously with artifact isolation.

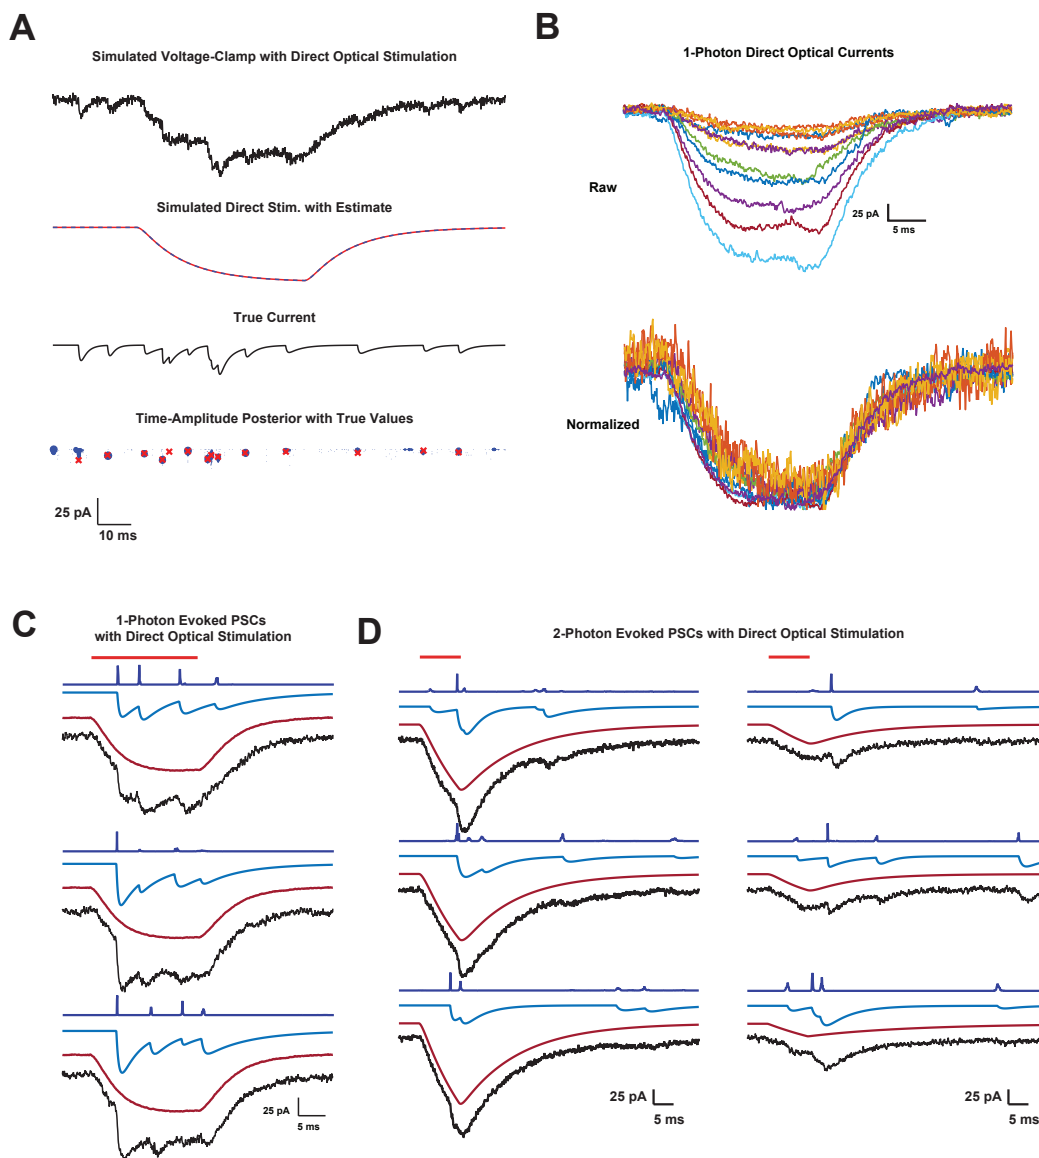


Figure 4.5: Removing direct stimulation artifacts from active mapping data. **A** Top: a simulated voltage-clamp recording that is contaminated by a large direct, optical current. Top middle: the true direct current (blue) with an estimate of the inferred direct current (dashed, red). Bottom middle: the true current used in the top trace. Bottom, the inferred amplitude-time posterior (blue) with true amplitude-time coordinates overlaid (red X's). **B** Top: direct optical currents evoked in a single cell by stimulating different locations with one-photon excitation and a DMD. Bottom: same as above but each trace is normalized to have the same amplitude. **C-D** PSC detection with direct optical stimulation on real mapping data. **C** Inference results for one-photon, DMD-based mapping data with direct stimulation contamination. All three trials are for a single-photon, DMD-based mapping site. Red line shows when the stimulating laser is on. Dark blue trace shows the posterior over event times. Light blue trace shows the MAP estimate for the synaptic currents. Maroon trace, MAP estimate for the direct stimulation. Black, raw voltage-clamp observation. **D** Inference results for two-photon, SLM-based mapping data with direct stimulation contamination. Left, three trials from a particular stimulation location which shows putative evoked EPSCs riding on top of direct stimulation with inference. Right, same as the results to the left but for a stimulation location further from the cell.

In Figure 4.5B-D, we present results on real data obtained by combining voltage-clamp recording with one-photon stimulation via a digital micromirror device (DMD) as well as two-photon, holographic stimulation of neurons. In 4.5B, we show direct currents at many different stimulation locations with one-photon excitation via a DMD. Under these conditions, the shape of the current is not well characterized by a parameterized template function so we must use an empirically derived template. When normalized, all of the currents have roughly similar shapes, validating our approximation of the artifact as a scaled template. In 4.5C we show PSC inference on three trials from a single stimulation location that has putative evoked PSCs overlapping the direct stimulation artifact. In Figure 4.5D, we show similar results except the stimulation is performed using two-photon excitation with a spatial-light modulator (SLM). We found that under these conditions, we were able to use the model in eqn. 4.9 and fit two time constants to the optical current. This allowed for a more flexible fit that could account for the differences in the direct current shape as the stimulation location varied. The two sets of traces and results in 4.5D are for three trials at two different locations. These locations were chosen because they contained putative evoked EPSCs.

4.3.4 Extension 2: Passive mapping experiment (simulation)

In addition to electrical recordings from a post-synaptic cell, we may also have calcium indicator available in pre-synaptic cells [1]. This allows for a passive mapping experiment from many pre-synaptic cells (optically imaged) to a single post-synaptic cell (patched). Here we provide an example of the usage of the joint inference procedure (Figure 4.6A). We have simulated a population of presynaptic cells which are observed via calcium imaging (SNR for these cells is plausible for a well-tuned setting, and bin size is 35ms). These cells drive post-synaptic events in the simulated patched cell (noise is biologically plausible, and bin size is .05ms). In 4.6A, 4 of 6 candidate presynaptic neurons are observed. Postsynaptic events evoked by unobserved neurons are inferred on a per event basis whereas events co-occurring with presynaptic events arise from a single variable (see model eqns. 4.12).

We see that we can recover events jointly from the calcium imaging and electrophysiology, with event identity successfully linked across the two modalities. In addition, the histogram in the right panel of Figure 4.6B indicates that combining electrophysiology and calcium imaging tends to yield increased temporal precision of event times compared to inference from calcium imaging alone (this intuitively follows from the fact that the electrophysiology has much higher temporal resolution and the Bayesian inference can combine information across modalities). This simulation serves as a proof-of-concept that this probabilistic approach can provide meaningful automatic analysis for passive mapping experiments.

4.4 Discussion

In this work we have presented a probabilistic formulation of the event detection problem for electrophysiological recordings with an emphasis on PSC detection. This method works on simulated and real data and its performance compares favorably relative to existing methods. We have also shown how the probabilistic model can be incorporated into two

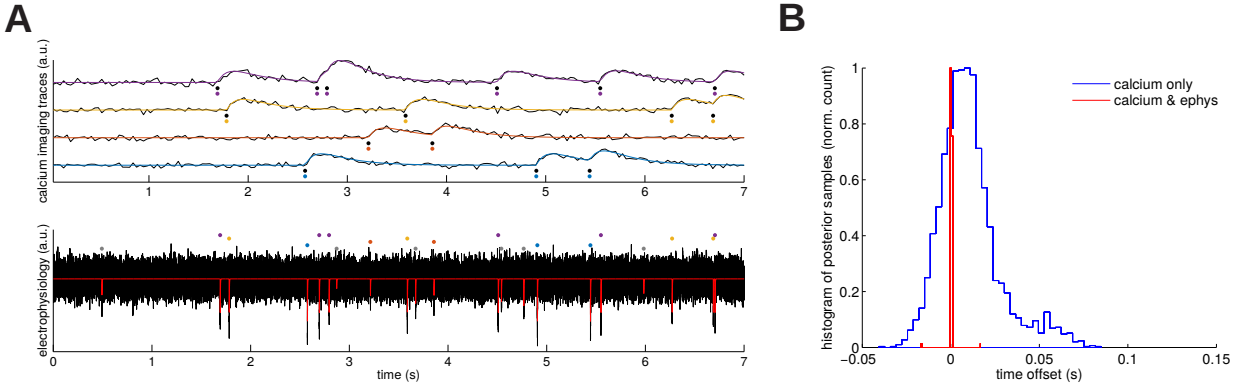


Figure 4.6: Joint inference over simulated calcium and voltage data. **A** Observed data and true/inferred event times for four simultaneously observed calcium traces (top) and (for simplicity) just a single voltage trace (bottom). Colored dots indicate to which trace the algorithm matched each event; gray dots in bottom correspond to voltage events that were detected but (correctly) not matched to any detected calcium events. In this example the algorithm correctly matched all calcium-observed events. **B** Histograms of temporal error in estimated event times. Given calcium only, the variance of inferred event times is large; incorporating voltage information drastically reduces this variance, indicating significant potential gains in temporal resolution from this Bayesian data fusion approach.

extensions applicable for mapping experiments which make use either of optical stimulation or multi-modal fusion of electrophysiology with calcium imaging.

The PSC detection problem explored in this paper is similar to the more general set of problems in neuroscience (and signal processing more broadly) having to do with inferring events in noisy timeseries. In neuroscience, recent work has proposed probabilistic, MCMC-based tools for analysis of calcium imaging data [172, 139]. However, we are not aware of similar tools yet being leveraged for analysis of electrophysiology data. The models developed in this paper are formally similar to those used in the calcium imaging setting, with appropriate modifications to allow for structured AR noise and distinct per event kinetics, as well as detailing the extensions related to the mapping setting.

While there are other discriminative methods that have been proposed for deconvolution (e.g. see [166] for alternative event-detection methods for calcium imaging), these approaches lack the clear probabilistic generative semantics which facilitate modularity and extension to hierarchical models such as those useful for mapping experiments. Our approach also complements related work inferring synaptic inputs in a probabilistic fashion using particle filtering [127] from voltage traces. The present approach focuses on current traces, allows for per event kinetics, is perhaps simpler to implement, and does not require temporal discretization.

The inference approach presented in this work makes most sense when PSC rates are relatively low such that overlap is limited. For very high levels of overlap (i.e. the high-rate case), we do not expect there is always hope to resolve precise timing of single events, and it is likely that approaches more like [171] or [127] may be appropriate. That said, we have found that our method can still give sensible results even when event rates are high (for example, see Figure 4.4I).

