Characterization of Response Properties and Connectivity in Mouse Visual Thalamus and Cortex

Simon Renner

Graduate School of
Systemic Neurosciences
LMU Munich

Dissertation der
Graduate School of Systemic Neurosciences
Ludwig-Maximilians-Universität München

January 2022

Characterization of Response Properties and Connectivity in Mouse Visual Thalamus and Cortex

Doctoral dissertation
by Simon Renner
Graduate School of Systemic Neurosciences
Ludwig-Maximilians-Universität München
January 2022

Supervisor
Prof. Dr. Laura Busse
Department Biologie II, Division Neurobiology
Ludwig-Maximilians-Universität München

First Reviewer:	Prof. Dr. Laura Busse
Second Reviewer:	Prof. Dr. Tatjana Tchumatchenko
Third Reviewer:	Prof. Dr. Matthias Kaschube

Date of submission:	January 20th 2022
Date of defense:	June 22nd 2022

Acknowledgements

I first want to thank three people without whom this thesis would never have started. Prof. Dr. Thorsten Stein first encouraged me to travel the path of neuroscience and Prof. Dr. Marco Taubert guided me during my first steps in this field. Prof. Dr. Stefan Glasauer gave me the chance to come to Munich and start my PhD at the GSN. Dear Thorsten, Marco, and Stefan, thank you for your trust and support, my path would have been very different without you.

Foremost, I want to thank Prof. Dr. Laura Busse for supervising and guiding me through my PhD. Dear Laura, thank you for your advice, your critique, your help, your strength, your understanding and your support in all matters; you have made me a better scientist and a better person.

I want to thank the members of my Thesis Advisory Committee – Prof. Dr. Tatjana Tchumatchenko, Prof. Dr. Thomas Wachtler, and Dr. Martin Stemmler – for the advice and fruitful discussions during my PhD.

I want to thank my collaborators – Prof. Dr. Tatjana Tchumatchenko and Dr. Nataliya Kraynyukova – for the chance to work with them and learn from them.

I want to thank the entire Neurobiology Department for three great years at the Biocenter; especially Yannik, who recruited me into the lab, Martin, who taught me how to perform experiments and be a better programmer, and my office mates Gregory and Felix, who made the bad times bearable and the good times even better. Thank you for the support, the discussions, and the many happy moments, I will miss you.

I also want to thank Davide and Lukas for proofreading my thesis, your keen eyes and helpful comments have greatly improved this text.

I want to thank the *SPP2041 Computational Connectomics* for funding my PhD and providing an embedding into the German neuroscience scene, the *RTG2175 Perception in Context and its Neural Basis* for interesting courses and wonderful retreats, and the GSN for a beautiful and supportive environment to study in.

My deepest gratitude goes towards my girlfriend Lisa, who has been with me through all the ups and downs of this journey, and my family, who has given me

unconditional support and freedom on my path towards this degree. Thank you for your love and support, I could not have done this without you.

Summary

How neuronal activity is shaped by circuit connectivity between neuronal populations is a central question in visual neuroscience. Combined with experimental data, computational models allow causal investigation and prediction of both how connectivity influences activity and how activity constrains connectivity. In order to develop and refine these computational models of the visual system, thorough characterization of neuronal response patterns is required. In this thesis, I first present an approach to infer connectivity from *in vivo* stimulus responses in mouse visual cortex, revealing underlying principles of connectivity between excitatory and inhibitory neurons. Second, I investigate suppressed-by-contrast neurons, which, while known since the 1960s, still remain to be included in standard models of visual function. I present a characterization of intrinsic firing properties and stimulus responses that expands the knowledge about this obscure neuron type.

Inferring the neuronal connectome from neural activity is a major objective of computational connectomics. Complementary to direct experimental investigation of connectivity, inference approaches combine simultaneous activity data of individual neurons with methods ranging from statistical considerations of similarity to large-scale simulations of neuronal networks. However, due to the mathematically ill-defined nature of inferring connectivity from *in vivo* activity, most approaches have to constrain the inference procedure using experimental findings that are not part of the neural activity data set at hand. Combining the stabilized-supralinear network model with response data from the visual thalamus and cortex of mice, my collaborators and I have found a way to infer connectivity from *in vivo* data *alone*. Leveraging a property of neural responses known as contrast-invariance of orientation tuning, our inference approach reveals a consistent order of connection strengths between cortical neuron populations as well as tuning differences between thalamic inputs and cortex.

Throughout the history of visual neuroscience, neurons that respond to a visual stimulus with an increase in firing have been at the center of attention. A different response type that decreases its activity in response to visual stimuli, however, has

been only sparsely investigated. Consequently, these suppressed-by-contrast neurons, while recently receiving renewed attention from researchers, have not been characterized in depth. Together with my collaborators, I have conducted a survey of SbC properties covering firing reliability, cortical location, and tuning to stimulus orientation. We find SbC neurons to fire less regularly than expected, be located in the lower parts of cortex, and show significant tuning to oriented gratings.

Contents

Summary **5**

1 Introduction **9**
1.1 Primer: Function and Connectivity 9
1.2 Emergence of Orientation Selectivity 11
1.3 Models of Orientation Selectivity 11
1.4 Connectivity 16
1.5 Inference of Connectivity 19
1.6 Functional Neuron Classification 21
1.7 Aims of this thesis 22

2 Model-based Connectivity Inference in V1 **23**
2.1 Summary 23
2.2 Contributions 23

3 Suppressed-by-Contrast Neurons in dLGN and V1 **59**
3.1 Summary 59
3.2 Contributions 59

4 Conclusion **79**
4.1 Discussion 79
4.2 Outlook 82

Bibliography **83**
List of Publications 101
Author Contributions 102

1 Introduction

1.1 Primer: Function and Connectivity

Our visual system helps us to extract relevant information from the constant stream of input that our surroundings bombard our eyes with. Illusions like the Rubin vase, which doubles as two faces in profile, or Daniel Simons' *monkey business*, where we fail to spot a Gorilla running through a scene we observe, demonstrate powerfully that this extraction of information is neither entirely truthful, nor is it trivial. It takes complex pattern recognition to dissect a visual image into meaningful components and represent them in conscious or unconscious perception. So how does the brain handle this tremendous challenge? Today, with public attention on deep learning and self-driving cars, it may seem obvious to think of the brain as a network of computational units that send signals back and forth to extract information patterns. But it was not always so. At the end of the 19^{th} century, when many believed the brain was working as one unified construct, Ramón y Cajal identified the individual neuron as the main functional unit of the brain. With Sherrington formulating the idea of individual synapses as connection points between these neurons and the discovery of the action potential as a communication signal passing along these synapses, the ingredients for a distributed information processing system were in place. In the following decades, pioneering work in the visual system led to the discovery of receptive fields, convergent connections, and the formulation of distributed coding principles in the brain. Equally supported by evidence about how the brain stores memories and navigates in space, the idea that neural networks can extract and store information because of their specific network connections was established as a fundamental principle in neuroscience.

Today, uncovering the direct relation between structure and function of brain networks lies at the heart of modern neuroscience and provides one of its biggest challenges: observing brain structure and function *at the same time*. As an experimental technique that can provide such a unified picture is still out-of-reach, any attempt at directly relating structure and function must use combinations of

different methods that capture aspects of one or the other, and carefully piece together the evidence. However, even with such an approach, structure and function are still measured *separately* both in time and space. Additionally, what can be measured about network structure depends highly on the method employed, with electron microscopy giving a much different picture than synaptic tracing or current stimulation techniques. But beyond the methodological differences, which range from anatomical descriptions to *in vitro* response propagation in neural tissue, lies a joint characteristic: the *computational* structure of the network. In it, all biological details about ion channels, membrane potentials, transmitters, and synaptic boutons are combined into one fundamental property: how information flows from one neuron to the next. This is an abstract property that can be implemented in biological structure and can even remain intact as the detailed physical structure changes over time. Motivated by such robustness, the question arises if the *computational* structure of a brain network could be determined *computationally*, that is to say, by combining experimental data with a mathematical model of brain function. In the first part of this thesis, I will present a study that uses the stabilized supralinear network (Ahmadian et al., 2013), a state-of-the-art cortical model, to infer connectivity in mouse visual cortex.

The stabilized supralinear network model used in the first part of the thesis highlights the interaction between populations of excitatory and inhibitory neurons in cortex. At the same time, it excludes suppressed-by-contrast (SbC) neurons (Rodieck, 1967), a neuronal response type that was discovered in the the 1960s but has remained on the fringes of scientific interest. Consequently, SbC neurons, which respond to visual stimuli with a decrease in activity rather than an increase in activity, have remained understudied and have not found their way into computational models of visual function. As SbC neurons nonetheless represent a sizeable fraction of visually responsive neurons and are likely to participate in visual feature processing, a better understanding of their response properties and role in the circuit is needed. In the second part of this thesis, I will present a study characterising both intrinsic and stimulus response properties of SbC neurons.

Before presenting the two manuscripts, I will cover the necessary background on response phenomena, computational models, and connectivity in the early visual system.

1.2 Emergence of Orientation Selectivity

When light hits the retina, it is transduced by photoreceptors into a bio-electrical signal and thus becomes accessible to the brain. The retina additionally starts to process the image signal, by, for example, extracting differences in contrast to reduce redundancy (Kuffler, 1953), before sending information on to the brain via retinal ganglion cells. While retinal ganglion cells project to multiple brain areas, like the hypothalamus, superior colliculus and pretectal region in mice, the strongest projection is towards the dorsolateral geniculate nucleus (dLGN) in the thalamus (Martersteck et al., 2017). This projection route is also the most important for conscious visual perception, as it further projects to the primary visual cortex (V1), the first of multiple cortical areas for visual processing (Kandel et al., 2000).

Early recordings of action potentials in the dLGN, emitted in response to flashes of light, have provided evidence that most neurons in dLGN possess center-surround receptive fields (Hubel, 1960). This means that light in a specific circular region of the visual field causes firing of a neuron at light onset (on-subfield), while light in a surrounding annulus inhibits firing (off-subfield), or vice versa. In V1 on the other hand, receptive fields are elongated into bar-like shapes, with different configurations of an on-subfield being flanked by an off-subfield, or vice versa (Hubel and Wiesel, 1962). Consequently, responses to an elongated stimulus are strongly dependent on the orientation of the stimulus, constituting one of the hallmarks of visual cortex: orientation selectivity. As a readily quantifiable property, the emergency of orientation selectivity in visual processing has served both as an experimental anchor to explore a multitude of other phenomena and as a foundation for modeling the visual system.

1.3 Models of Orientation Selectivity

The first conceptual model for the emergence of orientation selectivity during the transition from dLGN to V1 was provided by Hubel and Wiesel (1962), who proposed a feed-forward model where multiple dLGN neurons with spatially offset receptive field centers converge onto one cortical neuron, endowing that neuron with a ”simple” receptive field with an angled, elongated central region and two flanking regions of opposite polarity. These *simple cells* were in turn hypothesized to converge onto other neurons in V1 and thus form more ”complex” receptive fields that are independent of the exact location of the presented edge as long as it is oriented the correct way (Fig. 1.1). This simple model immediately offered predictions about

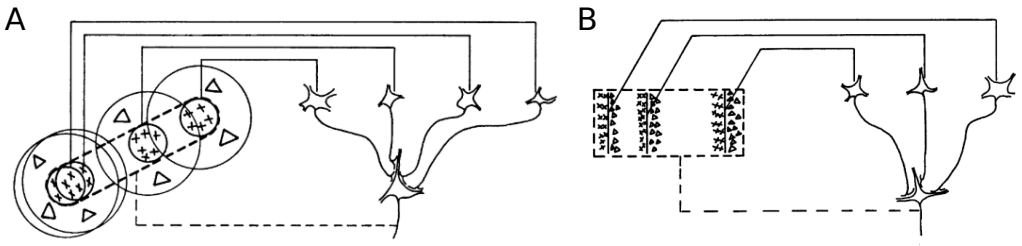

Figure 1.1: Convergence pattern in a feed-forward model. **A** Multiple dLGN neurons with center-surround receptive fields converge onto one cortical *simple cell*, forming an edge detector. **B** Multiple *simple cells* with vertical receptive fields converge onto a *complex cell*, forming an edge detector that is invariant to the spatial phase of the edge (Hubel and Wiesel, 1962, reproduced with permission from John Wiley and Sons).

how dLGN neurons and V1 neurons would connect: 1) Neurons with a central on-subfield in dLGN should be connected to neurons with a central on-subfield in V1. 2) dLGN neurons that project to V1 should be spatially aligned with their targets. Both predictions were experimentally verified in the following years (e.g. matching polarity: Reid and Alonso (1995), spatial alignment: Chapman et al. (1991)), solidifying the idea that feed-forward convergence contributes to orientation selectivity (Lien and Scanziani, 2013). However, besides the evidence that the circuit scheme connecting dLGN and V1 is actually much more complicated (Jürgens et al., 2012), there are also response phenomena that cannot be explained by such a simple model.

A surprisingly potent catalyst and checkpoint for models of orientation selectivity has been the experimentally observed contrast-invariance of the width of the orientation tuning curve, called *contrast-invariance* for short. It means that, during presentation of orientated gratings, the width of the orientation tuning curve of a neuron stays constant across different levels of contrast. The feed-forward model, however, would predict the so-called iceberg effect, which is a stark contradiction of contrast-invariance (Fig. 1.2). Experimentally confirmed in cats (Sclar and Freeman, 1982; Skottun et al., 1987), ferrets (Alitto and Usrey, 2004), primates (Carandini et al., 1997) and mice (Niell and Stryker, 2008), contrast-invariance has been shown to be a cornerstone of visual processing. As such, contrast-invariance has driven the creation of new computational models that can adequately explain how it is generated (Somers et al., 1995; Carandini et al., 1997), and continues to serve to test and refine models as they progress (Finn et al., 2007; Sadeh and Rotter, 2015).

One approach to formulating models that are consistent with contrast-invariance revolves around including cortical inhibitory neurons. Since orientation tuning was found to be highly dependent on cortical inhibition (Sillito et al., 1980), adding

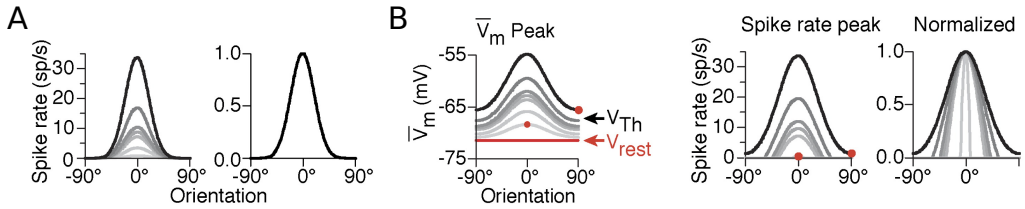

Figure 1.2: Contrast-invariance is violated in the feed-forward model. **A** *Left*: Contrast invariance means that the width of the orientation tuning curve remains constant as contrast increases. *Right*: Normalizing the curves leads to perfect superposition. **B** *Left*: In the feed-forward model, the baseline membrane potential of a V1 *simple cell* increases with contrast due to higher overall drive from dLGN inputs. The spiking threshold may be higher than the maximum membrane potential for some contrasts. *Right*: A linear spiking threshold for action potential generation leads to unequal tuning curve widths across contrast levels (Finn et al., 2007, reproduced with permission from Elsevier).

either divisive or subtractive inhibition to the feed-forward model has been proposed. Divisive inhibition normalizes the activity of a neuron by the total activity of a large group of neurons (Carandini and Heeger, 1994), while subtractive inhibition subtracts the activity of the surrounding neurons in a broadly-tuned fashion (Wörgötter and Koch, 1991). However, while these models can generate contrast-invariance, some of their properties spark controversy: divisive inhibition through changes in membrane conductance occurs in multiple contexts and brain regions (Carandini and Heeger, 2012) but seems not to be involved in orientation tuning (Douglas et al., 1988). Broad subtractive inhibition on the other hand, while being plausible from a connectivity perspective (see also section 1.4), strongly influences not only the width of a tuning curve but also its height, prompting the next step in model evolution: recurrent connections (Somers et al., 1995).

The neocortex consists of an intricately connected network of inhibitory and excitatory neurons (Binzegger, 2004), lending substantial biological plausibility to modeling cortical response phenomena with recurrently connected networks. However, the seemingly small step of adding recurrent connections propels a model into a different qualitative class: it is now a dynamical system. Dynamical systems offer great explanatory power for time-dependent processes like brain activity but they are also require complicated mathematical analysis tools to analyze and understand (Vyas et al., 2020). At the same time, their mathematics immediately introduce a property that is of vital importance in the brain: stability or homeostasis of firing activity. When the homeostasis of brain activity is strongly perturbed, for example by blocking inhibition on a large scale, the activity of excitatory neurons is no longer

balanced, which leads to runaway excitation and epileptic seizures (McCormick and Contreras, 2001). For dynamical systems models of cortex, this means that strong recurrent excitatory connections have to be balanced by inhibitory connections to form so-called *inhibition stabilized networks* (ISN) (Tsodyks et al., 1997). Exemplifying the sometimes counter-intuitive nature of dynamical systems, ISNs predict that increasing the external input to inhibitory neurons paradoxically decreases their firing activity, as more excitation is withdrawn via inhibition of the excitatory neurons than added via the external input (Sadeh and Clopath, 2020a). Underscoring the biological relevance of the ISN, there has been recent evidence of this paradoxical suppression across visual, somatosensory and motor cortices (Sanzeni et al., 2020). However, ISNs typically fail to reconcile contrast invariance with another experimentally observed property: contrast response functions that saturate earlier in V1 than in dLGN (Scholl et al., 2012). ISN models allow only one or the other, thus leaving gaps open for other modeling approaches (Finn et al., 2007; Persi et al., 2011).

Initially designed to combine sigmoidal contrast response functions with contrast-invariance of orientation tuning, the stabilized supralinear network (SSN) (Persi et al., 2011; Ahmadian et al., 2013) has proven so successful at modeling a variety of cortical phenomena, that it has been hailed as a *unifying circuit motif* (Rubin et al., 2015). The SSN is identical to the ISN in its connectivity assumptions and only differs in the activation function that its neurons use to convert input into firing activity. While classical ISN models use threshold-linear activation functions, the SSN uses a power-law activation function (Fig. 1.3), smoothing the transition from inactivity to firing (Persi et al., 2011). Using a power-law activation function has a legitimate basis in the experimental literature: in cats (Anderson et al., 2000; Priebe et al., 2004; Finn et al., 2007) and mice (Tan et al., 2011), trial-to-trial variability in firing threshold crossings of the membrane potential leads to a power-law relationship between membrane potential and firing rate. This relationship has been further supported by modeling work showing that the most-common spiking neuron models lead to a power-law activation function (Hansel and van Vreeswijk, 2002) and that a power-law activation function is necessary to convert contrast-invariance of the membrane potential to contrast-invariance of the firing rate (Miller and Troyer, 2002).

The strength of the SSN lies in its ability to describe multiple cortical phenomena. Its most characteristic - and eponymous - feature stems from the way it behaves when transitioning from weak to strong inputs (Ahmadian et al., 2013). As the authors describe, due to the recurrent connections in the SSN, a neuron's

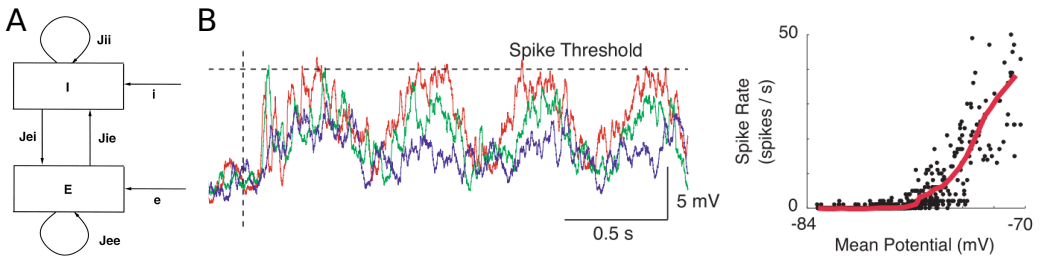

Figure 1.3: Connectivity and power-law activation function in the stabilized-supralinear network (SSN) **A** Like the inhibition-stabilized network (ISN), the SSN combines external input with recurrent connections between excitatory and inhibitory neurons. (Tsodyks et al., 1997). Copyright 1997 Society for Neuroscience. **B** *Left*: Membrane potential for low (blue), medium (green), and high (red) contrast drifting gratings. Noise fluctuations in the membrane potential cause threshold crossings of the membrane potential even if the mean is below threshold. *Right*: Threshold crossings smooth the activation function and lead to a power-low relationship between mean membrane potential and firing rate (From [Anderson et al. (2000)]. Reprinted with permission from AAAS.)

activity is fed back into itself. Mathematically, this means that the firing will be governed by the initial input term and infinitely many terms where the input term and the recurrent term are multiplied. For weak inputs, these infinitely many terms will converge to a constant value much smaller than the initial input term. Thus the network is essentially feed-forward-driven, with the input being amplified supralinearly by the power-law activation function. For a steady-state with large inputs on the other hand, the recurrent term must cancel the input term to maintain stability. This indeed happens as inhibition dynamically stabilizes the network when the input becomes stronger, leading to a sublinear dependence on the input. The network thus transitions from a supralinear into a sublinear regime while being stabilized by inhibition (Ahmadian et al., 2013).

Partly stemming from the characteristic transition from supralinear to sublinear regime, additional emergent properties give the SSN considerable explanatory power. As the contrast of a visual stimulus is increased, the activity of some neurons in V1 reaches a peak and then decreases again, a phenomenon termed supersaturation (Peirce, 2007), which can also occur in the SSN (Ahmadian et al., 2013). The SSN can also model surround suppression, a phenomenon where responses in V1 decrease when a stimulus is extended beyond a preferred size (Rubin et al., 2015). If two competing stimuli are presented, the SSN also correctly predicts that the response will depend on the relative strength of the stimuli, transitioning from summing the individual responses sublinearly to a winner-take-all response (Rubin et al., 2015). Additionally, presenting a stimulus quenches the variability of activity

in many brain areas (Churchland et al., 2010), which the SSN can model using only a single, stimulus-driven attractor (Hennequin et al., 2018). Remarkably, contrary to other ISNs (van Vreeswijk and Sompolinsky, 1996), the SSN can achieve these computations without requiring tight coupling between the network and external inputs due to its dynamic stabilization (Ahmadian et al., 2013; Rubin et al., 2015).

While most initial work involving the SSN has centered on V1, recent work shows that it can be both applied to other visual areas (Liu et al., 2018), as well as at different scales, by either modeling multiple layers of cortex (Obeid and Miller, 2019) or interactions between interneurons (Palmigiano et al., 2020). Furthermore, the SSN has been shown to allow multiple stable states, which could be the computational basis for working memory (Durstewitz et al., 2000), persistent activity, which is involved in decision-making (Curtis and Lee, 2010), and rhythm generation, which is a fundamental cortical property (Singer, 1993), demonstrating its applicability to many different brain functions (Kraynyukova and Tchumatchenko, 2018). Critically, the behavior of the SSN depends on the configuration of its connectivity parameters. For example, the transition to sublinear summation depends on the excitatory recurrent connection (Ahmadian et al., 2013) and the number of stable states the SSN can support is dependent on the determinant of the weight matrix (Kraynyukova and Tchumatchenko, 2018). To judge the applicability of the model and generate further predictions, it is thus vital to relate the model to the connectivity found in the brain networks of interest.

1.4 Connectivity

Since Ramón y Cajal's groundbreaking work showed how neurons as building blocks of the nervous system form specific connections with each other (Llinás, 2003), the notion that a *neuronal wiring diagram* forms the basis of brain function has dominated the field and even birthed a whole subfield of research intent on quantifying these connections, termed *connectomics* (Toga et al., 2012). However, scientists are heavily debating how much information is gained by studying the connectome (Morgan and Lichtman, 2013), with critics pointing out that even complete knowledge of the neuronal structure does not account for effects of neuronal dynamics or neuromodulation (Bargmann and Marder, 2013). The most prominent example of this controversy is the nematode *C. elegans*: its connectome has been completely mapped over 30 years ago (White et al., 1986), yet a complete functional understanding of the nematode's nervous system has not been achieved. Nonetheless, understanding the connectome of *C. elegans* has benefited research

immensely, opening new avenues for studying behavioral paradigms and genetic manipulation (Schafer, 2018). At the heart of the controversy around connectomics lies the question: what *is* the connectome? There are many ways of measuring and interpreting connectivity, ranging from electron microscopy to diffusion magnetic-resonance imaging (Rossini et al., 2019). Each method can be used to construct a network of nodes and edges, yielding static or dynamic connectivity measures between neurons or entire brain areas (Rossini et al., 2019). The connectome thus becomes a multi-faceted entity that depends on the method, species and brain area it is derived from, as well as the research question that motivated its investigation.

To model circuits of the visual cortex using the SSN, principled knowledge of neuron-type specific wiring as well as quantified connection probabilities and strengths are essential. In rodents, connections between excitatory neurons in V1 are generally sparse (Seeman et al., 2018) and highly non-random: strong connections tend to cluster together (Song et al., 2005), neurons with similar response properties tend to be connected (Cossell et al., 2015), and neurons that receive common thalamic input (Morgenstern et al., 2016) tend to be connected. Combining these results, a consensus has emerged that the rodent visual system, like the visual system of carnivores and primates, employs response-based connectivity that enhances feature coding in visual cortex (Harris and Mrsic-Flogel, 2013).

In recent years, and with the arsenal of genetic tools in mice expanding rapidly, inhibitory interneurons have come into the focus of circuit mapping. While excitatory neurons constitute the majority of neurons in cortex, inhibitory neurons come in a variety of anatomical and morphological subtypes (Defelipe et al., 2013). With the discovery that these subtypes selectively express neuropeptides like parvalbumin, calretinin, and calbindin, histological methods to investigate different inhibitory interneurons were developed around 1990 (Kawaguchi et al., 1987; Baimbridge et al., 1992). Once the expression of these proteins was leveraged for the expression of light-sensitive ion channels that allowed optical manipulation of specific cell types (Boyden et al., 2005; Deisseroth et al., 2006; Kuhlman and Huang, 2008), it was only a matter of time before large-scale investigations of inhibitory neurons would develop (de Vries et al., 2020). Such targeted investigations of inhibitory neurons have shown that they seem to have less functional preference than excitatory pyramidal (PYR) neurons and form denser connections to PYR neurons than PYR neurons amongst themselves (Packer and Yuste, 2011; Karnani et al., 2016).

Especially interesting has been the differentiation of the three major subtypes of inhibitory interneurons: parvalbumin-positive (PV), somatostatin-positive (SOM) and vasoactive-intestinal-peptide-positive (VIP) neurons (Rudy et al., 2011). Of the

three types, PV neurons have been studied most intensely and play an important role in oscillatory activity, feature processing, gain control, and plasticity (Atallah et al., 2012; Hu et al., 2014). SOM neurons on the other hand are thought to be involved in spatial processing (Adesnik, 2017) and interactions between neurons over longer distances (Veit et al., 2017; Lee et al., 2017), while VIP neurons have been connected to the enhancement of weak stimuli (Millman et al., 2020) and the effects of locomotion (Fu et al., 2014). Together, they seem to form a canonical circuit of cortical inhibition, where VIP neurons connect only to SOM neurons, PV neurons connect to excitatory pyramidal neurons as well as themselves, and SOM neurons connect to all (Pfeffer et al., 2013). Combined with strong input from thalamus (Cruikshank et al., 2007), this scheme puts PV neurons in the position to provide both feedforward and recurrent inhibition in network models of visual cortex.

Connectivity is typically quantified by measuring the connection probability and strength between different neurons. In an imaging setting for example, tracing synaptic boutons will yield an estimate of the probability to find a connection between two neuron types (Bock et al., 2011). During paired patch-clamp recordings (Edwards et al., 1989; Miles and Poncer, 1996), the current gold standard to measure functional connectivity, current is injected into one neuron, triggering an action potential, which then induces a detectable postsynaptic potential (PSP) in a second neuron. By measuring the strength of the PSP and defining a threshold for considering two neurons to be connected, this method yields connection probability and connection strength, making it an ideal candidate to supply connectivity estimations for modeling.

One major challenge when quantifying connectivity between different types of neurons, is to relate that connectivity to other functional response properties. Since patch-clamp techniques are carried out *in vitro*, responses of a patched neuron to e.g. visual stimuli cannot be obtained during the recording. Consequently, attempts to link stimulus responses to connectivity involve first probing neuronal responses to a set of stimuli while uniquely identifying the neurons under investigation, before re-identifying these neurons in the brain slice that is subsequently used for probing connectivity (Ko et al., 2011; Hofer et al., 2011; Cossell et al., 2015). The technical complexity and inevitably low through-put of such dual-technique approaches highlight both the immense value of these data sets and the need to refine complementary approaches to provide estimates of connectivity that can be captured on a larger scale.

1.5 Inference of Connectivity

In contrast to detailed connectivity measurements, which remain technically challenging and labor-intensive to obtain, the development of ever denser electrode arrays (Jia et al., 2019) and large-scale calcium imaging techniques (Jercog et al., 2016) has led to a drastic increase in the amount of neural response data gathered around the world. Utilizing available data sets and surrogate data, computational neuroscientists and modelers have put forth increasing efforts to infer network connectivity from large-scale neural activity recordings (Magrans de Abril et al., 2018). Two approaches can be distinguished: model-free and model-based inference methods. Model-free approaches link two neuronal activity traces based on statistical or information-theoretic estimations of the traces' similarity. While widely applicable, these approaches offer little explanatory power about the underlying network and cannot generate predictions that could be tested and verified. Model-based approaches on the other hand allow assumptions about the process that generated the activity data, such that models can be compared not just in terms of accurately predicting connectivity in ground-truth data but also in terms of how biologically plausible the underlying assumptions are. Consequently, model-based approaches can generate predictions, thus allowing to close the loop between theoretical and experimental work (Magrans de Abril et al., 2018).

However, connectivity inference faces multiple challenges, of both biological and computational nature. First, an inferred connection might not actually be a biological connection because of the common input problem: if a neuron that is not being observed drives two neurons that are being observed, they will appear connected. This problem extends to input from other brain areas, possibly causing entire populations of neurons to appear falsely connected. Researchers are trying to account for the common input problem by adapting experimental protocols (Soudry et al., 2013) or by estimating common input in network simulations (Ladenbauer et al., 2019), but a general-purpose solution is unlikely to be found. Second, the targeted biological networks consist of a multitude of neuron types with heterogeneous characteristics: neurons can be excitatory or inhibitory, they can produce burst firing (Stiefel et al., 2013), and their connectivity might be generally restricted, e.g. to certain cortical layers or neuron types (Jia et al., 2019). Additionally, signal integration at synapses can be non-linear (Nettleton and Spain, 2000), further complicating the already non-linear activation function of neurons (Anderson et al., 2000). Representing these biological traits in computational models rapidly increases model complexity, which leads to increased computational demands or impairs model in-

vertibility.

To cope with the above-mentioned challenges in connectivity inference, researchers typically have to make trade-offs in their methodology, which I will briefly illustrate using two recent studies. Ladenbauer et al. (2019) infer neuron-to-neuron connectivity strengths in randomly connected networks of leaky integrate-and-fire (LIF) neurons. Fitting networks of LIF neurons to simulated spiking data, they demonstrate improved inference compared to other methods, especially when the network is subsampled, mimicking the common input problem. However, the authors define fixed connection probabilities and only infer the corresponding connection strengths, which are additionally bounded by experimentally plausible values. Thus, they use key a-priori assumptions to enable their inference. Additionally, Ladenbauer et al. use a random-connectivity setting without stimulus tuning, making it hard to relate their work to the highly non-random, stimulus-related connectivity in the visual system. Baker et al. (2020), on the other hand, derive mathematical constraints for inferring connectivity in a linear model of neurons with continuous firing rates. They show that biological constraints such as neuron type, spatial distance or tuning can improve inference performance when using linear activation functions. However, the authors can only infer the existence of connections rather than their strength if the neurons are equipped with a non-linear activation function. When they apply their method to *in vivo* data, the authors are faced with the problem that the external inputs to the networks and their covariance structure is unknown. Without knowing the input covariances, the authors pivot to *fixing* the connectivity with experimentally derived values and inferring the covariance structure rather than inferring the connectivity.

The examples above show that model-based connectivity inference from *in vivo* data remains a challenge in neuroscience that has yet to be solved. In the first manuscript of this thesis, my collaborators and I present an approach to enable inference by utilizing inherent response properties of *in vivo* recordings. We show that contrast-invariance allows a transition from a two-population model describing the contrast response functions of the network to a ring model that additionally accounts for orientation selectivity. This transition means that response data can be used to infer connectivity in a two-step process *without* additional assumptions or parameter limitations. We then demonstrate contrast-invariance in electrophysiological recordings from mouse V1, enabling us to combine the data with our model. By using data from mouse V1 and dLGN, we finally infer connectivity profiles in our model, showing which connectivity regime is compatible with the experimental data. We find that connections in V1 follow a specific order, where strong connec-

tions correlate with broad connectivity profiles across different stimulus orientation preferences. We additionally compare these findings with connectivity measurements from the literature, which show considerable variations. However, when reanalyzing the data from the literature, we recover the same connectivity principles our inference procedure suggested, indicating that a consistent order of connection strengths and widths could indeed be a principle of connectivity in visual cortex.

1.6 Functional Neuron Classification

Computational models of neuronal function provide a simplified, *interpretable* perspective through the reduction of empirical complexity. The SSN model used in the first manuscript highlights connections between inhibitory and excitatory simple cells in V1, which constitute a specialized functional neuron type.

Classifying neurons by their responses to stimuli has shown that there exist multiple functional streams, reaching from the retina to visual cortex, that can be characterized by their receptive fields (center-surround vs other types) and their temporal response patterns (sustained vs transient) (Van Hooser, 2007). Furthermore, neurons can fire action potentials in regular intervals (tonic firing) or in short bursts (Grubb and Thompson, 2005), which shapes development of the visual system (Shatz, 1996) and is also thought play a more general role in plasticity (Letzkus et al., 2006). In visual cortex, neurons respond differently to stimulus orientation (simple and complex cells), levels of contrast, spatial and temporal frequency, and are even modulated by the stimulus' surrounding area, which in itself will not cause any response of the neuron (Van Hooser, 2007).

While scientists attempted to unite this functional variety in the theory of the linear receptive field, mounting failures of classical receptive field theory in explaining non-linear phenomena such as contrast gain adaptation (Baccus and Meister, 2002), suppression (Carandini, 2004), and surround modulation (Fitzpatrick, 2000), combined with an increasing awareness of biases neuron sampling and stimulus selection (Olshausen and Field, 2005), have diversified the approaches of functional classification. Recently, Baden et al. (2016) have used a collection of dynamic stimuli and unsupervised clustering to discover over 40 functional response types in the retina. In the dLGN, Piscopo et al. (2013) similarly used clustering techniques to identify functional neuron types beyond orientation selectivity, while Walker et al. (2019) have developed a closed-loop paradigm to optimize stimuli for individual neurons in V1, resulting in multitude of input profiles beyond classical receptive fields. Taken together, these results solidify the contemporary view that there exists

a multitude of functional neuron types in the visual system.

One such neuron type that has remained outside the broad consideration of visual processing, despite being discovered in the early years of visual neuroscience, consists of neurons that are suppressed by visual stimuli (Levick, 1967; Rodieck, 1967). Because they have been on the fringes of scientific attention, an extensive description of their intrinsic and functional properties outside the retina has been lacking. In the second manuscript of this thesis, my collaborators and I address this gap by presenting a survey of spiking variability, electrophysiological characteristics, and functional responses of suppressed-by-contrast (SbC) neurons in the dLGN and visual cortex. Contrary to previous investigations, we find that SbC neurons fire less regularly than non-SbC neurons and that this difference cannot be explained by burst firing or oscillations. We also find SbC neurons to be putatively excitatory, to be preferentially located in lower cortical layers, and to have higher response latencies than non-SbC neurons. Finally, we find SbC neurons to be more contrast sensitive than non-SbC neurons while showing significant stimulus tuning. These results suggest that SbC neurons actively contribute to visual processing beyond detection of luminance uniformity and that suppressed responses may be generated *de novo* at multiple processing stages in the visual system, expanding the significance of this understudied response type.