Aside from the extensions we considered, other sophisticated prior structure can be incorporated by making the model hierarchical. For example, in the single trace setting we might expect the events to cluster by pre-synaptic cell identity or cell type. Even without the additional observed traces employed in our mapping model, it would be conceivable to specify and infer latent source clusters based on structure in the distribution of shape or amplitude of post-synaptic events, depending on what were of most interest scientifically. In the simplest case, this would result in a mixture model for the events, with each event associated with a latent pre-synaptic source. We plan to pursue these model extensions in future work.

4.5 Supplemental Material

4.5.1 Supplemental Table 1: Inference Hyperparameters

	a_{min} (pA)	a_{max} (pA)	τ_{min}^r (msec)	τ_{max}^r (msec)	τ_{min}^d (msec)	τ_{max}^d (msec)	λ_n events/sec	# Gibbs Sweeps
Figure 2C-D	1.5	Inf	0.25	2.50	0.25	2.50	0.0013	1000
Figure 3	0.01	Inf	0.05	1.00	0.50	10.00	2.00	1000
Figure 4A-C	-Inf	Inf	0.25	3.00	1.00	30.00	0.02	2000
Figure 4G LII/III Pyr	0.5	Inf	0.25	1.50	1.00	5.00	2.00	2000
Figure 4G SOM+	0.5	Inf	0.25	1.50	1.00	5.00	2.00	2000
Figure 4G PV+	0.5	Inf	0.25	1.50	1.00	5.00	2.00	2000
Figure 4H	0.5	Inf	1.00	3.00	5.00	30.00	2.00	2000
Figure 4I	-Inf	Inf	0.05	1.00	5.00	7.50	2000	2000
Figure 5A	5.0	Inf	0.05	0.50	0.50	10.00	200	2000
Figure 5C	5.0	Inf	0.05	0.50	0.50	10.00	0.20	2000
Figure 5D	2.5	Inf	0.10	1.00	1.00	8.00	0.20	3000

4.5.2 Supplemental Figure 1: Timing Error Distributions

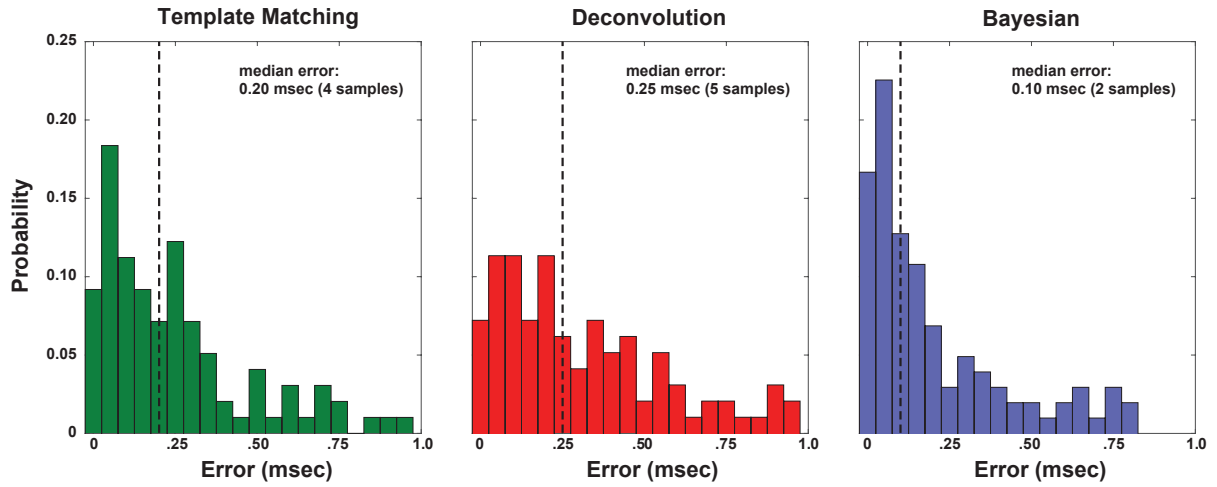


Figure 4.7: The distributions of timing errors from the results corresponding to the highlighted points in Figures 3B-D. The histograms have been normalized such that the areas for each histogram sum to one. The vertical dashed line shows the median value for each distribution. Although the Bayesian method models the timing of events in continuous time, the resolution of detection for all methods was set to the sampling resolution of the data (20 kHz) as a result of the peak detection process.

Bibliography

- [1] Gloster Aaron and Rafael Yuste. “Reverse optical probing (ROPING) of neocortical circuits”. In: *Synapse* 60.6 (2006), p. 437.
- [2] Hillel Adesnik and Massimo Scanziani. “Lateral competition for cortical space by layer-specific horizontal circuits.” In: *Nature* 464 (7292 Apr. 2010), pp. 1155–60. ISSN: 1476-4687. DOI: 10.1038/nature08935. URL: <http://dx.doi.org/10.1038/nature08935>%20<http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2908490&tool=pmcentrez&rendertype=abstract>.
- [3] Hillel Adesnik et al. “A neural circuit for spatial summation in visual cortex.” In: *Nature* 490 (7419 2012), pp. 226–31. ISSN: 1476-4687. DOI: 10.1038/nature11526. URL: <http://www.ncbi.nlm.nih.gov/pubmed/23060193>.
- [4] a. Amarasingham et al. “Conditional modeling and the jitter method of spike resampling”. In: *Journal of Neurophysiology* 107 (2 2012), pp. 517–531. ISSN: 0022-3077. DOI: 10.1152/jn.00633.2011. URL: <http://jn.physiology.org/cgi/doi/10.1152/jn.00633.2011>.
- [5] Paul G Anastasiades et al. “GABAergic interneurons form transient layer-specific circuits in early postnatal neocortex.” In: *Nature communications* 7 (2016), p. 10584. ISSN: 2041-1723. DOI: 10.1038/ncomms10584. URL: <http://www.nature.com/ncomms/2016/160204/ncomms10584/full/ncomms10584.html>.
- [6] Paul G. Anastasiades and Simon J. B. Butt. “A Role for Silent Synapses in the Development of the Pathway from Layer 2/3 to 5 Pyramidal Cells in the Neocortex”. In: *Journal of Neuroscience* 32.38 (2012), pp. 13085–13099. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.1262-12.2012. eprint: <https://www.jneurosci.org/content/32/38/13085.full.pdf>. URL: <https://www.jneurosci.org/content/32/38/13085>.
- [7] Paul G. Anastasiades, Andre Marques-Smith, and Simon J. B. Butt. “Studies of cortical connectivity using optical circuit mapping methods”. In: *The Journal of Physiology* 596.2 (2018), pp. 145–162. DOI: <https://doi.org/10.1113/JP273463>. eprint: <https://physoc.onlinelibrary.wiley.com/doi/pdf/10.1113/JP273463>. URL: <https://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/JP273463>.
- [8] N Ankri et al. “Automatic detection of spontaneous synaptic responses in central neurons”. In: *Journal of Neuroscience Methods* 52.1 (1994), pp. 87–100.

- [9] Alfonso J Apicella et al. “Laminarly orthogonal excitation of fast-spiking and low-threshold-spiking interneurons in mouse motor cortex.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 32 (20 May 2012), pp. 7021–33. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.0011-12.2012. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3377057&tool=pmcentrez&rendertype=abstract>.
- [10] Michael C. Ashby and John T. R. Isaac. “Maturation of a recurrent excitatory neocortical circuit by experience-dependent unsilencing of newly formed dendritic spines”. eng. In: *Neuron* 70.3 (May 2011). S0896-6273(11)00301-1[PII], pp. 510–521. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2011.02.057. URL: <https://doi.org/10.1016/j.neuron.2011.02.057>.
- [11] Bassam V Atallah et al. “Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli.” In: *Neuron* 73 (1 Jan. 2012), pp. 159–70. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2011.12.013. URL: <http://www.ncbi.nlm.nih.gov/pubmed/22243754>.
- [12] Craig A Atencio, Tatyana Sharpee, and Christoph E Schreiner. “Hierarchical computation in the canonical auditory cortical circuit.” In: *Proceedings of the National Academy of Sciences* 106 (51 2009), pp. 21894–21899. ISSN: 0027-8424. DOI: 10.1073/pnas.0908383106. URL: <http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=19918079&retmode=ref&cmd=prlinks%5Cnpapers3://publication/doi/10.1073/pnas.0908383106>.
- [13] Christopher A Baker et al. “Cellular resolution circuit mapping in mouse brain with temporal-focused excitation of soma-targeted channelrhodopsin”. In: *eLife* 5 (2016), pp. 1–15. ISSN: 2050-084X. DOI: 10.7554/eLife.14193. URL: <http://elifesciences.org/lookup/doi/10.7554/eLife.14193>.
- [14] Christopher A Baker et al. “Cellular resolution circuit mapping with temporal-focused excitation of soma-targeted channelrhodopsin”. In: *eLife* 5 (Aug. 2016). Ed. by Michael Häusser, e14193. ISSN: 2050-084X. DOI: 10.7554/eLife.14193. URL: <https://doi.org/10.7554/eLife.14193>.
- [15] Dan Bar-Yehuda and Alon Korngreen. “Space-Clamp Problems When Voltage Clamping Neurons Expressing Voltage-Gated Conductances”. In: *Journal of Neurophysiology* 99.3 (2008), pp. 1127–1136.
- [16] M Beierlein, J R Gibson, and B W Connors. “A network of electrically coupled interneurons drives synchronized inhibition in neocortex”. In: *Nature neuroscience* 3 (9 2000), pp. 904–10. ISSN: 1097-6256. DOI: 10.1038/78809. URL: <http://www.ncbi.nlm.nih.gov/pubmed/10966621>.
- [17] Michael Beierlein, Jay R Gibson, and Barry W Connors. “Two dynamically distinct inhibitory networks in layer 4 of the neocortex.” In: *Journal of neurophysiology* 90 (5 2003), pp. 2987–3000. ISSN: 0022-3077. DOI: 10.1152/jn.00283.2003. URL: <http://jn.physiology.org/content/90/5/2987.long>.

- [18] Y. Benjamini and Y. Hochberg. “Controlling the False Discovery Rate : A Practical and Powerful Approach to Multiple Testing Author (s): Yoav Benjamini and Yosef Hochberg Source : Journal of the Royal Statistical Society . Series B (Methodological), Vol . 57 , No . 1 Published by :” in: *J. R. Statist. Soc.* 57 (1 1995), pp. 289–300.
- [19] Thomas K Berger et al. “Brief bursts self-inhibit and correlate the pyramidal network.” In: *PLoS biology* 8 (9 Jan. 2010). ISSN: 1545-7885. DOI: 10.1371/journal.pbio.1000473. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2935452&tool=pmcentrez&rendertype=abstract>.
- [20] Thomas K Berger et al. “Frequency-dependent disynaptic inhibition in the pyramidal network: a ubiquitous pathway in the developing rat neocortex.” In: *The Journal of physiology* 587 (Pt 22 Nov. 2009), pp. 5411–25. ISSN: 1469-7793. DOI: 10.1113/jphysiol.2009.176552. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2793873&tool=pmcentrez&rendertype=abstract>.
- [21] Davi D Bock et al. “Network anatomy and in vivo physiology of visual cortical neurons.” In: *Nature* 471 (7337 Mar. 2011), pp. 177–82. ISSN: 1476-4687. DOI: 10.1038/nature09802. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3095821&tool=pmcentrez&rendertype=abstract>.
- [22] Dante S. Bortone, Shawn R. Olsen, and Massimo Scanziani. “Translaminar inhibitory cells recruited by layer 6 corticothalamic neurons suppress visual cortex”. In: *Neuron* 82 (2 2014), pp. 474–485. ISSN: 10974199. DOI: 10.1016/j.neuron.2014.02.021. URL: <http://dx.doi.org/10.1016/j.neuron.2014.02.021>.
- [23] Edward S. Boyden et al. “Millisecond-timescale, genetically targeted optical control of neural activity”. In: *Nature Neuroscience* 8.9 (Sept. 2005), pp. 1263–1268. ISSN: 1546-1726. DOI: 10.1038/nn1525. URL: <https://doi.org/10.1038/nn1525>.
- [24] J. Brill et al. “LSPS/Optogenetics to Improve Synaptic Connectivity Mapping: Unmasking the Role of Basket Cell Mediated Feed-Forward Inhibition”. In: *eNeuro* 3 (August 2016), ENEURO.0142–15.2016. ISSN: 2373-2822. DOI: 10.1523/ENEURO.0142-15.2016. URL: <http://eneuro.sfn.org/cgi/doi/10.1523/ENEURO.0142-15.2016>.
- [25] Katherine A. Buchanan et al. “Target-Specific Expression of Presynaptic NMDA Receptors in Neocortical Microcircuits”. In: *Neuron* 75 (3 2012), pp. 451–466. ISSN: 08966273. DOI: 10.1016/j.neuron.2012.06.017.
- [26] E. M. Callaway and L. C. Katz. “Photostimulation using caged glutamate reveals functional circuitry in living brain slices”. eng. In: *Proceedings of the National Academy of Sciences of the United States of America* 90.16 (Aug. 1993). PMC47202[pmcid], pp. 7661–7665. ISSN: 0027-8424. DOI: 10.1073/pnas.90.16.7661. URL: <https://doi.org/10.1073/pnas.90.16.7661>.
- [27] Edward M Callaway and Rafael Yuste. “Stimulating neurons with light”. In: *Current Opinion in Neurobiology* 12.5 (2002), pp. 587–592.

- [28] Edward M. Callaway et al. “A multimodal cell census and atlas of the mammalian primary motor cortex”. In: *Nature* 598.7879 (Oct. 2021), pp. 86–102. ISSN: 1476-4687. DOI: 10.1038/s41586-021-03950-0. URL: <https://doi.org/10.1038/s41586-021-03950-0>.
- [29] Luke Campagnola et al. “Local connectivity and synaptic dynamics in mouse and human neocortex”. In: *Science* 375.6585 (2022), eabj5861. DOI: 10.1126/science.abj5861. eprint: <https://www.science.org/doi/pdf/10.1126/science.abj5861>. URL: <https://www.science.org/doi/abs/10.1126/science.abj5861>.
- [30] M Carandini and Dj Heeger. “Normalization as a canonical neural computation.” In: *Nature Reviews Neuroscience* (November 2012), pp. 1–12. ISSN: 1471-003X. DOI: 10.1038/nrn3136. URL: <http://discovery.ucl.ac.uk/1332718/>.
- [31] I-Wen Chen et al. “In Vivo Submillisecond Two-Photon Optogenetics with Temporally Focused Patterned Light”. eng. In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 39.18 (May 2019). JNEUROSCI.1785-18.2018[PII], pp. 3484–3497. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.1785-18.2018. URL: <https://doi.org/10.1523/JNEUROSCI.1785-18.2018>.
- [32] Siddhartha Chib and Edward Greenberg. “Bayes inference in regression models with ARMA (p, q) errors”. In: *Journal of Econometrics* 64.1 (1994), pp. 183–206.
- [33] JD Clements and JM Bekkers. “Detection of spontaneous synaptic events with an optimally scaled template.” In: *Biophysical Journal* 73.1 (1997), p. 220.
- [34] Christine M Constantinople and Randy M Bruno. “Deep cortical layers are activated directly by thalamus.” In: *Science (New York, N.Y.)* 340 (6140 June 2013), pp. 1591–4. ISSN: 1095-9203. DOI: 10.1126/science.1236425. URL: <http://www.ncbi.nlm.nih.gov/pubmed/23812718>.
- [35] James C H Cottam, Spencer L Smith, and Michael Häusser. “Target-specific effects of somatostatin-expressing interneurons on neocortical visual processing.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 33 (50 Dec. 2013), pp. 19567–78. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.2624-13.2013. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3858626&tool=pmcentrez&rendertype=abstract>.
- [36] Robert D. Dayton, David B. Wang, and Ronald L. Klein. “The advent of AAV9 expands applications for brain and spinal cord gene delivery”. eng. In: *Expert opinion on biological therapy* 12.6 (June 2012). PMC3361729[pmcid], pp. 757–766. ISSN: 1744-7682. DOI: 10.1517/14712598.2012.681463. URL: <https://doi.org/10.1517/14712598.2012.681463>.
- [37] Karl Deisseroth. “Optogenetics: 10 years of microbial opsins in neuroscience”. In: *Nature Neuroscience* 18.9 (Sept. 2015), pp. 1213–1225. ISSN: 1546-1726. DOI: 10.1038/nn.4091. URL: <https://doi.org/10.1038/nn.4091>.

- [38] Winifried Denk, James H. Strickler, and Watt W. Webb. “Two-Photon Laser Scanning Fluorescence Microscopy”. In: *Science* 248.4951 (1990), pp. 73–76. DOI: 10.1126/science.2321027. eprint: <https://www.science.org/doi/pdf/10.1126/science.2321027>. URL: <https://www.science.org/doi/abs/10.1126/science.2321027>.
- [39] R J Douglas and K A Martin. “A functional microcircuit for cat visual cortex.” In: *The Journal of Physiology* 440.1 (1991), pp. 735–769. DOI: <https://doi.org/10.1113/jphysiol.1991.sp018733>. eprint: <https://physoc.onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1991.sp018733>. URL: <https://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1991.sp018733>.
- [40] Chee-Huat Linus Eng et al. “Transcriptome-scale super-resolved imaging in tissues by RNA seqFISH+”. In: *Nature* 568.7751 (Apr. 2019), pp. 235–239. ISSN: 1476-4687. DOI: 10.1038/s41586-019-1049-y. URL: <https://doi.org/10.1038/s41586-019-1049-y>.
- [41] Erika E Fanselow, Kristen A Richardson, and Barry W Connors. “Selective , State-Dependent Activation of Somatostatin-Expressing Inhibitory Interneurons in Mouse Neocortex”. In: *Journal of neurophysiology* 100 (2008), pp. 2640–2652. ISSN: 0022-3077. DOI: 10.1152/jn.90691.2008..
- [42] Ira C. Farber and Amiram Grinvald. “Identification of Presynaptic Neurons by Laser Photostimulation”. In: *Science* 222.4627 (1983), pp. 1025–1027. DOI: 10.1126/science.6648515. eprint: <https://www.science.org/doi/pdf/10.1126/science.6648515>. URL: <https://www.science.org/doi/abs/10.1126/science.6648515>.
- [43] Dirk Feldmeyer. “Excitatory neuronal connectivity in the barrel cortex.” In: *Frontiers in neuroanatomy* 6 (July Jan. 2012), p. 24. ISSN: 1662-5129. DOI: 10.3389/fnana.2012.00024. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3394394&tool=pmcentrez&rendertype=abstract>.
- [44] Lief Fenno, Ofer Yizhar, and Karl Deisseroth. “The Development and Application of Optogenetics”. In: *Annual Review of Neuroscience* 34.1 (2011), pp. 389–412.
- [45] Elodie Fino, Adam M Packer, and Rafael Yuste. “The logic of inhibitory connectivity in the neocortex.” In: *The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry* 19 (3 June 2013), pp. 228–37. ISSN: 1089-4098. DOI: 10.1177/1073858412456743. URL: <http://www.ncbi.nlm.nih.gov/pubmed/22922685>.
- [46] Elodie Fino and Rafael Yuste. “Dense Inhibitory Connectivity in Neocortex”. In: *Neuron* 69.6 (Mar. 2011), pp. 1188–1203. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2011.02.025. URL: <https://doi.org/10.1016/j.neuron.2011.02.025>.
- [47] Elodie Fino et al. “RuBi-Glutamate: Two-photon and visible-light photoactivation of neurons and dendritic spines”. In: *Frontiers in Neural Circuits* 3 (2009). ISSN: 1662-5110. DOI: 10.3389/neuro.04.002.2009. URL: <https://www.frontiersin.org/article/10.3389/neuro.04.002.2009>.
- [48] Yu Fu et al. “A cortical circuit for gain control by behavioral state”. In: *Cell* 156 (6 2014), pp. 1139–1152. ISSN: 10974172. DOI: 10.1016/j.cell.2014.01.050.

- [49] Andrew Gelman et al. *Bayesian data analysis*. Vol. 2. Taylor & Francis, 2014.
- [50] R. W. GERCHBERG and W. O. SAXTON. “PRACTICAL ALGORITHM FOR DETERMINATION OF PHASE FROM IMAGE AND DIFFRACTION PLANE PICTURES”. In: *OPTIK* 35.2 (1972), 237–&. ISSN: 0030-4026.
- [51] Stephen D. Glasgow et al. “Approaches and Limitations in the Investigation of Synaptic Transmission and Plasticity”. In: *Frontiers in Synaptic Neuroscience* 11 (2019). ISSN: 1663-3563. DOI: 10.3389/fnsyn.2019.00020. URL: <https://www.frontiersin.org/article/10.3389/fnsyn.2019.00020>.
- [52] Segundo J Guzman, Alois Schlögl, and Christoph Schmidt-Hieber. “Stimfit: quantifying electrophysiological data with Python”. In: *Frontiers in Neuroinformatics* 8.16 (2014), *online*.
- [53] Travis A Hage et al. “Synaptic connectivity to L2/3 of primary visual cortex measured by two-photon optogenetic stimulation”. In: *eLife* 11 (Jan. 2022). Ed. by Tianyi Mao and Ronald L Calabrese, e71103. ISSN: 2050-084X. DOI: 10.7554/eLife.71103. URL: <https://doi.org/10.7554/eLife.71103>.
- [54] Kenneth D Harris and Gordon M G Shepherd. “The neocortical circuit: themes and variations”. In: *Nature Neuroscience* 18 (2 2015), pp. 170–181. ISSN: 1097-6256. DOI: 10.1038/nn.3917. URL: http://www.nature.com/neuro/journal/v18/n2/full/nn.3917.html?WT.ec_id=NEURO-201502.
- [55] Alexis M Hattox and Sacha B Nelson. “Layer V neurons in mouse cortex projecting to different targets have distinct physiological properties.” In: *Journal of Neurophysiology* 98 (6 Dec. 2007), pp. 3330–3340. ISSN: 0022-3077. DOI: 10.1152/jn.00397.2007. URL: <http://www.ncbi.nlm.nih.gov/pubmed/17898147>.
- [56] Miao He et al. “Strategies and tools for combinatorial targeting of GABAergic neurons in mouse cerebral cortex”. In: *Neuron* (In Press 2016), pp. 1–16. ISSN: 08966273. DOI: 10.1016/j.neuron.2016.08.021.
- [57] M. Helmstaedter, B. Sakmann, and D. Feldmeyer. “Neuronal Correlates of Local, Lateral, and Translaminar Inhibition with Reference to Cortical Columns”. In: *Cerebral Cortex* 19 (4 2009), pp. 926–937. ISSN: 1047-3211. DOI: 10.1093/cercor/bhn141. URL: <http://www.cercor.oxfordjournals.org/cgi/doi/10.1093/cercor/bhn141>.
- [58] Oscar Hernandez et al. “Three-dimensional spatiotemporal focusing of holographic patterns”. In: *Nature Communications* 7.1 (June 2016), p. 11928. ISSN: 2041-1723. DOI: 10.1038/ncomms11928. URL: <https://doi.org/10.1038/ncomms11928>.
- [59] Sonja B. Hofer et al. “Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex.” In: *Nature neuroscience* 14 (8 2011), pp. 1045–1052. ISSN: 1097-6256. DOI: 10.1038/nn.2876.
- [60] B. M. Hooks et al. “Laminar Analysis of Excitatory Local Circuits in Vibrissal Motor and Sensory Cortical Areas”. In: *PLOS Biology* 9.1 (Jan. 2011), pp. 1–15. DOI: 10.1371/journal.pbio.1000572. URL: <https://doi.org/10.1371/journal.pbio.1000572>.