1.7 Aims of this thesis

In an attempt to connect a contemporary model of the visual system with experimental observations of neural activity, this thesis aims to:

- Provide a proof of concept, that connectivity in V1 and inputs from dLGN can be inferred via the SSN model
- Demonstrate that connectivity inference can yield insights into biological connectivity principles
- Highlight the diversity of functional neuronal response types
- Improve understanding of suppressed neurons in the visual system

2 Model-based Connectivity Inference in V1

2.1 Summary

In the cortex, feedforward inputs and recurrent connectivity provide the scaffold for neural computations. However, relating circuit computations to the underlying connectivity remains a fundamental, but as yet unsolved, problem. Here, we use a novel, theory-driven approach to deduce V1 connectivity from sensory responses. Combining *in vivo* extracellular activity from mouse visual thalamus and primary visual cortex (V1) with the stabilized supralinear network model (SSN), we determined the connectivity weights for V1's input and recurrent connections. These connectivity weights had a specific ascending order in their orientation-dependent profiles: the excitatory-to-inhibitory weight was the strongest and broadest, and the inhibitory-to-excitatory weight was the weakest and narrowest. Surprisingly, we discovered the same order within the diverse prior experimental connectivity measurements, in which the individual connection probability and the synaptic strengths varied by orders of magnitude across studies, hiding the universal order. Thus, we discovered robust wiring principles contained in seemingly incompatible connectivity estimates.

2.2 Contributions

The following authors contributed to this manuscript. Nataliya Kraynyukova (NK), Laura Busse (LB), and Tatjana Tchumatchenko (TT) conceived the study. NK developed the mathematical proof of how contrast-invariance constrains connectivity, and designed the connectivity inference method. Simon Renner (SR) recorded and curated V1 data in one mouse, developed reverse correlation analysis techniques for flashed gratings stimuli, optogenetic tagging analysis techniques, wave shape classification techniques, contrast-invariance analysis techniques, programmed the data

analysis software, and showed contrast-invariance of the recorded V1 data. SR performed all analyses in Figures 1, S1, S2, S3, S4 and made these Figures. NK and SR pre-processed data for model fits, where SR developed and tested two-dimensional response functions and constructed population responses. SR and NK reviewed and analyzed V1 connectivity measures available in the experimental literature in relation to the inferred connectivity with input by TT and LB. Here, SR selected appropriate studies from the literature based on methodological considerations and extracted connectivity parameters from publicly available data sets. Georgi Tushev contributed to the connectivity inference algorithm. Gregory Born, Yannik Bauer and Martin Spacek recorded and curated dLGN data and additional V1 data, and developed data analysis software. NK, SR, LB, and TT contributed conceptual ideas, discussed and coordinated the project at all stages, wrote and edited the manuscript. All authors contributed to the final version of the manuscript.

Copyright

In vivo recordings of V1 and thalamic extracellular activity reveal cortical connectivity rules

Nataliya Kraynyukova[1,5†], Simon Renner[2,3,†], Gregory Born[2,3], Yannik Bauer[2,3], Martin Spacek[2], Georgi Tushev[5], Laura Busse[2,4,‡], and Tatjana Tchumatchenko[1,5‡]

[1]Institute of Experimental Epileptology and Cognition Research, University of Bonn Medical Center
[2]Division of Neurobiology, Department Biology II, LMU Munich, Munich
[3]Graduate School of Systemic Neurosciences, LMU Munich, Munich
[4]Bernstein Center for Computational Neuroscience, Munich
[5]MPI for Brain Research, Theory of neural dynamics group, Frankfurt am Main
[†,‡]equal contribution

Abstract

In the cortex, feedforward inputs and recurrent connectivity provide the scaffold for neural computations. However, relating circuit computations to the underlying connectivity remains a fundamental yet unsolved problem. Here, we use a novel, theory-driven approach to deduce V1 connectivity from sensory responses. Combining *in vivo* extracellular activity from mouse visual thalamus and primary visual cortex (V1) with the stabilized supralinear network model (SSN), we determined the connectivity weights for V1's input and recurrent connections. These connectivity weights had a specific ascending order in their orientation-dependent profiles: the excitatory-to-inhibitory weight was the strongest and broadest, and the inhibitory-to-excitatory weight was the weakest and narrowest. Surprisingly, we discovered the same order within the diverse prior experimental connectivity measurements, in which the individual connection probability and the synaptic strengths varied by orders of magnitude across studies, hiding the universal order. Thus, we discovered robust wiring principles contained in seemingly incompatible connectivity estimates.

Introduction

In the brain, neuronal activity is constrained by the circuit's connectivity, yet a tractable and interpretable link between experimentally recorded *in vivo* activity and connectivity remains to be found. Efforts to establish such a link are currently fuelled by major progress in both measurements of the activity in neural circuits as well as analysis of circuit connectivity. For example, it is now possible to perform brain-wide recordings with single-cell resolution and targeted manipulations of neural activity[1–3]. At the same time, progress in imaging, tracing, and automated analysis has put the first brain-wide mammalian connectome within reach[4]. Despite this tremendous progress in both fields, deriving connectivity that is consistent with the recorded *in vivo* activity remains a challenge[5].

In an effort to provide a link between *in vivo* activity and connectivity, circuit models of varying complexity have been devised and refined over the years (reviewed in[6]). To specifically connect recent connectomics data with recorded activity, powerful circuit models and large-scale simulations have been proposed[7–12]. Yet, such circuit simulations are computationally demanding and experimental data concerning specific parameters in these highly heterogeneous networks is often scarce and diverse. Therefore, these models typically require additional constraints on the connectivity weights as well as assumptions about unknown external inputs to reproduce activity. Thus, inference of circuit connectivity based solely on *in vivo* activity is yet to be achieved.

Relating circuit connectivity to *in vivo* activity is arguably most realistic for primary visual cortex (area V1), because there, knowledge about cell types is rapidly advancing[13], local connectivity measurements are continuously being refined[14–18], and investigations of feature selectivity and computations have a long tradition[19]. However, even in this advantageous setting, linking V1 circuit connectivity to *in vivo* activity is hampered by the fact that cell-type specific connectivity measurements can vary by an order of magnitude across the available experimental studies[14–18]. Additionally, V1's characteristic property of orientation-selectivity[19] relies at least in part on V1's feature-specific connectivity[16–18,20]. Yet, quantification of this feature-specific connectivity is scarce or even completely absent, in particular between rarer neuron types and deeper layers of V1, as it requires technically challenging re-identification of *in vivo* functionally characterized neurons *in vitro*[16–18,20].

An alternative avenue for linking activity and connectivity is offered by population models, such as the stabilized

supralinear network (SSN) model[21,22]. Indeed, population models reduce the multitude of individual neurons found in large-scale, individual neuron simulations to few populations of similar neurons[23]. As a result, the connection weight matrices between the neuronal populations hold direct information about possible states of the underlying network, which can be meaningfully interpreted. As a population model of V1, the SSN captures non-linear neuronal response properties[24] and at the same time can account for multiple activity features observed in V1, including normalization[22], quenched variability at stimulus onset[25], and contextual modulation[26]. Importantly, the SSN is one of the few nonlinear network models for which the mathematical description of possible firing steady-state configurations is available[27], granting a unique opportunity to compare possible state configurations to observed ones.

Here, we show that combining *in vivo* extracellular recordings of visual responses in mouse dorsolateral geniculate nucleus (dLGN) of the thalamus and V1 with the stabilized supralinear network (SSN) model allows determining, without further assumptions, V1 connectivity weights and their orientation-dependent profiles. Exploiting contrast-invariance of orientation tuning in V1, we designed a novel algorithm that can predict recurrent cortical connections and thalamic input weights consistent with the specific firing rate responses recorded *in vivo*. We found robust relationships between the resulting recurrent connection weights. For example, the excitatory-to-inhibitory cortical connection was the strongest and depended the least on the similarity in the populations' preferred orientations. In contrary, the inhibitory-to-excitatory connection was the weakest and depended strongly on the similarity in the populations' preferred orientations. Despite large variability across previous experimental *in vitro* connectomics studies[14–18], we discovered that the connectivity weights had the same ascending order as our inferred weights, thus reconciling seemingly different experimental connectivity estimates. Additionally, the connectivity patterns we found support the role of V1 circuits in sharpening broadly orientation-tuned thalamic excitation, and their operation in an inhibition stabilized network (ISN) regime, proposed to be the operational state of the cortex[28]. Overall, determining network connectivity solely based on the recorded *in vivo* extracellular activity allowed us to find V1 connectivity principles underlying a variety of nonlinear cortical computations.

Results

Our goal was to understand how *in vivo* extracellular activity constrains the connectivity between neural populations in mouse V1. Here, we briefly illustrate our approach (Fig. 1). We combined the SSN model[21,22] with our recordings of dLGN and V1 responses (Fig. **1A, top**) to infer connectivity by exploiting contrast invariance[29–34], i.e. the finding that orientation tuning curves preserve their width across different levels of contrast (Fig. **1A, bottom**). Contrast invariance, which we confirmed in our data, allowed us to split the response into contrast and orientation components and develop a two-step inference procedure. In the first step of our connectivity inference, we concentrated on the recorded contrast responses (Fig. **1B, top, left**), which allowed us to fit a SSN model consisting of one excitatory (E) and one inhibitory (I) population that receive input from dLGN (Fig. **1B, top, middle**). The fit determined the dLGN input weights and the recurrent V1 connectivity weights between the E and I populations that best matched the recorded activity (Fig. **1B, top, right**). In the second step, we could then expand our inference to the orientation components of the cortical responses (Fig. **1B, bottom, left**) within an extended SSN model that took into account the orientation tuning of the cortical populations (Fig. **1B, bottom, middle**). Here, we determined the relative amount of orientation-specific input from each presynaptic source (Fig. **1B, bottom, right**). Because our connectivity predicting algorithm was unbiased and did not use any prior information about plausible connectivity regimes, we were curious how our results would relate to previous experimental measurements. We thus compared the inferred connectivity to recurrent and feedforward thalamic input connectivity extracted from previous connectomics results, which employed different techniques e.g. *in vitro* multi-patch recordings.

Mouse V1 has contrast-invariant orientation tuning

To obtain neuronal responses for the inference of connectivity, we performed extracellular recordings in awake, head-fixed mice (Fig. **2A**). We measured visual responses across all layers of V1 (Fig. **2B**), with the majority of recorded neurons located in granular and infragranular layers (see Methods for quantification). To obtain stimulus selectivities separately for the E and I V1 populations, we classified individual neurons (N = 242), based on the extracellular spike wave shape, as putative excitatory (broad-spiking, *orange*) and putative inhibitory (narrow-spiking, *teal*) (Fig. **2C**). Narrow spike wave shapes are indicative of parvalbumin-positive (PV+) inhibitory interneurons[35–37], which preferentially inhibit one another and pyramidal cells[38,39], and thus resemble the SSN I population.

During our recordings, we presented briefly flashed, static gratings (84 ms), in order to efficiently sample stimulus responses from the large parameter space spanned not only by orientation and contrast, but also spatial frequency

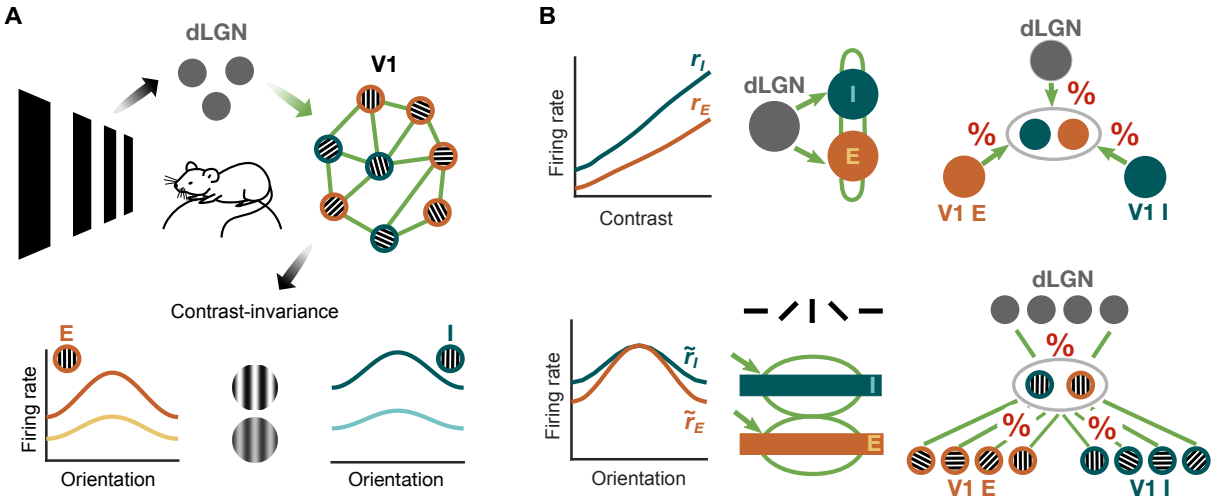

Fig. 1. Inference of thalamic input and V1 recurrent connectivity using *in vivo* extracellular recordings from mouse dLGN and V1, and a network model. A, *Top*: Neurons in area V1 integrate inputs from dLGN, and V1 excitatory (E) and inhibitory (I) neurons. *Bottom*: Contrast invariance of E and I V1 neurons means that the width of orientation responses is invariant with respect to contrast. **B**, Contrast invariance in V1 allows contrast responses (*top, left*) and orientation tuning (*bottom, left*) to be treated separately. Due to contrast invariance, the two-population SSN model used for the contrast response analysis (*top, middle*) can be embedded in the extended SSN model, which reproduces V1 responses as a function of stimulus contrast and orientation (*bottom, middle*). The shape of the contrast responses determines the connectivity weights to the E and I V1 populations, arising from dLGN inputs, and recurrent V1 E and V1 I connections (*top, right*). The orientation component of the response contains information about the relative amounts of orientation-specific inputs to a population with a particular orientation preference (*bottom, right*).

and phase (Fig. 2D). To obtain tuning curves from the acquired spike data, we first performed subspace reverse correlation analysis[44], yielding time-resolved responses (Fig. S1A), and then used each neuron's optimal response time (Fig. S1B) to extract two-dimensional response profiles covering orientation and contrast (Fig. 2E). For further analysis, we only considered visually responsive neurons with optimal response time greater than 25 ms (174/242, Fig. S1B).

We first asked whether responses of individual V1 neurons are contrast-invariant. We applied a singular value decomposition (SVD) to split response profiles into their largest contrast-invariant component and a residual component[45] (Fig. 2E). Since violations of contrast invariance would be visible as non-random spatial patterns in the residual, we used spatial autocorrelation analysis[46] to statistically assess the presence of spatial structure in the residuals (z-scored Gamma index, g_z). Additionally, we computed the power of the residual within the SVD as a measure of the strength of the remaining responses not captured by the separable component (see Methods). Similar to the example neuron in Fig. 2E, top ($g_z = -0.48$, power of residual = 4%), most neurons' residual lacked spatial structure and constituted only a minor part of the response profile. For other neurons, such as the example neuron in Fig. 2E, bottom, the SVD residual indeed contained spatial structure ($g_z = 8.21$), but its power was negligible (residual power = 2%). Across the population of recorded responses, we thus classified neurons as contrast-invariant if their residual SVD component's power was small ($< 5\%$) or if the residual did not contain a significant spatial pattern ($g_z < 1.96$, Fig. 2F). Consistent with previous studies[29–32], we classified the vast majority of V1 neurons as contrast-invariant (91%, 160/174). The observed violations of contrast invariance were restricted to excitatory neurons (E: 14/144, I: 0/30), but the low number of inhibitory neurons leaves open the possibility that the proportion of contrast invariance in the two populations was similar (Fisher-Yates-Test: $p = 0.13$).

Next, we investigated more closely the orientation tuning and contrast response curves of the recorded neurons. We fitted individual response profiles with a separable two-dimensional tuning model constructed from a hyperbolic ratio function[47] for contrast responses and a wrapped Gaussian[41] for orientation tuning (Fig. 2G). The separable tuning model described the responses of the majority of neurons well (mean $R^2 = 0.81$ for both excitatory and inhibitory neurons; Fig. 2H), further strengthening our assertion of contrast invariance. Among the well-fit neurons ($R^2 > 0.4$) used for further analysis, we found a sizeable fraction of neurons (19%) whose response was suppressed by contrast (SbC, Fig. S2). Since V1 SbC neurons might correspond to the vasoactive intestinal peptide-expressing (VIP) interneurons[48], and the function of SbC neurons is not well understood[49], we continued our analyses using only neurons whose responses were enhanced by contrast (101 excitatory, 30 inhibitory).

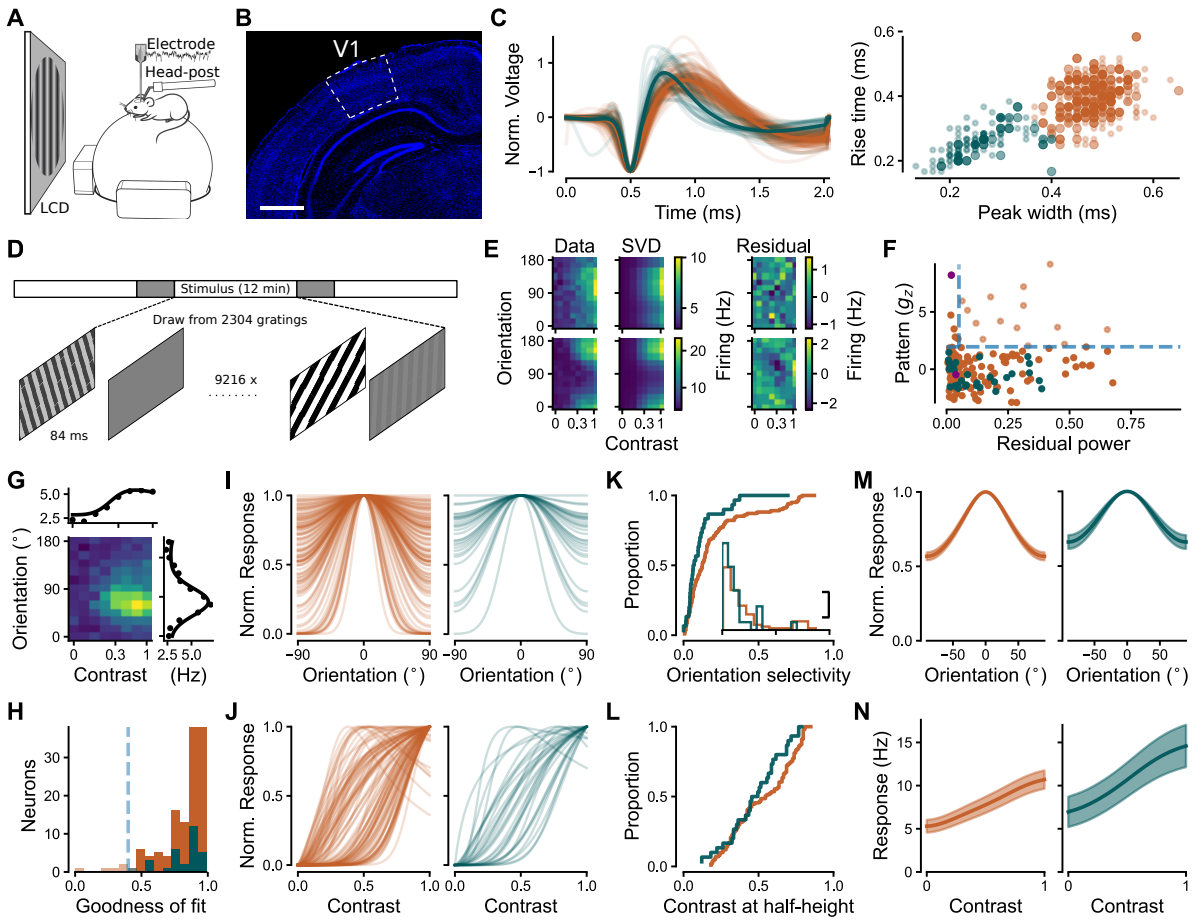

Fig. 2. Population tuning curves for excitatory and inhibitory neurons in V1 are contrast-invariant. **A**, Schematic of the setup for extracellular silicon probe recordings in V1. **B**, Histology of example V1 recording site. *Blue*: DAPI, *white outline*: V1, labeled according to[40]. Scale bar: 1 mm. Bregma: -2.7 mm. **C**, Classification based on extracellular wave shape into putative excitatory (broad-spiking, E, *orange*) and inhibitory (narrow-spiking, I, *teal*) neurons. *Left*: Normalized extracellular waveshapes. *Right*: Clustering based on wave shape parameters. *Large dots:* neurons recorded in sessions relevant to the current study, i.e. sessions containing the flashed grating stimulus, (N = 204 E + 38 I); *small dots:* V1 neurons recorded in other sessions, used to improve N for clustering. **D**, Flashed gratings stimulus paradigm consisting of three stimulus intervals containing a random sequence of gratings (*white bars*), with interleaved blank periods (*gray bars*). In some sessions, the blank periods were used for photoactivation of ChR2 expressed in V1 PV+ inhibitory interneurons (not used for this study). **E**, Responses of two V1 example neurons to combinations of orientation and contrast (*left*), first SVD component (*middle*) and SVD residual (*right*). Spatial patterns in the residual, such as evident for the lower example neuron, reveal violations of contrast invariance. **F**, Violations of contrast invariance were assessed by the power of the SVD residual ($>$ 5%) and significance of spatial autocorrelation ($g_z > 1.96$). *Light dots*: contrast-dependent neurons (14/144 E, 0/30 I neurons). *Solid dots*: contrast-invariant neurons, considered for further analysis. *Purple*: example neurons from (E). **G**, Two-dimensional tuning fit consisting of a product of a hyperbolic ratio function and a wrapped Gaussian[41] for an example V1 neuron. **H**, Distribution of fit quality across neurons. *Dashed line*: fit quality threshold (0.4). *Solid bars*: neurons considered for further analysis (125/130 E, 30/30 I neurons). **I**, Normalized orientation tuning component for E (*left*) and I (*right*) neurons. **K**, Cumulative distribution of orientation selectivity OSI;[42,43]. *Inset:* Density histogram of orientation selectivity. Same x-axis as cumulative distribution, y-scale bar represents 2 neurons per bin of OSI. **J, L**, Same as (I, K) for normalized contrast response component and cumulative distribution of contrast sensitivity (contrast at which the contrast response function reaches half height). **M,N**, Pooled population responses from V1 (mean ± sem). (C,F,H–N) *orange*: putative excitatory neurons, *teal*: putative inhibitory neurons.

To characterize selectivity for orientation and contrast, we next analyzed the individual, normalized tuning components of the fitted 2D tuning model (Fig. 2I–L). Focusing first on orientation tuning (Fig. 2I), visual inspection suggested – in the face of a broad range of tuning for both populations – that the inhibitory population had few neurons of narrow tuning. To quantify differences in tuning strength, we computed an orientation selectivity index (OSI)[42,43] (Fig. 2K). Consistent with previous studies[17,50–53], we found that excitatory neurons had overall stronger orientation selectivity than inhibitory neurons (E: 0.21 ± 0.02, I: 0.13 ± 0.03, mean $\pm$ sem; two-tailed Welch's t-Test: $t = 2.28$, $p = 0.03$). We then investigated the normalized contrast response component of the fitted 2D tuning functions (Fig. 2J), and found a broad range of contrast sensitivity. Inspecting the contrast at half-height of the normalized responses (Fig. 2L), we did not find a significant difference between excitatory and inhibitory neurons (E: 0.54 ± 0.02, I: 0.49 ± 0.02, mean contrast at half-height $\pm$ sem; two-tailed Welch's t-Test: $t = 1.07$, $p = 0.29$).

Finally, to test whether the average population response is also contrast-invariant, we pooled across individual contrast-invariant neurons by averaging their tuning curves, after aligning them to their preferred orientations. This is important because an average of contrast-invariant neurons does not guarantee contrast invariance of the population[54–56]. Analogous to the procedure for the individual neurons, we applied SVD to the population average, and studied the spatial autocorrelation of the residual in order to test for contrast invariance. This analysis revealed that the first SVD component accounted for 99.9% of the variance, albeit with a significant spatial pattern in the residual (Fig. S3A). Due to the small strength of the residual ($< 0.1\%$), we conclude that the pooled population response was also contrast-invariant, complementing the invariance of the individual neurons.

To allow connectivity inference in the SSN model (Fig. 1), another key ingredient besides the V1 E and I population responses, is the input to the cortical network. Instead of assuming the thalamic input to be linear or retrieving average values from the literature, we used the same stimulus and analysis framework to obtain the dLGN population response for contrast (72 neurons, Fig. S4). We found its shape to be sigmoidal, with firing rates lying in between those for V1 E and I neurons (Fig. S4, Fig. 3A). Taken together, our *in vivo* extracellular measurements of dLGN and V1 E and I population responses provided realistic activity for connectivity inference in the SSN model.

Contrast responses reveal consistent relations between connectivity and input weights

Having corroborated contrast-invariant tuning in the E and I V1 populations, we split the response components, and first focused on the recorded contrast responses in dLGN and V1 for inference of the thalamic feedforward and recurrent cortical connectivity weights. The thalamic contrast response T_{dLGN} (Fig. 3A, left) provided input to our V1 network, which we represented via a two-dimensional SSN model (Fig. 3A, middle). The SSN model translated the thalamic input into contrast responses of the E and I cortical populations (Fig. 3A, right). The four constants J_{XY} represented the connectivity weights from the cortical population Y to X, and the two constants g_X corresponded to thalamic input weights to the cortical population X (Fig. 3A, middle). The power-law transfer function $I_+^n = (\max\{I, 0\})^n$ with the exponent n described the non-linear relationship between the input and spike rate output of a V1 neuron (see Fig. 5J, K in[57], Fig. 3A, middle). Our goal was to determine distributions of J_{XY}, g_X and n, for which the SSN model reproduced the recorded thalamic and cortical contrast responses. More specifically, for each recorded contrast C, we required that the stable steady-states r_E and r_I and the input T_{dLGN} of the SSN model

$$r_X(C)^{\frac{1}{n}} = J_{XE} \cdot r_E(C) - J_{XI} \cdot r_I(C) + g_X \cdot T_{\text{dLGN}}(C), \quad X \in \{E, I\} \tag{1}$$

accurately approximated the mean recorded contrast responses of the V1 and dLGN populations, relative to the spontaneous activity at 0% contrast. To this end, we used the mean responses of the thalamic and V1 populations to the eight contrasts (Fig. 3A, Eq. S3). To provide an unbiased comparison with the direct, experimental measurements of connectivity, we computed the J_{XY} and g_X without prior assumptions about their possible experimental range. Our only experimentally motivated parameter constraint was the assumption that the power-law exponent n ranged between 1 and 5[58,59]. Plotting the V1 population contrast responses as a function of contrast response measured in dLGN (Fig. 3B), we found a supralinear relationship, which ruled out the possibility that the transfer function was linear ($n = 1$).

To obtains sets of J_{XY}, g_X and n, we generated a triplet of sample dLGN, V1 excitatory, and inhibitory contrast responses as random monotonically increasing functions within the $\pm$ sem areas of the measured contrast responses (Fig. 3C, left). For each such triplet, we kept n constant in the range between 1.1 to 5 and initially computed the six weights J_{XY}, g_X as a solution of the over-determined linear system in Eq. 1. We considered the weights J_{XY}, g_X to be computationally valid, if they were positive, and lead to a stable steady state of the SSN model (Eq. S4 – Eq. S7). In the next step, we minimized a score function (Eq. S8) in order to obtain fitted weights which accurately reproduced

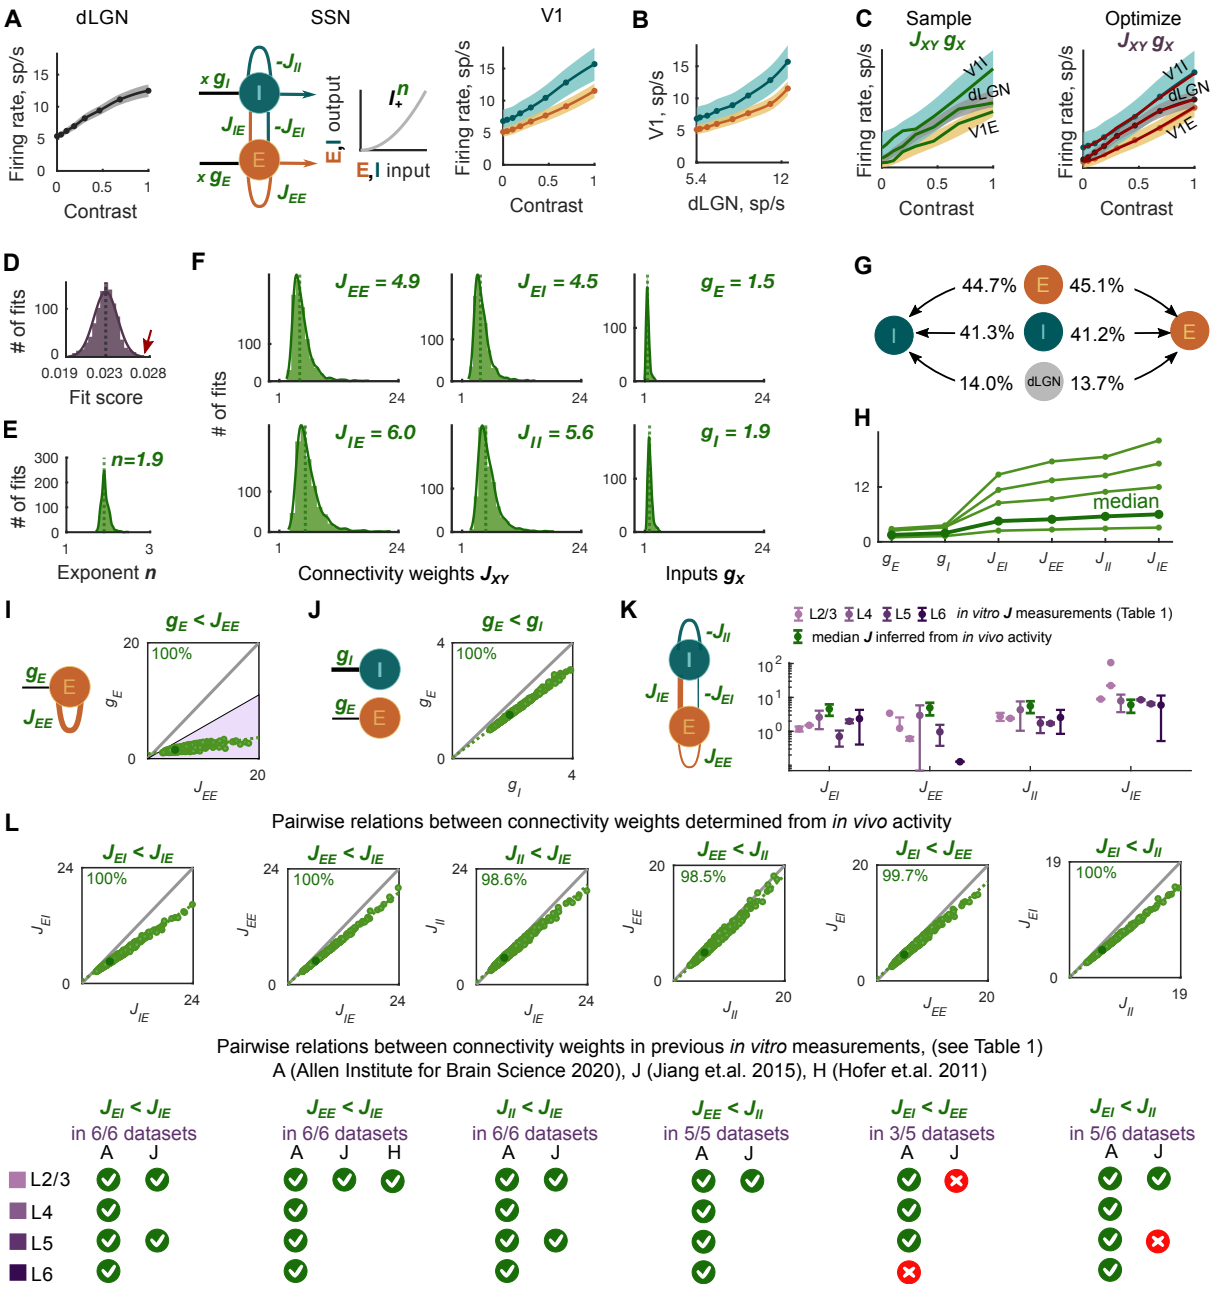

Fig. 3. Inference of cortical and thalamic input weights from contrast responses using the SSN model. A, The cortical network transforms the dLGN population contrast response (*left*) to the contrast responses of E and I V1 populations (*right*). The SSN model (*middle*) represents the cortical network. The weights $J_{XY} > 0$ denote the strength of the recurrent connections from cortical population Y to X, and $g_X > 0$ denote the strength of the thalamic input to cortical population X. The transfer function relating input to output of the cortical E and I populations is $I_+^n = (\max\{I, 0\})^n$. Shaded areas represent ±sem of the recorded responses. **B**, Supralinear thalamocortical mapping of contrast responses excludes the exponent $n = 1$ in the SSN model. **C**, *Left:* Randomly generated sample curves (*green lines*) within ±sem of the dLGN and V1 contrast responses (*shaded areas*) were substituted into the SSN model to yield the initial parameters J_{XY} and g_X for the fixed exponents n. *Right:* Subsequently, the SSN parameters J_{XY}, g_X and n were optimized to minimize the distance between the fit (*red*) and the mean of the recorded responses (*orange, gray, teal*). **D**, The distribution of fit scores of the final 10^3 fits has a mean of $2 \cdot 10^{-2}$. The arrow (*red*) points to the largest score value used to produce the example fit in (C, *red*). **E**, The distribution of the optimized power-law exponents n corresponding to the final 10^3 fits has a median of $n = 1.9$. **F**, The distributions of the optimized connectivity weights J_{XY} are broader than that of input weights g_X. The location of vertical dotted lines corresponds to the J_{XY} and g_X medians presented in the plot's upper right corners. **G**, The average ratios of connectivity weights from three input sources (local E, local I, external dLGN) to E cortical population are similar to those of the I population. **H**, In 99% of the inferred connectivity parameter sets, the recurrent and input connectivity weights, ordered according to the x-axis labels, build an increasing sequence. *Thin green lines*: four representative examples of fits, *bold green line*: median of parameter distributions. **I**, For cortical E neurons, the weight g_E of the dLGN input was lower than the recurrent cortical weight J_{EE} in 100% of the inferred connectivity parameter sets. The colored area (*purple*) comprises relations between J_{EE} and g_E for which thalamic contribution of excitatory input entering the E cortical population would

be smaller than the experimentally determined upper bound of $36 \pm 2\%$[60,61]. **J**, The thalamic input weights to E cortical population were lower than the I-weights in 100% of the inferred connectivity parameter sets. **K**, The connectivity weights between V1 populations computed from previously published *in vitro* measurements have largely the same increasing order as the connectivity parameters obtained from our SSN-based inference procedure based on *in vivo* activity. **L**, Most of pairwise relations between connectivity weights (H, *x*-axis labels) inferred from *in vivo* activity were consistent with relations computed from previous *in vitro* measurements (**Table 1**).

the mean contrast responses (Fig. **3C, right**). The resulting optimal weights yielded excellent fits with low scores (Eq. **S8**; Fig. **3D**; to illustrate the accuracy of these fits we show the worst resulting fit in Fig. **3C, right**). Considering the resulting optimal fits, we found that the distribution of the power-law exponents n peaked at $n = 1.9$ (Fig. **3E**), which is in line with previous experimental and theoretical evidence[33,34,58,59].

Next, we highlight three characteristic features of the resulting connectivity distributions. First, the recorded thalamic and cortical responses were supported by connectivity weights of diverse magnitudes (Fig. **3F**). Specifically, we found that the distributions of the inferred thalamic input weights g_E and g_I were largely limited to the range 1 to 5, while distributions of recurrent connectivity weights J_{XY} were broader and spanned almost a 5-fold range, with values from 1 to 24 (Fig. **3F**). Second, the connectivity weights to the I population were stronger than the weights to the E population in all inferred sets of parameters (Fig. **3F**, *top*: weights to E, *bottom*: weights to I). However, the relative contributions of the weights from all three input sources (E, I, dLGN) targeting the E population were remarkably similar to the relative contributions of the weights targeting the I population (Fig. **3G**). Third, even though J_{XY} and g_X were distributed broadly, the connectivity weights in 98.5% of the 10^3 final parameter sets represented an ascending sequence following the order $g_E < g_I < J_{EI} < J_{EE} < J_{II} < J_{IE}$ (Fig. **3H**).

Having discovered this consistent order within the inferred connectivity weights, we next focused on their pairwise relations and tested if they could also be found in previous direct, experimental measurements of V1 connectivity. We first studied the relation between the thalamic feedforward (g_E) and recurrent (J_{EE}) weights. We did this for both our inferred connectivity weights and for experimental measurements obtained in Lien et al.[60], who performed *in vivo* recordings of excitatory postsynaptic currents in V1 L4 pyramidal cells and found that thalamic excitation, isolated by optogenetic silencing of V1, to be $36 \pm 2\%$ of the total V1 L4 pyramidal cell excitation. To understand whether this relative contribution to overall excitation was compatible with our inferred connectivity, we first estimated thalamic and cortical input to V1 E neurons as a product of dLGN and V1 population firing rates and the inferred connectivity weights (see Methods). The maximal relative contribution of the thalamic input of $36 \pm 2\%$ compared to the total input[60] yielded an approximate upper bound for the relation g_E/J_{EE} of 0.55, suggesting that g_E was smaller than J_{EE} in direct experimental measurements (see Methods). The relationship $g_E < J_{EE}$ was preserved across 100% of our inferred sets of connectivity weights (Fig. **3I**). This emphasizes the predominant role of recurrent excitatory cortical connections over thalamic input weights in shaping visual contrast responses of the V1 E population.

In the next step, we addressed the relationship between the thalamic input weights g_I and g_E. Recently, an extensive *in vitro* study reported connection probabilities and response amplitudes across V1 layers and cell types[62]. Connection probabilities combined with corresponding response amplitudes provide a measure for a connectivity weight between cortical populations. Specifically, Ji et al.[62] demonstrate that the thalamic projections to V1 excitatory cells had either similar (layers 2/3 and 6) or slightly lower (layers 4 and 5) connection probabilities than thalamic projections to PV+ V1 neurons; the response amplitudes, however, were substantially lower in V1 excitatory cells than in PV+ neurons in all V1 layers. These data suggest that the connectivity strength of thalamocortical projections to excitatory cells is lower than to PV+ V1 neurons. Consistent with this report, our inferred input weights to E cortical neurons were lower (median $g_E = 1.5$) than the thalamic input weights to the I neurons (median $g_I = 1.9$) across all inferred parameter sets (Fig. **3J**).

We now turned to the cortical circuit and compared the inferred recurrent connectivity weights J_{XY} with reports from *in vitro* whole-cell patch-clamp recordings in V1 excitatory and inhibitory PV+ neurons (**Table 1**). For each experimental source, we extracted the connectivity weights J_{XY} (*gray* cells in **Table 1**) as a product of the measured connection probability (CP), the measured strength of the postsynaptic potential (PSP) (*white* cells in **Table 1**), and the percentage of neurons in the source population Y of the connection XY. Since our network represents the circuit motif between V1 pyramidal cells and PV+ inhibitory neurons, we made the following assumptions to determine the percentage of neurons in the source population: we set the fraction of excitatory neurons to be 80% of all V1 neurons[63] and the fraction of PV+ neurons to 50% of the inhibitory V1 population in layers 2/3–6[64] (Methods). Together, this resulted in 89% of neurons in the E SSN population and 11% of neurons in the I SSN population (Methods). Despite CPs and PSPs varying across studies by an order of magnitude, the computed connectivity estimates from the experimental

sources in **Table 1** had a surprisingly similar order as that observed in our inferred connectivity sets (Fig. **3K**, Fig. **S5B**). For example, the connectivity strength J_{IE} had consistently the largest value among all four J_{XY} in all cortical layers and studies (**Table 1**, Fig. **3K**).

We therefore asked next whether the pairwise relations found in our *in vivo*-inferred connectivity weights (Fig. **3L, top**) were consistent with the relations between the extracted *in vitro* connectivity weights (**Table 1**). We constrained our analysis to experimental sources in **Table 1**, which contained measurements of both connections in each relation of interest (Fig. **3L, top**). Remarkably, four out of six relations ($J_{EI} < J_{IE}$, $J_{EE} < J_{IE}$, $J_{II} < J_{IE}$, and $J_{EE} < J_{II}$) could also be computed from the experimental sources and across V1 layers (**Table 1**, Fig. **3L, bottom**). Specifically, J_{IE} was always the strongest connectivity weight in all cortical layers and in our inferred connectivity sets (**Table 1**, Fig. **3L**). The remaining two pairwise relations found in our inferred connectivity sets ($J_{EI} < J_{EE}$ and $J_{EI} < J_{II}$) held true for the majority of cortical layers and experimental sources. Specifically, we found three exceptions from the above rules: the relation $J_{EI} < J_{EE}$ did not hold in layer 6, and two available experimental reports for the relation (J_{EE}, J_{EI}) in layer 2/3 and (J_{II}, J_{EI}) in layer 5 did not lead to consistent conclusions (Fig. **3L, bottom**).

Finally, we analyzed the inferred connectivity weights with respect to possible computational regimes. First, we asked if the resulting circuit was an inhibition stabilized network (ISN)[28,65]. We found that all inferred connectivity weights were consistent with the ISN regime, starting from the smallest contrast value of 4% recorded in our experiments (Fig. **S5C**). In other words, for the contrast values exceeding 4%, an increase of external input to the I population of the SSN model would lead to a paradoxical decrease in the activity of both E and I populations - an observation reported in previous V1 experiments[65–67]. Second, we analyzed whether the resulting inferred circuit can have different levels of activity for the same input, i.e. whether the SSN model with the inferred connectivity weights can have multiple stable steady states. From our previous work[27] we know that for specific connectivity regimes, the SSN model can have multiple stable steady states. However, we found that for 100% of the inferred connectivity weights, the SSN model circuit had a unique stable steady state, ruling out that the network was able to generate different levels of activity for the same input.

Taken together, with minimal prior assumptions and solely based on the recorded *in vivo* extracellular data, our SSN-based inference method found a consistent ascending order between the inferred cortical connectivity weights ($J_{EI} < J_{IE}$, $J_{EE} < J_{IE}$, $J_{II} < J_{IE}$, and $J_{EE} < J_{II}$). Remarkably, despite substantial variability in the absolute values of connectivity measurements in previous *in vitro* connectomics estimates (*white* rows in **Table 1**), the inferred pairwise relations between connectivity weights were preserved across experimental sources and cortical layers and matched the relations identified through our inference method.

Determining connectivity and input profiles from contrast and orientation responses

Since it is known that V1 neurons form fine-scale subnetworks according to orientation preference and general response similarity[16–18,20,70,71], our next goal was to infer orientation-dependent connectivity patterns in a network of orientation-selective cortical neurons. To reproduce the recorded dLGN contrast responses and V1 responses to orientations and contrasts (Fig. **4A,B**), we assumed that the cortical firing rates are the steady states R_E and R_I of the SSN model

$$R_X(\psi-\theta,C) = \Big(\int_{-\pi/2}^{\pi/2} \Big(W_{XE}(\theta-\theta')R_E(\psi-\theta',C)d\theta' - W_{XI}(\theta-\theta')R_I(\psi-\theta',C) \Big) d\theta' + T_{\mathrm{dLGN}}(C)L_X(\psi-\theta) \Big)_+^n. \quad (2)$$

Here, C represents the stimulus contrast, n is the exponent of the transfer function, θ is the preferred orientation of a cortical population X, $X \in \{E, I\}$, and $R_X(\psi-\theta, C)$ represents the firing rate of the population X with the preferred orientation θ in response to a grating stimulus having orientation ψ. The connectivity profiles $W_{XY}(\theta-\theta')$ correspond to the strength of connection from population Y to population X that differ in their preferred orientations by $|\theta-\theta'|$. The X population with the preferred orientation θ receives the thalamic input $T_{\mathrm{dLGN}}(C)L_X(\psi-\theta)$.

Next, we used the experimental observation of contrast invariance as a constraint for the steady-state solutions R_E and R_I of the SSN model in Eq. **2**. Specifically, we analyzed how contrast invariance constrains the recurrent and input connectivity profiles W_{XY} and L_X in Eq. **2**. Contrast invariance means that for a fixed grating orientation, the response can be represented as product of contrast and orientation components, namely, the firing rates R_E and R_I follow the relations[33,34]

$$R_E(\psi-\theta,C) = r_E(C)\widetilde{r}_E(\psi-\theta), \quad R_I(\psi-\theta,C) = r_I(C)\widetilde{r}_I(\psi-\theta). \quad (3)$$

Layer 2/3							
	Hofer et al.[17]	Ko et al.[16]	Cossell et al.[18]	Seeman et al.[14]	Karnani et al.[68]	Allen Institute for Brain Science[69]	Jiang et al.[15]
CP_{EI}						17/49	27/83
PSP_{EI}						0.31±0.20	0.42±0.03
J_{EI}						1.18±0.76	1.50±0.11
CP_{EE}	45/235	43/222	75/520	13/130		5/80	1/50
PSP_{EE}	0.2		0.45±0.68	0.34±0.32		0.22±0.24	0.34±0.08
J_{EE}	3.41		5.78±8.73	3.03±2.85		1.22±1.34	0.61±0.14
CP_{II}					13/32	36/97	28/78
PSP_{II}					0.58±0.10	0.66±0.74	0.61±0.10
J_{II}					2.59±0.45	2.69±3.02	2.41±0.39
CP_{IE}	36/41					19/50	13/83
PSP_{IE}	1.36					0.27±0.21	1.6±0.23
J_{IE}	106.28					9.13±7.10	22.30±3.21

Layer 4	
	Allen Institute for Brain Science[69]
CP_{EI}	7/34
PSP_{EI}	1.15±0.64
J_{EI}	2.60±1.45
CP_{EE}	17/236
PSP_{EE}	0.46±0.45
J_{EE}	2.95±2.88
CP_{II}	26/55
PSP_{II}	0.83±0.63
J_{II}	4.32±3.28
CP_{IE}	4/33
PSP_{IE}	0.73±0.39
J_{IE}	7.88±4.21

Layer 5		
	Allen Institute for Brain Science[69]	Jiang et al.[15]
CP_{EI}	10/72	8/32
PSP_{EI}	0.46±0.35	0.72±0.09
J_{EI}	0.7±0.53	1.98±0.25
CP_{EE}	13/553	0/12
PSP_{EE}	0.46±0.59	
J_{EE}	0.96±1.23	
CP_{II}	27/126	9/48
PSP_{II}	0.74±0.86	0.83±0.12
J_{II}	1.74±2.03	1.71±0.25
CP_{IE}	7/74	3/36
PSP_{IE}	1.01±0.41	0.87±0.07
J_{IE}	8.50±3.45	6.45±0.52

Layer 6	
	Allen Institute for Brain Science[69]
CP_{EI}	16/69
PSP_{EI}	0.92±0.76
J_{EI}	2.35±1.94
CP_{EE}	3/443
PSP_{EE}	0.21
J_{EE}	0.13
CP_{II}	41/118
PSP_{II}	0.67±0.45
J_{II}	2.56±1.72
CP_{IE}	11/76
PSP_{IE}	0.46±0.42
J_{IE}	5.93±5.41

Table 1: **Connectivity matrix for pyramidal and PV+ neurons in layers 2/3, 4, 5, and 6, summarized from recent *in vitro* studies.** The entries of the connectivity matrix J_{XY} are computed for each layer based on the experimentally measured connection probability (CP) and amplitude of the postsynaptic potential in mV (PSP) using the formula $J_{XY} = \mathrm{CP} * \mathrm{PSP} * N_Y/N * 100\%$, where N_Y/N is the proportion of neurons in the population N_Y. Here we use $N_E/N = 0.89$ and $N_I/N = 0.11$ [64] (see Methods).

Here, r_E and r_I are the contrast response functions studied in the previous section, and $\widetilde{r}_E$ and $\widetilde{r}_I$ are the peak-aligned and normalized orientation tuning curves (Fig. **4B, left**). The property of contrast invariance is computationally advantageous because it allowed us to reduce the two-dimensional product of its one-dimensional contrast and orientation components and study each component independently. We found that contrast responses r_X, T_{dLGN} (Fig. **4B, right**) and orientation tuning curves $\widetilde{r}_X$ (Fig. **4B, left**) provide direct access to the connectivity W_{XY} and input tuning profiles L_X as expressed by the following relations

$$\int_{-\pi/2}^{\pi/2} W_{XY}(\theta-\theta')\widetilde{r}_Y(\psi-\theta')d\theta' = J_{XY}(\widetilde{r}_X(\psi-\theta))^{1/n}, \quad L_X(\psi-\theta) = g_X(\widetilde{r}_X(\psi-\theta))^{1/n}, \quad X,Y \in \{E,I\}, \quad (4)$$

where the connectivity and input weights J_{XY} and g_X correspond to the cortical and thalamic contrast response functions studied in the previous section (for full derivation, see Methods and Eq. **S9** – Eq. **S18**). Relations in Eq. **4** generalize the result in Persi et al.[72], which showed that the wrapped Gaussian approximation makes it possible to relate orientation tuning curves and connectivity profiles.

In order to understand how our inferred connectivity between populations depends on the similarity of their orientation

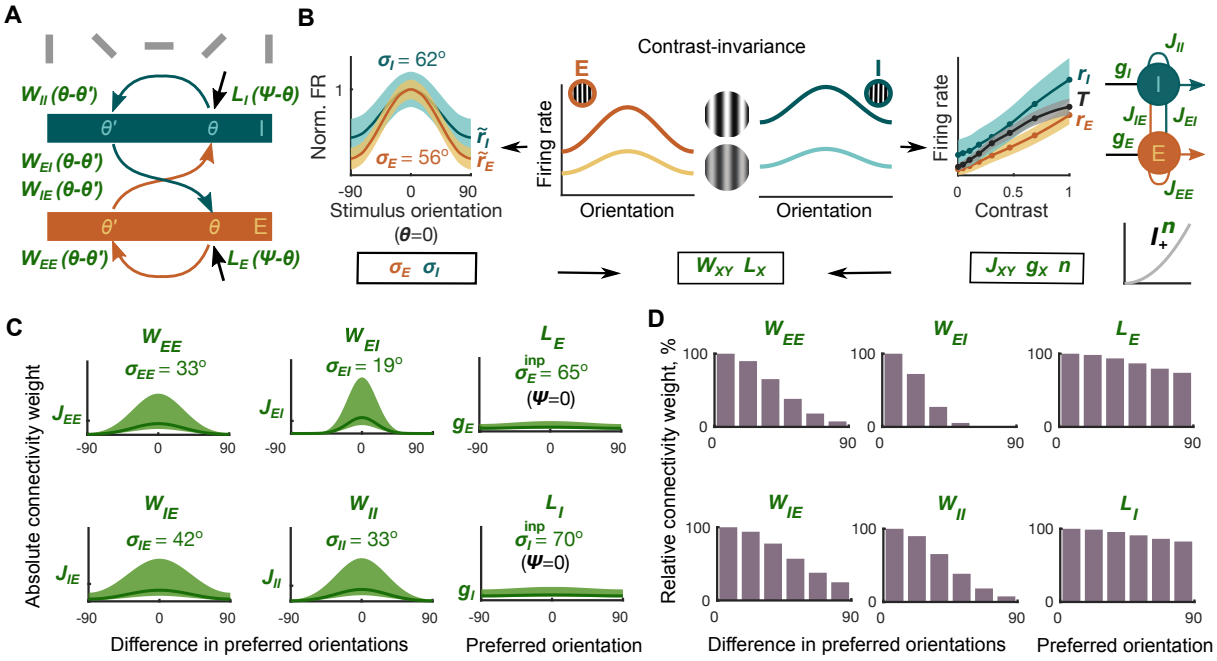

Fig. 4. Inferring orientation dependence of recurrent and feedforward connectivity. A, The connectivity profile $W_{XY}(\theta - \theta')$ describes the connection strength from a population Y to X with a difference in their preferred orientations of $|\theta - \theta'|$. The input profile $L_X(\psi - \theta)$ represents the thalamic input to a population with the preferred orientation θ. **B**, The width of the normalized orientation tuning curve $\widetilde{r}_E$ is smaller than $\widetilde{r}_I$, wrapped Gaussian fit $\sigma_E = 56° \pm 2°$, $\sigma_I = 62° \pm 3°$ mean±sem (*Left*). The orientation-dependent connectivity (W_{XY}) and input connectivity profiles (L_X) are determined using the widths of the recorded orientation tuning curves σ_E and σ_I and the connectivity weights J_{XY}, g_X, n derived from the contrast responses in Fig. 3. **C**, The connectivity profiles W_{XY} (*left and middle columns*) can be calculated using σ_X, J_{XY}, and n. The input profiles L_X (*right column*) depend on σ_X, g_X, and n. *Line*: connectivity and input profiles for median J_{XY}, g_X, and n. *Shaded area*: distributions of connectivity and input profiles corresponding to J_{XY}, g_X, and n distributions in Fig. 3E,F. **D**, Normalized connection strengths between populations with different preferred orientations. The orientations are binned in 15° steps to aid comparison with experimental findings.

preference, we next determined the widths of the recurrent connectivity profiles W_{XY} (Fig. 4C, left and middle) and the input profiles L_X (Fig. 4C, right; Eq. S20 and Eq. S22). Interestingly, inspecting equations in Eq. S20 and Eq. S22, we realized that the recorded contrast responses corresponding to J_{XY} and g_X mostly determine the amplitudes of the connectivity and input profiles, and represent connectivity strength between populations with similar preferred orientations (Fig. 4C). In contrast, the widths of orientation tuning curves σ_X together with the power-law exponent n characterizing the neuronal transfer properties determine the widths of connectivity profiles. The relative widths of connectivity and input profiles are best visible after normalizing their amplitudes (Fig. 4D), where we combined orientations in 15° bins. We found that the widths of connectivity profiles in each parameter set followed the order $\sigma_{EI} < \sigma_{EE} \approx \sigma_{II} < \sigma_{IE}$. This order of connectivity profiles' widths resembled the ascending order of connectivity weights $J_{EI} < J_{EE} < J_{II} < J_{IE}$ discussed in the previous paragraph.

How do our inferred connectivity profiles W_{XY} and tuning properties of the thalamic inputs L_X relate to previous experimental results? We first compared the width of the inferred thalamic input profiles L_X with the width of V1 orientation tuning curves (Fig. 4B, left), a relation addressed in prior experimental reports (e.g.[60,73–76]). We found that the widths of the inferred orientation-specific thalamic inputs to cortical neurons (mean $\sigma_E^{inp} = 65°$, $\sigma_I^{inp} = 70°$) have a broader tuning than the average V1 orientation tuning curves themselves (mean $\sigma_E = 56°$, $\sigma_I = 62°$). This is consistent with a wide range of experimental studies across cortical layers, showing broader dLGN input than V1 orientation tuning curves[60,73–76].

We also compared our derived cortical connectivity profiles W_{XY} with previous direct connectivity measurements. To the best of our knowledge, information on how connection probability and the relative connection strength of I-to-E, and I-to-I projections depend on the difference in the populations' preferred orientations does not seem to be available (but see[20]). Hence, for our comparison of inferred and experimentally measured orientation-dependent profiles, we concentrated on the direct measurements currently available for E-to-E and E-to-PV+ connections in L2/3[16,17]. The orientation-dependence of our inferred connectivity profiles (mean $\sigma_{EE} = 33°$, $\sigma_{IE} = 42°$ Eq. 4 and Eq. S22) is consistent with direct experimental measurements of connection probability[16,17], demonstrating that the connectivity

profile between pairs of E neurons was sharper than the E-to-I connectivity profile. However, both connectivity profiles E-to-E and E-to-PV+ measured experimentally[16,17] were flatter than our inferred profiles. This might be related to overall sharper orientation tuning in V1 L2/3[32,75] compared to the lower layers, which are over-represented in our recorded data (see Discussion).

In summary, we found that the widths of the connectivity profiles we derived from *in vivo* extracellular activity followed the order $\sigma_{EI} < \sigma_{EE} \approx \sigma_{II} < \sigma_{IE}$, resembling the order of the connectivity weights $J_{EI} < J_{EE} < J_{II} < J_{IE}$. Consistent with experimental reports across cortical layers, we found that both E and I input profiles L_X are broader than the average cortical orientation tuning curves. For future experiments, our results predict that W_{EI} is the narrowest and W_{IE} the broadest profile of all four projections. In our model, this results from the width of I orientation tuning profile, σ_I, exceeding that of the E population σ_E, and the width σ_E^{inp} of the E input profile L_E exceeding the width of orientation tuning curve, σ_I (see Methods).

Discussion

Here, we combined *in vivo* extracellular dLGN and V1 responses with the stabilized supralinear network (SSN) model to infer feedforward and recurrent connectivity weights of mouse primary visual cortex, and their orientation-dependent connectivity profiles. In a novel, 2-step inference procedure based on the separation of contrast and orientation tuning of V1 responses, we identified consistent relationships between the input and recurrent cortical connectivity weights, $g_E < g_I$, $g_E < J_{EE}$ and $J_{EI} < J_{EE} < J_{II} < J_{IE}$. These relations were hidden in experimental connectomics results, and had to be revealed using a model-based approach. Our results demonstrate a single connectomic origin underlying cortical signal amplification, contrast invariance and inhibition stabilisation. These results provide evidence that key computations of visual processing are imprinted into the underlying connectivity and can be discovered and linked via a circuit model.

Inferred connectivity reveals principles consistent with previous connectomics estimates

Cell-type specific connectivity in mouse V1 has been measured by a number of connectomics studies[15,17,69]. These studies reported amplitudes of postsynaptic potentials and connection probabilities which can differ by an order of magnitude across experiments (**Table 1**). This variability is likely due to differences in experimental conditions, e.g. slice thickness or distance between cell bodies of interest. Despite this variability, we found that the pairwise relations between connectivity weights within studies were remarkably consistent. Specifically, considering relations between pairs of connectivity weights – a procedure similar to that applied to single slices to remove experimental variability (e.g.[62,77]) – revealed that the relative magnitude of the recurrent weights followed the order $J_{EI} < J_{EE} < J_{II} < J_{IE}$ across the majority of experimental studies and across the different cortical layers. Intriguingly, this order was also contained in the connectivity we derived from our *in vivo* data.

Our results support several experimental and theoretical findings related to the computational regime of the cortex. First, the strong E-to-I connection has been linked in a previous computational study to image discriminability[78]. Second, we found that 100% of the inferred feedforward and recurrent connectivity weights were consistent with the inhibition stabilized regime[66,67]. This regime has been associated with the paradoxical firing rate reduction in response to increased excitation of the I population and fast, flexible stabilization mechanism that balances otherwise unstable networks[28,66]. Such a paradoxical response of cortical populations seems to be present in both superficial and deep layers of mouse V1[67]. Interestingly, although previous studies hypothesized that a strong connectivity weight J_{EE} might be a requirement for inhibitory stabilization[28,66,78], the ISN regime arises in our network even though the connectivity weights J_{EE} and J_{EI} were the smallest entries in the connectivity matrix (Fig. **3**). Lastly, our inference method exploited the property of contrast invariance[33,34], which we confirmed at the level of single neurons as well as cortical excitatory and inhibitory populations. This indicates that contrast invariance and inhibitory stabilization can be supported jointly by a connectivity pattern consistent with recorded cortical and thalamic activity.

Furthermore, we determined the widths of the orientation-dependent connectivity profiles showing that they follow the order $\sigma_{EI} < \sigma_{EE} \approx \sigma_{II} < \sigma_{IE}$. Specifically, we have shown that this ascending order of the connectivity profiles' widths is a direct consequence of the contrast-invariance property combined with the observation that the recorded inhibitory orientation tuning curve is broader than its excitatory counterpart. In line with our inferred relation $\sigma_{EI} < \sigma_{EE}$, direct measurements of orientation-dependent connectivity profiles report that E-to-E connection is sharper than that of E-to-PV+ in V1 L2/3[16,17]. However, both profiles W_{EE} and W_{IE} which we inferred from our *in vivo* recorded data were sharper than respective profiles reported in the direct connectivity measurements[16,17]. Investigating our network

model and its implications, we found that broader profiles W_{EE} and W_{IE} correlate with narrower orientation tuning in the excitatory population. Interestingly, V1 L2/3 pyramidal neurons targeted in the previous connectivity measurements[16,17] are known to display sharper orientation tuning (e.g.[32,75]) compared to lower layers, which dominate in our experimental data. Thus, the overall broader orientation tuning curves in deeper cortical layers might cause the sharper tuning of connectivity profiles W_{EE} and W_{IE} we find in our recorded data.

Previous work reported that E cortical neurons receive primarily untuned input from their local inhibitory afferents[79]. We emphasize that our highly selective W_{EI} represents the connectivity profile and not the recurrent input current profile measured in these experiments. Indeed, our theory predicts that the input arriving at E cortical population from their local I afferents is a convolution of the functions W_{EI} with the I orientation tuning curve. Consequently, we show that the resulting width of the recurrent input profiles from each of the local afferents is equal to the widths of the corresponding external thalamic input profiles, which we found to be almost flat. Additionally, although direct measurements of highly selective connectivity profile W_{EI} do not seem to be available at present, the overall similarity of neuronal feature selectivity found for this connection in Znamenskiy et al.[20] can serve as a predictor of connectivity in mouse V1 L2/3.

We studied not only the recurrent connections in cortex, but also considered simultaneously the thalamic input strengths and their profiles. Recurrent connections and feedforward input weights are rarely measured at the same time[60,61,76,80], therefore it is often difficult to study their relative strength. Consistent with previous experimental results, we found cortical connections to be stronger than the thalamic input weights, indicating that only a fraction of the excitation in cortex is due to the thalamic connections[60,61,76]. At the same time, the tuning of thalamocortical afferents and their interplay with intra-cortical connections in the emergence of orientation selectivity have been under intense discussion. Yet, there is no consensus as to how much tuning the afferents carry or how much their tuning contributes to cortical orientation selectivity. Our model results provide evidence that, if thalamic afferents, cortical inputs, membrane transfer functions, and stimulus responses are considered together, orientation tuning can be generated with a combination of weaker, broadly tuned thalamic input and strong, sharply tuned cortical connections[60,61,76].

Theory-based inference of connectivity from *in vivo* responses complements existing *in vitro* approaches

Inferring network connectivity using a theory-driven network model and *in vivo* responses complements existing *in vitro* approaches based on paired whole-cell recordings[14–18,39,68,69], photo-stimulation[81] or glutamate uncaging[82], gold standard methods for assessing circuit-level connectivity. One strength of our method is that it is based on functional measurements in the intact, living brain, which has the potential to overcome limitations imposed by the unavoidable truncation of axo-dendritic branches in slices, where connectivity measurements constitute a lower bound. In addition, since our *in vivo* activity recordings are obtained under thalamo-cortical operating regimes established by local and long-range activity, they reflect, for instance, neuromodulatory input[83,84], specific short-term synaptic dynamics[85], and background synaptic bombardment, rarely present under *in vitro* conditions[86]. Finally, our inference method rests on stimulus-driven responses of neuronal populations and can thus yield estimates of connectivity in the context of functional response properties. This is important because fine-scale specificity of connectivity with respect to visual tuning similarity is a prominent motif in primary visual cortex[16–18,71], but typically requires technically challenging experiments involving a re-identification of neurons *in vitro* after their visual response properties were first characterized *in vivo*[16–18,20,71,87–89]. Such mapping is currently only performed by a few labs world-wide, which have the appropriate technical resources and broad methodological expertise.

Our model-based connectivity inference connects a mathematically interpretable, generative model of brain activity with the underlying circuit. In contrast to direct connectivity reconstruction[9,14,90] or inference approaches using large spiking-networks[10–12], which focus on detailed neuron-to-neuron connectivity, we are able to deduce cortical firing regimes and possible network states directly from the inferred weight matrix. Since the inferred connectivity is connected to a generative model, it can additionally be used to generate predictions about network activity that can then be tested experimentally. Other groups have recently made complementary progress in inferring connectivity from constrained models, focusing on spatial integration and behavioral state[91], inhibitory cell types[26] or response perturbations[92], showing that model-based connectivity inference can be used flexibly to elucidate multiple aspects of cortical information processing. Yet, our approach is unique in that it succeeds in inferring remarkably accurate cortical connectivity features directly and exclusively from dLGN and V1 response data, without imposing connectivity constraints taken from the literature.

Future directions

Our results can be extended in several directions. Conceptual advances on the experimental side demonstrate that not only orientation preference but general similarity in stimulus selectivity can influence V1 connectivity[18,20]. Future modeling efforts could thus expand the computational framework we present here to other aspects of feature selectivity, such as receptive field location, spatial and temporal frequency preferences, and similarity of responses to temporally varying stimuli. More substantially, while V1 neurons can be broadly classified into excitatory and inhibitory types, there are many known subtypes for both excitatory and in particular for inhibitory V1 neurons[13,15,15,39,48,93–97]. Future work, both on the recording and modeling front, could therefore consider more neuronal subtypes in relation to specific computations. For instance, interneuron-specific 2-photon calcium imaging[26,91,92] or optogenetic perturbations[67,92] would have the potential to inform multi-dimensional SSN-type models. Future work could also consider cortical layers separately, providing insights into the pronounced differences in connectivity and potentially operating regime observed experimentally[28,98,99]. For instance, L2/3 is well-known for its long-range connectivity between similarly tuned pyramidal cells[16,70], strong lateral inhibition and the sparsity of responses[51,98], while L5 coding is considered dense, with higher firing rates and broader stimulus selectivity[98,100]. Finally, future studies could investigate how recurrent V1 and thalamic input connectivity change as a function of behavioral state of the animal. Indeed, previous work has suggested that effective connectivity might change with locomotion and stimulus context[91,101]. Exploiting *in vivo* recordings, which can encompass the full spectrum of behavioral state-related neuronal modulations, and our model-based inference framework to study circuit connectivity, promises to generate novel insights into the potentially highly dynamic relationship between connectivity and computations.

Acknowledgements

This research was supported by the German Research Foundation (DFG) SPP2041 (TT, LB), DFG BU 1808/5-1 (LB) and TC 67/4-1 (TT), the Max Planck Society, University of Bonn Medical Center, DFG SFB 1233, Robust Vision: Inference Principles and Neural Mechanisms, TP 13, project number: 276693517 (LB), and by an add-on fellowship of the Joachim Herz Stiftung (GB). LB thanks M. Sotgia for lab management and support with animal handling and histology, A.H. Kotkat and X. Liu for contributions with spike sorting, S. Schörnich for IT support, B. Grothe for providing excellent research infrastructure. TT thanks L. Rebelo for assistance with literature review and L. Bernaez Timon for feedback on an earlier version of the manuscript. TT and LB thank all our group members for fruitful discussions.

Author contributions

N.K., L.B., and T.T. conceived the study. N.K. developed the mathematical proof of how contrast invariance constrains network connectivity, and designed the connectivity inference method. S.R. recorded and curated V1 data, developed the data analysis software, and showed contrast invariance of the recorded V1 data. N.K and S.R. pre-processed data for model fits, reviewed and analyzed V1 connectivity measures available in the experimental literature in relation to the inferred connectivity with input by T.T. and L.B.. G.T. contributed to the connectivity inference algorithm. G.B., Y.B. and M.S. recorded and curated dLGN data and additional V1 data, developed data analysis software. N.K., S.R., L.B., and T.T. contributed conceptual ideas, discussed and coordinated the project at all stages, wrote and edited the manuscript. All authors contributed to the final version of the manuscript.

References

1. Ahrens, M. B., Orger, M. B., Robson, D. N., Li, J. M. & Keller, P. J. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature Methods* **10**, 413–420 (2013).

2. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).

3. Zhang, Z., Russell, L. E., Packer, A. M., Gauld, O. M. & Häusser, M. Closed-loop all-optical interrogation of neural circuits in vivo. *Nature Methods* **15**, 1037–1040 (2018).

4. Abbott, L. F. *et al.* The mind of a mouse. *Cell* **182** (2020).

5. Bassett, D. S. & Sporns, O. Network neuroscience. *Nature Neuroscience* **20**, 353–364 (2017).

6. Vogels, T. P., Rajan, K. & Abbott, L. Neural network dynamics. *Annual Review of Neuroscience* **28**, 357–376 (2005).

7. Arkhipov, A. *et al.* Visual physiology of the layer 4 cortical circuit in silico. *PLoS Computational Biology* **14**, e1006535 (2018).

8. Billeh, Y. N. *et al.* Systematic Integration of Structural and Functional Data into Multi-scale Models of Mouse Primary Visual Cortex. *Neuron* **106**, 388–403.e18 (2020).

9. Markram, H. *et al.* Reconstruction and Simulation of Neocortical Microcircuitry. *Cell* **163**, 456–492 (2015).

10. Pillow, J. W. *et al.* Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–999 (2008).

11. Stevenson, I. H. *et al.* Bayesian inference of functional connectivity and network structure from spikes. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **17**, 203–213 (2008).

12. Ladenbauer, J., McKenzie, S., English, D. F., Hagens, O. & Ostojic, S. Inferring and validating mechanistic models of neural microcircuits based on spike-train data. *Nature Communications* **10**, 1–17 (2019).

13. Gouwens, N. W. *et al.* Classification of electrophysiological and morphological neuron types in the mouse visual cortex. *Nature Neuroscience* **22**, 1182–1195 (2019).

14. Seeman, S. C. *et al.* Sparse recurrent excitatory connectivity in the microcircuit of the adult mouse and human cortex. *Elife* **7**, e37349 (2018).

15. Jiang, X. *et al.* Principles of connectivity among morphologically defined cell types in adult neocortex. *Science* **350**, aac9462 (2015).

16. Ko, H. *et al.* Functional specificity of local synaptic connections in neocortical networks. *Nature* **473**, 87–91 (2011).

17. Hofer, S. B. *et al.* Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nature Neuroscience* **14**, 1045–52 (2011).

18. Cossell, L. *et al.* Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* **518**, 399–403 (2015).

19. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology* **160**, 106–154 (1962).

20. Znamenskiy, P. *et al.* Functional selectivity and specific connectivity of inhibitory neurons in primary visual cortex. *bioRxiv* (2018).

21. Ahmadian, Y., Rubin, D. B. & Miller, K. D. Analysis of the Stabilized Supralinear Network. *Neural Computation* **25**, 1994–2037 (2013). 1803.01446.

22. Rubin, D. B., Van Hooser, S. D. & Miller, K. D. The stabilized supralinear network: a unifying circuit motif underlying multi-input integration in sensory cortex. *Neuron* **85**, 402–417 (2015).

23. Gerstner, W., Sprekeler, H. & Deco, G. Theory and simulation in neuroscience. *Science* **338**, 60–65 (2012).

24. Anderson, J. S., Lampl, I., Gillespie, D. C. & Ferster, D. The Contribution of Noise to Orientation Tuning in Cat Visual Cortex. *Science* **290**, 1968–1972 (2000).

25. Hennequin, G., Ahmadian, Y., Rubin, D. B., Lengyel, M. & Miller, K. D. The Dynamical Regime of Sensory Cortex: Stable Dynamics around a Single Stimulus-Tuned Attractor Account for Patterns of Noise Variability. *Neuron* **98**, 846–860.e5 (2018).

26. Keller, A. J. *et al.* A Disinhibitory Circuit for Contextual Modulation in Primary Visual Cortex. *Neuron* **108**, 1181–1193.e8 (2020).

27. Kraynyukova, N. & Tchumatchenko, T. Stabilized supralinear network can give rise to bistable, oscillatory, and persistent activity. *PNAS* **115**, 3464–3469 (2018).