- [61] H. Hu and A. Agmon. “Differential Excitation of Distally versus Proximally Targeting Cortical Interneurons by Unitary Thalamocortical Bursts”. In: *Journal of Neuroscience* 36 (26 2016), pp. 6906–6916. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.0739-16.2016. URL: <http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.0739-16.2016>.
- [62] D. H. HUBEL and T. N. WIESEL. “Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex”. eng. In: *The Journal of physiology* 160.1 (Jan. 1962). PMC1359523[pmcid], pp. 106–154. ISSN: 0022-3751. DOI: 10.1113/jphysiol.1962.sp006837. URL: <https://doi.org/10.1113/jphysiol.1962.sp006837>.
- [63] Thomas N Hwang and David R Copenhagen. “Automatic detection, characterization, and discrimination of kinetically distinct spontaneous synaptic events”. In: *Journal of Neuroscience Methods* 92.1 (1999), pp. 65–73.
- [64] Mercè Izquierdo-Serra et al. “Two-Photon Optogenetic Mapping of Excitatory Synaptic Connectivity and Strength”. In: *iScience* 8 (Oct. 2018), pp. 15–28. ISSN: 2589-0042. DOI: 10.1016/j.isci.2018.09.008. URL: <https://doi.org/10.1016/j.isci.2018.09.008>.
- [65] X.-y. Ji et al. “Thalamocortical Innervation Pattern in Mouse Auditory and Visual Cortex: Laminar and Cell-Type Specificity”. In: *Cerebral Cortex* (2015), bhv099-. ISSN: 1047-3211. DOI: 10.1093/cercor/bhv099. URL: <http://cercor.oxfordjournals.org/content/early/2015/05/14/cercor.bhv099.abstract>.
- [66] X. Jiang et al. “Principles of connectivity among morphologically defined cell types in adult neocortex”. In: *Science* 350 (6264 2015), aac9462–aac9462. ISSN: 0036-8075. DOI: 10.1126/science.aac9462. URL: <http://www.sciencemag.org/cgi/doi/10.1126/science.aac9462>.
- [67] Xiaolong Jiang et al. “Principles of connectivity among morphologically defined cell types in adult neocortex”. In: *Science* 350.6264 (2015), aac9462. DOI: 10.1126/science.aac9462. eprint: <https://www.science.org/doi/pdf/10.1126/science.aac9462>. URL: <https://www.science.org/doi/abs/10.1126/science.aac9462>.
- [68] Xiaolong Jiang et al. “The organization of two new cortical interneuronal circuits.” In: *Nature neuroscience* 16 (2 Feb. 2013), pp. 210–8. ISSN: 1546-1726. DOI: 10.1038/nn.3305. URL: <http://dx.doi.org/10.1038/nn.3305>.
- [69] Peter Jonas, Guy Major, and Bert Sakmann. “Quantal components of unitary EPSCs at the mossy fibre synapse on CA3 pyramidal cells of rat hippocampus.” In: *The Journal of Physiology* 472.1 (1993), pp. 615–663.
- [70] Christoph Kapfer et al. “Supralinear increase of recurrent inhibition during sparse activity in the somatosensory cortex.” In: *Nature neuroscience* 10 (6 June 2007), pp. 743–53. ISSN: 1097-6256. DOI: 10.1038/nn1909. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3518866&tool=pmcentrez&rendertype=abstract>.

- [71] Mahesh M Karnani, Masakazu Agetsuma, and Rafael Yuste. “A blanket of inhibition: functional inferences from dense inhibitory connectivity.” In: *Current opinion in neurobiology* 26 (June 2014), pp. 96–102. ISSN: 1873-6882. DOI: 10.1016/j.conb.2013.12.015. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4024353&tool=pmcentrez&rendertype=abstract>.
- [72] Mahesh M M. Karnani et al. “Cooperative Subnetworks of Molecularly Similar Interneurons in Mouse Neocortex”. In: *Neuron* 90 (1 2016), pp. 86–100. ISSN: 10974199. DOI: 10.1016/j.neuron.2016.02.037. URL: <http://dx.doi.org/10.1016/j.neuron.2016.02.037>.
- [73] Lawrence C. Katz and Matthew B. Dalva. “Scanning laser photostimulation: a new approach for analyzing brain circuits”. In: *Journal of Neuroscience Methods* 54.2 (1994). Imaging Techniques in Neurobiology, pp. 205–218. ISSN: 0165-0270. DOI: [https://doi.org/10.1016/0165-0270\(94\)90194-5](https://doi.org/10.1016/0165-0270(94)90194-5). URL: <https://www.sciencedirect.com/science/article/pii/0165027094901945>.
- [74] Dennis Katznel et al. “The columnar and laminar organization of inhibitory connections to neocortical excitatory cells”. In: *Nature Neuroscience* 14.1 (Jan. 2011), pp. 100–107.
- [75] Y Kawaguchi. “Groupings of nonpyramidal and pyramidal cells with specific physiological and morphological characteristics in rat frontal cortex.” In: *Journal of neurophysiology* 69 (2 1993), pp. 416–431. ISSN: 0022-3077.
- [76] Adam Kepecs and Gordon Fishell. “Interneuron cell types are fit to function”. In: *Nature* 505 (7483 2014), pp. 318–326. ISSN: 0028-0836. DOI: 10.1038/nature12983. URL: <http://www.nature.com/doifinder/10.1038/nature12983>.
- [77] Dohoung Kim et al. “Distinct Roles of Parvalbumin- and Somatostatin- Expressing Interneurons in Working Memory Article Distinct Roles of Parvalbumin- and Somatostatin-Expressing Interneurons in Working Memory”. In: *Neuron* 92 (4 2016), pp. 1–14. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2016.09.023. URL: <http://dx.doi.org/10.1016/j.neuron.2016.09.023>.
- [78] Euseok J Kim et al. “Three Types of Cortical Layer 5 Neurons That Differ in Brain-wide Connectivity and Function”. In: *Neuron* 88 (January 2015), pp. 1–15. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2015.11.002. URL: <http://dx.doi.org/10.1016/j.neuron.2015.11.002>.
- [79] Amanda K Kinnischtzke et al. “Postnatal maturation of somatostatin-expressing inhibitory cells in the somatosensory cortex of GIN mice.” In: *Frontiers in neural circuits* 6 (May 2012), p. 33. ISSN: 1662-5110. DOI: 10.3389/fncir.2012.00033. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3364579&tool=pmcentrez&rendertype=abstract>.
- [80] Thomas Knöpfel and Chenchen Song. “Optical voltage imaging in neurons: moving from technology development to practical tool”. In: *Nature Reviews Neuroscience* 20.12 (Dec. 2019), pp. 719–727. ISSN: 1471-0048. DOI: 10.1038/s41583-019-0231-4. URL: <https://doi.org/10.1038/s41583-019-0231-4>.

- [81] Ho Ko et al. “Functional specificity of local synaptic connections in neocortical networks”. In: *Nature* 473.7345 (May 2011), pp. 87–91. ISSN: 1476-4687. DOI: 10.1038/nature09880. URL: <https://doi.org/10.1038/nature09880>.
- [82] Christian Koelbl et al. “A barrel-related interneuron in layer 4 of rat somatosensory cortex with a high intrabarrel connectivity”. In: *Cerebral Cortex* 25 (3 2015), pp. 713–725. ISSN: 14602199. DOI: 10.1093/cercor/bht263.
- [83] Aaron T. Kuan et al. “Synaptic wiring motifs in posterior parietal cortex support decision-making”. In: *bioRxiv* (2022). DOI: 10.1101/2022.04.13.488176. eprint: <https://www.biorxiv.org/content/early/2022/04/14/2022.04.13.488176.full.pdf>. URL: <https://www.biorxiv.org/content/early/2022/04/14/2022.04.13.488176>.
- [84] Suguru N Kudoh and Takahisa Taguchi. “A simple exploratory algorithm for the accurate and fast detection of spontaneous synaptic events”. In: *Biosensors and Bioelectronics* 17.9 (2002), pp. 773–782.
- [85] D Kvitsiani et al. “Distinct behavioural and network correlates of two interneuron types in prefrontal cortex.” In: *Nature* 498 (7454 June 2013), pp. 363–6. ISSN: 1476-4687. DOI: 10.1038/nature12176. URL: <http://www.ncbi.nlm.nih.gov/pubmed/23708967>.
- [86] David C. Lay. *Linear Algebra and Its Applications*. 4th ed. Pearson, 2012.
- [87] Seung-Hee Lee et al. “Activation of specific interneurons improves V1 feature selectivity and visual perception”. In: *Nature* 488 (7411 Aug. 2012), pp. 379–83. ISSN: 0028-0836. DOI: 10.1038/nature11312. URL: <http://www.nature.com/doifinder/10.1038/nature11312%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3422431&tool=pmcentrez&rendertype=abstract>.
- [88] Soohyun Lee et al. “A disinhibitory circuit mediates motor integration in the somatosensory cortex.” In: *Nature neuroscience* 16 (11 Nov. 2013), pp. 1662–70. ISSN: 1546-1726. DOI: 10.1038/nn.3544. URL: <http://www.ncbi.nlm.nih.gov/pubmed/24097044>.
- [89] Sandrine Lefort et al. “The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex”. In: *Neuron* 61.2 (2009), pp. 301–316. ISSN: 0896-6273. DOI: <https://doi.org/10.1016/j.neuron.2008.12.020>. URL: <https://www.sciencedirect.com/science/article/pii/S0896627308010921>.
- [90] Erich Leo Lehmann and George Casella. *Theory of point estimation*. Vol. 31. Springer Science & Business Media, 1998.
- [91] Robert B Levy and Alex D Reyes. “Spatial profile of excitatory and inhibitory synaptic connectivity in mouse primary auditory cortex.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 32 (16 Apr. 2012), pp. 5609–19. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.5158-11.2012. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3359703&tool=pmcentrez&rendertype=abstract>.