28. Sadeh, S. & Clopath, C. Inhibitory stabilization and cortical computation. *Nature Reviews Neuroscience* **22**, 21–37 (2021).

29. Carandini, M., Heeger, D. J. & Movshon, J. A. Linearity and normalization in simple cells of the macaque primary visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **17**, 8621–8644 (1997).

30. Sclar, G. & Freeman, R. D. Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental Brain Research* **46**, 457–461 (1982).

31. Skottun, B. C., Bradley, A., Sclar, G., Ohzawa, I. & Freeman, R. D. The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *Journal of Neurophysiology* **57**, 773–786 (1987).

32. Niell, C. M. & Stryker, M. P. Highly Selective Receptive Fields in Mouse Visual Cortex. *Journal of Neuroscience* **28**, 7520–7536 (2008).

33. Miller, K. D. & Troyer, T. W. Neural noise can explain expansive, power-law nonlinearities in neural response functions. *Journal of Neurophysiology* **87**, 653–659 (2002).

34. Hansel, D. & Van Vreeswijk, C. How noise contributes to contrast invariance of orientation tuning in cat visual cortex. *Journal of Neuroscience* **22**, 5118–5128 (2002).

35. Kawaguchi, Y., Katsumaru, H., Kosaka, T., Heizmann, C. W. & Hama, K. Fast spiking cells in rat hippocampus (CA1 region) contain the calcium-binding protein parvalbumin. *Brain Research* **416**, 369–374 (1987).

36. Celio, M. R. Parvalbumin in most γ-aminobutyric acid-containing neurons of the rat cerebral cortex. *Science* **231**, 995–997 (1986).

37. Kawaguchi, Y. Physiological subgroups of nonpyramidal cells with specific morphological characteristics in layer II/III of rat frontal cortex. *Journal of Neuroscience* **15**, 2638–2655 (1995).

38. Tremblay, R., Lee, S. & Rudy, B. GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits. *Neuron* **91**, 260–292 (2016).

39. Pfeffer, C. K., Xue, M., He, M., Huang, Z. J. & Scanziani, M. Inhibition of inhibition in visual cortex: The logic of connections between molecularly distinct interneurons. *Nature Neuroscience* **16**, 1068–1076 (2013).

40. Paxinos, G. & Franklin, K. B. *Paxinos and Franklin's the mouse brain in stereotaxic coordinates* (Academic press, 2019).

41. Swindale, N. V. Orientation tuning curves: Empirical description and estimation of parameters. *Biological Cybernetics* **78**, 45–56 (1998).

42. Bonhoeffer, T., Kim, D.-S., Malonek, D., Shoham, D. & Grinvald, A. Optical Imaging of the Layout of Functional Domains in Area 17 and Across the Area 17/18 Border in Cat Visual Cortex. *European Journal of Neuroscience* **7**, 1973–1988 (1995).

43. Olsen, S. R., Bortone, D. S., Adesnik, H. & Scanziani, M. Gain control by layer six in cortical circuits of vision. *Nature* **483**, 47–54 (2012).

44. Ringach, D. L., Sapiro, G. & Shapley, R. A subspace reverse-correlation technique for the study of visual neurons. *Vision Research* **37**, 2455–2464 (1997).

45. Mazer, J. A., Vinje, W. E., McDermott, J., Schiller, P. H. & Gallant, J. L. Spatial frequency and orientation tuning dynamics in area V1. *Proceedings of the National Academy of Sciences of the United States of America* **99**, 1645–1650 (2002).

46. Hubert, L. J., Golledge, R. G. & Costanzo, C. M. Generalized procedures for evaluating spatial autocorrelation. *Geographical Analysis* **13**, 224–233 (1981).

47. Peirce, J. W. The potential importance of saturating and supersaturating contrast response functions in visual cortex. *Journal of Vision* **7**, 1–10 (2007).

48. Millman, D. J. *et al.* VIP interneurons in mouse primary visual cortex selectively enhance responses to weak but specific stimuli. *eLife* **9**, e55130 (2020).

49. Masland, R. H. & Martin, P. R. The unsolved mystery of vision. *Current Biology* **17**, R577–R582 (2007).

50. Kerlin, A. M., Andermann, M. L., Berezovskii, V. K. & Reid, R. C. Broadly Tuned Response Properties of Diverse Inhibitory Neuron Subtypes in Mouse Visual Cortex. *Neuron* **67**, 858–871 (2010).

51. Niell, C. M. & Stryker, M. P. Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. *Neuron* **65**, 472–479 (2010).

52. Sohya, K., Kameyama, K., Yanagawa, Y., Obata, K. & Tsumoto, T. GABAergic Neurons Are Less Selective to Stimulus Orientation than Excitatory Neurons in Layer II/III of Visual Cortex, as Revealed by In Vivo Functional Ca2+ Imaging in Transgenic Mice. *J Neurosci* **27**, 2145–2149 (2007).

53. Liu, B.-h. *et al.* Visual Receptive Field Structure of Cortical Inhibitory Neurons Revealed by Two-Photon Imaging Guided Recording. *Journal of Neuroscience* **29**, 10520–10532 (2009).

54. Busse, L., Wade, A. R. & Carandini, M. Representation of Concurrent Stimuli by Population Activity in Visual Cortex. *Neuron* **64**, 931–942 (2009).

55. Lee, S., Park, J. & Smirnakis, S. M. Internal Gain Modulations, But Not Changes in Stimulus Contrast, Preserve the Neural Code. *Journal of Neuroscience* **39**, 1671–1687 (2019).

56. Tring, E. & Ringach, D. L. On the Subspace Invariance of Population Responses. *arXiv* (2018). 1811.03251.

57. Priebe, N. J. & Ferster, D. Inhibition, spike threshold, and stimulus selectivity in primary visual cortex. *Neuron* **57**, 482–497 (2008).

58. Priebe, N. J., Mechler, F., Carandini, M. & Ferster, D. The contribution of spike threshold to the dichotomy of cortical simple and complex cells. *Nature Neuroscience* **7**, 1113–1122 (2004).

59. Tan, A. Y., Brown, B. D., Scholl, B., Mohanty, D. & Priebe, N. J. Orientation selectivity of synaptic input to neurons in mouse and cat primary visual cortex. *Journal of Neuroscience* **31**, 12339–12350 (2011).

60. Lien, A. D. & Scanziani, M. Tuned thalamic excitation is amplified by visual cortical circuits. *Nature Neuroscience* **16**, 1315–1323 (2013).

61. Reinhold, K., Lien, A. D. & Scanziani, M. Distinct recurrent versus afferent dynamics in cortical visual processing. *Nature Neuroscience* **18**, 1789–1797 (2015).

62. Ji, X.-y. *et al.* Thalamocortical innervation pattern in mouse auditory and visual cortex: laminar and cell-type specificity. *Cerebral Cortex* **26**, 2612–2625 (2016).

63. Markram, H. *et al.* Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience* **5**, 793–807 (2004).

64. Gonchar, Y., Wang, Q. & Burkhalter, A. H. Multiple distinct subtypes of gabaergic neurons in mouse visual cortex identified by triple immunostaining. *Frontiers in neuroanatomy* **2**, 3 (2008).

65. Tsodyks, M. V., Skaggs, W. E., Sejnowski, T. J. & McNaughton, B. L. Paradoxical effects of external modulation of inhibitory interneurons. *Journal of neuroscience* **17**, 4382–4388 (1997).

66. Ozeki, H., Finn, I. M., Schaffer, E. S., Miller, K. D. & Ferster, D. Inhibitory stabilization of the cortical network underlies visual surround suppression. *Neuron* **62**, 578–592 (2009).

67. Sanzeni, A. *et al.* Inhibition stabilization is a widespread property of cortical networks. *eLife* **9**, e54875 (2020).

68. Karnani, M. M. *et al.* Cooperative subnetworks of molecularly similar interneurons in mouse neocortex. *Neuron* **90**, 86–100 (2016).

69. Allen Institute for Brain Science. Synaptic physiology coarse matrix dataset (2019). Available from: https://brain-map.org/explore/connectivity/synaptic-physiology.

70. Bosking, W. H., Zhang, Y., Schofield, B. & Fitzpatrick, D. Orientation Selectivity and the Arrangement of Horizontal Connections in Tree Shrew Striate Cortex. *J Neurosci* **17**, 2112–2127 (1997).

71. Bock, D. D. *et al.* Network anatomy and in vivo physiology of visual cortical neurons. *Nature* **471**, 177–184 (2011).

72. Persi, E., Hansel, D., Nowak, L., Barone, P. & Van Vreeswijk, C. Power-law input-output transfer functions explain the contrast-response and tuning properties of neurons in visual cortex. *PLoS Comput Biol* **7**, e1001078 (2011).

73. Scholl, B., Tan, A. Y. Y., Corey, J. & Priebe, N. J. Emergence of Orientation Selectivity in the Mammalian Visual Pathway. *Journal of Neuroscience* **33**, 10616–10624 (2013).

74. Kondo, S. & Ohki, K. Laminar differences in the orientation selectivity of geniculate afferents in mouse primary visual cortex. *Nature Neuroscience* **19**, 316–319 (2016).

75. Sun, W., Tan, Z., Mensh, B. D. & Ji, N. Thalamus provides layer 4 of primary visual cortex with orientation-and direction-tuned inputs. *Nature Neuroscience* **19**, 308 (2016).

76. Li, Y. T., Ibrahim, L. A., Liu, B. H., Zhang, L. I. & Tao, H. W. Linear transformation of thalamocortical input by intracortical excitation. *Nature Neuroscience* **16**, 1324–1330 (2013).

77. Young, H., Belbut, B., Baeta, M. & Petreanu, L. Laminar-specific cortico-cortical loops in mouse visual cortex. *eLife* **10**, e59551 (2021).

78. Sadeh, S. & Clopath, C. Theory of neuronal perturbome in cortical networks. *Proceedings of the National Academy of Sciences* **117**, 26966–26976 (2020).

79. Karnani, M. M., Agetsuma, M. & Yuste, R. A blanket of inhibition: functional inferences from dense inhibitory connectivity. *Current opinion in Neurobiology* **26**, 96–102 (2014).

80. Morgenstern, N. A., Bourg, J. & Petreanu, L. Multilaminar networks of cortical neurons integrate common inputs from sensory thalamus. *Nature Neuroscience* **19**, 1034–1040 (2016).

81. Petreanu, L., Mao, T., Sternson, S. M. & Svoboda, K. The subcellular organization of neocortical excitatory connections. *Nature* **457**, 1142–1145 (2009).

82. Callaway, E. M. & Katz, L. C. Photostimulation using caged glutamate reveals functional circuitry in living brain slices. *Proceedings of the National Academy of Sciences* **90**, 7661–7665 (1993).

83. Jacob, S. N. & Nienborg, H. Monoaminergic Neuromodulation of Sensory Processing. *Frontiers in Neural Circuits* **12**, 51 (2018).

84. Thiele, A. Muscarinic Signaling in the Brain. *Annual Review of Neuroscience* **36**, 271–294 (2013).

85. Motanis, H., Seay, M. J. & Buonomano, D. V. Short-Term Synaptic Plasticity as a Mechanism for Sensory Timing. *Trends in neurosciences* **41**, 701–711 (2018).

86. Destexhe, A., Rudolph, M. & Paré, D. The high-conductance state of neocortical neurons in vivo. *Nature Reviews Neuroscience* **4**, 739–751 (2003).

87. Weiler, S. *et al.* High-yield in vitro recordings from neurons functionally characterized in vivo. *Nature Protocols* **13**, 1275–1293 (2018).

88. Weiler, S. *et al.* Relationship between input connectivity, morphology and orientation tuning of layer 2/3 pyramidal cells in mouse visual cortex. *bioRxiv* (2020).

89. Lee, W.-C. A. *et al.* Anatomy and function of an excitatory network in the visual cortex. *Nature* **532**, 370–374 (2016).

90. Reimann, M. W., King, J. G., Muller, E. B., Ramaswamy, S. & Markram, H. An algorithm to predict the connectome of neural microcircuits. *Frontiers in Computational Neuroscience* **9**, 1–18 (2015).

91. Dipoppa, M. *et al.* Vision and Locomotion Shape the Interactions between Neuron Types in Mouse Visual Cortex. *Neuron* **98**, 602–615.e8 (2018).

92. Palmigiano, A. *et al.* Structure and variability of optogenetic responses identify the operating regime of cortex. *bioRxiv* 2020.11.11.378729 (2021).

93. Gouwens, N. W. *et al.* Integrated Morphoelectric and Transcriptomic Classification of Cortical GABAergic Cells. *Cell* **183**, 935–953.e19 (2020).

94. Rudy, B., Fishell, G., Lee, S. & Hjerling-Leffler, J. Three groups of interneurons account for nearly 100% of neocortical GABAergic neurons. *Developmental Neurobiology* **71**, 45–61 (2011).

95. Peron, S. *et al.* Recurrent interactions in local cortical circuits. *Nature* **579**, 1–4 (2020).

96. Okun, M. *et al.* Diverse coupling of neurons to populations in sensory cortex. *Nature* **521**, 511–515 (2015).

97. Muñoz, W., Tremblay, R., Levenstein, D. & Rudy, B. Layer-specific modulation of neocortical dendritic inhibition during active wakefulness. *Science* **355**, 954–959 (2017).

98. Harris, K. D. & Mrsic-Flogel, T. D. Cortical connectivity and sensory coding. *Nature* **503**, 51–58 (2013).

99. Rao, R. P. N. & Ballard, D. H. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* **2**, 79–87 (1999).

100. Sakata, S. & Harris, K. D. Laminar Structure of Spontaneous and Sensory-Evoked Population Activity in Auditory Cortex. *Neuron* **64**, 404–418 (2009).

101. Pakan, J. M. *et al.* Behavioral-state modulation of inhibition is context-dependent and cell type specific in mouse visual cortex. *eLife* **5**, e14985 (2016).

102. Spacek, M. A. *et al.* Robust effects of corticothalamic feedback during naturalistic visual stimulation. *bioRxiv* (2021).

103. Rueden, C. T. *et al.* Imagej2: Imagej for the next generation of scientific image data. *BMC Bioinformatics* **18**, 529 (2017).

104. Schindelin, J. *et al.* Fiji: an open-source platform for biological-image analysis. *Nature Methods* **9**, 676–682 (2012).

105. Pachitariu, M., Steinmetz, N. A., Kadir, S. N., Carandini, M. & Harris, K. D. Fast and accurate spike sorting of high-channel count probes with KiloSort. *Advances in Neural Information Processing Systems* **29**, 4448–4456 (2016).

106. Spacek, M. A., Blanche, T. & Swindale, N. Python for large-scale electrophysiology. *Frontiers in Neuroinformatics* **2**, 9 (2009).

107. Van Rossum, G. & Drake, F. L. *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).

108. Yatsenko, D., Walker, E. Y. & Tolias, A. S. Datajoint: a simpler relational data model. *arXiv preprint arXiv:1807.11104* (2018).

109. Mitzdorf, U. Current source-density method and application in cat cerebral cortex: Investigation of evoked potentials and EEG phenomena. *Physiological Reviews* **65**, 37–100 (1985).

110. Pettersen, K. H., Devor, A., Ulbert, I., Dale, A. M. & Einevoll, G. T. Current-source density estimation based on inversion of electrostatic forward solution: Effects of finite extent of neuronal activity and conductivity discontinuities. *Journal of Neuroscience Methods* **154**, 116–133 (2006).

111. Denker, M., Yegenoglu, A. & Grün, S. Collaborative HPC-enabled workflows on the HBP Collaboratory using the Elephant framework. In *Neuroinformatics 2018*, P19 (2018).

112. Heumann, D., Leuba, G. & Rabinowicz, T. Postnatal development of the mouse cerebral neocortex. II. Quantitative cytoarchitectonics of visual and auditory areas. *Journal Fur Hirnforschung* **18**, 483–500 (1977).

113. Virtanen, P. *et al.* Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature Methods* **17**, 261–272 (2020).

114. Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**, 357–362 (2020).

115. Rey, S. J. & Anselin, L. PySAL: A Python Library of Spatial Analytical Methods. *The Review of Regional Studies* **37**, 5–27 (2007).

116. Mardia, K. & Jupp, P. *Directional Statistics*. Wiley Series in Probability and Statistics (Wiley, 2009).

Methods

Data and code availability

Any additional information required to reanalyze the data reported in this paper is available from the lead author upon request. During the review process, the reviewers can find the computer code underlying our data analysis here. www.tchumatchenko.de/CodeForReviewers.zip, We will publish the code upon the publication of the manuscript.

Experiments and data analysis

All procedures complied with the European Communities Council Directive 2010/63/EU and the German Law for Protection of Animals, and were approved by local authorities, following appropriate ethics review.

Animals
Recordings were performed in 4 adult male Ntsr1-Cre mice (3 hemizygous Tg, 1 negative control, median age at first recording session: 24.4 weeks; B6.FVB(Cg)-Tg(Ntsr1-cre)GN220Gsat/Mmcd; MMRRC, #030648-UCD) and 2 (1 male, 1 female) PV-Cre mice (median age: 17.9 weeks; B6.129P2-Pvalb$^{tm1(cre)Arbr}$/J; Jackson Laboratory, #017320).

Surgery
The surgical procedures are described in detail in[102]. In brief: mice were administered an analgesic (Metamizole, 200 mg/kg, sc, MSD Animal Health, Brussels, Belgium) and put under isoflurane anesthesia (5% in oxygen at start, then lowered to 0.5%–2% in oxygen, CP-Pharma, Burgdorf, Germany), the depth of which was constantly monitored. After shaving and disinfecting the scalp, a skin incision was performed and the skull cleaned of any remaining tissue. Upon positioning the head in a skull-flat position, a custom lightweight aluminium head bar with an opening over dLGN and V1 was placed on the skull and fixated using dental cement. For V1 recordings and optogenetic stimulation unrelated to this study in PV-Cre mice, a small craniotomy above V1 was performed and $\sim$ 0.2μL of pAAV9/1.EF1a.DIO.hChR2(H134R)-eYFP.WPRE.hGH (Addgene, #20298-AAV9) dyed with fast-green (Sigma-Aldrich, St. Louis, USA) was injected through the entire depth of the cortex. In the Ntsr1-Cre mice used for additional V1 and dLGN recordings, a similar craniotomy was performed and $\sim$ 0.35μL of stGtACR2 (rAAV2/1-pAAV-hSyn1-SIO-stGtACR2-FusionRed, Addgene, #105677) were injected in the infragranular layers of cortex for experiments with suppression of corticothalamic feedback unrelated to the current study. Post-injection, the opening was filled with Kwik-Cast (WPI Germany, Berlin, Germany). Long-term analgesic (Meloxicam, 2 mg/kg, sc, Böhringer Ingelheim, Ingelheim, Germany) was administered and continued to be administered for 3 consecutive days. After at least 1 week of recovery, animals were gradually habituated to the experimental setup, by first handling them and then simulating the experimental procedure. To allow for virus expression, neural recordings started no sooner than 3 weeks after injection. On the day prior to the first day of recording, mice were fully anesthetized using the same procedures as for the initial surgery, and a craniotomy (ca. 1.5 mm^2) was performed over dLGN and/or V1, and re-sealed with Kwik-Cast. As long as the animals did not show signs of discomfort, the long-term analgesic Metacam was administered only once at the end of surgery, to avoid any confounding effect on experimental results. Recordings were performed daily and continued for as long as the quality of the electrophysiological signals remained high.

Experimental setup
Our experimental configuration for *in-vivo* recordings is described in detail in[102]. In brief: mice were head-fixed and could run freely on an air-suspended styrofoam ball while stimuli were presented to the right visual field on a gamma-corrected LCD screen. Extracellular neural signals were recorded with 32-channel silicon probes (Neuronexus, A1x32Edge-5mm-20-177-A32, Ann Arbor, USA) for the 4 Ntsr1-Cre mice, a 32-channel silicon probe for one PV-Cre mouse (A1x32-Edge-5mm-20-177-A32 and A1x32Edge-5mm-20-177-A32), and a 64-channel silicon probe (A1x64-Poly2-6mm-23s-160-A64) for the other PV-Cre mouse. Ball movements were registered at 90 Hz by two optical mice connected to an Arduino-type microcontroller. Eye movements were monitored under infrared light illumination.

For photostimulation of V1 PV+ inhibitory interneurons, an optic fiber (910 μm diameter, Thorlabs, Newton, USA) was coupled to a light-emitting diode (LED, center wavelength 470 nm, M470F1, Thorlabs, Newton, USA) and positioned with a micromanipulator less than 1 mm above the exposed surface of V1. A black metal foil surrounding the tip of the head bar holder prevented the photostimulation light from reaching the animal's eyes.

Perfusion and histology
After the final recording session, mice were first administered an analgesic (Metamizole, 200 mg/kg, sc, MSD Animal Health, Brussels, Belgium) and following a 30 min wait period were transcardially perfused under deep anesthesia using a cocktail of Medetomidin (0.5 mL/kg), Midazolam (1 mL/kg), and Fentanyl (1 mL/kg) (ip). Perfusion was first done with Ringer's lactate solution followed by 4% paraformaldehyde (PFA) in 0.2 M sodium phosphate buffer (PBS).

To verify recording site and virus expression, we performed histological analyses. Brains were removed, postfixed in PFA for 24 h, and then rinsed with and stored in PBS at 4°C. Slices (40 μm) were cut using a vibrotome (Leica VT1200 S, Leica, Wetzlar, Germany), mounted on glass slides with Vectashield DAPI (Vector Laboratories, Burlingame, USA), and coverslipped. A fluorescent microscope (BX61 Systems Microscope, Olympus, Tokyo, Japan) was used to inspect slices for the presence of yellow fluorescent protein (eYFP) and DiI. Recorded images were processed using FIJI[103,104].

Stimulus
We used custom software (EXPO, https://sites.google.com/a/nyu.edu/expo/home) to present visual stimuli on a gamma-calibrated liquid crystal display (LCD) monitor (Samsung SyncMaster 2233RZ; mean luminance 50 cd/m^2, 60 Hz) at 25 cm distance to the animal's right eye (spanning $\sim$ 108x66°, small angle approximation). Mice were presented with three 12 min random sequences of briefly flashed (84 ms), full-screen grating stimuli. The random sequences were drawn from 2304 unique gratings covering 12 orientations ($0, 15, 30, 45, 60, 75, 90, 105, 120, 135, 150, 165°$), 8 contrasts (0, 0.04, 0.10, 0.19, 0.30, 0.46, 0.69, 1), 6 spatial frequencies (0.01, 0.02, 0.06, 0.14, 0.33, 0.80 cyc/°) and 4 spatial phases ($0, 90, 180, 270°$). One sequence consisted of 9216 gratings. Between the sequences, a blank gray screen was displayed for 1 min. For V1 recordings in PV-Cre mice expressing ChR2, light pulses (10 Hz, 1 ms pulses) were delivered from the optical fiber during these periods; analyses of the blank screen responses or photostimulation effects were not included in the current study. Typically, the stimulus sequence was presented once per electrode penetration, except in two cases, where the sequence was run twice during one electrode penetration but data from each run was analyzed separately.

Data analysis
Wideband extracellular signals were digitized at 30 kHz (Blackrock microsystems, Blackrock Microsystems Europe GmbH, Hannover DE). To obtain single unit activity from extracellular recordings, the open source, Matlab-based, automated spike sorting toolbox Kilosort[105] was used. Resulting clusters were manually refined using Spyke[106], a Python application that allows the selection of channels and time ranges around clustered spikes for realignment, as well as representation in 3D space using dimension reduction (multichannel PCA, ICA, and/or spike time). Exhaustive pairwise comparisons of similar clusters allowed merging of potentially over-clustered units. All further analyses were performed using an SQL data base and a custom-made analysis pipeline programmed in `python`[107] and managed via `datajoint`[108].

Spike waveshape analysis
From the mean waveform of the maximum-response electrode channel of each single unit, the time between trough and peak (rise time) and the half-width at half-height of the peak were calculated. Exploiting the waveshapes of all V1 units processed using the same pipeline (N = 428 from 10 mice), a k-means algorithm was used to cluster the data into 2 populations.

Laminar location
We used current source density (CSD) analysis[109] for recordings in area V1 to determine the laminar position of electrode contacts. To obtain the LFP, we first down-sampled the signal to 1 kHz before applying a bandpass filter (3–90 Hz, second-order Butterworth filter). We computed the CSD using the iCSD method[110] implemented in `elephant` (RRID:SCR_003833)[111]. We assigned the base of layer 4 to the contact that was closest to the earliest CSD polarity inversion. The remaining contacts were assigned to layers based on relative layer thickness reported by[112], assuming a thickness of 1.2 mm for mouse visual cortex. The resulting distribution of neurons across layers was: L2/3 - 10/131 (7.6%), L4 - 29/131 (22.1%), L5 - 47/131 (35.9%), L6 - 45/131 (34.4%).

Temporal response kernels via reverse correlation
After eliminating units with overall very low firing rate (<0.1 Hz), the probability of a stimulus preceding a spike by a time δt from -50 to 350 ms was computed for each unique grating stimulus. This was done by binning spike times in 1 ms windows and then counting how often a specific stimulus occurred δt before a spike. After normalizing this histogram by dividing by the total number of spikes, the posterior distribution was calculated according to Bayes'

theorem by multiplying with the probability of a spike occurring and dividing by the probability of the stimulus occurring:

$$
\begin{aligned}
P(\text{Spike}|\text{Grating}) &= \frac{P(\text{Grating}|\text{Spike}) \cdot P(\text{Spike})}{P(\text{Grating})} \\
&= \frac{\frac{\text{Grating And Spike Bins}}{\text{Total Spike Bins}} \cdot \frac{\text{Total Spike Bins}}{\text{Total Bins}}}{\frac{\text{Grating Bins}}{\text{Total Bins}}} \\
&= \frac{\text{Grating And Spike Bins}}{\text{Grating Bins}}
\end{aligned}
$$

By dividing by the 1 ms bin duration, the resulting probabilities could be directly converted to firing rates and thus give a temporal response kernel for each unique grating stimulus.

Determining visual responsiveness

To eliminate non-responsive or noise-dominated units, the variance of the temporal kernels across stimuli was calculated and tested for non-randomness using the Wald-Wolfowitz-Test (WWT). The WWT uses the distribution of consecutive ones and zeros in a binary sequence, which should follow a normal distribution in a random sequence, to statistically determine randomness of the sequence. The test can be applied to a non-binary sequence by converting it to a binary sequence via a threshold criterion, typically the mean or median. Before applying the WWT, global trends in the sequence should be removed by either filtering or applying an approximation of the derivative. As the response to grating contrast is a robust indicator of visual responsiveness, the analysis was performed using the aggregate variance across grating contrast, which was computed by averaging the kernels across grating orientation, spatial frequency and spatial phase, before computing the variance across the resulting contrast kernels. To further increase signal-to-noise ratio, the partial variance was squared before computing the differences across time to remove any global trends. On the resulting sequence, the WWT was performed using the median as cutoff criterion. Because the WWT can miss narrow peaks, even if they are high, a second WWT was computed on the absolute values of the sequence. A unit was classified as visually responsive if one of the WWT's was significant and the test statistic of both WWT's was negative, indicating fewer sign changes than expected by chance.

Determining optimal time point

The time point of optimal response was determined via the peak of the summed aggregate variances across stimuli. First, the partial variances were computed for all four stimulus parameters as described above for grating contrast. The resulting partial variances were then summed and the time point of the first peak exceeding half the modulation depth of the result was selected as the optimal response time point δt_{opt} peak detection using `scipy`,[113].

Response profiles

In a 20 ms window around δt_{opt}, responses were averaged over grating spatial phase and frequency, resulting in two-dimensional response matrices covering grating orientation and contrast.

Contrast-invariance

Contrast-invariance of a unit was assessed by applying a singular value decomposition (SVD) to its response matrix, separating the SVD's principle component and residual, and computing the Gamma index of spatial autocorrelation on the residual[46]. The Gamma index is computed by computing a similarity matrix for all data points and then masking the similarity matrix with a contiguity matrix that considers data points that share an edge to be neighbors. The index itself is the sum of all entries in the masked similarity matrix. Patterns are detected by randomly shuffling data points and comparing the original index against the resulting distribution. The strength of the first SVD component was assessed as its power in the SVD: the quotient between the squared first singular value and the sum of squares over all singular values. Residual strength was then computed as one minus first component power. The SVD was calculated using `numpy`[114] and the Gamma index using `pysal`[115]. Neurons with a z-scored Gamma index $g_z > 1.96$ and residual strength $p > 0.05$ were classified as contrast-dependent.

Tuning model

Contrast invariant units were fitted with a two-dimensional tuning model consisting of a hyperbolic ratio function with supersaturation parameter[47] and a wrapped Gaussian[41]:

$$
r(c, \psi, \theta) = r_0 + (r_{max} - r_0) \cdot \frac{c^n}{c_{50}^{s \cdot n} + c^{s \cdot n}} \cdot \sum_{m=-\infty}^{m=\infty} \exp\left\{\frac{-(\psi - \theta + \pi m)^2}{2\sigma^2}\right\}, \tag{S1}
$$

where c is stimulus contrast, ψ is the stimulus orientation, θ the preferred orientation, r_0 the baseline response and r_{max} the peak response of the neuron. The tuning model was fitted to each neuron's response matrix in a least-

squared sense using the `scipy.optimize` library. Quality of fit was assessed via the coefficient of determination (r^2) and only units with $r^2 > 0.4$ were used for further analysis.

Width parameters of orientation tuning functions, σ in our case, scale non-linearly when tuning is either strong or weak, depending on the specific function. When quantitatively analysing data that contains a broad spectrum of tuning, it is thus advisable to use measures that are not distorted by such non-linear scaling. Accordingly, orientation selectivity[42,43] was quantified as

$$\mathrm{OSI} = \frac{\sqrt{(\sum R_k \sin(2\theta_k))^2 + (\sum R_k \cos(2\theta_k))^2}}{\sum R_k} \tag{S2}$$

where R_k is the response to the kth direction given by θ_k.

Contrast sensitivity was quantified as contrast at half height of the contrast response function. Since some of our recorded contrast response functions did not saturate, even for full contrast, we preferred this measure as opposed to the parameter c_{50} of the hyperbolic ratio function.

To construct population tuning curves for orientation and contrast, individual orientation and contrast tuning curves were averaged, after aligning individual neurons to their preferred orientation. Population contrast-invariance was assessed by applying the above explained SVD and spatial autocorrelation analysis to the population tuning curves.

Statistics
All statistics were performed using functions from `scipy.stats` and `statsmodels`.

Theory and model

Determining populations' responses from the recorded data for the SSN model fit

The hyperbolic ratio function used to describe and quantify the recorded contrast responses imposes a sigmoidal shape on the contrast response. However, the output of the SSN model itself can explain how S-shaped contrast responses arise from a recurrent network wiring. Therefore, to avoid an additional fitting bias, we didn't use the hyperbolic ratio function to represent the recorded populations' responses before we fitted the SSN model to the data. To determine E, I and thalamic population responses to a stimulus of orientation ψ, the function

$$R(c_i, \psi - \theta) = r_0 + (r(c_i) - r_0) \cdot \sum_{m=-\infty}^{\infty} \exp\left\{\frac{-(\psi - \theta + \pi m)^2}{2\sigma^2}\right\} \tag{S3}$$

was fitted to the two-dimensional contrast and orientation responses of individual units to determine the contrast response functions $r(c_i)$ at eight contrast values c_i as well as the width σ of orientation tuning curves (see Swindale[41] for justification of wrapped Gaussian fit). Then the responses of the units were aligned such that their preferred orientations θ coincided with 0°. The E, I, and thalamic population contrast responses at each contrast value c_i in (Fig. 3) were computed as an average $r(c_i)$ in the corresponding population. The population orientation tuning widths were computed as an average σ over the corresponding population.

SSN model with two populations, stability of steady states

The two population SSN model is given by the equations

$$\tau_X \cdot \frac{dr_X(t,C)}{dt} + r_X(t,C) = \Big(J_{XE} \cdot r_E(t,C) - J_{XI} \cdot r_I(t,C) + T_{\mathrm{dLGN}}(C) \cdot g_X\Big)_+^n, \quad X \in \{E,I\}. \tag{S4}$$

The steady states $r_E(C)$ and $r_I(C)$ defined by the equations ($dr_X/dt = 0$)

$$r_X(C) = \Big(J_{XE} \cdot r_E(C) - J_{XI} \cdot r_I(C) + T_{\mathrm{dLGN}}(C) \cdot g_X\Big)_+^n, \quad X \in \{E,I\}. \tag{S5}$$

are stable exactly when the inequalities

$$J_{EE} r_E^{1-1/n} - (n \cdot \det J \cdot r_E^{1-1/n} + J_{II}) r_I^{1-1/n} < 1/n \tag{S6}$$

and

$$\tau_E + \tau_I + \tau_E J_{II} n r_I^{1-1/n} - \tau_I J_{EE} n r_E^{1-1/n} > 0 \tag{S7}$$

are fulfilled[27]. To guarantee stability of the fitted firing rates, we incorporated the inequality in Eq. S6 in the parameter inference algorithm. We note that the second inequality in Eq. S7 can always be fulfilled if we choose sufficiently large τ_E and/or small τ_I.

Determining the two-population SSN parameters from contrast responses

The SSN model with the initially unknown parameters J_{XY}, g_X, and n was required to generate stable steady states $r_X(C)$ (Eq. S5 - Eq. S7), which closely approximated the average recorded cortical and thalamic contrast responses (Eq. S3). For each fixed n and eight contrasts C, Eq. 1 represented an over-determined system of 16 linear equations with six unknown connectivity weights J_{XY}, g_X, which always has a unique solution. We called this solution valid, if additionally, the constants J_{XY}, g_X were positive and lead to a stable steady state of the SSN model. We note that the weights computed directly from the average V1 and thalamic contrast responses did not lead to any valid solutions for the exponents n ranging from 1.1 to 5. Therefore, we randomly generated triplets of V1 and thalamic contrast response curves within ±sem error bar areas of the contrast responses and computed corresponding sets of the connectivity weights J_{XY}, g_X as solutions of the over-determined linear system in Eq. 1 for each triplet. Overall, the fraction of valid connectivity weights was less than 0.1% for all n, and was a monotonically increasing function of n with few valid fits found for n close to 1 (Fig. S5A). Since the initial SSN parameters J_{XY}, g_X, and n were computed for random response triplets and not for the average contrast responses, we optimized them to closely approximate the average responses by minimizing the score function

$$\text{Score}_{\text{fit}}(J, g, n) = \frac{1}{F(n)} \sum_{i=1}^{8} \frac{(r_E^{\text{fit}}(C_i) - r_E^{\text{av}}(C_i))^2}{\sigma_E^2(C_i)} + \frac{(r_I^{\text{fit}}(C_i) - r_I^{\text{av}}(C_i))^2}{\sigma_I^2(C_i)} + \frac{(T_{\text{dLGN}}^{\text{fit}}(C_i) - T_{\text{dLGN}}^{\text{av}}(C_i))^2}{\sigma_{\text{dLGN}}^2(C_i)}. \tag{S8}$$

Based on the score function (Eq. S8), the contrast responses with smaller standard deviation σ_X were approximated with higher precision than those with larger standard deviation. We note that as expected, lower firing rates had lower variability in our recordings. We divided the difference between the fit and the recorded average by the fraction of valid fits F as a function of n to reinforce the exponents n leading to a larger fraction of initial valid fits. Each parameter set J_{XY}, g_X, n in the final distribution of 10^3 fits was a parameter set with the best score out of 10^4 optimized randomly generated valid initial fits.