- [92] Guo-Hua Li, Michael F Jackson, and John F MacDonald. “Weighted least squares fitting with multiple templates for detection of small spontaneous signals”. In: *Journal of Neuroscience Methods* 164.1 (2007), pp. 139–148.
- [93] Jeff W. Lichtman, Hanspeter Pfister, and Nir Shavit. “The big data challenges of connectomics”. In: *Nature Neuroscience* 17.11 (Nov. 2014), pp. 1448–1454. ISSN: 1546-1726. DOI: 10.1038/nn.3837. URL: <https://doi.org/10.1038/nn.3837>.
- [94] Michael Z. Lin and Mark J. Schnitzer. “Genetically encoded indicators of neuronal activity”. eng. In: *Nature neuroscience* 19.9 (Aug. 2016). nn.4359[PII], pp. 1142–1153. ISSN: 1546-1726. DOI: 10.1038/nn.4359. URL: <https://doi.org/10.1038/nn.4359>.
- [95] Ashok Litwin-Kumar, Robert Rosenbaum, and Brent Doiron. “Inhibitory stabilization and visual coding in cortical circuits with multiple interneuron subtypes.” In: *Journal of neurophysiology* (2016), jn.00732.2015. ISSN: 1522-1598. DOI: 10.1152/jn.00732.2015. URL: <http://jn.physiology.org/content/early/2016/01/04/jn.00732.2015.full-text.pdf+html>.
- [96] Wen-pei Ma et al. “Visual representations by cortical somatostatin inhibitory neurons—selective but with weak and delayed responses.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 30 (43 Oct. 2010), pp. 14371–9. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.3248-10.2010. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3001391&tool=pmcentrez&rendertype=abstract>.
- [97] Yunyong Ma et al. “Distinct subtypes of somatostatin-containing neocortical interneurons revealed in transgenic mice.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 26 (19 May 2006), pp. 5069–82. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.0661-06.2006. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2020857&tool=pmcentrez&rendertype=abstract>.
- [98] Hiroshi Makino and Takaki Komiyama. “Learning enhances the relative impact of top-down processing in the visual cortex”. In: *Nature Neuroscience* 18 (8 July 2015), pp. 1116–1122. ISSN: 1097-6256. DOI: 10.1038/nn.4061. URL: <http://www.nature.com/doi/10.1038/nn.4061>.
- [99] A. R. Mardinly et al. “Sensory experience regulates cortical inhibition by inducing IGF1 in VIP neurons”. In: *Nature* 531.7594 (Mar. 2016), pp. 371–375.
- [100] Alan R. Mardinly et al. “Precise multimodal optical control of neural ensemble activity”. eng. In: *Nature neuroscience* 21.6 (June 2018). PMC5970968[pmcid], pp. 881–893. ISSN: 1546-1726. DOI: 10.1038/s41593-018-0139-8. URL: <https://doi.org/10.1038/s41593-018-0139-8>.
- [101] Henry Markram et al. “Interneurons of the neocortical inhibitory system.” In: *Nature reviews. Neuroscience* 5 (10 Oct. 2004), pp. 793–807. ISSN: 1471-003X. DOI: 10.1038/nrn1519. URL: <http://www.ncbi.nlm.nih.gov/pubmed/15378039>.
- [102] Henry Markram et al. “Reconstruction and Simulation of Neocortical Microcircuitry”. In: *Cell* 163.2 (Oct. 2015), pp. 456–492. ISSN: 0092-8674. DOI: 10.1016/j.cell.2015.09.029. URL: <https://doi.org/10.1016/j.cell.2015.09.029>.

- [103] Andre Marques-Smith et al. “A Transient Translaminar GABAergic Interneuron Circuit Connects Thalamocortical Recipient Layers in Neonatal Somatosensory Cortex”. In: *Neuron* 89.3 (Feb. 2016), pp. 536–549. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2016.01.015. URL: <https://doi.org/10.1016/j.neuron.2016.01.015>.
- [104] L. M. Martinez et al. “Receptive field structure varies with layer in the primary visual cortex”. In: *Nat Neurosci* 8 (3 2005), pp. 372–9. ISSN: 1097-6256 (PRINT) 1097-6256 (LINKING). DOI: 10.1038/nn1404. URL: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15711543.
- [105] Laura M McGarry et al. “Quantitative classification of somatostatin-positive neocortical interneurons identifies three interneuron subtypes.” In: *Frontiers in neural circuits* 4 (May Jan. 2010), p. 12. ISSN: 1662-5110. DOI: 10.3389/fncir.2010.00012. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2896209&tool=pmcentrez&rendertype=abstract>.
- [106] Manuela Medelin et al. *Altered development in GABA co-release shapes glycinergic synaptic currents in cultured spinal slices of the SOD1G93A mouse model of ALS*. The Journal of Physiology. 2016.
- [107] Gonzalo Mena et al. “Large-scale Multi Electrode Array Spike Sorting Algorithm Introducing Concurrent Recording and Stimulation”. In: *NIPS workshop on Statistical Methods for Understanding Neural Systems*. 2015.
- [108] J. Merel et al. “Bayesian methods for event analysis of intracellular currents”. In: *Journal of Neuroscience Methods* 269 (2016). ISSN: 1872678X. DOI: 10.1016/j.jneumeth.2016.05.015.
- [109] Kenneth D. Miller. “Canonical computations of cerebral cortex”. In: *Current Opinion in Neurobiology* 37 (2016), pp. 75–84. ISSN: 18736882. DOI: 10.1016/j.conb.2016.01.008. URL: <http://dx.doi.org/10.1016/j.conb.2016.01.008>.
- [110] Jesper Moller and Rasmus Plenge Waagepetersen. *Statistical inference and simulation for spatial point processes*. CRC Press, 2004.
- [111] Nicolás A. Morgenstern, Jacques Bourg, and Leopoldo Petreanu. “Multilaminar networks of cortical neurons integrate common inputs from sensory thalamus.” In: *Nature neuroscience* 19 (8 2016), pp. 1034–40. ISSN: 1546-1726. DOI: 10.1038/nn.4339. URL: <http://www.nature.com/neuro/research/index.html%5C%5Cnhttp://www.nature.com/doifinder/10.1038/nn.4339%5C%5Cnhttp://www.ncbi.nlm.nih.gov/pubmed/27376765>.
- [112] William Muñoz, Robin Tremblay, and Bernardo Rudy. “Channelrhodopsin-Assisted Patching: InVivo Recording of Genetically and Morphologically Identified Neurons throughout the Brain”. In: *Cell Reports* 9 (6 2014), pp. 2304–2316. ISSN: 22111247. DOI: 10.1016/j.celrep.2014.11.042.
- [113] William Muñoz et al. “Layer-specific modulation of neocortical dendritic inhibition during active wakefulness”. In: *Science* 355 (6328 Mar. 2017), pp. 954–959. ISSN: 0036-8075. DOI: 10.1126/science.aag2599. URL: <http://www.sciencemag.org/lookup/doi/10.1126/science.aag2599>.

- [114] Masanori Murayama et al. “Dendritic encoding of sensory stimuli controlled by deep cortical interneurons.” In: *Nature* 457 (February 2009), pp. 1137–1141. ISSN: 0028-0836. DOI: 10.1038/nature07663.
- [115] Miho Nakajima, Andreas Görlich, and Nathaniel Heintz. “Oxytocin modulates female sociosexual behavior through a specific class of prefrontal cortical interneurons”. In: *Cell* 159 (2 2014), pp. 295–305. ISSN: 10974172. DOI: 10.1016/j.cell.2014.09.020.
- [116] Radford M Neal. *Probabilistic inference using Markov chain Monte Carlo methods*. Department of Computer Science, University of Toronto Toronto, CA, 1993.
- [117] Volodymyr Nikolenko et al. “SLM Microscopy: Scanless Two-Photon Imaging and Photostimulation with Spatial Light Modulators”. eng. In: *Frontiers in neural circuits* 2 (Dec. 2008). PMC2614319[pmcid], pp. 5–5. ISSN: 1662-5110. DOI: 10.3389/neuro.04.005.2008. URL: <https://doi.org/10.3389/neuro.04.005.2008>.
- [118] a a Oliva et al. “Novel hippocampal interneuronal subtypes identified using transgenic mice that express green fluorescent protein in GABAergic interneurons.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 20 (9 2000), pp. 3354–3368. ISSN: 1529-2401. DOI: 0270-6474/00/203354-15\$15.00/0.
- [119] Shawn R. Olsen et al. “Gain control by layer six in cortical circuits of vision”. In: *Nature* 483.7387 (Mar. 2012), pp. 47–52. ISSN: 1476-4687. DOI: 10.1038/nature10835. URL: <https://doi.org/10.1038/nature10835>.
- [120] Dan Oron and Yaron Silberberg. “Temporal Focusing Microscopy”. In: *Cold Spring Harbor Protocols* 2015.2 (2015), pdb.top085928. DOI: 10.1101/pdb.top085928. eprint: <http://cshprotocols.cshlp.org/content/2015/2/pdb.top085928.full.pdf+html>. URL: <http://cshprotocols.cshlp.org/content/2015/2/pdb.top085928.abstract>.
- [121] Takeshi Otsuka and Yasuo Kawaguchi. “Cortical inhibitory cell types differentially form intralaminar and interlaminar subnetworks with excitatory neurons.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 29 (34 Aug. 2009), pp. 10533–40. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.2219-09.2009. URL: <http://www.ncbi.nlm.nih.gov/pubmed/19710306>.
- [122] Adam M Packer and Rafael Yuste. “Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: a canonical microcircuit for inhibition?” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 31 (37 2011), pp. 13260–13271. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.3131-11.2011.
- [123] Adam M. Packer and Rafael Yuste. “Dense, Unspecific Connectivity of Neocortical Parvalbumin-Positive Interneurons: A Canonical Microcircuit for Inhibition?” In: *Journal of Neuroscience* 31.37 (2011), pp. 13260–13271. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.3131-11.2011. eprint: <https://www.jneurosci.org/content/31/37/13260.full.pdf>. URL: <https://www.jneurosci.org/content/31/37/13260>.
- [124] Adam M. Packer et al. “Two-photon optogenetics of dendritic spines and neural circuits”. In: *Nature Methods* 9.12 (Dec. 2012), pp. 1202–1205. ISSN: 1548-7105. DOI: 10.1038/nmeth.2249. URL: <https://doi.org/10.1038/nmeth.2249>.

- [125] Janelle M P Pakan et al. “Behavioral-state modulation of inhibition is context-dependent and cell type specific in mouse visual cortex”. In: *eLife* 5 (AUGUST 2016), pp. 1–18. ISSN: 2050084X. DOI: 10.7554/eLife.14985.
- [126] Francisco Palma-Cerda et al. “New caged neurotransmitter analogs selective for glutamate receptor sub-types based on methoxynitroindoline and nitrophenylethoxycarbonyl caging groups”. In: *Neuropharmacology* 63.4 (2012), pp. 624–634. ISSN: 0028-3908. DOI: <https://doi.org/10.1016/j.neuropharm.2012.05.010>. URL: <https://www.sciencedirect.com/science/article/pii/S0028390812001955>.
- [127] Liam Paninski et al. “Inferring synaptic inputs given a noisy voltage trace via sequential Monte Carlo methods”. In: *Journal of Computational Neuroscience* 33.1 (2012), pp. 1–19.
- [128] Eirini Papagiakoumou et al. “Patterned two-photon illumination by spatiotemporal shaping of ultrashort pulses”. In: *Opt. Express* 16.26 (Dec. 2008), pp. 22039–22047. DOI: 10.1364/OE.16.022039. URL: <http://opg.optica.org/oe/abstract.cfm?URI=oe-16-26-22039>.
- [129] Eirini Papagiakoumou et al. “Scanless two-photon excitation of channelrhodopsin-2”. In: *Nature Methods* 7.10 (Oct. 2010), pp. 848–854. ISSN: 1548-7105. DOI: 10.1038/nmeth.1505. URL: <https://doi.org/10.1038/nmeth.1505>.
- [130] Nicolas C. Pégard et al. “Three-dimensional scanless holographic optogenetics with temporal focusing (3D-SHOT)”. In: *Nature Communications* 8.1 (Oct. 2017), p. 1228. ISSN: 2041-1723. DOI: 10.1038/s41467-017-01031-3. URL: <https://doi.org/10.1038/s41467-017-01031-3>.
- [131] Hanchuan Peng et al. “Morphological diversity of single neurons in molecularly defined cell types”. In: *Nature* 598.7879 (Oct. 2021), pp. 174–181. ISSN: 1476-4687. DOI: 10.1038/s41586-021-03941-1. URL: <https://doi.org/10.1038/s41586-021-03941-1>.
- [132] Alejandro Javier Pernia-Andrade et al. “A deconvolution-based method with high sensitivity and temporal resolution for detection of spontaneous synaptic currents in vitro and in vivo”. In: *Biophysical Journal* 103.7 (2012), pp. 1429–1439.
- [133] Carl C H Petersen and Bert Sakmann. “The excitatory neuronal network of rat layer 4 barrel cortex.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 20 (20 2000), pp. 7579–7586. ISSN: 1529-2401.
- [134] Leopoldo Petreanu et al. “Activity in motor–sensory projections reveals distributed coding in somatosensation”. In: *Nature* (Aug. 2012), pp. 0-7. ISSN: 0028-0836. DOI: 10.1038/nature11321. URL: <http://www.nature.com/doifinder/10.1038/nature11321>.
- [135] Leopoldo Petreanu et al. “Channelrhodopsin-2–assisted circuit mapping of long-range callosal projections”. In: *Nature Neuroscience* 10.5 (May 2007), pp. 663–668. ISSN: 1546-1726. DOI: 10.1038/nn1891. URL: <https://doi.org/10.1038/nn1891>.

- [136] Carsten K Pfeffer et al. “Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons.” In: *Nature neuroscience* 16 (8 Aug. 2013), pp. 1068–76. ISSN: 1546-1726. DOI: 10.1038/nn.3446. URL: <http://dx.doi.org/10.1038/nn.3446>.
- [137] Hyun-Jae Pi et al. “Cortical interneurons that specialize in disinhibitory control.” In: *Nature* 503 (7477 Nov. 2013), pp. 521–4. ISSN: 1476-4687. DOI: 10.1038/nature12676. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4017628&tool=pmcentrez&rendertype=abstract>.
- [138] Scott Pluta et al. “A direct translaminar inhibitory circuit tunes cortical output”. In: *Nature Neuroscience* 18.11 (Nov. 2015), pp. 1631–1640.
- [139] Eftychios A Pnevmatikakis et al. “Bayesian spike inference from calcium imaging data”. In: *Asilomar Conference on Signals, Systems and Computers*. IEEE. 2013, pp. 349–353.
- [140] Pierre-Olivier Polack and Diego Contreras. “Long-range parallel processing and local recurrent activity in the visual cortex of the mouse.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 32 (32 Aug. 2012), pp. 11120–31. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.6304-11.2012. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3437251&tool=pmcentrez&rendertype=abstract>.
- [141] Cindy Poo and Jeffrey S. Isaacson. “Odor Representations in Olfactory Cortex: ”Sparse” Coding, Global Inhibition, and Oscillations”. In: *Neuron* 62 (6 2009), pp. 850–861. ISSN: 08966273. DOI: 10.1016/j.neuron.2009.05.022. URL: <http://dx.doi.org/10.1016/j.neuron.2009.05.022>.
- [142] Frédéric Pouille et al. “Input normalization by global feedforward inhibition expands cortical dynamic range.” In: *Nature neuroscience* 12 (12 Dec. 2009), pp. 1577–85. ISSN: 1546-1726. DOI: 10.1038/nn.2441. URL: <http://www.ncbi.nlm.nih.gov/pubmed/19881502%20http://dx.doi.org/10.1038/nn.2441>.
- [143] Rajesh Ranganath, Sean Gerrish, and David Blei. “Black Box Variational Inference”. In: *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*. Ed. by Samuel Kaski and Jukka Corander. Vol. 33. Proceedings of Machine Learning Research. Reykjavik, Iceland: PMLR, 22–25 Apr 2014, pp. 814–822. URL: <https://proceedings.mlr.press/v33/ranganath14.html>.
- [144] Jacob Reimer et al. “Pupil Fluctuations Track Fast Switching of Cortical States during Quiet Wakefulness”. In: *Neuron* 84 (2 2014), pp. 355–362. ISSN: 08966273. DOI: 10.1016/j.neuron.2014.09.033. URL: <http://www.sciencedirect.com/science/article/pii/S0896627314008915>.
- [145] Magnus JE Richardson and Gilad Silberberg. “Measurement and analysis of postsynaptic potentials using a novel voltage-deconvolution method”. In: *Journal of Neurophysiology* 99.2 (2008), pp. 1020–1031.
- [146] John Peter Rickgauer, Karl Deisseroth, and David W. Tank. “Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields”. In: *Nature Neuroscience* 17.12 (Dec. 2014), pp. 1816–1824.

- [147] Jeffrey S Rosenthal. “Optimal proposal distributions and adaptive MCMC”. In: *Handbook of Markov Chain Monte Carlo* (2011), pp. 93–112.
- [148] Bernardo Rudy et al. “Three groups of interneurons account for nearly 100% of neocortical GABAergic neurons.” In: *Developmental neurobiology* 71 (1 Jan. 2011), pp. 45–61. ISSN: 1932-846X. DOI: 10.1002/dneu.20853. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3556905&tool=pmcentrez&rendertype=abstract>.
- [149] Emilio Salinas and Peter Thier. “Gain Modulation : A Major Computational Principle of the Central Nervous System”. In: *Neuron* 27 (x 2000), pp. 15–21. ISSN: 08966273. DOI: 10.1016/S0896-6273(00)00004-0.
- [150] Massimo Scanziani and Michael Häusser. “Electrophysiology in the age of light”. In: *Nature* 461.7266 (2009), pp. 930–939.
- [151] Benjamin Scholl et al. “Local Integration Accounts for Weak Selectivity of Mouse Neocortical Parvalbumin Interneurons”. In: *Neuron* 87 (2 2015), pp. 424–437. ISSN: 10974199. DOI: 10.1016/j.neuron.2015.06.030. URL: <http://dx.doi.org/10.1016/j.neuron.2015.06.030>.
- [152] D Schubert et al. “Layer-specific intracolumnar and transcolumnar functional connectivity of layer V pyramidal cells in rat barrel cortex.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 21 (10 2001), pp. 3580–3592. ISSN: 1529-2401. DOI: 21/10/3580[pii].
- [153] Murat Sekerli et al. “Estimating action potential thresholds from neuronal time-series: New metrics and evaluation of methodologies”. In: *IEEE Transactions on Biomedical Engineering* 51 (9 2004), pp. 1665–1672. ISSN: 00189294. DOI: 10.1109/TBME.2004.827531.
- [154] Bryan A Seybold et al. “Inhibitory Actions Unified by Network Integration”. In: *Neuron* 87 (6 2015), pp. 1181–1192. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2015.09.013. URL: <http://dx.doi.org/10.1016/j.neuron.2015.09.013>.
- [155] Ben Shababo et al. “Bayesian inference and online experimental design for mapping neural microcircuits”. In: *Advances in Neural Information Processing Systems (NIPS)* (2013), pp. 1304–1312.
- [156] Or A. Shemesh et al. “Temporally precise single-cell-resolution optogenetics”. In: *Nature Neuroscience* 20.12 (Dec. 2017), pp. 1796–1806. ISSN: 1546-1726. DOI: 10.1038/s41593-017-0018-8. URL: <https://doi.org/10.1038/s41593-017-0018-8>.
- [157] Gordon M. G. Shepherd and Karel Svoboda. “Laminar and Columnar Organization of Ascending Excitatory Projections to Layer 2/3 Pyramidal Neurons in Rat Barrel Cortex”. In: *Journal of Neuroscience* 25.24 (2005), pp. 5670–5679. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.1173-05.2005. eprint: <https://www.jneurosci.org/content/25/24/5670.full.pdf>. URL: <https://www.jneurosci.org/content/25/24/5670>.
- [158] Gordon M. G. Shepherd et al. “Geometric and functional organization of cortical circuits”. In: *Nature Neuroscience* 8.6 (June 2005), pp. 782–790.

- [159] Gordon M.G Shepherd, Thomas A Pologruto, and Karel Svoboda. “Circuit Analysis of Experience-Dependent Plasticity in the Developing Rat Barrel Cortex”. In: *Neuron* 38.2 (2003), pp. 277–289. ISSN: 0896-6273. DOI: [https://doi.org/10.1016/S0896-6273\(03\)00152-1](https://doi.org/10.1016/S0896-6273(03)00152-1). URL: <https://www.sciencedirect.com/science/article/pii/S0896627303001521>.
- [160] Yulin Shi, Zoran Nenadic, and Xiangmin Xu. “Novel use of matched filtering for synaptic event detection and extraction”. In: *PLoS One* 5.11 (2010), e15517.
- [161] Gilad Silberberg and Henry Markram. “Disynaptic inhibition between neocortical pyramidal cells mediated by Martinotti cells.” In: *Neuron* 53 (5 Mar. 2007), pp. 735–46. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2007.02.012. URL: <http://www.ncbi.nlm.nih.gov/pubmed/17329212>.
- [162] V.Y.F. Tan and V.K. Goyal. “Estimating Signals With Finite Rate of Innovation From Noisy Samples: A Stochastic Algorithm”. In: *IEEE Transactions on Signal Processing* 56.10 (Oct. 2008), pp. 5135–5146.
- [163] Jonathan C Y Tang et al. “Cell type-specific manipulation with GFP-dependent Cre recombinase”. In: *Nature Neuroscience* 18 (9 2015), pp. 1334–1341. ISSN: 1097-6256. DOI: 10.1038/nn.4081. URL: <http://dx.doi.org/10.1038/nn.4081>.
- [164] Hiroki Taniguchi et al. “A Resource of Cre Driver Lines for Genetic Targeting of GABAergic Neurons in Cerebral Cortex”. In: *Neuron* 71 (6 Sept. 2011), pp. 995–1013. ISSN: 08966273. DOI: 10.1016/j.neuron.2011.07.026. URL: <http://www.sciencedirect.com/science/article/pii/S0896627311006799> <http://linkinghub.elsevier.com/retrieve/pii/S0896627311006799>.
- [165] Bosiljka Tasic et al. “Adult mouse cortical cell taxonomy revealed by single cell transcriptomics”. In: *Nature Neuroscience* advance on (January 2016), pp. 1–37. ISSN: 1097-6256. DOI: 10.1038/nn.4216. URL: <http://dx.doi.org/10.1038/nn.4216>.
- [166] L. Theis et al. “Benchmarking spike rate inference in population calcium imaging”. In: *Neuron* 90.3 (May 2016), pp. 471–482.
- [167] Alex Thomson and Christophe Lamy. “Functional maps of neocortical local circuitry”. In: *Frontiers in Neuroscience* 1 (2007). ISSN: 1662-453X. DOI: 10.3389/neuro.01.1.1.002.2007. URL: <https://www.frontiersin.org/article/10.3389/neuro.01.1.1.002.2007>.
- [168] Yu Toyoshima et al. “Accurate Automatic Detection of Densely Distributed Cell Nuclei in 3D Space”. In: *PLOS Computational Biology* 12.6 (June 2016), pp. 1–20. DOI: 10.1371/journal.pcbi.1004970. URL: <https://doi.org/10.1371/journal.pcbi.1004970>.
- [169] Robin Tremblay, Soohyun Lee, and Bernardo Rudy. “GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits”. In: *Neuron* 91 (2 2016), pp. 260–292. ISSN: 08966273. DOI: 10.1016/j.neuron.2016.06.033. URL: <http://dx.doi.org/10.1016/j.neuron.2016.06.033>.