Comparison of the inferred connectivity parameters with direct connectivity measures: upper bound for the relation g_E/J_{EE}

Lien et al.[60] reported that the upper bound for the contribution of thalamic inputs compared to the total postsynaptic charge of the E cortical neurons was 36 + 2% for full screen, 100% contrast drifting gratings see also[61,76]. We used the upper bound of 38% to estimate an experimentally plausible region for the relation g_E/J_{EE} in Fig. 3I. To this end, we assumed that the relative contribution of the thalamic input $g_E \cdot T_{\text{dLGN}}$ to the E population with respect to the total input to E population $J_{EE} \cdot r_E - J_{EI} \cdot r_I + g_E \cdot T_{\text{dLGN}}$ was smaller than 38%. Using the estimate

$$\frac{g_E \cdot T_{\text{dLGN}}}{J_{EE} \cdot r_E + g_E \cdot T_{\text{dLGN}}} < \frac{g_E \cdot T_{\text{dLGN}}}{J_{EE} \cdot r_E - J_{EI} \cdot r_I + g_E \cdot T_{\text{dLGN}}} < 0.38,$$

we computed the approximate upper bound for the relation g_E/J_{EE}

$$\frac{g_E}{J_{EE}} < \frac{0.38}{1 - 0.38} \cdot \frac{r_E}{T_{\text{dLGN}}} \approx 0.55.$$

Here we used the firing rates r_E and T_{dLGN} recorded for 100% of contrast.

Relation between parameters g_I and g_E

We determined the relation between the parameters g_E and g_I based on the measurements published in Ji et al.[62]. The V1 E neurons received direct thalamic input with the probabilities 15/19 in layer 2/3, 19/19 in layer 4, 8/8 in layer 5, and 7/9 in layer 6[62]. The V1 I neurons received direct thalamic input with the probabilities 14/17 in layer 2/3, 15/15 in layer 4, 15/15 in layer 5, and 9/11 in layer 6[62]. The adjusted peak amplitudes of postsynaptic potentials amounted to $190 \pm 78pA$ for E and $475 \pm 178pA$ for I V1 populations in layer 2/3, $430 \pm 97pA$ for E and $1111 \pm 260pA$ for I cortical populations in layer 4, $190 \pm 73pA$ for E and $596 \pm 178pA$ for I cortical populations in layer 5, and $160 \pm 49pA$ for E

and $412 \pm 167pA$ for I cortical populations in layer 6[62]. In total, the experimentally measured g_I was higher than g_E in all layers 2/3, 4, 5 and 6 (Fig. 3J).

Relations between connectivity constants J_{XY}

The weights of the network connectivity matrix J_{XY} were computed as a product of connection probability (CP), strength of postsynaptic potential (PSP), and the fraction of neurons in the source population with respect to the total number of neurons included in the network model. The data was extracted from experimental sources introduced in the first rows of Table 1 for the layers 2/3, 4, 5, and 6. We note that only two experimental reports contained complete information on both connectivity measures for all four V1 connections in layers 2/3 and 5[15,69], and only one source on connectivity measures in layers 4 and 6[69]. We assumed that our network contained 89% of excitatory and 11% of PV+ neurons, based on the following calculation: The V1 network consists of approximately 80% of excitatory and 20% of inhibitory neurons. A survey of inhibitory subpopulations in V1[64] reported that PV+ neurons constitute 37, 49, 53, and 42% of inhibitory neurons in layers 2/3, 4, 5, and 6, respectively. 10 out of 131 recorded neurons belonged to the layer 2/3, 29 to layer 4, 47 to layer 5, and 45 to layer 6. We computed the fraction of PV+ neurons in inhibitory population as the weighted percentage of the recorded neurons $(10 \cdot 0.37 + 29 \cdot 0.49 + 47 \cdot 0.53 + 45 \cdot 0.42)/131 \approx 0.47$. Thus, the percentage of PV+ neurons in our cortical network is $0.47 \cdot 20/(80 + 0.47 \cdot 20) \cdot 100\% \approx 11\%$.

Contrast invariance constrains connectivity and input profiles

We denote $\phi \equiv \psi - \theta$, $\phi' \equiv \psi - \theta'$ and show that the property of contrast invariance (Eq. 3)

$$R_X(\phi, C) = r_X(C)\widetilde{r}_X(\phi), \quad X \in \{E, I\} \tag{S9}$$

combined with steady state equations Eq. 2 leads to equations in Eq. 4, which relate orientation tuning curves $\widetilde{r}_X$ with connectivity and input profiles W_{XY} and L_X.

The steady-state equations of the extended SSN model in Eq. 2 are given by

$$R_X(\phi, C) = \Big(\int_{-\pi/2}^{\pi/2} W_{XE}(\phi - \phi') R_E(\phi', C) d\phi' - \int_{-\pi/2}^{\pi/2} W_{XI}(\phi - \phi') R_I(\phi', C) d\phi' + T_{\mathrm{dLGN}}(C) L_X(\phi) \Big)_+^n. \tag{S10}$$

We insert the contrast invariant representation of steady states (Eq. S9) into (Eq. S10) and divide Eq. S10 by $\widetilde{r}_X(\phi)$ to obtain

$$r_X(C) = \Big(J_{XE}(\phi) r_E(C) - J_{XI}(\phi) r_I(C) + T_{\mathrm{dLGN}}(C) g_X(\phi) \Big)_+^n, \tag{S11}$$

where $J_{XY}(\phi) = \int_{-\pi/2}^{\pi/2} W_{XY}(\phi - \phi') \widetilde{r}_Y(\phi') d\phi' / (\widetilde{r}_X(\phi))^{1/n}$ and $g_X(\phi) = L_X(\phi)/(\widetilde{r}_X(\phi))^{1/n}$.

Now we show that g_X are constants independent of ϕ, then we show that J_{XY} are constants provided the contrast response functions are not exactly linearly dependent, i.e. r_E and r_I do not satisfy $r_E(C) = a \cdot r_I(C)$ for all contrasts C with some constant a. We prove this statement by contradiction, i.e. we assume there is at least one parameter J_{XY} or g_X such that $J_{XY}(\phi) \neq J_{XY}(\phi')$ or $g_X(\phi) \neq g_X(\phi')$ for some $\phi \neq \phi'$, and derive a contradiction.

First, we show that g_X are independent of ϕ. We substitute $S = T_{\mathrm{dLGN}}(C)$ into Eq. S11

$$\widehat{r}_X(S) = \Big(J_{XE}(\phi) \widehat{r}_E(S) - J_{XI}(\phi) \widehat{r}_I(S) + S g_X(\phi) \Big)_+^n. \tag{S12}$$

Here, $\widehat{r}_X(S) = r_X(C) = r_X(T_{\mathrm{dLGN}}^{-1}(S))$. Since $\widehat{r}_X$ are non-negative, the content of the bracket on the right side of the equations is positive and we can remove the sign +. Next, we apply the exponent $1/n$ to both sides of Eq. S12 to obtain

$$(\widehat{r}_X(S))^{1/n} = J_{XE}(\phi) \widehat{r}_E(S) - J_{XI}(\phi) \widehat{r}_I(S) + S g_X(\phi). \tag{S13}$$

Now we denote $\widehat{J}_{XY}(\phi) = J_{XY}(\phi) - J_{XY}(\phi')$ and $\widehat{g}_X(\phi) = g_X(\phi) - g_X(\phi')$ and subtract from Eq. S13 the same equation with ϕ substituted by ϕ' to obtain

$$0 = \widehat{J}_{XE}(\phi) \widehat{r}_E(S) - \widehat{J}_{XI}(\phi) \widehat{r}_I(S) + S \widehat{g}_X(\phi). \tag{S14}$$

Next, we compute a derivative of Eq. S14 with respect to S and set $S = 0$. We obtain

$$0 = \widehat{J}_{XE}(\phi)\widehat{r}'_E(S) - \widehat{J}_{XI}(\phi)\widehat{r}'_I(S) + \widehat{g}_X(\phi). \tag{S15}$$

We note that $\widehat{r}'_X(0) = 0$ always holds for the zero steady state of Eq. S12 corresponding to $S = 0$ input. Then we obtain $\widehat{g}_X(\phi) = 0$ from Eq. S15, which by definition implies $g_X(\phi) = g_X(\phi')$ for all ϕ. We have shown that g_E and g_I are constants independent of ϕ.

Now we show when J_{XY} are independent of ϕ. From $\widehat{g}_E(\phi) = \widehat{g}_I(\phi) = 0$, we obtain that Eq. S14 is equivalent to

$$\widehat{r}_E(S)/\widehat{r}_I(S) = \widehat{J}_{EI}(\phi)/\widehat{J}_{EE}(\phi). \tag{S16}$$

Since the left side of Eq. S16 depends only on S and the right side only on ϕ, both sides are equal to the same constant which we denote by a. In particular, this last observation implies exact linear dependence of the contrast response functions: $r_E(C) = a \cdot r_I(C)$.

We have shown that the property of contrast-invariance Eq. 3 restricts the shape of the input functions L_E and L_I and the interaction profiles W_{EE}, W_{EI}, W_{IE} and W_{II} accordingly to the following relations

$$L_E(\phi) = g_E \cdot (\widetilde{r}_E(\phi))^{1/n}, \quad L_I(\phi) = g_I \cdot (\widetilde{r}_I(\phi))^{1/n}, \tag{S17}$$

$$\begin{aligned} &\int_{-\pi/2}^{\pi/2} W_{EE}(\phi-\phi')\widetilde{r}_E(\phi')d\phi' = J_{EE} \cdot (\widetilde{r}_E(\phi))^{1/n}, \quad \int_{-\pi/2}^{\pi/2} W_{EI}(\phi-\phi')\widetilde{r}_I(\phi')d\phi' = J_{EI} \cdot (\widetilde{r}_E(\phi))^{1/n}, \\ &\int_{-\pi/2}^{\pi/2} W_{IE}(\phi-\phi')\widetilde{r}_E(\phi')d\phi' = J_{IE} \cdot (\widetilde{r}_I(\phi))^{1/n}, \quad \int_{-\pi/2}^{\pi/2} W_{II}(\phi-\phi')\widetilde{r}_I(\phi')d\theta' = J_{II} \cdot (\widetilde{r}_I(\phi))^{1/n}, \end{aligned} \tag{S18}$$

where the constants g_X and J_{XY} depend on the shape of contrast responses r_E and r_I that are steady states of the two population SSN model.

Determining connectivity and input profiles of the extended SSN model

To determine the input and connectivity profiles L_X and W_{XY} from Eq. S17 and Eq. S18, we used the wrapped Gaussian approximation of orientation tuning curves. The wrapped Gaussian function is given by

$$G_{(\sigma)}(\phi) = \sum_{m=-\infty}^{\infty} \exp\left\{-\frac{(\phi+\pi m)^2}{2\sigma^2}\right\}.$$

The widths of the orientation tuning curves $\widetilde{r}_E$ and $\widetilde{r}_I$ were $\sigma_E \approx 0.31\pi \approx 56^\circ$ and $\sigma_I \approx 0.34\pi \approx 62^\circ$ (Fig. 4B), the orientation tuning curves are represented by

$$\widetilde{r}_E(\phi) = G_{(\sigma_E)}(\phi)/\max_{\phi} G_{(\sigma_E)}, \quad \widetilde{r}_I(\phi) = G_{(\sigma_I)}(\phi)/\max_{\phi} G_{(\sigma_I)}. \tag{S19}$$

To derive the input profiles L_E and L_I using Eq. S17, we fitted normalized wrapped Gaussian functions to the power-law transformations of orientation tuning curves $(\widetilde{r}_E)^{1/n}$ and $(\widetilde{r}_I)^{1/n}$, where n were the power-law exponents inferred from the recorded contrast responses. We obtained that the mean widths of the curves $(\widetilde{r}_E)^{1/n}$ and $(\widetilde{r}_I)^{1/n}$ were $\sigma_E^{\text{inp}} \approx 0.36\pi \approx 65^\circ$ and $\sigma_I^{\text{inp}} \approx 0.39\pi \approx 70^\circ$, respectively, and the input profiles L_E and L_I (Fig. 4C) were represented by

$$L_E(\phi) = g_E \cdot G_{(\sigma_E^{\text{inp}})}(\phi)/\max_{\phi} G_{(\sigma_E^{\text{inp}})}, \quad L_I(\phi) = g_I \cdot G_{(\sigma_I^{\text{inp}})}(\phi)/\max G_{(\sigma_I^{\text{inp}})}. \tag{S20}$$

Our next goal was to determine the connectivity profiles W_{XY} using Eq. S18. To this end, we used the formula for the convolution of two wrapped Gaussian functions [116]

$$\frac{\sigma_2}{\sqrt{2\pi}\sigma_1(\sigma_2^2-\sigma_1^2)^{1/2}} \int_{-\pi/2}^{\pi/2} G_{((\sigma_2^2-\sigma_1^2)^{1/2})}(\phi-\phi')G_{(\sigma_1)}(\phi')d\phi' = G_{(\sigma_2)}(\phi). \tag{S21}$$

Next, we combined Eq. S18 and Eq. S21 to obtain the wrapped Gaussian representation of W_{XY}. Using

$$(\widetilde{r}_X(\phi))^{1/n} = G_{(\sigma_X^{\text{inp}})}(\phi)/\max_{\phi} G_{(\sigma_X^{\text{inp}})}, \quad \widetilde{r}_Y(\phi) = G_{(\sigma_Y)}(\phi)/\max_{\phi} G_{(\sigma_Y)},$$

Eq. S18, and Eq. S21, we obtained

$$W_{XY}(\phi) = \frac{J_{XY} \cdot \sigma_X^{\text{inp}}}{\sqrt{2\pi} \cdot \sigma_Y \cdot ((\sigma_X^{\text{inp}})^2 - \sigma_Y^2)^{1/2}} \cdot \frac{\max_\phi G_{(\sigma_Y)}}{\max_\phi G_{(\sigma_X^{\text{inp}})}} \cdot G_{(((\sigma_X^{\text{inp}})^2 - \sigma_Y^2)^{1/2})}(\phi) \tag{S22}$$

and

$$\sigma_{XY} = ((\sigma_X^{\text{inp}})^2 - \sigma_Y^2)^{1/2}. \tag{S23}$$

For our recorded data, we obtained $\sigma_{EE} = 33°$, $\sigma_{EI} = 19°$, $\sigma_{IE} = 42°$, and $\sigma_{II} = 33°$.

Ascending order between the widths of connectivity profiles

The widths of connectivity profiles follow the order $\sigma_{EI} < \sigma_{EE} < \sigma_{II} < \sigma_{IE}$ in each inferred parameter set. However, this result would also follow for a specific order between only σ_E, σ_I and σ_E^{inp}. Here we show that the above derivations constrain the possible order of connectivity widths: independently of the exact values of σ_E and σ_I, the assumptions

$$\sigma_E < \sigma_I, \quad \sigma_I < \sigma_E^{\text{inp}}, \quad n > 1 \tag{S24}$$

always imply the relations

$$\sigma_{EI} < \sigma_{EE} < \sigma_{IE} \quad \sigma_{EI} < \sigma_{II} < \sigma_{IE}.$$

Indeed, since $\sigma_E < \sigma_I$, we always have $\sigma_E^{\text{inp}} < \sigma_I^{\text{inp}}$ for $n > 1$. Based on this inequality and Eq. S24, we obtain $\sigma_E^2 < \sigma_I^2 < (\sigma_E^{\text{inp}})^2 < (\sigma_I^{\text{inp}})^2$. Next, we obtain

$$(\sigma_E^{\text{inp}})^2 - \sigma_I^2 < (\sigma_E^{\text{inp}})^2 - \sigma_E^2 < (\sigma_I^{\text{inp}})^2 - \sigma_E^2$$

and

$$(\sigma_E^{\text{inp}})^2 - \sigma_I^2 < (\sigma_I^{\text{inp}})^2 - \sigma_I^2 < (\sigma_I^{\text{inp}})^2 - \sigma_E^2.$$

Together with Eq. S23, these chains of inequalities are equivalent to the relations $\sigma_{EI} < \sigma_{EE} < \sigma_{IE}$ and $\sigma_{EI} < \sigma_{II} < \sigma_{IE}$. In particular, inequalities in Eq. S24 always imply that σ_{EI} is the smallest and σ_{IE} the largest connectivity profile width, while σ_{EE} and σ_{II} are constrained between σ_{EI} and σ_{IE}.

Supplementary Information

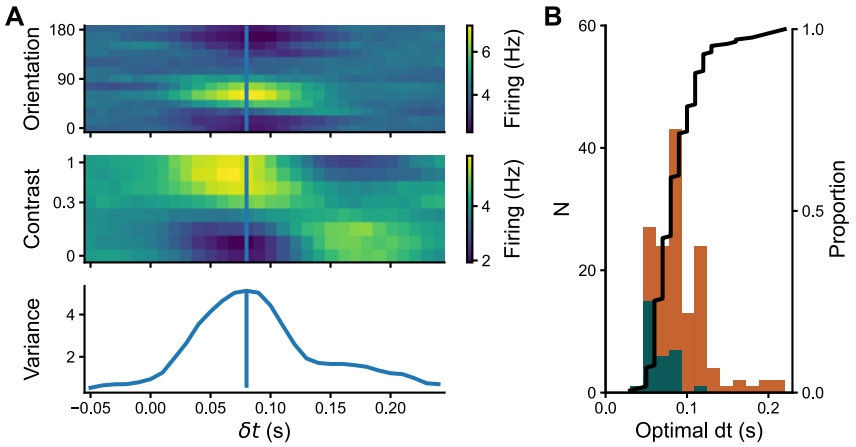

Fig. S1. Reverse correlation and optimal response time point (related to Fig. 2). **A**, Reverse correlation analysis for an example V1 neuron. Reverse correlation computes the firing rate at a time point δt relative to stimulus occurrence, yielding temporal kernels for each stimulus combination (*Middle*: orientation; *bottom*: contrast; average across all other stimulus dimensions for visualization). The optimal response time was calculated by using the sum of the aggregated variances in firing rate across stimulus conditions (*top*, see Methods) and selecting its peak as the latency of optimal response (*vertical line*)[45]. **B**, Distribution of optimal response times (*teal*: inhibitory, *orange*: excitatory, *black*: cumulative distribution). Inhibitory neurons had a significantly lower latency of optimal response time (69 ± 3 ms) than excitatory neurons (89 ± 3 ms; mean ± sem; two-tailed Welch's t-Test: t=5.18, $p < .001$).

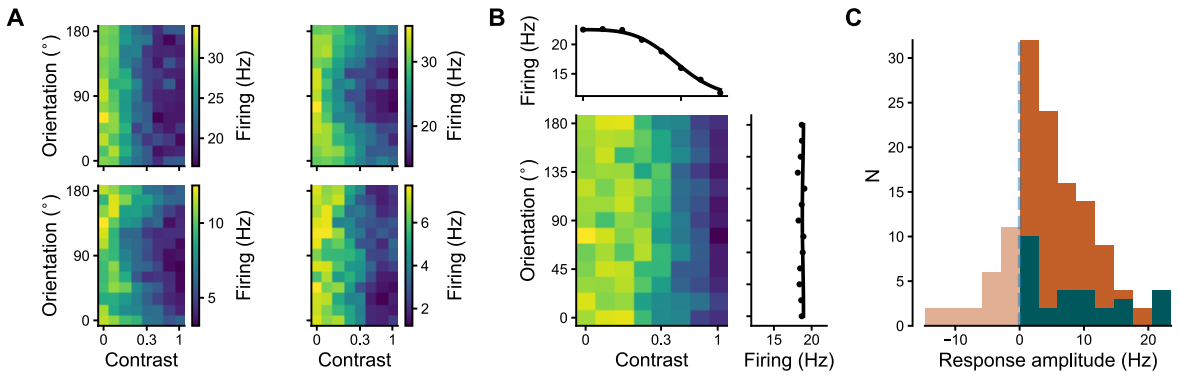

Fig. S2. Suppressed-by-contrast neurons (SbC) (related to Fig. 2). **A**, Four example SbC neurons, characterized by their stronger response to low contrast and an increasing suppression with higher contrasts. **B**, Two-dimensional tuning model fitted to an SbC neuron. **C**, Distribution of response amplitudes, defined as the difference between responses to 100% and 0% contrast. A considerable fraction of broad-spiking, putative excitatory neurons (24/125) were suppressed by contrast (negative response amplitude, *transparent*). Since SbC neurons might correspond to VIP interneurons[48], SbC neurons were excluded from further analysis.

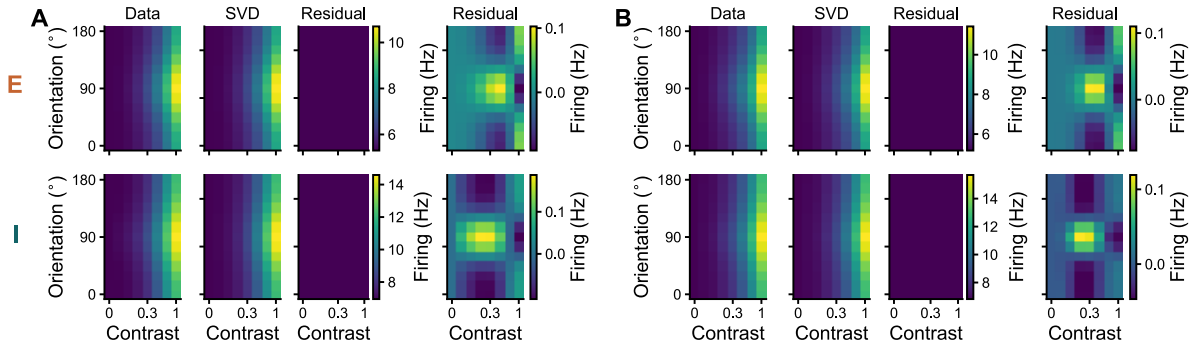

Fig. S3. Contrast-invariance of V1 population response (related to Fig. 2, Fig. 3). *Top:* excitatory, *bottom* inhibitory. **A**, Population response for two-dimensional tuning model using hyperbolic ratio function and wrapped Gaussian. Residuals are shown once on the same scale as the data and once on a separate scale. The residuals show a significant, but very weak pattern (E: $g_z = 10.41$, residual strength $< 0.1\%$; I: $g_z = 14.31$, residual strength $= 0.1\%$). **B**, Same as (A) for two-dimensional tuning model using model-free contrast response and wrapped Gaussian, as used in Fig. 3 (E: $g_z = 11.58$, residual strength $< 0.1\%$; I: $g_z = 15.07$, residual strength $< 0.1\%$).

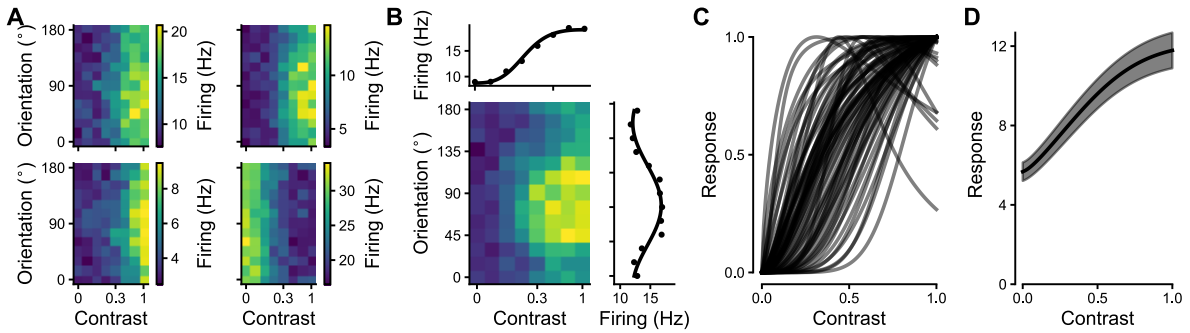

Fig. S4. Responses in dLGN (related to Fig. 2). **A**, Four example neurons from dLGN. As for the V1 data, SbC neurons were excluded from further analysis (17/89). **B**, Two-dimensional tuning model fitted to a dLGN neuron. **C**, Normalized contrast response functions. **D**, Population response curve obtained by averaged individual contrast response functions.

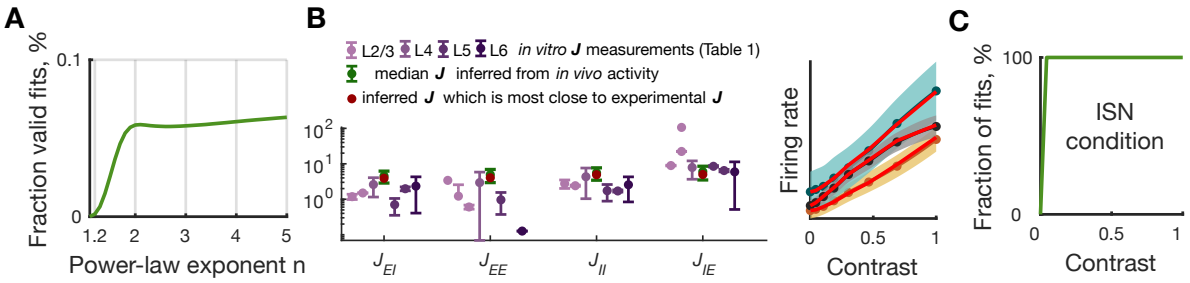

Fig. S5. Inference of connectivity parameters from contrast responses (related to Fig. 3). **A**, The fraction of valid initial connectivity weights is an increasing function of n with few valid fits found for n close to 1. **B**, *Left:* The presented inferred set of connectivity weights (red) is the closest to the weighted average of connectivity constants across layers (purple) extracted from experimental sources Table 1. *Right:* The inferred parameters set (red, left) leads to a close approximation of the average recorded contrast responses. **C**, Starting from the smallest measured contrast value of 4%, the SSN model represented an inhibition stabilized network for all inferred connectivity weights.

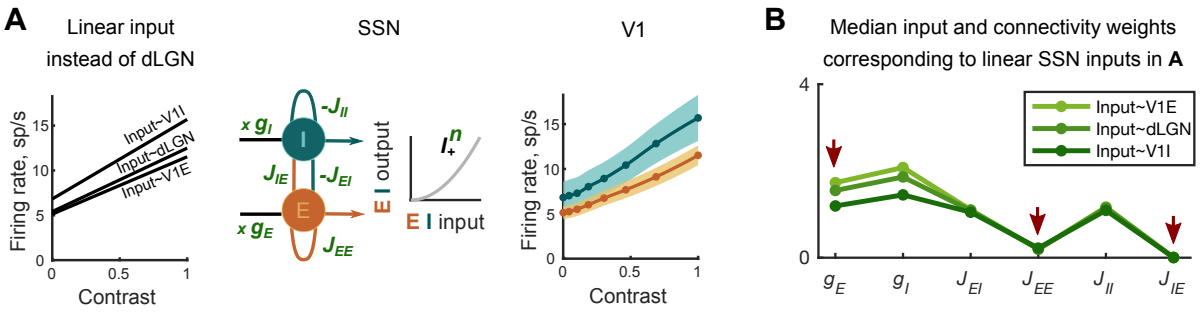

Fig. S6. Inference of connectivity parameters for linear inputs instead of recorded dLGN input (related to Fig. 3). **A**, Linear inputs (*left*), instead of recorded dLGN input introduced in Fig. 3A, Fig. S4, were used for the inference of input and connectivity weights g_X and J_{XY} of the SSN model (*middle*). The linear inputs (*left*) have the same values for the contrast $C = 0$ and $C = 100\%$ as the recorded contrast response functions of the E, dLGN, and I populations, respectively. The connectivity parameters were computed using the same method as in Fig. 3C to generate the thalamo-cortical mapping of linear inputs to the recorded V1 contrast responses (*right*). **B**, The inferred parameters g_X and J_{XY} followed a different order than the order found for the parameters determined from the recorded dLGN contrast response. Inconsistent with direct experimental measurements, J_{EE} was smaller than g_E and J_{IE} was the smallest connectivity weight for all three linear inputs (*red arrows*).

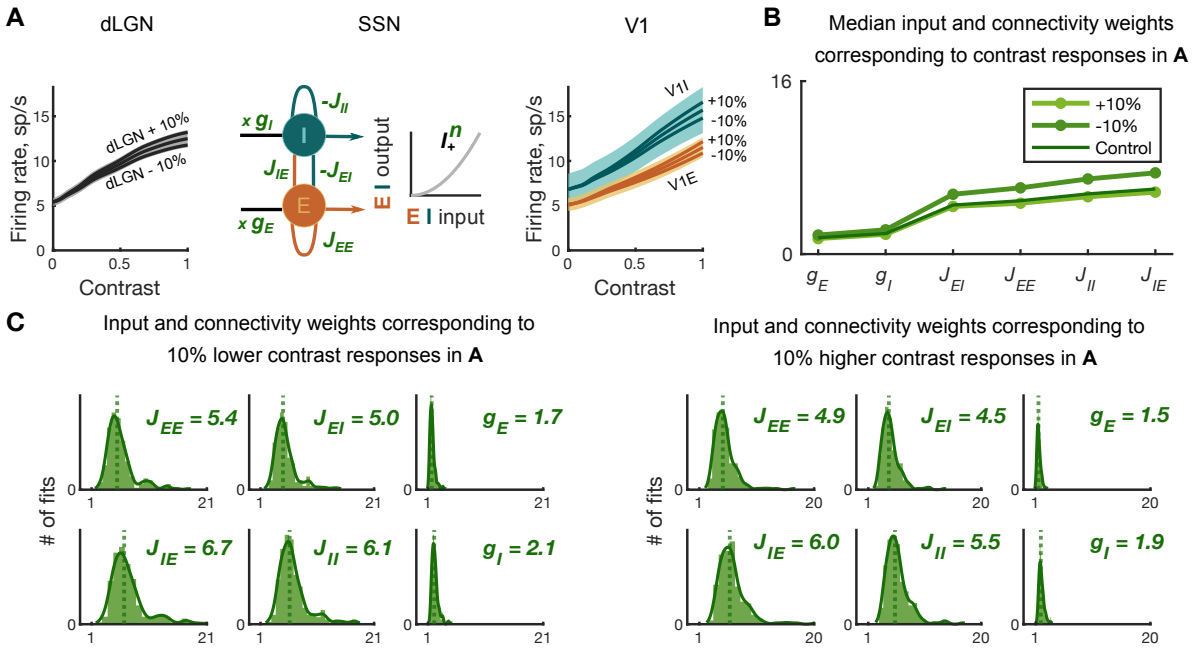

Fig. S7. Inference of connectivity parameters for modified contrast responses (related to Fig. 3). **A**, To compute connectivity and input weights for slight modifications of recorded responses, all three average contrast responses (dLGN, V1E, and V1I) were either increased or decreased by 10% accordingly to the relations (FR-FR(0))*1.1 and (FR-FR(0))*0.9, where FR denotes firing rate. **B**, In both cases most of the inferred weights followed the order found for average contrast responses. **C**, While medians of weights corresponding to decreased contrast responses remained almost the same, the medians of weights computed from increased responses were up to 13% higher.

3 Suppressed-by-Contrast Neurons in dLGN and V1

3.1 Summary

Neurons whose firing is suppressed by almost any visual stimulus exist along the early processing stages of the mammalian visual system: retina, dorsolateral geniculate nucleus (dLGN) and primary visual cortex (V1). However, the role of these suppressed-by-contrast (SbC) neurons in visual processing remains unclear. In mice, recent studies have investigated their connection to behavioral state and molecular markers, but a dedicated multifaceted survey linking past and present findings has been largely lacking. Here, we investigate intrinsic firing characteristics, functional aspects, and stimulus response properties of SbC neurons in the dLGN and V1 of mice. We find SbC neurons to fire less regularly and to be located in lower cortical layers than non-SbC neurons. Additionally, we find SbC neurons to be broad-spiking and to be tuned to the presented stimulus. Our results suggest that SbC responses are generated *de-novo* at multiple processing stages along the visual stream and highlight the diversity that underlies the SbC response type, which likely spans multiple known cell types.

3.2 Contributions

The following authors contributed to this manuscript. Laura Busse (LB) and Simon Renner (SR) conceived the study. SR recorded and curated data V1 data in one mouse and developed the data analysis software. Gregory Born, Yannik Bauer and Martin Spacek recorded and curated dLGN data and additional V1 data. Emma Müller-Seydlitz developed analysis techniques for spontaneous activity data related to Figure 2. SR expanded and finalized analysis techniques for variability and oscillations of spontaneous activity. SR performed data analysis for all figures, including identification of SbC neurons, wave shape classification, layer location analysis,

and stimulus tuning analysis. SR made the figures. SR and LB contributed conceptual ideas, discussed and coordinated the project at all stages, wrote and edited the manuscript.

Suppressed-by-contrast neurons in mouse dLGN and V1 have distinct circuit and tuning properties

Simon Renner[1,2], **Emma Müller-Seydlitz**[1], **and Laura Busse**[1,3,*]

[1]Division of Neurobiology, Department Biology II, LMU Munich, 82152 Munich, Germany
[2]Graduate School of Systemic Neurosciences (GSN), LMU Munich, 82152 Munich, Germany
[3]Bernstein Center for Computational Neuroscience, 82152 Munich, Germany
[*]busse@bio.lmu.de

ABSTRACT

Neurons whose firing is suppressed by almost any visual stimulus exist all along the early processing stages of the mammalian visual system. However, the role of these suppressed-by-contrast (SbC) neurons in visual processing remains unclear. In mice, recent studies have investigated their connection to behavioral state and molecular markers, but a dedicated multifaceted survey linking past and present findings has been largely lacking. Here, we investigate intrinsic spike characteristics, functional aspects and stimulus response properties of SbC neurons in the dorsolateral geniculate nucleus (dLGN) and primary visual cortex (V1) of mice. We find SbC neurons to differ from non-SbC neurons in firing regularity, cortical layer location, to be broad-spiking and to be tuned to the presented stimulus. Our results suggest that SbC responses are generated at multiple processing stations along the visual stream and highlight the diversity that underlies the SbC response type, which likely spans multiple known cell types.

Introduction

Classically, visual neuroscience has studied action potentials evoked by visual stimuli. By probing neurons with stimuli ranging from small spots of light to gratings and naturalistic images, studies have uncovered functional response types from simple center-surround, Gabor-like [1] to, more recently, complex activation patterns [2, 3]. However, while action potentials elicited by visual stimuli have been the subject of hundreds of studies and therefore have dominated the field, neurons which respond with suppression to visual stimuli have received far less attention [4]. Despite this lack of attention, a small but steady flow of studies has discovered examples of neurons that are suppressed by visual stimuli along all processing steps of the early visual system [5, 6, 7]. Aided by large-scale recording techniques and genetic tools, recent investigation have probed the role of suppressed neurons in different behavioral states as well as their relation to specific inhibitory neurons [8, 9]. In most aspects, however, the knowledge of suppressed neurons' computational and circuit role still remains scarce.

Neurons with high spontaneous activity, which is suppressed by visual stimuli, were first discovered in the cat and rabbit retina during the late 1906s [5, 10]. Because these retinal ganglion cells were suppressed by virtually any stimulus that changed a uniform blank screen, they were hypothesized to function as "uniformity detectors" [5, 10, 11]. With vision research shifting to different species, similar neurons were found in the dorsolateral geniculate nucleus (dLGN) of the thalamus in primates and mice [12, 6] as well as the primary visual cortex (V1) in mice [7, 13]. This indicates that "suppressed-by-contrast" (SbC) neurons might be an integral part of visual coding, even in higher processing centers and across different species.

Investigations of SbC neurons have been distributed across time, different labs and research contexts, resulting in spotty and diverse evidence of their role in visual processing. Despite being suppressed by almost any stimulus, retinal SbC neurons already show diversity in their processing of stimulus onset and offset as well as their temporal dynamics and spike behavior [10, 5, 14, 15, 2]. This diversity continues in dLGN, where SbC neurons have been observed to have different firing rates across species [6, 12], fire more regularly than the rest of the population [12], and exist as different subtypes, potentially within a projection pathway from the retina [16]. In visual cortex, SbC neurons have been implied to correspond to different populations of inhibitory neurons [7, 9], to be involved in firing rate modulations during transitions between a quiescent and more active brain state [13, 7], as well as to be part of a corticothalamic feedback channel [8]. In light of this scattered evidence, there is no consensus on the functional role of SbC neurons. They could serve to mainly signal uniformity, perhaps in the context of saccadic eye movements [15] and aid prey capture [17], they could regulate overall inhibitory drive to balance activity in different behavioral states [7], perhaps in a feedback pathway [8], or they could provide disinhibition to selectively enhance activity [9].

As knowledge about SbC neurons in dLGN and V1 is still limited and most studies have been carried out with different

angles of attack, we here provide a wide screening of SbC and non-SbC neurons in mouse dLGN and V1, gathering evidence on their firing patterns, cell type (excitatory vs. inhibitory), layer location and tuning properties. We find SbC neurons to be mainly excitatory, be located in lower cortical layers, show tuning to multiple stimulus parameters, and their firing patterns to vary between brain areas. As such, we offer a survey investigation that connects previous studies and provides details about SbC neurons' circuit and computational role.