- [170] Joanna Urban-Ciecko and Alison L. Barth. “Somatostatin-expressing neurons in cortical networks.” In: *Nature reviews. Neuroscience* 17 (7 2016), pp. 401–9. ISSN: 1471-0048. DOI: 10.1038/nrn.2016.53. URL: <http://www.nature.com/doifinder/10.1038/nrn.2016.53%5C%5Cnhttp://www.ncbi.nlm.nih.gov/pubmed/27225074>.
- [171] Joshua T Vogelstein et al. “Fast nonnegative deconvolution for spike train inference from population calcium imaging”. In: *Journal of Neurophysiology* 104.6 (2010), pp. 3691–3704.
- [172] Joshua T Vogelstein et al. “Spike inference from calcium imaging using sequential Monte Carlo methods”. In: *Biophysical Journal* 97.2 (2009), pp. 636–655.
- [173] Yun Wang et al. “Anatomical, physiological and molecular properties of Martinotti cells in the somatosensory cortex of the juvenile rat.” In: *The Journal of physiology* 561 (Pt 1 Nov. 2004), pp. 65–90. ISSN: 0022-3751. DOI: 10.1113/jphysiol.2004.073353. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1665344&tool=pmcentrez&rendertype=abstract>.
- [174] Norbert Wiener. *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*. Technology press books in science and engineering. Technology Press of the Massachusetts Institute of Technology, 1949.
- [175] Stephen R Williams and Simon J Mitchell. “Direct measurement of somatic voltage clamp errors in central neurons.” In: *Nature neuroscience* 11 (7 July 2008), pp. 790–8. ISSN: 1097-6256. DOI: 10.1038/nn.2137. URL: <http://www.ncbi.nlm.nih.gov/pubmed/18552844>.
- [176] Nathan R. Wilson et al. “Division and subtraction by distinct cortical inhibitory networks in vivo”. In: *Nature* 488 (7411 Aug. 2012), pp. 1–6. ISSN: 0028-0836. DOI: 10.1038/nature11347. URL: <http://www.nature.com/doifinder/10.1038/nature11347%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3653570&tool=pmcentrez&rendertype=abstract>.
- [177] V. C. Wimmer et al. “Dimensions of a Projection Column and Architecture of VPM and POm Axons in Rat Vibrissal Cortex”. In: *Cerebral Cortex* 20 (10 2010), pp. 2265–2276. ISSN: 1047-3211. DOI: 10.1093/cercor/bhq068. URL: <http://www.cercor.oxfordjournals.org/cgi/doi/10.1093/cercor/bhq068>.
- [178] Chaowen Wu et al. “rAAV-Mediated Subcellular Targeting of Optogenetic Tools in Retinal Ganglion Cells In Vivo”. In: *PLOS ONE* 8.6 (June 2013), pp. 1–10. DOI: 10.1371/journal.pone.0066332. URL: <https://doi.org/10.1371/journal.pone.0066332>.
- [179] Z. Xiang, John R Huguenard, and David a Prince. “Cholinergic Switching Within Neocortical Inhibitory Networks”. In: *Science* 281 (5379 1998), pp. 985–988. DOI: 10.1126/science.281.5379.985. URL: <http://www.sciencemag.org/cgi/doi/10.1126/science.281.5379.985>.
- [180] Han Xu et al. “Neocortical Somatostatin-Expressing GABAergic Interneurons Disinhibit the Thalamorecipient Layer 4”. In: *Neuron* 77 (1 2013), pp. 155–167. ISSN: 08966273. DOI: 10.1016/j.neuron.2012.11.004. URL: <http://dx.doi.org/10.1016/j.neuron.2012.11.004>.

- [181] Xiangmin Xu and Edward M Callaway. “Laminar specificity of functional input to distinct types of inhibitory cortical neurons.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 29 (1 Jan. 2009), pp. 70–85. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.4104-08.2009. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2656387&tool=pmcentrez&rendertype=abstract>.
- [182] Xiangmin Xu and Edward M. Callaway. “Laminar Specificity of Functional Input to Distinct Types of Inhibitory Cortical Neurons”. In: *Journal of Neuroscience* 29.1 (2009), pp. 70–85. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.4104-08.2009. eprint: <https://www.jneurosci.org/content/29/1/70.full.pdf>. URL: <https://www.jneurosci.org/content/29/1/70>.
- [183] Zizhen Yao et al. “A transcriptomic and epigenomic cell atlas of the mouse primary motor cortex”. In: *Nature* 598.7879 (Oct. 2021), pp. 103–110. ISSN: 1476-4687. DOI: 10.1038/s41586-021-03500-8. URL: <https://doi.org/10.1038/s41586-021-03500-8>.
- [184] Iryna Yavorska and Michael Wehr. “Somatostatin-Expressing Inhibitory Interneurons in Cortical Circuits”. In: *Frontiers in Neural Circuits* 10 (September 2016), pp. 1–18. ISSN: 1662-5110. DOI: 10.3389/fncir.2016.00076. URL: <http://journal.frontiersin.org/article/10.3389/fncir.2016.00076>.
- [185] Yumiko Yoshimura and Edward M Callaway. “Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity.” In: *Nature neuroscience* 8 (11 Nov. 2005), pp. 1552–9. ISSN: 1097-6256. DOI: 10.1038/nn1565. URL: <http://www.ncbi.nlm.nih.gov/pubmed/16222228>.
- [186] Rafael Yuste. “Circuit neuroscience: the road ahead”. eng. In: *Frontiers in neuroscience* 2.1 (Aug. 2008). PMC2570080[pmcid], pp. 6–9. ISSN: 1662-453X. DOI: 10.3389/neuro.01.017.2008. URL: <https://doi.org/10.3389/neuro.01.017.2008>.

Chapter 5

Appendix

Further details on the model and algorithm from Chapter 2.

5.1 Technical Details for Gaussian Processes Learning

In this section, we discuss the procedure for learning the mean and variance functions for the Gaussian process prior on the shape function $s_j(\cdot)$. For reference, we introduced the the optical response experiments to collect pilot data in Section 2.5.2. We discuss how to normalize the measured current with respect to power and each cell's unique optical gain in Section 5.1.1. We explain the model for the variability in the induced current in Section 5.1.2. We explain how to estimate the mean and variance functions in Section 5.1.3.

Recall from Section 2.5.2 that we need the mean function $\mu(\cdot)$, the variance function $\sigma^2(\cdot)$, and τ in the kernel $K(\cdot, \cdot)$ to define the Gaussian process. We use the mean and variance functions estimated from the response data and set the kernel parameter $\tau = 7.5, 7.5$, or 20 for the three dimensions, respectively.

5.1.1 Normalization of measured current

To extract the function of interest $s_j(\cdot)$ from the measured current (2.1), we normalize the raw measurements $m_{i,k}$ by the power level as well as the maximum current for that cell. To be specific, we first adjust the raw measurement by the power in the corresponding trial

$$\hat{m}_{i,k} \equiv m_{j,i}/I_{j,i}, \quad (5.1)$$

and then by the maximum power-adjusted current of this cell

$$\tilde{m}_{j,i} \equiv \hat{m}_{j,i}/\max_k(\bar{m}_{j,k}), \quad (5.2)$$

where $\bar{m}_{j,k} = \sum_{i:v_i=v_k} \hat{m}_{j,i}/\text{card}(\{i : v_i = v_k\})$.

Plugging (2.1) into (5.1), we can see that

$$\hat{m}_{j,i} = s_j(v_i)\phi_j + \frac{c_{j,i}\epsilon_{j,i}}{I_{j,i}}, \quad (5.3)$$

and

$$\bar{m}_{j,k} = s_j(v_k)\phi_j + \sum_{i:v_i=v_k} \frac{c_{j,i,\epsilon}\epsilon_{j,i}}{I_{j,i}\text{card}(\{i : v_i = v_k\})}. \quad (5.4)$$

Therefore, the normalized current is

$$\tilde{m}_{j,i} = \frac{s_j(v_i)\phi_j}{\max_k \bar{m}_{j,k}} + \frac{c_{j,i,\epsilon}\epsilon_{j,i}}{I_{j,i}\max_k \bar{m}_{j,k}}. \quad (5.5)$$

Recalling that $\max_v s_j(v) \equiv 1$, we have

$$\tilde{m}_{j,i} \approx s_j(v_i) + \frac{c_{j,i,\epsilon}\epsilon_{j,i}}{I_{j,i}\phi_j}, \quad (5.6)$$

where we use the approximation that $\max_k \bar{m}_{j,k} \approx \phi_j$

Recall that the shape function should satisfy that $\max_i (s_j(x_i, 0, 0)) = 1$. Thus,

$$\tilde{m}_{j,i,X} = s_j(x_i, 0, 0) + c_{j,i,\epsilon}\epsilon_{j,i}/\{I_{j,i}\max_i (s_j(x_i, 0, 0))\phi_j\}. \quad (5.7)$$

Therefore, when learning the GP using $\tilde{m}_{j,i,X}$, we need to account for the change of variance as shown in (5.7). Specifically, let $\hat{c}_{j,i,\epsilon}$ be an estimate of the standard deviation, we have $\tilde{c}_{j,i,\epsilon} = \hat{c}_{j,i,\epsilon}/\{I_{j,i}\max_k(\bar{x}_{j,k,X})\}$.

5.1.2 Estimating the scaled variance

We assume that the standard deviation $c_{j,i,\epsilon}$ in (2.15) is a linear function of the mean current. In other words, $c_{j,i,\epsilon} = a\mathbb{E}[m_{j,i}] + b$. Using the repeated measurements at the same stimulation location, we have

$$\bar{m}_{j,k,X} = \sum_{i \in S_{i'}} \tilde{m}_{j,i,X} / \text{card}(S_{i'}), \quad (5.8)$$

and

$$\hat{c}_{j,k,\epsilon}^2 = \frac{1}{\text{card}(S_{i'}) - 1} \sum_{i \in S_{i'}} (\tilde{m}_{j,i,X} - \bar{m}_{j,k,X})^2, \quad (5.9)$$

$\{x_{j,i'}^u\}_{i'=1}^{n'}$ is the set of unique stimulation locations for cell j , $S_{i'} \equiv \{i : x_{j,i} = x_{j,i'}^u\}$ is the index set of $x_{j,i}$ that equals $x_{j,k}^u$, and $\text{card}(S_{i'})$ is the number of indices in $S_{i'}$.

5.1.3 Estimating the mean and variance function with boundary conditions

Our goal here is to obtain a population mean shape given the noisy measurements of multiple cell shapes. To this end, we use a Gaussian process $GP(\mu^0(\cdot), \Sigma^0(\cdot, \cdot))$ as our smoother. We set the mean function to be zero, i.e., $\mu^0(\cdot) \equiv 0$, and set the kernel parameter τ to be a constant (7.5 for $l = 1, 2$ and 20 for $l = 3$).

To enforce the boundary condition, that the mean function is zero beyond some boundaries, we set the standard deviation function $\sigma^0(\cdot)$ as

$$\sigma^0(v) = a \left[\frac{1}{1 + \exp(-B - v)} \frac{1}{1 + \exp(v - B)} \right], \quad (5.10)$$

Figure 5.1: Posterior mean of the GP fits on the mean function (left) and the prior variance for the GP. The red dots represent the normalized current $\tilde{m}_{j,i}$, the blue curve represent the noninformative variance $\sigma^0(\cdot)$, and the black curve represents the posterior mean, i.e, the estimated mean function $\tilde{\mu}(\cdot)$ in Section 2.4.

where $a = 1$, $B = 70$ for $l = 1, 2$, or $B = 120$ for $l = 3$. The posterior mean is our estimated mean function $\tilde{\mu}_X(\cdot)$.

Given the estimated mean function $\hat{\mu}(\cdot)$, we can calculate the squared deviations $[\tilde{m}_{j,i} - \hat{\mu}(v_{j,i})]^2$. We can then take a similar strategy to learn a smooth variance function $\tilde{\sigma}(\cdot)$ with a GP prior $GP(\mu^0(\cdot), \Sigma^0(\cdot, \cdot))$, and a crude observation model

$$[\tilde{m}_{j,i} - \hat{\mu}(v_{j,i})]^2 = \sigma^2(v_{j,i}) + \dot{\epsilon}_{j,i}, \quad (5.11)$$

where $\dot{\epsilon}_{j,i} \sim \mathcal{N}(0, 1)$.