Results

To identify SbC neurons in dLGN and V1, we performed silicon probe recordings in awake, head-fixed mice (**Figure 1A**) while presenting full-field flashed gratings of different contrasts, orientations, spatial phase and spatial frequency (**Figure 1B**). Using reverse correlation analysis [18, 19], we calculated temporal response profiles for each unique stimulus and extracted each neuron's optimal response latency as the time of maximal variance across stimuli [20] (**Figure 1B**). We then computed their modulation strength by normalizing their responses to their firing rate during 0% contrast stimuli. We classified neurons as SbC, if their mean modulation strength was below -0.1 i.e. if their average firing response was at least 10% lower than the response to 0% contrast stimuli (**Figure 1C, inset**). Similarly, we only classified neurons as non-SbC, if their mean modulation strength was larger than 0.1. Requiring modulation stronger than 10% effectively filtered out noisy neurons whose responses fluctuated around their baseline responses (**Figure 1D**). Using this classification, we found 16% SbC neurons in dLGN (17/108) and 13% (31/239) in V1 (χ^2-test: $\chi^2 = 0.22$, p=.640).

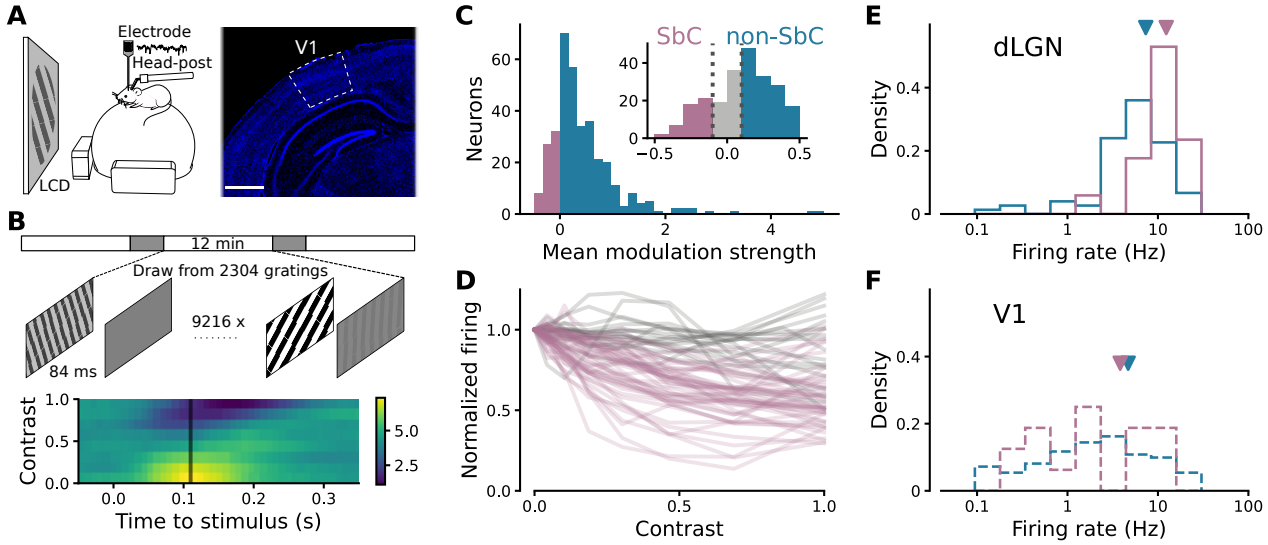

Figure 1. Identifying SbC neurons. (**A**) *Left*: Experimental recording setup for awake, head-fixed mice. *Right*: Histological identification of V1. Scale bar: 1 mm. Bregma: -2.70 mm. (**B**) *Top*: Flashed gratings stimulus. During gray periods, a median luminance gray screen was shown. *Bottom*: Example contrast response profile retrieved via reverse correlation. (**C**) Distribution of mean contrast modulation strength (mean of responses normalized to firing at 0%-contrast stimuli) across neurons used to classify neurons as SbC (*pink*) and non-SbC (*turquoise*). *Inset*: neurons with small modulation strength were not associated with a response type (*gray*). (**D**) Contrast response functions of SbC neurons (*pink*) and neurons with mean modulation strength between -0.1 and 0 from panel C (*gray*). (**E**) Distribution of spontaneous firing rate in dLGN (*solid lines*) for SbC (*pink*) and non-SbC (*turquoise*) neurons. Arrows indicate means. (**F**) Same as E for V1 (*dashed lines*).

We first explored whether firing rates during spontaneous activity, a simpler criterion than suppression by sensory stimuli, could equally well distinguish SbC from non-SbC neuron types. Indeed, a previous study in the retina [11] reported that the high baseline firing rates of SbC neurons were distinct enough to constitute a criterion sufficient for identification, and Piscopo et al. [6] observed a 4-fold difference in spontaneous firing rates between SbC and non-SbC neurons in dLGN that might have served as a predictor for identifying SbC neurons. We therefore focused on periods of spontaneous activity in response to a median luminance gray screen both during blank periods in our stimulus (**Figure 1B, top**) and during separate measurements where available (see Methods). Comparing distributions of baseline firing rates in SbC and non-SbC neurons, we found the distributions to be largely overlapping in both brain regions (**Figure 1E, F**). Consistent with previous studies [12, 7], we found that mean firing rates were overall lower in V1 compared to dLGN (5.4 ± 7.5 vs. 8.7 ± 6.4, two-way ANOVA, main effect of

region: $F = 21.17, p < 0.001$). More importantly, SbC neurons overall tended to have higher firing rates than non-SbC neurons (8.2 ± 7.3 vs. 5.8 ± 6.6, main effect of response type: $F = 2.92, p = 0.088$). This trend arose because firing rate differences between SbC and non-SbC neurons were larger in dLGN (dLGN SbC: 12.3 ± 7.0 vs non-SbC 7.3 ± 5.3) than in V1 (V1 SbC: 3.8 ± 4.5 vs non-SbC 4.7 ± 7.2, interaction effect: $F = 5.84, p = 0.016$, **Figure 1E, F**).

To better understand to which degree these differences in mean spontaneous firing rates might support differentiation between SbC and non-SbC neurons, we next applied an ideal observer analysis. From the distributions of spontaneous firing rates for SbC and non-SbC neurons, we determined the area under the receiver operating characteristics (ROC) curve, and found, in line with the substantial overlap between the firing rate distributions (**Figure 1F**), classification performance for V1 to be poor ($AUC = 0.51$, Mann-Whitney-U test: $U = 872.0, p = 0.455$, **Figure S1**). In dLGN, classification accuracy was different from chance (Mann-Whitney-U test: $U = 327.0, p = .001$; **Figure S1**), yet the overall low performance ($AUC = 0.74$) did not not permit reliably identifying SbC neurons based on firing rate alone.

Firing properties of SbC and non-SbC neurons differ between brain areas

We next turned to firing regularity, bursting, and rhythmicity in activity of SbC and non-SbC neurons, to gain insights into intrinsic properties of these neuron classes and how they might be embedded in their local network [21]. These investigations were further motivated by previous studies, which have reported considerable differences in spike train regularity and burst firing between SbC and non-SbC neurons [14, 12]. To measure firing regularity, we focused again on spontaneous activity and computed the coefficient of variation (CV), i.e. the standard deviation (SD) of the spike interval distribution divided by its mean, with $CV > 1$ representing a spike train that is less regular than a Poisson process with the same firing rate. Comparing CVs across brain regions and response types (**Figure 2A**), we found that the CV depended on both response type and region (two-way ANOVA interaction, $F = 4.81, p = .029$): while both response types had similar CVs in dLGN (1.48 ± 0.33 vs. 1.49 ± 0.28), post-hoc Welch's t-test: $t = 0.09, p = 0.931$; **Figure 2A, top**), SbC neurons tended to have higher CVs than non-SbC neurons in V1 (1.72 ± 0.77 vs. 1.34 ± 0.52, post-hoc Welch's t-test: $t = 1.91, p = 0.073$; **Figure 2A, bottom**). Thus, contrary to a previous report in dLGN of anesthetized macaques, which found that SbC neurons had lower CVs (i.e. more regular spike trains, more rhythmic activity) [12], SbC neurons in awake mouse seem to have similarly regular spike trains as non-SbC neurons in dLGN, and more irregular spike trains than non-SbC neurons in V1.

Since higher CV's can be driven by burst firing, we next computed, for each neuron, the ratio of burst spikes in the spike train. Relying on definitions for bursts driven by thalamic low-threshold calcium spikes [22], we identified bursts as groups of spikes that were preceded by at least 100 ms of quiescence and had inter-spike intervals below 4 ms. Comparing again across brain regions and response types, we found – consistent with the well-known propensity of thalamus to fire in bursts [23] – higher mean burst ratios in dLGN than in V1 (0.088 ± 0.088 vs. 0.012 ± 0.036, two-way ANOVA, main effect of region: $F = 28.57, p < .001$). In addition, we found a lower mean burst ratio for SbC neurons than for non-SbC neurons (0.028 ± 0.050 vs. 0.047 ± 0.078, main effect of response type: $F = 5.38, p = .021$) **Figure 2B**). This result indicates that burst firing in SbC neurons is unlikely to be the underlying reason for higher CVs in SbC neurons, particularly in V1.

To corroborate this finding and investigate more directly to which degree burst firing determined the differences in firing regularity between SbC and non-SbC neurons, we re-computed the CV while considering only tonic spikes (**Figure S2**) and compared the CV across response types, brain regions and firing modes (full vs tonic-only). We found the CV to depend on all three factors (three-way mixed ANOVA, region vs response type vs firing mode interaction effect, $F = 4.13, p = 0.043$). Post-hoc tests (Bonferroni-adjusted: $\alpha = 0.00625$) revealed that the removal of burst spikes did not affect either response types in V1 (CV for all vs tonic-only spikes: non-SbC: $1.34 \pm 0.0.52$ vs 1.33 ± 0.52, Welch's t-test: $t = 0.23, p = 0.819$; SbC: 1.72 ± 0.77 vs. 1.72 ± 0.76, Welch's t-test: $t = 0.03, p = 0.98$) nor SbC units in dLGN (1.48 ± 0.33 vs. 1.41 ± 0.34, Welch's t-test: $t = 0.59, p = 0.557$). Instead, removing burst spikes exclusively lowered the CV of non-SbC neurons in dLGN (1.49 ± 0.28 vs. 1.38 ± 0.30, Welch's t-test: $t = 3.22, p = 0.00156$). We thus conclude that burst firing cannot explain differences in firing regularity between SbC and non-SbC neurons in V1, but might rather mask additional differences in firing regularity in dLGN.

Having observed characteristic differences between SbC and non-SbC neurons in the relative variability of inter-spike intervals, we next turned to spike autocorrelograms (**Figure 2C**). While the CV only considers neighboring spikes, autocorrelograms take into account more distant spikes in the spike train, whose dynamics can reflect some of the intrinsic properties of the neurons, ranging from firing modes to morphology [24]. Since the parameters used to compute an autocorrelogram influence which frequencies can be detected, we used an oscillation score (OS) [25] to extract the strongest autocorrelogram oscillation frequency (PF) and its relative strength of oscillation within specific frequency bands (θ: 4–8 Hz, α: 8–12 Hz, β_-: 12–20 Hz, β^+: 20-30 Hz, γ_-: 30-50 Hz, γ^+: 50-80 Hz). To optimally capture oscillations in these bands, the method computes a separate autocorrelogram with adjusted maximal time lag and binning for each band before extracting OS and PF from the Fourier spectrum of each autocorrelogram. For instance, oscillations in the θ-band can be mathematically optimally extracted using time lags up to ≈ 1000 ms, while oscillations in the γ_+-band, can be optimally extracted using time lags up to ≈ 250 ms

(**Figure 2C, top left**). To compare oscillation properties between brain regions and response types across the entire frequency spectrum, we smoothed the obtained $PF - OS$ diagrams via kernel convolution and extracted frequency ranges that differed between SbC and non-SbC neurons or areas by computing their impact on the correlation between the spectra (see Methods).

We found the smoothed spectra of SbC and non-SbC neurons in V1 to be very similar (Pearson correlation, $\rho = 0.986$, $p < 0.001$), lacking continuous frequency ranges impacting the correlation and prominent peaks (data not shown). In dLGN (**Figure 2C, top right**), the spectra of SbC and non-SbC neurons were also similar (Pearson correlation, $\rho = 0.966$, $p < 0.001$), both having prominent peaks at 12 and 59 Hz. The latter peak is consistent with a conspicuous narrow gamma-band oscillation with a subcortical origin reported in previous studies [26, 27]. Interestingly, dLGN SbC neurons had lower *OS* than dLGN non-SbC neurons in lower frequency ranges (4-18 Hz, impact 0.021). We also compared response types between dLGN and V1 (**Figure 2C, bottom**). For non-SbC neurons, we found they had higher *OS* in dLGN compared to V1 in the γ^+-band (SbC: 56-62 Hz, impact 0.034; non-SbC: 56-61 Hz, impact 0.011) as well as across the α and β_--band (SbC: 10-14 Hz, impact 0.125; non-SbC: 9-15 Hz, impact 0.122) (**Figure 2C, bottom left**). For SbC neurons, we found that they had lower *OS* in dLGN than V1 in the θ-band (4-10 Hz, impact 0.233) and higher *OS* across the α and β_--band (11-14 Hz, impact 0.017) (**Figure 2C, bottom right**). Performing the same oscillation analyses without burst spikes yielded qualitatively similar results. In summary, we found SbC neurons in dLGN to have a unique *OS* spectrum, with low *OS* in lower frequency ranges. Furthermore, our analyses show that dLGN SbC and non-SbC neurons equally participate in the prominent narrow-band gamma oscillation known to exist in mouse dLGN.

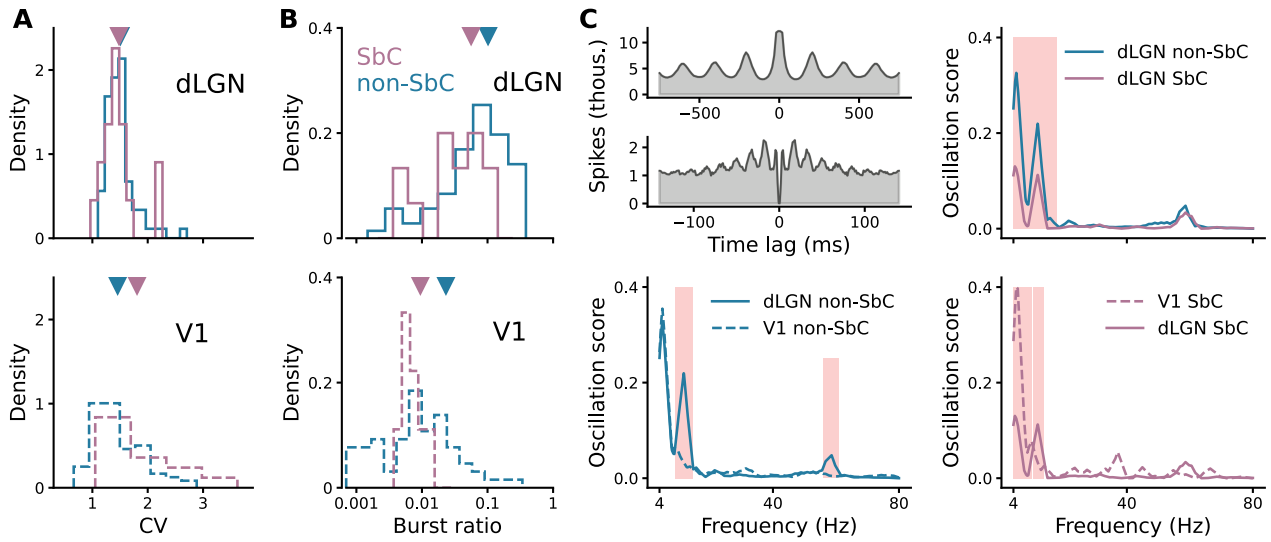

Figure 2. Firing regularity, burst firing, and autocorrelogram oscillations for dLGN and V1 during spontaneous activity. (**A**) Coefficient of variation (*CV*) of SbC and non-SbC neurons. (**B**) Ratio between burst spikes and total spikes (burst ratio) of non-SbC and SbC neurons. (**C**) Differences in autocorrelogram oscillations between regions and response types. *Top left*: Example autocorrelogram of one non-SbC neuron in V1 (*top*) with high oscillation score ($OS = 55.5$, $PF = 4.9$ Hz) in the θ-band (4–8 Hz) and one SbC neuron in dLGN (*bottom*) with high oscillation score ($OS = 5.2$, $PF = 58.6$ Hz) in the γ_+-band (50–80 Hz). Binning adjusted for optimal visibility of oscillation. *Bottom left, top right, bottom right*: Normalized, smoothed oscillation score values across the tested frequency spectrum for different neuron populations. Shaded areas indicate ranges with significant differences between compared populations. (A-C) Arrows indicate means. *Pink*: SbC; *turquoise*: non-SbC; *solid lines*: dLGN; *dashed lines*: V1.

SbC neurons are mainly broad-spiking and have longer response latencies

In addition to differences in their spike statistics, SbC neurons might differ from non-SbC neurons with regards to their specific neuronal type (e.g., excitatory vs inhibitory), or holding a specific place in the circuitry with regards to layers and the associated projection patterns. For example, SbC neurons in mouse dLGN and V1 could be part of an excitatory projection pathway signaling uniformity from the retina to the cortex [14, 16], provide complementary inhibition to balance overall firing rates [8] and counteract effects of arousal [7], or influence the gain of pyramidal cells via disinhibition [9].

To determine whether SbC neurons belonged to the group of broad-spiking, putative excitatory neurons, or narrow-spiking, putative inhibitory neurons, we first concentrated on primary visual cortex, for which the extracellular waveshape is known

to serve as a proxy for cell type [28]. From the extracellular waveshape of 540 neurons recorded in V1 (**Figure 3A, top**), we extracted characteristic parameters (peak width, trough-to-peak time) and identified two well-separated clusters (Mean silhouette coefficient=0.64, **Figure 3B, top**). Remarkably, when we then labeled neurons recorded specifically for this study according to their responses to contrast, we found a strong trend towards fewer SbC neurons in the narrow-spiking cluster (3%, 1/34) than in the broad-spiking cluster (20%, 26/131, exact Fisher-test: $oddsratio = 0.15, p = 0.053$, **Figure 3B, top**). This suggests that the V1 SbC neurons in their majority do not correspond to the class of fast-spiking, inhibitory basket cells. Applying clustering to the dLGN wave shapes (**Figure 3A, bottom**) yielded, as expected, a single cluster, in which the SbC and non-SbC response types overlapped (**Figure 3B, bottom**). This indicates that the same waveshape parameters as used for V1 do not predict SbC vs. non-SbC neuron type in dLGN. In both V1 and dLGN, we also found a small group of neurons with positive-polarity wave shapes (V1: 21/519, dLGN: 42/512, **Figure 3A**), which we did not include in the clustering. Positive-polarity wave shapes in V1 have recently been associated with axonal recordings of thalamic inputs [29], which makes them good candidates to investigate the propagation of SbC responses between dLGN and V1. We did indeed find SbC neurons among the positive-polarity V1 waveshapes recorded for this study (40%, 4/10), suggesting that SbC feature selectivity could be transmitted to cortex. We did not find any SbC neurons among the small number of dLGN positive-polarity neurons recorded for this study (0/5).

To investigate the distribution of SbC neurons across V1 layers, we performed current-source density (CSD) analysis [30, 31], and found that SbC neurons were more preferentially located in deeper layers. When we compared our recorded neurons' relative depths across response types, we observed SbC neurons and non-SbC neurons across all V1 layers (**Figure 3C, top**), with little difference in average relative depth (SbC: 116 ± 233 µm, non-SbC: 191 ± 211 µm, Student's t-test: $t = 1.44, p = 0.151$). After assigning the neurons to their cortical layers, however, we found a trend for SbC neurons to be more strongly biased towards infragranular layers (83%, 20/24 neurons) compared to non-SbC neurons (63%, 76/120 neurons, exact Fisher-test: $oddsratio = 2.89, p = 0.062$, **Figure 3C, bottom**).

As layer IV of primary visual cortex receives the majority of thalamocortical axons [32], finding a higher proportion of SbC neurons in lower layers might indicate that at least some part of contrast suppression is generated *de-novo* or enhanced in V1, in addition to potentially being inherited from dLGN. We further investigated this hypothesis by comparing optimal response latencies (i.e., time of peak variance across stimuli, **Figure 1B**) between SbC and non-SbC neurons as well as brain regions. As expected, latencies were overall higher in V1 than in dLGN (87 ± 28 vs. 70 ± 24 ms, two-way ANOVA, main effect of region: $F = 24.7, p < .001$). At the same time, SbC neurons had a significantly larger mean latency (103 ± 36 ms) than non-SbC neurons (77 ± 24 ms, main effect of response type: $F = 33.0, p < .001$, **Figure 3C**). We found no significant interaction effect ($F = 1.84, p = 0.176$), meaning the latency of SbC neurons in V1 was not additionally altered compared to dLGN. Since cortical depth and optimal response latency were significantly, but only weakly correlated ($\rho = 0.17, p = 0.047$), and explanatory power of cortical depth for optimal response latency was low ($R^2 = 0.027$), our data cannot discern if the higher latencies of SbC neurons in V1 could come from a slower direct projection pathway or because suppression is generated *de-novo* at each processing station.

SbC neurons have stimulus selectivity

Classically, SbC cells have been reported to respond similarly to almost any kind of stimulus [5, 12] and not show specific selectivity to stimulus properties like spatial phase, orientation or direction [14, 6, 7]. Recently however, there have been reports of direction and phase tuning in SbC neurons [9]. Since these sensory response properties can reflect underlying connectivity and computations (e.g., simple vs. complex cells in the case of phase tuning), we more closely investigated the tuning of SbC and non-SbC neurons.

We first investigated the contrast response functions of SbC and non-SbC neurons in more detail, and found that SbC neurons had higher contrast-sensitivity than non-SbC neurons (**Figure 4A, B**). We investigated the shape of the contrast response functions by computing the contrast at half height (CHH) (**Figure 4A, B; right**). In addition to dLGN being overall more contrast-sensitive than V1 (0.38 ± 0.18 vs 0.48 ± 0.20, main effect of region: $F = 7.3, p = 0.007$), SbC neurons across both V1 and dLGN showed earlier saturation (0.30 ± 0.14) than non-SbC neurons (0.47 ± 0.19; two-way ANOVA, main effect of response type: $F = 28.9, p < .001$). Strikingly, while almost half of the non-SbC neurons were late-saturating ($CHH > 0.5$; 102/244), the SbC population only contained 12.5% of late-saturating cells with (6/48 neurons; χ^2-test: $\chi^2 = 7.09, p = 0.008$). So, while SbC neurons seem to reach their peak response later in time (see above), they seem to respond at lower stimulus intensities.

Spatial frequency tuning of neurons in V1 and dLGN typically falls into three broad categories: low-pass, high-pass, and band-pass [33]. SbC neurons, on the other hand, are expected to have a U-shaped (convex) band-pass profile [6], being suppressed by medium spatial frequencies and returning to their baseline firing once frequencies are too high to be resolved. We used unbiased clustering to confirm this finding by recovering a convex response type consisting exclusively of SbC neurons. More specifically, we first used k-means clustering on the individual normalized spatial frequency responses and embedded

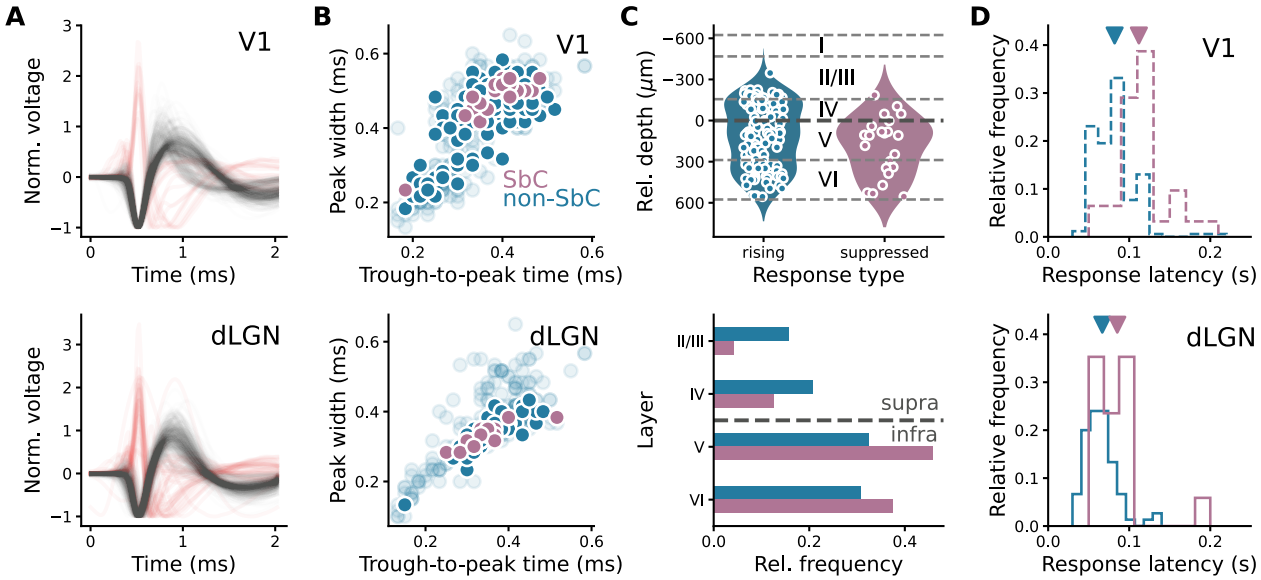

Figure 3. Comparison of functional type, layer location and response latencies in dLGN and V1. (**A**) Extracellular wave shapes of neurons in V1 and dLGN. *Red:* non-classical, positive-polarity wave shapes. (**B**) Classification based on wave shape parameters. *Light dots*: neurons recorded in other experiments to improve power for clustering, *solid dots*: neurons recorded for the current study. (**C**) *Top*: cortical depth relative to the bottom of layer IV, separately for rising and contrast-suppressed neurons. *Bottom*: layer affiliation based on layer border estimation. Thick line indicates split into supra- and infragranular layers used for statistical analysis. (**D**) Response latency distributions. Arrows indicate means. *solid lines*: dLGN; *dashed lines*: V1. (B-D) *Pink*: SbC; *turquoise*: non-SbC.

the result in a 2D space for visualization (**Figure 4C, D, top left**). We found that, for both brain regions, four clusters best represented the data (maximum Variance Ratio Criterion, dLGN: 78.4, V1: 94.5). In dLGN, these four cluster corresponded closely to the previously reported low-pass, band-pass and high-pass types and an additional convex-shaped type (**Figure 4C, bottom left**), which exclusively contained SbC neurons (15 neurons, **Figure 4C, top right**). Only two SbC neurons were assigned to a different cluster. In V1, we also found a convex cluster consisting exclusively of SbC neurons (30 neurons, **Figure 4D, top right**) with only one SbC neuron being assigned to a different cluster. This close correspondence between unsupervised clustering and hypothesis-driven classification of SbC neurons across two different response modalities (contrast, spatial frequency) accentuates SbC neurons as an independent response type.

Next, we computed the mean spatial frequency tuning function for SbC and non-SbC neurons (**Figure 4C, D, bottom right**). While non-SbC neurons had a concave tuning function (negative slope of a parabola fitted to mean tuning function, dLGN: -1.29, V1: -1.48), SbC neurons had a convex tuning function (positive parabola slope, dLGN: 1.78, V1: 1.39). Interestingly, the preferred spatial frequency of the two populations (peak of fitted parabola) were similar for SbC and non-SbC neurons, especially in dLGN (dLGN: 0.087 vs 0.083 cyc/°, V1: 0.111 vs 0.072 cyc/°). This mirrored tuning could be an indication of a homeostatic coupling between SbC and non-SbC firing that serves to keep overall firing in balance [8].

Following reports that SbC cells are insensitive to the spatial phase of a presented grating [12], we investigated possible differences in phase tuning between SbC and non-SbC cells. To this end, we determined the preferred grating orientation for each neuron and compared responses to the four spatial phases at this preferred orientation. To quantify selectivity for spatial phase, we took advantage of spatial phase being a circular property and used a vector averaging method, which yields values between 0 (no tuning) and 1 (maximum tuning) [34]. While the majority of neurons for both response types showed little sensitivity for spatial phase (phase selectivity index (PSI) $<< 0.1$) (**Figure 4E,F, right**), the top 30% of neurons in each cell type did show tuning to spatial phase (dLGN: SbC: 0.11 ± 0.03, non-SbC: 0.21 ± 0.09, V1: SbC: 0.17 ± 0.06, non-SbC: 0.20 ± 0.09, **Figure 4E,F, left**). Mean tuning strength for spatial phase in this subpopulation was similar between brain regions (two-way ANOVA, main effect of region: $F = 0.88, p = 0.351$), but SbC neurons showed weaker selectivity for phase than non-SbC neurons (two-way ANOVA, main effect of response type: $F = 5.40, p = 0.025$) (**Figure 4E,F**). We thus conclude that, contrary to previous reports, SbC neurons can be sensitive to spatial phase, albeit less so than non-SbC cells.

In addition to reporting a lack of spatial phase tuning, previous studies in monkey dLGN [12] and mouse V1 [7, 8] have

found SbC cells to show no or only weak selectivity for grating orientation. Analogous to our analyses of spatial phase tuning, we extracted the optimal spatial phase for each neuron and quantified orientation tuning at this spatial phase using an orientation selectivity index (*OSI*) (**Figure 4G, H, right**). Focusing again on the 30% most selective neurons, we found orientation selectivity to depend on brain region and response type (**Figure 4G, H, left**). As expected, mean orientation tuning in this subpopulation was weaker in dLGN (0.12 ± 0.04) compared to V1 (0.22 ± 0.11, two-way ANOVA, main effect of region: $F = 12.11, p < 0.001$). Additionally, SbC neurons were less orientation-tuned than non-SbC neurons (0.11 ± 0.06 vs. 0.20 ± 0.11, main effect of response type: $F = 9.53, p = 0.003$). In dLGN, SbC neurons were, on average, at most suppressed to $74 \pm 5\%$ of their baseline firing, which represents a negligible amount of tuning. In V1 on the other hand, tuning was stronger, with SbC neurons being, on average, maximally suppressed to $53 \pm 22\%$ of baseline firing. Thus, while non-SbC neurons respond stronger overall to grating orientation, we discovered previously unreported orientation tuning in SbC neurons.

Discussion

In this work, we provide a description of basic functional properties and feature-selectivity of SbC neurons in mouse dLGN and V1. We found SbC neurons in V1 to fire less regularly, SbC neurons in dLGN to have weaker oscillations, and SbC neurons in both brain areas to have fewer bursts than non-SbC neurons. A waveshape analysis for SbC neurons in V1 revealed that SbC neurons were mainly broad-spiking, indicating that they likely do not correspond to fast-spiking basket cells, but rather to other types of inhibitory interneurons or excitatory neurons. Furthermore, V1 SbC neurons were located preferentially in lower cortical layers and had longer peak response latencies than non-SbC neurons. Lastly, we found sensitivity to spatial phase and orientation of gratings that could allow SbC neurons a computational role beyond uniformity detection. Together, our findings highlight that the sizeable fraction of SbC neurons is positioned to play a multifaceted role in visual processing along the thalamo-cortical pathway.

Basic response properties

Our results regarding firing regularity of SbC neurons during spontaneous activity is different from previous results in dLGN of anesthetized macaques, which had shown that SbC neurons fire more regularly than the rest of the population [12]. In light of this previous result, we were surprised to find no differences between SbC and non-SbC neurons in spike train regularity in dLGN, and even a lower regularity for SbC neuron in V1. These findings regarding spike train regularity could not be accounted for by differences in burst-firing or autocorrelogram oscillations. Firing regularity is known to vary between brain areas [35] and neuron type (excitatory vs. inhibitory) [36], both of which could, in principle, underlie the origin of the regularity bias between SbC and non-SbC neurons we find in our results. Indeed, one study conducted in monkey temporal cortex found broad-spiking neurons to fire less regularly than narrow-spiking neurons [37]. Since our SbC neurons overlapped almost entirely with the broad-spiking cluster, a low firing regularity might thus be expected. It is also noteworthy, that firing regularity is strongly influenced by the state of the animal [38] and seems to be particularly low during awake states, thus possibly explaining the reports of highly regular SbC spiking in dLGN of anesthetized macaques [12].

Oscillatory activity in the thalamo-cortical system of mice is influenced by specific stimuli, behavioral state, and top-down processing [39], and can globally influence firing regularity. When analyzing autocorrelogram oscillations, we found that SbC neurons in dLGN had weaker 3–5 Hz oscillations compared to the non-SbC populations. Oscillatory activity in the 3–5 Hz range within the mouse thalamo-cortical system is interesting, because this frequency has been linked to the primate alpha rhythm [40, 41, 42], which in turn is widely known to signal mental relaxation and low-arousal states. Our results indicate that SbC neurons in dLGN might be less affected by processes underlying the generation of the 3–5 Hz oscillation, which interestingly seem to involve low-threshold calcium bursts [42]. Consistent with this interpretation, we found that SbC neurons in dLGN displayed a lower burst ratio compared to the rest of the population. The 3–5 Hz rhythm in the mouse thalamo-cortical system is also remarkable since it is driven by a cyclical interaction between dLGN and V1, where temporary blocking of either dLGN or V1 abolishes the oscillation [42]. Together with our result that a similar difference in the strength of the 3–5 Hz oscillations between SbC and non-SbC neurons in V1 did not exist, this calls to question how much the responses of SbC neurons simply propagate in feedforward (or feedback) ways through the early visual system of mice, as speculated previously [15].

Interestingly, we found the largest differences in autocorrelogram oscillations between SbC and non-SbC neurons in dLGN, where firing regularity of SbC and non-SbC neurons was equal. Conversely, in V1, oscillation profiles were similar, while spike train regularity differed between response types. This suggests that oscillations in the autocorrelogram and spike train regularity are mediated, at least partially, by different mechanisms, such as interactions between different neuronal populations [43] vs. synaptic background activity [44, 45], respectively. These mechanisms, in turn, might thus impinge differentially on SbC and non-SbC neurons.

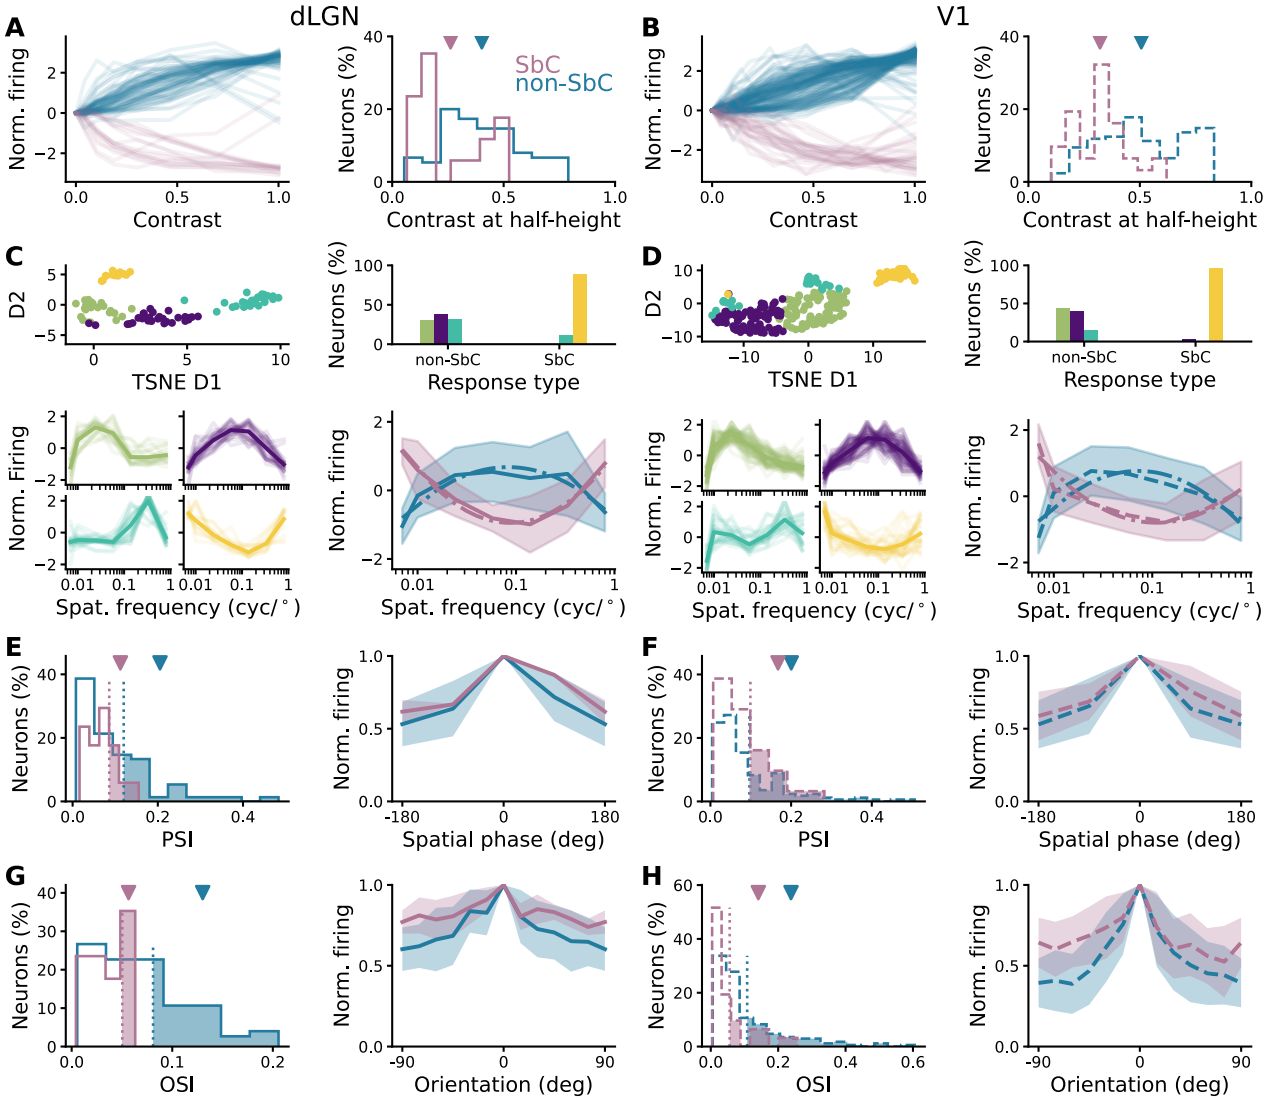

Figure 4. Tuning properties of SbC-cells in dLGN (A, C, E, G) and V1 (B, D, F, H). (**A, B**) *Left*: Z-scored contrast response functions of SbC cells. *Right*: Contrast-at-half-height distributions. (**C, D**) *Top left*: Embedded visualization and clustering of tuning curves reveals four distinct response groups. *Bottom left*: Response groups match previously reported response types (*green*: low-pass, dLGN: 23/92, V1: 75/200 neurons; *purple*: band-pass, dLGN: 28/92, V1: 70/200 neurons; *teal*: high-pass, dLGN: 28/92, V1: 26/200 neurons; *yellow*: band-stop, dLGN: 15/92, V1: 30/200 neurons). *Top right*: Unsupervised clustering matches supervised response type identification. *Bottom right*: Mean spatial frequency tuning of SbC neurons (*solid lines*). *Dashed lines*: fitted parabola to determine curvature. (**E, F**) SbC and non-SbC cells show similar mean tuning to spatial phase. *Left*: Distribution of phase selectivity index (*PSI*). Dotted lines and shaded areas indicate 30% most tuned neurons. *Right*: mean response at preferred orientation of the top 30% most tuned neurons, aligned to preferred spatial phase. Shaded areas correspond to standard deviation. (**G, H**) Comparison of orientation tuning between SbC and non-SbC neurons. *Left*: Distribution of orientation selectivity index (*OSI*). Dotted lines and shaded areas indicate 30% most tuned neurons. *Right*: Mean response of top 30% most tuned neurons, aligned to preferred orientation. Shaded areas correspond to standard deviation. (A-H) *Pink*: SbC; *turquoise*: non-SbC; *solid lines*: dLGN; *dashed lines*: V1.

Functional role in the circuit

Perhaps one of the most pressing questions around SbC neurons is if they represent a distinct response type that can also be connected to a distinct functional cell type. The evidence on this matter is currently mixed. Niell and Stryker [7] identify layer 2/3 SbC neurons as behavioral state dependent, fast-spiking inhibitory neurons that have been interpreted as PV+ interneurons.

Millman et al. [9], on the other hand, show that SbC neurons can be found in all cortical layers with the largest group being VIP neurons in layer 2/3. Interestingly, their group of VIP neurons differs from our SbC neurons in their tuning properties: for low contrasts, they display higher than baseline activity and strong orientation tuning that almost vanishes for high contrasts [9]. Additionally, the relatively tight distribution of waveshapes for SbC neurons (**Figure 3B**) we observe in our data seems different from that of vasoactive intestinal peptide-expressing (VIP) interneurons in other sources [24]. A recent study investigating both *in vivo* response properties and transcriptomic profiles of inhibitory V1 neurons identified another class of SbC neurons [46]. The SbC neurons in that study were synuclein-gamma-expressing (Sncg) cells, which responded to gratings, in particular of large size, with suppression [46]. Other results, however, point to excitatory neurons also belonging to the family of SbC neurons, including layer 2/3 pyramidal neurons [9]. Curiously, these layer 2/3 pyramidal cells with SbC response characteristics better matched the tuning properties of SbC neurons recorded in our study, but lacked the characteristically high baseline activity of SbC neurons. The finding of SbC neuron among L6 cortico-thalamic pyramidal cells [8] is well in line with our observed strong presence of SbC neurons in lower layers. Taken together, it seems save to say that, similar to the retina [2], there are likely multiple subtypes of SbC neurons in cortex, which do not belong to the same functional class or share the same genetic marker. Nonetheless, creating transgenic mouse lines that can selectively influence SbC neurons would be a prime way to gain information about their identity and function, as has been demonstrated recently in the retina [15, 17].

The finding of SbC neurons along all stages of the early visual pathway warrants the question whether suppression represents a projecting pathway, is generated de-novo multiple times, or a mixture of both. In the retina, SbC responses are generated via inhibition from amacrine cells that inhibit the responses of retinal ganglion cells with high spontaneous firing rates. The question is now whether involving inhibition to create SbC responses is also involved in other processing stages or if a mere "lack of activation" is projected along the pathway. In the case of an inhibitory mechanism, since there are no inhibitory forward projecting neurons from the retina or dLGN, local interneurons or maybe long-range inhibitory connections would have to be involved, resulting in one (or more) additional synapses in an SbC pathway compared to the classical non-SbC pathway. This would likely lead to a difference in response latency between SbC and non-SbC neurons when presented with abrupt stimuli. Our flashed gratings paradigm seems to indicate such an effect in both dLGN and V1, as mean latencies of SbC neurons are larger than non-SbC neurons. At the same time, the latency distributions still show considerable overlap, indicating that a mechanism involving multiple synapses and thus longer latencies may not be the only way that SbC responses are generated. Unfortunately, validating our results against other data is difficult; while [12] report equal or larger latencies for SbC neurons, we could not find additional reports of SbC response latencies. Re-analysis of existing datasets or consideration of latency as a relatively easily extracted indicator of projection route in future studies could lead to further insight on this matter. Additionally, genetic targeting of candidate neurons as mentioned above mentioned would allow causal investigations of the projecting pathway that leads to SbC responses.

Even though SbC neurons have been known for more than a half-century, their functional role is still not clear, partly because the existing evidence points in different directions. The most concrete investigations of high-level SbC functioning point towards involvement in saccade signaling in the retina [15] or prey capture [17]. Other studies, especially in dLGN and V1, have either observed SbC neurons as part of a larger population [9, 17], or observed the influence of variables like behavioral state *on* SbC neurons [7, 8], rather than the function *of* SbC Neurons. As such, these studies provide important guidance to direct future research towards causal investigations. One reason that the functional nature of SbC neurons has not yet been pinpointed is that they could be involved in almost anything: from a coding perspective, any feature that is coded in non-SbC fashion by an increase in activity, could also be coded inversely via a decrease of activity in SbC neurons, barring possibly higher metabolic costs. As such, one would expect diverse response features across brain areas and involvement in multiple contexts, as we currently see in the literature. The similar tuning properties we find in SbC neurons compared to non-SbC neurons similarly point towards the capability of SbC neurons to contribute to diverse coding processes.

Methods

All procedures complied with the European Communities Council Directive 2010/63/EU and the German Law for Protection of Animals, and were approved by local authorities, following appropriate ethics review.

Animals: Recordings were performed in 6 adult Ntsr1-Cre mice (5 male, 1 female, 5 hemizygous Tg, 1 negative control, median age at first recording session: 23.9 weeks; B6.FVB(Cg)-Tg(Ntsr1-cre)GN220Gsat/Mmcd; MMRRC, #030648-UCD) and 3 (2 male, 1 female) PV-Cre mice (median age: 14.1 weeks; B6.129P2-Pvalb$^{tm1(cre)Arbr}$/J; Jackson Laboratory, #017320).

Surgery: The surgical procedures are described in detail in [47]. In brief: mice were administered an analgesic (Metamizole, 200 mg/kg, sc, MSD Animal Health, Brussels, Belgium) and put under isoflurane anesthesia (5% in oxygen at start, then lowered to 0.5%–2% in oxygen, CP-Pharma, Burgdorf, Germany), the depth of which was constantly monitored. After shaving and disinfecting the scalp, a skin incision was performed and the skull cleaned of any remaining tissue. Upon positioning the head in a skull-flat position, a custom lightweight aluminium head bar with an opening over dLGN and V1 was placed on the skull and fixated using dental cement. For V1 recordings and optogenetic stimulation unrelated to this study in PV-Cre mice, a

small craniotomy above V1 was performed and $\sim 0.2\mu$L of pAAV9/1.EF1a.DIO.hChR2(H134R)-eYFP.WPRE.hGH (Addgene, #20298-AAV9) dyed with fast-green (Sigma-Aldrich, St. Louis, USA) was injected through the entire depth of the cortex. In the Ntsr1-Cre mice used for additional V1 and dLGN recordings, a similar craniotomy was performed and $\sim 0.35\mu$L of stGtACR2 (rAAV2/1-pAAV-hSyn1-SIO-stGtACR2-FusionRed, Addgene, #105677) were injected in the infragranular layers of cortex for experiments with suppression of corticothalamic feedback unrelated to the current study. Post-injection, the opening was filled with Kwik-Cast (WPI Germany, Berlin, Germany). Long-term analgesic (Meloxicam, 2 mg/kg, sc, Böhringer Ingelheim, Ingelheim, Germany) was administered and continued to be administered for 3 consecutive days. After at least 1 week of recovery, animals were gradually habituated to the experimental setup, by first handling them and then simulating the experimental procedure. To allow for virus expression, neural recordings started no sooner than 3 weeks after injection. On the day prior to the first day of recording, mice were fully anesthetized using the same procedures as for the initial surgery, and a craniotomy (ca. 1.5 mm^2) was performed over dLGN and/or V1, and re-sealed with Kwik-Cast. As long as the animals did not show signs of discomfort, the long-term analgesic Metacam was administered only once at the end of surgery, to avoid any confounding effect on experimental results. Recordings were performed daily and continued for as long as the quality of the electrophysiological signals remained high.

Experimental setup: Our experimental configuration for *in-vivo* recordings is described in detail in [47]. In brief: mice were head-fixed and could run freely on an air-suspended styrofoam ball while stimuli were presented to the right visual field on a gamma-corrected LCD screen. Extracellular neural signals were recorded with 32-channel silicon probes (Neuronexus, A1x32Edge-5mm-20-177-A32, Ann Arbor, USA) for the 4 Ntsr1-Cre mice, a 32-channel silicon probe for one PV-Cre mouse (A1x32-Edge-5mm-20-177-A32 and A1x32Edge-5mm-20-177-A32), and a 64-channel silicon probe (A1x64-Poly2-6mm-23s-160-A64) for the other PV-Cre mouse. Ball movements were registered at 90 Hz by two optical mice connected to an Arduino-type microcontroller. Eye movements were monitored under infrared light illumination.

For photostimulation of V1 PV+ inhibitory interneurons, unrelated to this study, an optic fiber (910 µm diameter, Thorlabs, Newton, USA) was coupled to a light-emitting diode (LED, center wavelength 470 nm, M470F1, Thorlabs, Newton, USA) and positioned with a micromanipulator less than 1 mm above the exposed surface of V1. A black metal foil surrounding the tip of the head bar holder prevented the photostimulation light from reaching the animal's eyes.

Perfusion and histology: After the final recording session, mice were first administered an analgesic (Metamizole, 200 mg/kg, sc, MSD Animal Health, Brussels, Belgium) and following a 30 min wait period were transcardially perfused under deep anesthesia using a cocktail of Medetomidin (0.5 mL/kg), Midazolam (1 mL/kg), and Fentanyl (1 mL/kg) (ip). Perfusion was first done with Ringer's lactate solution followed by 4% paraformaldehyde (PFA) in 0.2 M sodium phosphate buffer (PBS).

To verify recording site and virus expression, we performed histological analyses. Brains were removed, postfixed in PFA for 24 h, and then rinsed with and stored in PBS at 4°C. Slices (40 µm) were cut using a vibrotome (Leica VT1200 S, Leica, Wetzlar, Germany), mounted on glass slides with Vectashield DAPI (Vector Laboratories, Burlingame, USA), and coverslipped. A fluorescent microscope (BX61 Systems Microscope, Olympus, Tokyo, Japan) was used to inspect slices for the presence of yellow fluorescent protein (eYFP) and DiI. Recorded images were processed using FIJI [48, 49].

Stimulus: We used custom software (EXPO, https://sites.google.com/a/nyu.edu/expo/home) to present visual stimuli on a gamma-calibrated liquid crystal display (LCD) monitor (Samsung SyncMaster 2233RZ; mean luminance 50 cd/m^2, 60 Hz) at 25 cm distance to the animal's right eye (spanning $\sim$ 108x66°, small angle approximation). Mice were presented with three 12 min random sequences of briefly flashed (84 ms), full-screen grating stimuli. The random sequences were drawn from 2304 unique gratings covering 12 orientations ($0, 15, 30, 45, 60, 75, 90, 105, 120, 135, 150, 165°$), 8 contrasts (0, 0.04, 0.10, 0.19, 0.30, 0.46, 0.69, 1), 6 spatial frequencies (0.01, 0.02, 0.06, 0.14, 0.33, 0.80 cyc/°) and 4 spatial phases ($0, 90, 180, 270°$). One sequence consisted of 9216 gratings. Between the sequences, a blank gray screen was displayed for 1 min. For V1 recordings in PV-Cre mice expressing ChR2, light pulses (10 Hz, 1 ms pulses) were delivered from the optical fiber during these periods; analyses of the blank screen responses or photostimulation effects were not included in the current study. Typically, the stimulus sequence was presented once per electrode penetration, except in two cases, where the sequence was run twice during one electrode penetration but data from each run was analyzed separately.

Data analysis: Wideband extracellular signals were digitized at 30 kHz (Blackrock microsystems, Blackrock Microsystems Europe GmbH, Hannover DE). To obtain single unit activity from extracellular recordings, the open source, Matlab-based, automated spike sorting toolbox Kilosort [50] was used. Resulting clusters were manually refined using Spyke [51], a Python application that allows the selection of channels and time ranges around clustered spikes for realignment, as well as representation in 3D space using dimension reduction (multichannel PCA, ICA, and/or spike time). Exhaustive pairwise comparisons of similar clusters allowed merging of potentially over-clustered units. All further analyses were performed using an SQL data base and a custom-made analysis pipeline programmed in `python` [52] and managed via `datajoint` [53].

Spike waveshape analysis: From the mean waveform of the maximum-response electrode channel of each single unit, the time between trough and peak (trough-to-peak time) and the full-width at half-height of the peak (peak width) were calculated. Exploiting the waveshapes of all V1 units processed using the same pipeline (N = 540 from 11 mice), a k-means algorithm was

used to cluster the data into 2 populations.

Laminar location: We used current source density (CSD) analysis [30] for recordings in area V1 to determine the laminar position of electrode contacts. To obtain the LFP, we first down-sampled the signal to 1 kHz before applying a bandpass filter (3–90 Hz, second-order Butterworth filter). We computed the CSD using the iCSD method [31] implemented in `elephant` [(RRID:SCR_003833) 54]. We assigned the base of layer 4 to the contact that was closest to the earliest CSD polarity inversion. The remaining contacts were assigned to layers based on relative layer thickness reported by [55], assuming a thickness of 1.2 mm for mouse visual cortex.

Temporal response kernels via reverse correlation: After eliminating units with overall very low firing rate (<0.1 Hz), the probability of a stimulus preceding a spike by a time δt from -50 to 350 ms was computed for each unique grating stimulus. This was done by binning spike times in 1 ms windows and then counting how often a specific stimulus occurred δt before a spike. After normalizing this histogram by dividing by the total number of spikes, the posterior distribution was calculated according to Bayes' theorem by multiplying with the probability of a spike occurring and dividing by the probability of the stimulus occurring:

$$\begin{aligned} P(\text{Spike}|\text{Grating}) &= \frac{P(\text{Grating}|\text{Spike}) \cdot P(\text{Spike})}{P(\text{Grating})} \\ &= \frac{\frac{\text{Grating And Spike Bins}}{\text{Total Spike Bins}} \cdot \frac{\text{Total Spike Bins}}{\text{Total Bins}}}{\frac{\text{Grating Bins}}{\text{Total Bins}}} \\ &= \frac{\text{Grating And Spike Bins}}{\text{Grating Bins}} \end{aligned}$$

By dividing by the 1 ms bin duration, the resulting probabilities could be directly converted to firing rates and thus give a temporal response kernel for each unique grating stimulus.

Determining visual responsiveness: To eliminate non-responsive or noise-dominated units, the variance of the temporal kernels across stimuli was calculated and tested for non-randomness using the Wald-Wolfowitz-Test (WWT). The WWT uses the distribution of consecutive ones and zeros in a binary sequence, which should follow a normal distribution in a random sequence, to statistically determine randomness of the sequence. The test can be applied to a non-binary sequence by converting it to a binary sequence via a threshold criterion, typically the mean or median. Before applying the WWT, global trends in the sequence should be removed by either filtering or applying an approximation of the derivative. As the response to grating contrast is a robust indicator of visual responsiveness, the analysis was performed using the aggregate variance across grating contrast, which was computed by averaging the kernels across grating orientation, spatial frequency and spatial phase, before computing the variance across the resulting contrast kernels. To further increase signal-to-noise ratio, the partial variance was squared before computing the differences across time to remove any global trends. On the resulting sequence, the WWT was performed using the median as cutoff criterion. Because the WWT can miss narrow peaks, even if they are high, a second WWT was computed on the absolute values of the sequence. A unit was classified as visually responsive if one of the WWT's was significant and the test statistic of both WWT's was negative, indicating fewer sign changes than expected by chance.

Determining optimal time point: The time point of optimal response was determined via the peak of the summed aggregate variances across stimuli. First, the partial variances were computed for all four stimulus parameters as described above for grating contrast. The resulting partial variances were then summed and the time point of the first peak exceeding half the modulation depth of the result was selected as the optimal response latency δt_{opt} [peak detection using `scipy`, 56]. Responses were then averaged over a 20 ms window around δt_{opt}, yielding a 4 dimensional response matrix spanning contrast, spatial frequency, spatial phase and orientation.

Contrast response functions: Contrast response functions were computed by averaging over the 3 other stimulus dimensions: spatial frequency, spatial phase and orientation. Contrast sensitivity was quantified as contrast at half height of the contrast response function.

Response type classification: First, contrast response functions were converted to modulation strength by dividing by the response at 0 contrast. Neurons were classified as SbC, if their mean modulation strength was below -0.1. For the two cases where the experiment was performed twice during one electrode penetration, if the response classification of the experiments conflicted with one another (SbC vs non-SbC), the neuron was discarded from further analysis; in all other cases data from both experiments were used for further analysis.

Spike statistics during spontaneous activity: For 6 of the 9 mice involved in this study, we were able to use spontaneous activity during the grey-screen periods of the flashed gratings stimulus. For 1 of these mice and 1 of the remaining mice, we were able to use spontaneous activity during other grey-screen periods lasting multiple minutes, translating to 7 out of 9 mice contributing spontaneous activity data to the analysis. We computed mean firing rate, *CV* and burst ratio from this data and compared them as described above. In this analysis, we performed an additional two-way ANOVA on the rank-transformed

data in case of violations of the normality assumption. This yielded the same results.

Oscillation score: The *OS* analysis is described in detail in [25]. In brief: for a selected frequency band, a spike autocorrelogram is computed with specifically chosen time delays and binning. It is smoothed with a Gaussian kernel to eliminate high-frequency noise. The central peak of the autocorrelogram introduces strong low-frequency artifacts in the spectrum, so the peak is cut from the autocorrelogram. Next, a Fourier transform is applied to compute a power spectrum. The frequency of maximal power in the desired frequency band is then extracted and the *OS* is computed by dividing the maximal power in the desired band by the mean power of the spectrum. To recover an interpretable spectrum from the frequency-*OS* pairs calculated for each neuron and frequency band, we applied a triangular smoothing kernel of 3 Hz width. To find ranges of difference between two spectra, we iteratively calculated the Pearson correlation between the spectra, replaced one point, calculated the correlation again and interpreted the difference in ρ as the impact of the point. By always replacing the point of largest difference between the spectra, they became more and more similar until their correlation was larger than 0.999. Summing the impact of adjacent points yielded ranges of difference.

Tuning functions: Spatial frequency tuning curves were calculated by averaging over the other stimulus dimensions: contrast, spatial phase and orientation.

Spatial phase: Spatial phase tuning curves were calculated by selecting responses at the preferred stimulus orientation and then averaging over the remaining stimulus dimensions: contrast and spatial frequency. The preferred orientation was calculated by fitting a wrapped Gaussian function to the orientation tuning curve (averaged over the remaining stimulus dimensions), extracting the parameter of preferred orientation and selecting the stimulus orientation closest to this value. The then used vector averaging methods is described below.

Orientation selectivity: Orientation tuning curves were calculated by selecting responses at the preferred spatial phase and then averaging over the remaining stimulus dimensions: contrast and spatial frequency. The preferred spatial phase was selected as the phase of maximal response in the phase tuning curves (averaged over the remaining stimulus dimensions).

When analyzing tuning strength, we employed a vector averaging method with the following reasoning: width parameters of orientation tuning functions, σ in our case, scale non-linearly when tuning is either strong or weak, depending on the specific function. When quantitatively analysing data that contains a broad spectrum of tuning, it is thus advisable to use measures that are not distorted by such non-linear scaling. Accordingly, orientation selectivity [57, 58] was quantified as

$$\mathrm{OSI} = \frac{\sqrt{(\sum R_k \sin(2\theta_k))^2 + (\sum R_k \cos(2\theta_k))^2}}{\sum R_k} \tag{1}$$

where R_k is the response to the kth direction given by θ_k.

Statistics: All statistics were performed using functions from `scipy.stats` and `statsmodels` except for the three-way mixed ANOVA's used in the oscillation score analyses, which were implemented in `R`.

References

1. David H. Hubel and Thorsten N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160(1):106–154, 1962. doi: 10.1113/jphysiol.1962.sp006837.
2. Tom Baden, Philipp Berens, Katrin Franke, Miroslav Román Rosón, Matthias Bethge, and Thomas Euler. The functional diversity of retinal ganglion cells in the mouse. *Nature*, 529(7586):345–350, 2016. doi: 10.1038/nature16468.
3. Edgar Y. Walker, Fabian H. Sinz, Erick Cobos, Taliah Muhammad, Emmanouil Froudarakis, Paul G. Fahey, Alexander S. Ecker, Jacob Reimer, Xaq Pitkow, and Andreas S. Tolias. Inception loops discover what excites neurons most using deep predictive models. *Nature Neuroscience*, 22(12):2060–2065, 2019. doi: 10.1038/s41593-019-0517-x.
4. Richard H. Masland and Paul R. Martin. The unsolved mystery of vision. *Current Biology*, 17(15):577–582, 2007. doi: 10.1016/j.cub.2007.05.040.
5. Robert W. Rodieck. Receptive Fields in the Cat Retina: A New Type. *Science*, 157(3784):90–92, 1967. doi: 10.1126/science.157.3784.90.
6. Denise M. Piscopo, Rana N. El-Danaf, Andrew D. Huberman, and Cristopher M. Niell. Diverse Visual Features Encoded in Mouse Lateral Geniculate Nucleus. *Journal of Neuroscience*, 33(11):4642–4656, 2013. doi: 10.1523/JNEUROSCI.5187-12.2013.
7. Cristopher M. Niell and Michael P. Stryker. Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. *Neuron*, 65(4):472–479, 2010. doi: 10.1016/j.neuron.2010.01.033.
8. Sigita Augustinaite and Bernd Kuhn. Complementary Ca^{2+} Activity of Sensory Activated and Suppressed Layer 6 Corticothalamic Neurons Reflects Behavioral State. *Current Biology*, 30(20):3945–3960.e5, 2020. doi: 10.1016/j.cub.2020.07.069.
9. Daniel J Millman, Gabriel Koch Ocker, Shiella Caldejon, India Kato, Josh D Larkin, Eric Kenji Lee, Jennifer Luviano, Chelsea Nayan, Thuyanh V Nguyen, Kat North, Sam Seid, Cassandra White, Jerome Lecoq, Clay Reid, Michael A Buice,

and Saskia EJ de Vries. VIP interneurons in mouse primary visual cortex selectively enhance responses to weak but specific stimuli. *eLife*, 9:e55130, 2020. doi: 10.7554/eLife.55130.

10. William R. Levick. Receptive fields and trigger features of ganglion cells in the visual streak of the rabbit's retina. *The Journal of Physiology*, 188(3):285–307, 1967. doi: 10.1113/jphysiol.1967.sp008140.

11. Franklin R. Amthor, Ellen S. Takahashi, and Clyde W. Oyster. Morphologies of rabbit retinal ganglion cells with concentric receptive fields. *Journal of Comparative Neurology*, 280(1):72–96, 1989. doi: 10.1002/cne.902800107.

12. Chris Tailby, Samuel G. Solomon, Neel T. Dhruv, Najib J. Majaj, Sach H. Sokol, and Peter Lennie. A new code for contrast in the primate visual pathway. *Journal of Neuroscience*, 27(14):3904–3909, 2007. doi: 10.1523/JNEUROSCI.5343-06.2007.

13. Séverine Durand, Ramakrishnan Iyer, Kenji Mizuseki, Saskia De Vries, Stefan Mihalas, and R. Clay Reid. A comparison of visual response properties in the lateral geniculate nucleus and primary visual cortex of awake and anesthetized mice. *Journal of Neuroscience*, 36(48):12144–12156, 2016. doi: 10.1523/JNEUROSCI.1741-16.2016.

14. Benjamin Sivyer, W. Rowland Taylor, and David I. Vaney. Uniformity detector retinal ganglion cells fire complex spikes and receive only light-evoked inhibition. *Proceedings of the National Academy of Sciences of the United States of America*, 107(12):5628–5633, 2010. doi: 10.1073/pnas.0909621107.

15. Nai Wen Tien, James T. Pearson, Charles R. Heller, Jay Demas, and Daniel Kerschensteiner. Genetically identified suppressed-by-contrast retinal ganglion cells reliably signal self-generated visual stimuli. *Journal of Neuroscience*, 35(30):10815–10820, 2015. doi: 10.1523/JNEUROSCI.1521-15.2015.

16. Miroslav Román Rosón, Yannik Bauer, Ann H. Kotkat, Philipp Berens, Thomas Euler, and Laura Busse. Mouse dLGN Receives Functional Input from a Diverse Population of Retinal Ganglion Cells with Limited Convergence. *Neuron*, 102 (2), 2019. doi: 10.1016/j.neuron.2019.01.040.

17. Keith P. Johnson, Michael J. Fitzpatrick, Lei Zhao, Bing Wang, Sean McCracken, Philip R. Williams, and Daniel Kerschensteiner. Cell-type-specific binocular vision guides predation in mice. *Neuron*, 109(9):1527–1539.e4, 2021. doi: 10.1016/j.neuron.2021.03.010.

18. Dario Ringach and Robert Shapley. Reverse correlation in neurophysiology. *Cognitive Science*, 28(2):147–166, 2004. doi: 10.1016/j.cogsci.2003.11.003.

19. Elaine Tring and Dario L. Ringach. On the Subspace Invariance of Population Responses. *arXiv*, 2018.

20. James A. Mazer, William E. Vinje, Josh McDermott, Peter H. Schiller, and Jack L. Gallant. Spatial frequency and orientation tuning dynamics in area V1. *Proceedings of the National Academy of Sciences of the United States of America*, 99(3):1645–1650, 2002. doi: 10.1073/pnas.022638499.

21. Michael N. Shadlen and William T. Newsome. The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *The Journal of Neuroscience*, 18(10):3870–3896, 1998. doi: 10.1523/JNEUROSCI.18-10-03870.1998.

22. Wei Wang, Helen E Jones, Ian M Andolina, Thomas E Salt, and Adam M Sillito. Functional alignment of feedback effects from visual cortex to thalamus. *Nature Neuroscience*, 9(10):1330–1336, 2006.

23. S. Murray Sherman. Tonic and burst firing: Dual modes of thalamocortical relay. *Trends in Neurosciences*, 24(2):122–126, 2001.

24. Adrien Peyrache and Alain Destexhe. Electrophysiological monitoring of inhibition in mammalian species, from rodents to humans. *Neurobiology of Disease*, 130:104500, 2019. doi: 10.1016/j.nbd.2019.104500.

25. Raul C. Mureşan, Ovidiu F. Jurjuţ, Vasile V. Moca, Wolf Singer, and Danko Nikolić. The oscillation score: An efficient method for estimating oscillation strength in neuronal activity. *Journal of Neurophysiology*, 99(3):1333–1353, 2008. doi: 10.1152/jn.00772.2007.

26. Aman B. Saleem, Anthony D. Lien, Michael Krumin, Bilal Haider, Miroslav Román Rosón, Asli Ayaz, Kimberly Reinhold, Laura Busse, Matteo Carandini, and Kenneth D. Harris. Subcortical Source and Modulation of the Narrowband Gamma Oscillation in Mouse Visual Cortex. *Neuron*, 93(2):315–322, 2017. doi: 10.1016/j.neuron.2016.12.028.

27. Marius Schneider, Ana Clara Broggini, Benjamin Dann, Athanasia Tzanou, Cem Uran, Swathi Sheshadri, Hansjörg Scherberger, and Martin Vinck. A mechanism for inter-areal coherence through communication based on connectivity and oscillatory power. *Neuron*, page S0896627321007108, 2021. doi: 10.1016/j.neuron.2021.09.037.

28. Cristopher M. Niell and Michael P. Stryker. Highly Selective Receptive Fields in Mouse Visual Cortex. *Journal of Neuroscience*, 28(30):7520–7536, 2008. doi: 10.1523/jneurosci.0623-08.2008.

29. Shi H. Sun, Ali Almasi, Molis Yunzab, Syeda Zehra, Damien G. Hicks, Tatiana Kameneva, Michael R. Ibbotson, and Hamish Meffin. Analysis of extracellular spike waveforms and associated receptive fields of neurons in cat primary visual cortex. *Journal of Physiology*, 8:2211–2238, 2021. doi: 10.1113/JP280844.

30. Ulla Mitzdorf. Current source-density method and application in cat cerebral cortex: Investigation of evoked potentials and EEG phenomena. *Physiological Reviews*, 65(1):37–100, 1985. doi: 10.1152/physrev.1985.65.1.37.

31. Klas H. Pettersen, Anna Devor, Istvan Ulbert, Anders M. Dale, and Gaute T. Einevoll. Current-source density estimation based on inversion of electrostatic forward solution: Effects of finite extent of neuronal activity and conductivity discontinuities. *Journal of Neuroscience Methods*, 154(1-2):116–133, 2006. doi: 10.1016/j.jneumeth.2005.12.005.

32. Rodney J. Douglas and Kevan A.C. Martin. Neuronal circuits of the neocortex. *Annual Review of Neuroscience*, 27: 419–451, 2004. doi: 10.1146/annurev.neuro.27.070203.144152.

33. James H. Marshel, Marina E. Garrett, Ian Nauhaus, and Edward M. Callaway. Functional specialization of seven mouse visual cortical areas. *Neuron*, 72(6):1040–1054, 2011. doi: 10.1016/j.neuron.2011.12.004.

34. Nicholas V. Swindale. Orientation tuning curves: Empirical description and estimation of parameters. *Biological Cybernetics*, 78(1):45–56, 1998. doi: 10.1007/s004220050411.

35. Gaby Maimon and John A. Assad. Beyond Poisson: Increased Spike-Time Regularity across Primate Parietal Cortex. *Neuron*, 62(3):426–440, 2009. doi: 10.1016/j.neuron.2009.03.021.

36. Klaus M. Stiefel, Bernhard Englitz, and Terrence J. Sejnowski. Origin of intrinsic irregular firing in cortical interneurons. *Proceedings of the National Academy of Sciences of the United States of America*, 110(19):7886–7891, 2013. doi: 10.1073/pnas.1305219110.

37. Shigeru Shinomoto, Youichi Miyazaki, Hiroshi Tamura, and Ichiro Fujita. Regional and laminar differences in in vivo firing patterns of primate cortical neurons. *Journal of Neurophysiology*, 94(1):567–575, 2005. doi: 10.1152/jn.00896.2004.

38. Ida E.J. Aasebø, Mikkel E. Lepperød, Maria Stavrinou, Sandra Nøkkevangen, Gaute Einevoll, Torkel Hafting, and Marianne Fyhn. Temporal processing in the visual cortex of the awake and anesthetized rat. *eNeuro*, 4(4), 2017. doi: 10.1523/ENEURO.0059-17.2017.

39. Yuan Zhang and Xiaohui Zhang. Portrait of visual cortical circuits for generating neural oscillation dynamics. *Cognitive Neurodynamics*, 15(1):3–16, 2021. doi: 10.1007/s11571-020-09623-4.

40. Yuta Senzai, Antonio Fernandez-Ruiz, and György Buzsáki. Layer-specific physiological features and interlaminar interactions in the primary visual cortex of the mouse. *Neuron*, 101(3):500–513, 2019.

41. David A. McCormick, Dennis B. Nestvogel, and Biyu J. He. Neuromodulation of Brain State and Behavior. *Annual Review of Neuroscience*, 43:391–415, 2020. doi: 10.1146/annurev-neuro-100219-105424.

42. Dennis B Nestvogel and David A McCormick. Visual Thalamocortical Mechanisms of Waking State Dependent Activity and Alpha Oscillations. *bioRxiv*, 2021.

43. Guang Chen, Yuan Zhang, Xiang Li, Xiaochen Zhao, Qian Ye, Yingxi Lin, Huizhong W. Tao, Malte J. Rasch, and Xiaohui Zhang. Distinct Inhibitory Circuits Orchestrate Cortical beta and gamma Band Oscillations. *Neuron*, 96(6):1403–1418.e6, 2017. doi: 10.1016/j.neuron.2017.11.033.

44. Charles F. Stevens and Anthony M. Zador. Input synchrony and the irregular firing of cortical neurons. *Nature Neuroscience*, 1(3):210–217, 1998. doi: 10.1038/659.

45. Jean-Marc Fellous, Michelle Rudolph, Alain Destexhe, and Terrence J. Sejnowski. Synaptic background noise controls the input/output characteristics of single cells in an in vitro model of in vivo activity. *Neuroscience*, 122(3):811–829, 2003. doi: 10.1016/j.neuroscience.2003.08.027.