5.2 Implementation Details of Algorithm 1

5.2.1 Logit-normal Distributions

For logit-normal distributions, take ϕ_j for instance, we have

$$\bar{\phi}_j \equiv \log \left(\frac{\phi_j - \phi_{\text{low}}}{\phi_{\text{up}} - \phi_j} \right) \sim \mathcal{N}(\alpha_{\phi,j,1}, \exp(\alpha_{\phi,j,2})),$$

and $\phi_j = [\exp(\bar{\phi}_j)\phi_{\text{up}} + \phi_{\text{low}}]/[\exp(\bar{\phi}_j) + 1]$. As a result, the density function is

$$q(\phi_j | \alpha_{\phi,j,1}, \alpha_{\phi,j,2}) = \frac{1}{\exp(\alpha_{\phi,j,2}/2)\sqrt{2\pi}} \exp \left[-\frac{(\bar{\phi}_j - \alpha_{\phi,j,1})^2}{2\exp(\alpha_{\phi,j,2})} \right] \frac{\phi_{\text{up}} - \phi_{\text{low}}}{(\phi_j - \phi_{\text{low}})(\phi_{\text{up}} - \phi_j)}. \quad (5.12)$$

And thus the logarithm of $q(\phi_j | \alpha_{\phi,j,1}, \alpha_{\phi,j,2})$ takes the form

$$\log q(\phi_j | \alpha_{\phi,j,1}, \alpha_{\phi,j,2}) = C - \frac{1}{2}\alpha_{\phi,j,2} - \frac{(\bar{\phi}_j - \alpha_{\phi,j,1})^2}{2\exp(\alpha_{\phi,j,2})} - \log(\phi_j - \phi_{\text{low}}) - \log(\phi_{\text{up}} - \phi_j). \quad (5.13)$$

The partial derivatives are

$$\frac{\partial}{\partial \alpha_{\phi,j,1}} \log q(\phi_j | \alpha_{\phi,j,1}, \alpha_{\phi,j,2}) = \frac{\bar{\phi}_j - \alpha_{\phi,j,1}}{\exp(\alpha_{\phi,j,2})}, \quad \frac{\partial}{\partial \alpha_{\phi,j,2}} \log q(\phi_j | \alpha_{\phi,j,1}, \alpha_{\phi,j,2}) = -0.5 + \frac{(\bar{\phi}_j - \alpha_{\phi,j,1})^2}{2\exp(\alpha_{\phi,j,2})}.$$

5.2.2 Multivariate Normal Distributions

We approximate the posterior distribution of $s_j(\mathbf{L} - \mathbf{w}_j)$ using a multivariate Gaussian distribution $\mathcal{N}(\tilde{\boldsymbol{\mu}}_j, \tilde{\boldsymbol{\Sigma}}_j)$, where $\tilde{\boldsymbol{\Sigma}}_j^{-1} = \boldsymbol{\Sigma}_j^{-1} + \mathbf{D}_j^{-1}$ with \mathbf{D}_j being a diagonal matrix. We slightly abuse the notation and denote that $\mathbf{s}_j \equiv s_j(\mathbf{L} - \mathbf{w}_j)$ and assume that the $\mathbf{L}_j \in \mathbb{R}^{n \times 3}$.

The variational parameters to learn are $\boldsymbol{\mu}_j \equiv (\mu_{j,1}, \dots, \mu_{j,n})$ and $\mathbf{D}_j \equiv \text{diag}(d_{j,1}, \dots, d_{j,n})$. We drop the subscript j for simplicity in this section.

The log-likelihood of \mathbf{s} is

$$\begin{aligned} \log q(\mathbf{s} \mid \boldsymbol{\mu}, \mathbf{D}) &\equiv = \frac{1}{2} \log \det(2\pi \tilde{\boldsymbol{\Sigma}})^{-1} - \frac{1}{2} (\mathbf{s} - \boldsymbol{\mu})^T \tilde{\boldsymbol{\Sigma}}^{-1} (\mathbf{s} - \boldsymbol{\mu}) \\ &= \frac{n}{2} \log(2\pi) + \frac{1}{2} \log \det(\tilde{\boldsymbol{\Sigma}}^{-1}) - \frac{1}{2} (\mathbf{s} - \boldsymbol{\mu})^T \tilde{\boldsymbol{\Sigma}}^{-1} (\mathbf{s} - \boldsymbol{\mu}) \\ &= \frac{n}{2} \log(2\pi) + \frac{1}{2} \log \det(\boldsymbol{\Sigma}^{-1} + \mathbf{D}^{-1}) - \frac{1}{2} (\mathbf{s} - \boldsymbol{\mu})^T (\boldsymbol{\Sigma}^{-1} + \mathbf{D}^{-1}) (\mathbf{s} - \boldsymbol{\mu}) \end{aligned} \quad (5.14)$$

The derivatives w.r.t. μ_i and d_i are as follows

$$\begin{aligned} \frac{\partial}{\partial \mu_i} \log q(\mathbf{s} \mid \boldsymbol{\mu}, \mathbf{D}) &= -\frac{1}{2} \frac{\partial}{\partial \mu_i} \left[(\mathbf{s} - \boldsymbol{\mu})^T \tilde{\boldsymbol{\Sigma}}^{-1} (\mathbf{s} - \boldsymbol{\mu}) \right] \\ &= \left(\tilde{\boldsymbol{\Sigma}}^{-1} \right)_{i, \cdot} (\boldsymbol{\mu} - \mathbf{s}), \end{aligned} \quad (5.15)$$

and

$$\begin{aligned} \frac{\partial}{\partial d_i} \log q(\mathbf{s} \mid \boldsymbol{\mu}, \mathbf{D}) &= \frac{1}{2} \text{tr} \left(\tilde{\boldsymbol{\Sigma}} \frac{\partial}{\partial d_i} (\boldsymbol{\Sigma}^{-1} + \mathbf{D}^{-1}) \right) - \frac{1}{2} \frac{\partial}{\partial d_i} \left[(\mathbf{s} - \boldsymbol{\mu})^T (\boldsymbol{\Sigma}^{-1} + \mathbf{D}^{-1}) (\mathbf{s} - \boldsymbol{\mu}) \right] \\ &= \frac{1}{2} \text{tr} \left(\tilde{\boldsymbol{\Sigma}} \frac{\partial}{\partial d_i} (\mathbf{D}^{-1}) \right) + \frac{1}{2} (s_i - \mu_i)^2 d_i^{-2} \\ &= -\frac{1}{2} \tilde{\sigma}_{i,i} d_i^{-2} + \frac{1}{2} (s_i - \mu_i)^2 d_i^{-2}, \end{aligned} \quad (5.16)$$

where $\tilde{\sigma}_{i,i}$ is the i th diagonal element of $\tilde{\boldsymbol{\Sigma}}$.

5.2.3 Multivariate Normal Distributions with Constraints on Parameters

We want to make sure that the variational distribution of the shape values falls in the plausible region given our prior knowledge on $\mathbf{s}(\cdot)$. Thus, we set the variational parameters as

$$\mu_i = \frac{\exp(\alpha_{i,1})}{\exp(\alpha_{i,1}) + 1} (\mu_{i,\max} - \mu_{i,\min}) + \mu_{i,\min}, \quad d_i = \exp(\alpha_{i,2}), \quad (5.17)$$

where $\mu_{i,\max}$ and $\mu_{i,\min}$ are the upper and lower bound of the shape value at this location from the prior distribution.

The derivatives w.r.t. $\alpha_{i,1}$ and $\alpha_{i,2}$ are as follows

$$\frac{\partial}{\partial \alpha_{i,1}} \log q(\mathbf{s} \mid \boldsymbol{\mu}, \mathbf{D}) = -\frac{\exp(\alpha_{i,1}) (\mu_{i,\max} - \mu_{i,\min})}{[\exp(\alpha_{i,1}) + 1]^2} \left(\tilde{\boldsymbol{\Sigma}}^{-1} \right)_{i, \cdot} (\boldsymbol{\mu} - \mathbf{s}), \quad (5.18)$$

and

$$\frac{\partial}{\partial \alpha_{i,2}} \log q(\mathbf{s} \mid \boldsymbol{\mu}, \mathbf{D}) = -\frac{1}{2} \tilde{\sigma}_{i,i} \exp(-\alpha_{i,2}) + \frac{1}{2} (s_i - \mu_i)^2 \exp(-\alpha_{i,2}). \quad (5.19)$$

5.3 Table of Notation

Variable	Definition
\mathbf{w}_k	Spatial coordinates of the nucleus of neuron k
\mathbf{x}_i	Spatial coordinates of the stimulation in the i th trial
I_i	Intensity of optogenetic stimulation in the i th trial
$J_i(t)$	Postsynaptic current in the i th trial
\mathcal{T}_i	Set of detected postsynaptic events $\{t_{i,j} : j = 1, \dots, n_i\}$ in the i th trial
M	Number of trials
p	Number of neurons
$m_{i,k}$	Magnitude of elicited current in neuron k in the i th trial
$s_k(\cdot)$	Spatial distribution of opsin expression level of neuron k
ϕ_k	Overall opsin expression level of neuron k
$g(m)$	Mean first spike time given the magnitude of induced current m
$h^2(m)$	Variance of first spike time given the magnitude of induced current m
$\hat{\lambda}_{i,k}$	Intensity of the first spike time of neuron k in the i th trial
$\psi(\cdot)$	Density function of a standard normal distribution
d_k	Mean transmission delay of neuron k
$e_{i,k}$	Random transmission delay of neuron k in the i th trial
δ_k^2	Variance of transmission delay of neuron k
t_{i,j_k}	Postsynaptic event time induced by a spike of neuron k in the i th trial
γ_k	Probability of transmission of neuron k
\mathcal{D}_i	Data recorded in the i th trial
\mathcal{D}	Data recorded in the mapping experiment
$\lambda_{i,k}$	Intensity of postsynaptic event induced by neuron k in the i th trial
Ψ	Cumulative distribution function of a standard normal
$\mathcal{A}(p, n_i)$	Set of all length- n_i permutations out of the p neurons
$\mathbf{v}_{l,k}$	Relative position of the l th unique stimulation location to the nucleus of neuron k
L	Number of unique stimulation locations
\mathbf{V}_k	Matrix of $(\mathbf{v}_{1,k}, \dots, \mathbf{v}_{L,k})$
Θ_k	Set of unknown parameters related to neuron k , i.e., $\{\gamma_k, d_k, s_k(\cdot), \phi_k, \delta_k\}$
Θ	Set of all unknown parameters, i.e., $\Theta = \{\Theta_k : k = 1, \dots, p\}$
$\mathcal{L}(\mathcal{D}; \Theta)$	Likelihood function
θ	Variable in Θ
$p_0(\theta; \alpha_\theta^0, \beta_\theta^0, \mathbf{b}_\theta)$	Prior distribution of θ
$\alpha_\theta^0, \beta_\theta^0$	Parameters in the prior distribution of θ
\mathbf{b}_θ	Bounds of θ , $\mathbf{b}_\theta = (b_{\theta,1}, b_{\theta,2})^T$
$f(\theta; \mathbf{b}_\theta)$	Logit transformation of θ with bounds \mathbf{b}_θ
$\mu(\cdot)$	Mean function of the Gaussian process for the prior of $s_k(\cdot)$
$\Sigma(\cdot, \cdot)$	Variance function of the Gaussian process for the prior of $s_k(\cdot)$
α, β	Variational parameters
$\text{ELBO}(\alpha, \beta)$	Evidence lower bound
$q(\theta; \alpha_\theta, \beta_\theta, \mathbf{b}_\theta)$	Variational distribution for θ
\mathbf{B}	Diagonal matrix in the variational distribution for $(s_k(\mathbf{v}_{1,k}), \dots, s_k(\mathbf{v}_{L,k}))^T \in \mathbb{R}^L$
$\Theta[s]$	Monte-Carlo sample of Θ drawn from the variational distribution