46. Stephane Bugeon, Joshua Duffield, Mario Dipoppa, Anne Ritoux, Isabelle Prankerd, Dimitris Nicolout-sopoulos, David Orme, Maxwell Shinn, Han Peng, Hamish Forrest, Aiste Viduolyte, Charu Bai Reddy, Yoh Isogai, Matteo Carandini, and Kenneth D. Harris. A transcriptomic axis predicts state modulation of cortical interneurons. *bioRxiv*, October 2021. doi: 10.1101/2021.10.24.465600.

47. Martin A. Spacek, Gregory Born, Davide Crombie, Yannik Bauer, Xinyu Liu, Steffen Katzner, and Laura Busse. Robust effects of corticothalamic feedback during naturalistic visual stimulation. *bioRxiv*, 2021. doi: 10.1101/776237.

48. Curtis T Rueden, Johannes Schindelin, Mark C Hiner, Barry E DeZonia, Alison E Walter, Ellen T Arena, and Kevin W Eliceiri. Imagej2: Imagej for the next generation of scientific image data. *BMC Bioinformatics*, 18(1):529, 2017.

49. Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pietzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin Schmid, et al. Fiji: an open-source platform for biological-image analysis. *Nature Methods*, 9(7):676–682, 2012.

50. Marius Pachitariu, Nicholas A. Steinmetz, Shabnam N. Kadir, Matteo Carandini, and Kenneth D. Harris. Fast and accurate spike sorting of high-channel count probes with KiloSort. *Advances in Neural Information Processing Systems*, 29: 4448–4456, 2016.

51. Martin Spacek, Tim Blanche, and Nicholas V. Swindale. Python for large-scale electrophysiology. *Frontiers in Neuroinformatics*, 2(JAN):1–10, 2009. doi: 10.3389/neuro.11.009.2008.

52. Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.

53. Dimitri Yatsenko, Edgar Y. Walker, and Andreas S. Tolias. DataJoint: A Simpler Relational Data Model. *arXiv*, 2018.

54. Michael Denker, Alper Yegenoglu, and Sonja Grün. Collaborative HPC-enabled workflows on the HBP Collaboratory using the Elephant framework. In *Neuroinformatics 2018*, page P19, 2018. doi: 10.12751/incf.ni2018.0019.

55. Didier Heumann, Genevieve Leuba, and Theodore Rabinowicz. Postnatal development of the mouse cerebral neocortex. II. Quantitative cytoarchitectonics of visual and auditory areas. *Journal fur Hirnforschung*, 18(6):483–500, 1977.

56. Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R.J. Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, Aditya Vijaykumar, Alessandro Pietro Bardelli, Alex Rothberg, Andreas Hilboll, Andreas Kloeckner, Anthony Scopatz, Antony Lee, Ariel Rokem, C. Nathan Woods, Chad Fulton, Charles Masson, Christian Häggström, Clark Fitzgerald, David A. Nicholson, David R. Hagen, Dmitrii V. Pasechnik, Emanuele Olivetti, Eric Martin, Eric Wieser, Fabrice Silva, Felix Lenders, Florian Wilhelm, G. Young, Gavin A. Price, Gert Ludwig Ingold, Gregory E. Allen, Gregory R. Lee, Hervé Audren, Irvin Probst, Jörg P. Dietrich, Jacob Silterra, James T. Webber, Janko Slavič, Joel Nothman, Johannes Buchner, Johannes Kulick, Johannes L. Schönberger, José Vinícius de Miranda Cardoso, Joscha Reimer, Joseph Harrington, Juan Luis Cano Rodríguez, Juan Nunez-Iglesias, Justin Kuczynski, Kevin Tritz, Martin Thoma, Matthew Newville, Matthias Kümmerer, Maximilian Bolingbroke, Michael Tartre, Mikhail Pak, Nathaniel J. Smith, Nikolai Nowaczyk, Nikolay Shebanov, Oleksandr Pavlyk, Per A. Brodtkorb, Perry Lee, Robert T. McGibbon, Roman Feldbauer, Sam Lewis, Sam Tygier, Scott Sievert, Sebastiano Vigna, Stefan Peterson, Surhud More, Tadeusz Pudlik, Takuya Oshima, Thomas J. Pingel, Thomas P. Robitaille, Thomas Spura, Thouis R. Jones, Tim Cera, Tim Leslie, Tiziano Zito, Tom Krauss, Utkarsh Upadhyay, Yaroslav O. Halchenko, and Yoshiki Vázquez-Baeza. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3): 261–272, 2020. doi: 10.1038/s41592-019-0686-2.

57. Tobias Bonhoeffer, Dae-Shik Kim, Dov Malonek, Doron Shoham, and Amiram Grinvald. Optical Imaging of the Layout of Functional Domains in Area 17 and Across the Area 17/18 Border in Cat Visual Cortex. *European Journal of Neuroscience*, 7(9):1973–1988, 1995. doi: 10.1111/j.1460-9568.1995.tb00720.x.

58. Shawn R. Olsen, Dante S. Bortone, Hillel Adesnik, and Massimo Scanziani. Gain control by layer six in cortical circuits of vision. *Nature*, 483(7387):47–54, 2012. doi: 10.1038/nature10835.

Acknowledgements

This research was supported by the German Research Foundation (DFG) SPP2041 (LB), DFG BU 1808/5-1 (LB). Laura Busse (LB) thanks M. Sotgia for lab management and support with animal handling and histology, S. Schörnich for IT support, B. Grothe for providing excellent research infrastructure. LB and Simon Renner (SR) thank all our group members for fruitful discussions.

Author contributions statement

LB, and SR conceived the study. SR recorded and curated data and developed the data analysis software. Emma Müller-Seydlitz developed analysis techniques for spontaneous activity data. SR performed data analysis. SR and LB contributed conceptual ideas, discussed and coordinated the project at all stages, wrote and edited the manuscript.

Additional information

We have no competing interests to declare.

Supplements

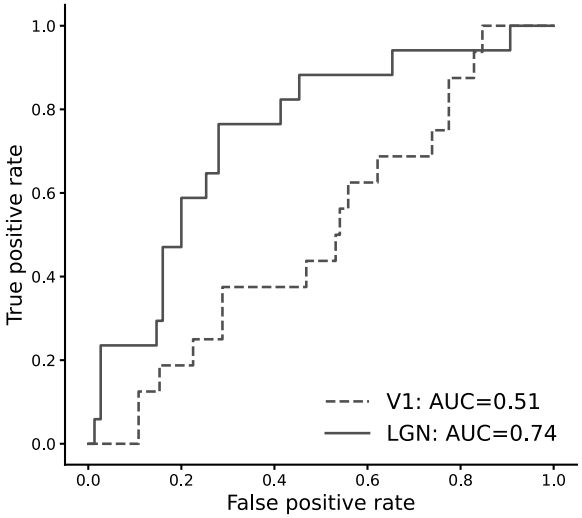

Figure S1. Receiver-operating-characteristics for spontaneous firing rates.

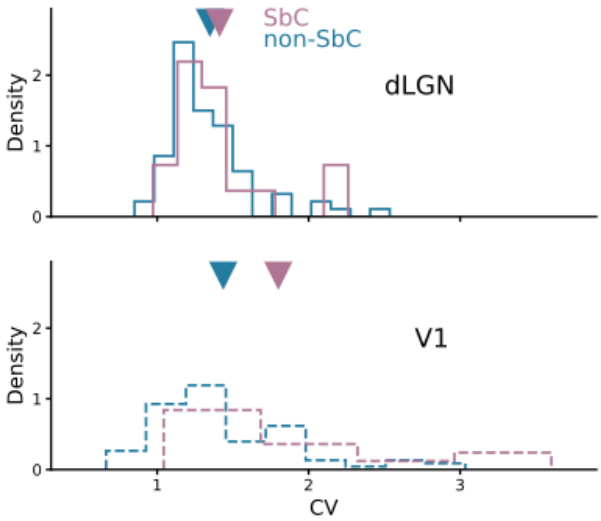

Figure S2. CV distributions without burst spikes. *Pink*: SbC; *turquoise*: non-SbC; *solid lines*: dLGN; *dashed lines*: V1.

4 Conclusion

In the two manuscripts presented in this thesis, my collaborators and I investigated how stimulus responses in dLGN and V1 constrain connectivity in V1, and how intrinsic properties and response characteristics differ between SbC and non-SbC neurons in dLGN and V1. We found that *in vivo* responses imply a specific order of connection strength and width, which we were able to recover from previous studies. Additionally, we found that SbC neurons fire less regularly than non-SbC neurons, are located in deeper cortical layers, have larger response latencies, and show significant stimulus tuning, suggesting the role of SbC neurons in visual processing goes beyond luminance uniformity detection.

4.1 Discussion

How and when biological complexity can be reflected in computational models has been a central question underlying the work presented in this thesis. Without employing detailed single neurons or intricate circuit details, our SSN model nonetheless produced predictions about connectivity that had *not* been suggested before and still could be recovered from the literature. At a time when many see large-scale network analyses as the way forward in neuroscience (Saxena and Cunningham, 2019), this underscores the power of distilling *interpretable* information from biological data sets as well as modeling approaches. At the same time, the increasing ability to classify cell types (Zeng and Sanes, 2017) presents an imperative to extend models while maintaining interpretability. Additionally, the differences to earlier studies we find in our investigation of SbC neurons, such as similar baseline activity as non-SbC neurons (Piscopo et al., 2013), irregular firing (Tailby et al., 2007), and significant stimulus tuning (Niell and Stryker, 2010), highlight how difficult it can be to distill interpretable information from experimental results. While some of these differences might be explained by species effects or behavioral state, the present picture of heterogeneity suggests that suppression is generated differently across brain regions and animal models. It thus seems clear that additional inves-

tigations are needed to uncover if SbC effects are bound to specific cell types or if they emerge from properties of the underlying network.

Reasons for separating neuronal response types

Manuscripts one and two are connected by the presence of SbC neurons in the data, which were deliberately excluded from modeling and analyses in manuscript one and studied in detail in manuscript two. The reasons for the exclusion in manuscript one are twofold: first, their role in visual processing is not well understood, making it unclear how they should (or could) be included in our model. To my knowledge, no study has investigated the thalamocortical projection patterns or polarity of SbC neurons in dLGN. It is therefore unclear whether they are excitatory or inhibitory, if their responses are processed purely locally, or if they project to cortex. Consequently, including them as inputs in our model would have been speculative and might have introduced biases that were not present in earlier studies. In visual cortex, SbC neurons can be excitatory pyramidal neurons (Augustinaite and Kuhn, 2020) but have also been associated with different subtypes of inhibitory neurons (Niell and Stryker, 2010; Millman et al., 2020; Bugeon et al., 2021), which have specific wiring patterns and functionality (Pfeffer et al., 2013). This heterogeneity precludes the integration of SbC neurons in a two-population model until their influences on other neurons and wiring patterns have been mapped more extensively. Modeling SbC neurons as separate populations on the other hand, would require additional assumptions or render the inference procedure mathematically invalid (see below for more details). Second, as SbC neurons have most likely been absent from many data sets that led to the conceptualizations of cortical function, their inclusion would likely lead to contradictions with the existing literature. Until automated, high-throughput methods were developed around the turn of the millennium (Gray et al., 1995; Hulata et al., 2002), neurons were often hand-mapped with high-contrast stimuli before recording actual data to optimize yield from labor-intensive experiments (Hubel and Wiesel, 1962; Alitto and Usrey, 2004), introducing a bias towards non-SbC neurons in recorded data and models constructed to fit the data (Wörgötter and Koch, 1991; Carandini and Heeger, 1994). Including SbC neurons in our population responses would significantly alter the shape of the contrast response functions, which has strong influence on the behavior of the SSN (Persi et al., 2011). Thus, we only used non-SbC neurons to construct the population response.

Relation to contemporary modeling approaches

The two-population SSN model, positioned between large-scale spiking networks and multi-population rate networks, offers mathematical tractability and could be extended to incorporate different (sub-)types of neurons. Since plasticity effects or individual spike timing were not relevant for our goal of inferring connectivity, we used a mean-field approximation model (Renart et al., 2003) similar to previous versions of the SSN (Ahmadian et al., 2013), rather than a detailed spiking network model (Maass, 1997), to describe cortical activity. This approach allowed us to work with tractable differential equations that only permitted the use of contrast-invariance to simplify the model. As the roles of inhibitory interneuron subtypes are currently being dissected meticulously (de Vries et al., 2020; Bugeon et al., 2021), other contemporary studies have focused on extending two-population models to include multiple interneuron populations (Litwin-Kumar et al., 2016; Lee et al., 2017; Bryson et al., 2021; Romero-Sosa et al., 2021). These approaches have typically used Wilson-Cowan-type models (Wilson and Cowan, 1972), which do not include the power-law activation function of the SSN, making them easier to handle mathematically but also omitting known properties like contrast-invariance. One recent study however, has used an SSN model with multiple inhibitory populations to predict perturbation effects of optogenetic stimulation (Palmigiano et al., 2020), paving the way to expanding the SSN to incorporate multiple subpopulations. In principle, the mathematics of our mapping procedure should be expandable to include more subtypes by applying the separation of contrast-dependent and orientation-dependent factors to the additional terms. However, this would require contrast-invariance in the responses of all subtypes, the presence of which is currently debated (Li et al., 2012) and would have to be verified in the process. What complicates this assessment is that contrast-invariance would not only have to hold for individual neurons of the populations but also for each population as a whole, which relies heavily on the distribution of contrast response functions between the neurons (see below). As more populations get added to the model and their properties become more distinct, this assumption becomes increasingly unlikely. Considering the difficulties around incorporating contrast-invariant subpopulations, one could try to relax this condition and only require invariance for specific subpopulations. The inference procedure developed in manuscript one would be rendered impossible in this case, but additional restrictions, like requiring weak orientation tuning or approximating a non-invariant subpopulation as a linear combination of invariant parts, could still permit connectivity inference. Additionally, one could mathematically derive or numerically test which connectivity regimes in a multi-

population model would result in contrast-invariant or contrast-dependent responses across subpopulations, thereby verifying if biological connectivity regimes rely on contrast-invariance in multiple populations.

Contrast-invariance of the population response depends on the sampled distribution of the response parameters of individual neurons. We combined individual, contrast-invariant responses into a population response by aligning the responses to the preferred orientation and then averaging the tuning curves. This procedure hinges on two assumptions: first, orientation tuning parameters and contrast response parameters need to be uncorrelated. If orientation tuning width and semi-saturation contrast of individual neurons were correlated, the average width would change between the most strongly represented saturation and the other parts of the distribution, breaking contrast-invariance. Recordings with multi-electrode arrays have found tuning width and semi-saturation contrast to be uncorrelated (Busse et al., 2009), matching our findings of a contrast-invariant population response. Second, contrast response parameters need to be uniformly distributed. Even without correlations between tuning width and semi-saturation contrast, over-representation of a particular semi-saturation contrast will break contrast-invariance because inactive neurons contribute flat tuning to the mean. As visible in the cumulative distribution of contrast sensitivity, our data met this criterion. Since evidence remains for violations of contrast-invariance in cortex (Li et al., 2012; Tring and Ringach, 2018; Lee et al., 2019), it is important to keep these confounders in mind.

4.2 Outlook

How the brain is affected by arousal and behavioral state is currently under intense discussion in the field (Froudarakis et al., 2019; McCormick et al., 2020). In the visual system, locomotion modulates the gain of firing in dLGN and V1 (Niell and Stryker, 2010; Erisken et al., 2014) and affects oscillatory activity (Nestvogel and McCormick, 2021). Circuit investigations have demonstrated the importance of interneuron projections in these phenomena (Polack et al., 2013; Dipoppa et al., 2018), raising the question if the underlying connectivity between excitatory and inhibitory populations also depends on behavioral state. Similarly, inhibitory neurons and behavioral state have been shown to be strongly related to suppression in the visual pathway (Augustinaite and Kuhn, 2020; Niell and Stryker, 2010), suggesting that SbC neurons in our data set might also be dependent on behavioral state. Since locomotion and arousal information are collected in our paradigm, splitting connectivity and SbC analyses into locomotion and quiescence states would be a natural

extension of our investigations.

Another source of untapped potential in the currently available data is temporal information. There is evidence that contrast responses, spatial frequency tuning and orientation selectivity are non-stationary over multiple time scales (Ringach et al., 1997; Mazer et al., 2002; Hu et al., 2011), which has been argued to also affect contrast-invariance (Nowak and Barone, 2009). However, a consensus on this matter has not been reached, with evidence for stable tuning (Gillespie et al., 2001) and evidence for sharpening over time (Nowak and Barone, 2009) both being present. Our stimulus protocol offers the possibility to contribute to this discussion by analyzing the temporal dynamics depending on multiple stimulus parameters. Additionally, synaptic plasticity has been shown to change connectivity on short time scales while depending strongly on different neuron types (Tan et al., 2008; Beierlein et al., 2003; Mongillo et al., 2018). If contrast or orientation responses in our data were time-dependent, inferred connectivity would likely follow, offering yet another perspective on network dynamics. Time-resolved analysis of our inference paradigm thus represents a promising step for future work.

Finally, probing the effects of activity perturbations on inferred cortical connectivity, as well as on different neuronal response types, could reveal details of the underlying circuits. Due to its sweeping implications for modeling and conceptualizing the cortex, the cortical operating regime and the underlying connectivity are a focal point of the current discussion in neuroscience (Sadeh and Clopath, 2020b; Sanzeni et al., 2020). Importantly, recent experimental tools have allowed *in vivo* perturbation of cortical activity to be established as a new gold standard for testing circuit models (Sadeh and Clopath, 2020a). As the SSN predicts severely altered response states for different connectivity regimes (Kraynyukova and Tchumatchenko, 2018), manipulating cortical responses is likely to strongly affect inferred connectivity. Investigating whether or not these effects work to re-balance cortical firing (Moore et al., 2018) would greatly benefit our understanding of cortex. As SbC neurons might be closely related to inhibitory firing (Millman et al., 2020), optogenetic tagging or perturbations of inhibitory firing could allow direct testing of the overlap between neuron types and response types and thus give insights about the role of suppression in the cortical circuit. Combined, these prospects make the investigation of perturbation effects a top priority for the future.

With such promising avenues to continue to impact the discussion about connectivity, inhibition, and cortical function in visual neuroscience and beyond, I hope the work presented in this thesis finds its way to strengthening the foundation that researchers can stand on to see further than we can imagine today.

Bibliography

Hillel Adesnik. Synaptic Mechanisms of Feature Coding in the Visual Cortex of Awake Mice. *Neuron*, 95(5):1147–1159.e4, 2017. doi: 10.1016/j.neuron.2017.08.014.

Yashar Ahmadian, Daniel B. Rubin, and Kenneth D. Miller. Analysis of the Stabilized Supralinear Network. *Neural Computation*, 25(8):1994–2037, 2013. doi: 10.1162/NECO_a_00472.

Henry J. Alitto and W. Martin Usrey. Influence of Contrast on Orientation and Temporal Frequency Tuning in Ferret Primary Visual Cortex. *Journal of Neurophysiology*, 91(6):2797–2808, 2004. doi: 10.1152/jn.00943.2003.

Jeffrey S. Anderson, Ilan Lampl, Deda C. Gillespie, and David Ferster. The Contribution of Noise to Contrast Invariance of Orientation Tuning in Cat Visual Cortex. *Science*, 290(5498):1968–1972, 2000. doi: 10.1126/science.290.5498.1968.

Bassam V. Atallah, William Bruns, Matteo Carandini, and Massimo Scanziani. Parvalbumin-Expressing Interneurons Linearly Transform Cortical Responses to Visual Stimuli. *Neuron*, 73(1):159–170, 2012. doi: 10.1016/j.neuron.2011.12.013.

Sigita Augustinaite and Bernd Kuhn. Complementary Ca^{2+} Activity of Sensory Activated and Suppressed Layer 6 Corticothalamic Neurons Reflects Behavioral State. *Current Biology*, 30(20):3945–3960.e5, 2020. doi: 10.1016/j.cub.2020.07.069.

Stephen A. Baccus and Markus Meister. Fast and slow contrast adaptation in retinal circuitry. *Neuron*, 36(5):909–919, 2002. doi: 10.1016/S0896-6273(02)01050-4.

Tom Baden, Philipp Berens, Katrin Franke, Miroslav Román Rosón, Matthias Bethge, and Thomas Euler. The functional diversity of retinal ganglion cells in the mouse. *Nature*, 529(7586):345–350, 2016. doi: 10.1038/nature16468.

Kenneth G. Baimbridge, Marco R. Celio, and John H. Rogers. Calcium-binding proteins in the nervous system. *Trends in Neurosciences*, 15(8):303–308, 1992. doi: 10.1016/0166-2236(92)90081-I.

Cody Baker, Emmanouil Froudarakis, Dimitri Yatsenko, Andreas S. Tolias, and Robert Rosenbaum. Inference of synaptic connectivity and external variability in neural microcircuits. *Journal of Computational Neuroscience*, 48(2):123–147, 2020. doi: 10.1007/s10827-020-00739-4.

Cornelia I. Bargmann and Eve Marder. From the connectome to brain function. *Nature Methods*, 10(6):483–490, 2013. doi: 10.1038/nmeth.2451.

Michael Beierlein, Jay R. Gibson, and Barry W. Connors. Two Dynamically Distinct Inhibitory Networks in Layer 4 of the Neocortex. *Journal of Neurophysiology*, 90(5):2987–3000, 2003. doi: 10.1152/jn.00283.2003.

Tom Binzegger. A Quantitative Map of the Circuit of Cat Primary Visual Cortex. *Journal of Neuroscience*, 24(39), 2004. doi: 10.1523/JNEUROSCI.1400-04.2004.

Davi D. Bock, Wei Chung Allen Lee, Aaron M. Kerlin, Mark L. Andermann, Greg Hood, Arthur W. Wetzel, Sergey Yurgenson, Edward R. Soucy, Hyon Suk Kim, and R. Clay Reid. Network anatomy and in vivo physiology of visual cortical neurons. *Nature*, 471(7337):177–184, 2011. doi: 10.1038/nature09802.

Edward S. Boyden, Feng Zhang, Ernst Bamberg, Georg Nagel, and Karl Deisseroth. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*, 8(9):1263–1268, 2005. ISSN 10976256. doi: 10.1038/nn1525.

Alexander Bryson, Samuel F. Berkovic, Steven Petrou, and David B. Grayden. State transitions through inhibitory interneurons in a cortical network model. *PLoS Computational Biology*, 17(10):1–29, 2021. doi: 10.1371/journal.pcbi.1009521.

Stephane Bugeon, Joshua Duffield, Mario Dipoppa, Isabelle Prankerd, Anne Ritoux, Dimitris Nicoloutsopoulos, David Orme, Maxwell Shinn, Han Peng, Hamish Forrest, Aiste Viduolyte, Charu Bai Reddy, Yoh Isogai, and Kenneth D Harris. A transcriptomic axis predicts state modulation of cortical interneurons. *bioRxiv*, 2021. doi: 10.1101/2021.10.24.465600.

Laura Busse, Alex R. Wade, and Matteo Carandini. Representation of Concurrent Stimuli by Population Activity in Visual Cortex. *Neuron*, 64(6):931–942, 2009. doi: 10.1016/j.neuron.2009.11.004.

Matteo Carandini. Receptive fields and suppressive fields in the early visual system. *The Cognitive Neurosciences*, 3:313–326, 2004.

Matteo Carandini and David J. Heeger. Summation and division by neurons in primate visual cortex. *Science*, 264(5163):1333–1336, 1994. doi: 10.1126/science.8191289.

Matteo Carandini and David J. Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51–62, 2012. doi: 10.1038/nrn3136.

Matteo Carandini, David J. Heeger, and J. Anthony Movshon. Linearity and normalization in simple cells of the macaque primary visual cortex. *Journal of Neuroscience*, 17(21):8621–8644, 1997. doi: 10.1523/JNEUROSCI.17-21-08621.1997.

Barbara Chapman, Kathleen R. Zahs, and Michael P. Stryker. Relation of cortical cell orientation selectivity to alignment of receptive fields of the geniculocortical afferents that arborize within a single orientation column in ferret visual cortex. *The Journal of Neuroscience*, 11(5):1347–58, 1991. doi: doi.org/10.1523/JNEUROSCI.11-05-01347.1991.

Mark M. Churchland, Byron M. Yu, John P. Cunningham, Leo P. Sugrue, Marlene R. Cohen, Greg S. Corrado, William T. Newsome, Andrew M. Clark, Paymon Hosseini, Benjamin B. Scott, David C. Bradley, Matthew A. Smith, Adam Kohn, J. Anthony Movshon, Katherine M. Armstrong, Tirin Moore, Steve W. Chang, Lawrence H. Snyder, Stephen G. Lisberger, Nicholas J. Priebe, Ian M. Finn, David Ferster, Stephen I. Ryu, Gopal Santhanam, Maneesh Sahani, and Krishna V. Shenoy. Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nature Neuroscience*, 13(3):369–378, 2010. doi: 10.1038/nn.2501.

Lee Cossell, Maria Florencia Iacaruso, Dylan R. Muir, Rachael Houlton, Elie N. Sader, Ho Ko, Sonja B. Hofer, and Thomas D. Mrsic-Flogel. Functional organization of excitatory synaptic strength in primary visual cortex. *Nature*, 518 (7539):399–403, 2015. doi: 10.1038/nature14182.

Scott J. Cruikshank, Timothy J. Lewis, and Barry W. Connors. Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex. *Nature Neuroscience*, 10(4):462–468, 2007. doi: 10.1038/nn1861.

Clayton E. Curtis and Daeyeol Lee. Beyond working memory: the role of persistent activity in decision making. *Trends in Cognitive Sciences*, 14(5):216–222, 2010. doi: 10.1016/j.tics.2010.03.006.

Saskia E.J. de Vries, Jerome A. Lecoq, Michael A. Buice, Peter A. Groblewski, Gabriel K. Ocker, Michael Oliver, David Feng, Nicholas Cain, Peter Ledochowitsch, Daniel Millman, Kate Roll, Marina Garrett, Tom Keenan, Leonard Kuan, Stefan Mihalas, Shawn Olsen, Carol Thompson, Wayne Wakeman, Jack Waters, Derric Williams, Chris Barber, Nathan Berbesque, Brandon Blanchard, Nicholas Bowles, Shiella D. Caldejon, Linzy Casal, Andrew Cho, Sissy Cross, Chinh Dang, Tim Dolbeare, Melise Edwards, John Galbraith, Nathalie Gaudreault, Terri L. Gilbert, Fiona Griffin, Perry Hargrave, Robert Howard, Lawrence Huang, Sean Jewell, Nika Keller, Ulf Knoblich, Josh D. Larkin, Rachael Larsen, Chris Lau, Eric Lee, Felix Lee, Arielle Leon, Lu Li, Fuhui Long, Jennifer Luviano, Kyla Mace, Thuyanh Nguyen, Jed Perkins, Miranda Robertson, Sam Seid, Eric Shea-Brown, Jianghong Shi, Nathan Sjoquist, Cliff Slaughterbeck, David Sullivan, Ryan Valenza, Casey White, Ali Williford, Daniela M. Witten, Jun Zhuang, Hongkui Zeng, Colin Farrell, Lydia Ng, Amy Bernard, John W. Phillips, R. Clay Reid, and Christof Koch. A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature Neuroscience*, 23(1): 138–151, 2020. doi: 10.1038/s41593-019-0550-9.

Javier Defelipe, Pedro L. López-Cruz, Ruth Benavides-Piccione, Concha Bielza, Pedro Larrañaga, Stewart Anderson, Andreas Burkhalter, Bruno Cauli, Alfonso Fairén, Dirk Feldmeyer, Gord Fishell, David Fitzpatrick, Tamás F. Freund, Guillermo González-Burgos, Shaul Hestrin, Sean Hill, Patrick R. Hof, Josh Huang, Edward G. Jones, Yasuo Kawaguchi, Zoltán Kisvárday, Yoshiyuki Kubota, David A. Lewis, Oscar Marín, Henry Markram, Chris J. McBain, Hanno S. Meyer, Hannah Monyer, Sacha B. Nelson, Kathleen Rockland, Jean Rossier, John L.R. Rubenstein, Bernardo Rudy, Massimo Scanziani, Gordon M. Shepherd, Chet C. Sherwood, Jochen F. Staiger, Gábor Tamás, Alex Thomson, Yun Weng, Rafael Yuste, and Giorgio A. Ascoli. New insights into the classification and nomenclature of cortical GABAergic interneurons. *Nature Reviews Neuroscience*, 2013. doi: 10.1038/nrn3444.

Karl Deisseroth, Guoping Feng, Ania K. Majewska, Gero Miesenböck, Alice Ting, and Mark J. Schnitzer. Next-generation optical technologies for illuminating genetically targeted brain circuits. *Journal of Neuroscience*, 26(41):10380–10386, 2006. ISSN 02706474. doi: 10.1523/JNEUROSCI.3863-06.2006.

Mario Dipoppa, Adam Ranson, Michael Krumin, Marius Pachitariu, Matteo Carandini, and Kenneth D. Harris. Vision and Locomotion Shape the Interactions between Neuron Types in Mouse Visual Cortex. *Neuron*, 98(3):602–615.e8, 2018. doi: 10.1016/j.neuron.2018.03.037.

Rodney J. Douglas, Kevan A.C. Martin, and David Whitteridge. Selective responses of visual cortical cells do not depend on shunting inhibition. *Nature*, 332(6165): 642–644, 1988. doi: 10.1038/332642a0.

Daniel Durstewitz, Jeremy K. Seamans, and Terrence J. Sejnowski. Neurocomputational Models of Working Memory. *Nature Neuroscience*, 3(11s):1184–1191, 2000. doi: 10.1038/81460.

Franics A. Edwards, Arthur Konnerth, Bert Sakmann, and Tomoyuki Takahashi. A thin slice preparation for patch clamp recordings from neurones of the mammalian central nervous system. *Pflügers Archiv European Journal of Physiology*, 414(5):600–612, 1989. doi: 10.1007/BF00580998.

Sinem Erisken, Agne Vaiceliunaite, Ovidiu Jurjut, Matilde Fiorini, Steffen Katzner, and Laura Busse. Effects of locomotion extend throughout the mouse early visual system. *Current Biology*, 24(24):2899–2907, 2014. doi: 10.1016/j.cub.2014.10.045.

Ian M. Finn, Nicholas J. Priebe, and David Ferster. The Emergence of Contrast-Invariant Orientation Tuning in Simple Cells of Cat Visual Cortex. *Neuron*, 54 (1):137–152, 2007. doi: 10.1016/j.neuron.2007.02.029.

David Fitzpatrick. Seeing beyond the receptive field in primary visual cortex. *Current Opinion in Neurobiology*, 10(4):438–443, 2000. doi: 10.1016/s0959-4388(00)00113-6.

Emmanouil Froudarakis, Paul G. Fahey, Jacob Reimer, Stelios M. Smirnakis, Edward J. Tehovnik, and Andreas S. Tolias. The Visual Cortex in Context. *Annual Review of Vision Science*, 5:317–339, 2019. doi: 10.1146/annurev-vision-091517-034407.

Yu Fu, Jason M. Tucciarone, J. Sebastian Espinosa, Nengyin Sheng, Daniel P. Darcy, Roger A. Nicoll, Z. Josh Huang, and Michael P. Stryker. A cortical circuit for gain control by behavioral state. *Cell*, 156(6):1139–1152, 2014. doi: 10.1016/j.cell.2014.01.050.

Deda. C. Gillespie, Ilan Lampl, Jeffrey S. Anderson, and David Ferster. Dynamics of the orientation-tuned membrane potential response in cat primary visual cortex. *Nature Neuroscience*, 4(10):1014–1019, 2001. doi: 10.1038/nn731.

Charles M. Gray, Pedro E. Maldonado, Mathew Wilson, and Bruce McNaughton. Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *Journal of Neuroscience Methods*, 63(1-2):43–54, 1995.

Matthew S. Grubb and Ian D. Thompson. Visual response properties of burst and tonic firing in the mouse dorsal lateral geniculate nucleus. *Journal of Neurophysiology*, 93:3224–3247, 2005. doi: 10.1152/jn.00445.2004.

David Hansel and Carl van Vreeswijk. How Noise Contributes to Contrast Invariance of Orientation Tuning in Cat Visual Cortex. *Journal of Neuroscience*, 22 (12):5118–5128, 2002. doi: 10.1523/jneurosci.22-12-05118.2002.

Kenneth D. Harris and Thomas D. Mrsic-Flogel. Cortical connectivity and sensory coding. *Nature*, 503(7474):51–58, 2013. doi: 10.1038/nature12654.

Guillaume Hennequin, Yashar Ahmadian, Daniel B. Rubin, Máté Lengyel, and Kenneth D. Miller. The Dynamical Regime of Sensory Cortex: Stable Dynamics around a Single Stimulus-Tuned Attractor Account for Patterns of Noise Variability. *Neuron*, 98(4):846–860.e5, 2018. doi: 10.1016/j.neuron.2018.04.017.

Sonja B. Hofer, Ho Ko, Bruno Pichler, Joshua Vogelstein, Hana Ros, Hongkui Zeng, Ed Lein, Nicholas A. Lesica, and Thomas D. Mrsic-Flogel. Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nature Neuroscience*, 14(8):1045–1052, 2011. doi: 10.1038/nn.2876.

Hua Hu, Jian Gan, and Peter Jonas. Fast-spiking, parvalbumin + GABAergic interneurons: From cellular design to microcircuit function. *Science*, 345(6196), 2014. doi: 10.1126/science.1255263.

Ming Hu, Yong Wang, and Yi Wang. Rapid dynamics of contrast responses in the cat primary visual cortex. *PLoS ONE*, 6(10), 2011. doi: 10.1371/journal.pone.0025410.

David H. Hubel. Single unit activity in lateral geniculate body and optic tract of unrestrained cats. *The Journal of Physiology*, 150(1):91–104, 1960. doi: 10.1113/jphysiol.1960.sp006375.

David H. Hubel and Thorsten N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160 (1):106–154, 1962. doi: 10.1113/jphysiol.1962.sp006837.

Eyal Hulata, Ronen Segev, and Eshel Ben-Jacob. A method for spike sorting and detection based on wavelet packets and shannon's mutual information. *Journal of Neuroscience Methods*, 117(1):1–12, 2002.

Pablo Jercog, Thomas Rogerson, and Mark J. Schnitzer. Large-scale fluorescence calcium-imaging methods for studies of long-term memory in behaving mammals. *Cold Spring Harbor Perspectives in Biology*, 8(5), 2016. doi: 10.1101/cshperspect.a021824.

Xiaoxuan Jia, Joshua H. Siegle, Corbett Bennett, Samuel D. Gale, Daniel J. Denman, Christof Koch, and Shawn R. Olsen. High-density extracellular probes reveal dendritic backpropagation and facilitate neuron classification. *Journal of Neurophysiology*, 121(5):1831–1847, 2019. doi: 10.1152/jn.00680.2018.

Chris W.D. Jürgens, Karen A. Bell, A. Rory McQuiston, and William Guido. Optogenetic Stimulation of the Corticothalamic Pathway Affects Relay Cells and GABAergic Neurons Differently in the Mouse Visual Thalamus. *PLoS ONE*, 7 (9), 2012. doi: 10.1371/journal.pone.0045717.

Eric R. Kandel, James H. Schwartz, Thomas M. Jessell, Steven Siegelbaum, A. James Hudspeth, and Sarah Mack. *Principles of neural science*, volume 4. McGraw-hill New York, 2000.

Mahesh M. Karnani, Jesse Jackson, Inbal Ayzenshtat, Jason Tucciarone, Kasra Manoocheri, William G. Snider, and Rafael Yuste. Cooperative Subnetworks of Molecularly Similar Interneurons in Mouse Neocortex. *Neuron*, 90(1):86–100, 2016. doi: 10.1016/j.neuron.2016.02.037.

Yasuo Kawaguchi, Hironobu Katsumaru, Toshio Kosaka, Claus W. Heizmann, and Kiyoshi Hama. Fast spiking cells in rat hippocampus (CA1 region) contain the calcium-binding protein parvalbumin. *Brain Research*, 416(2):369–374, 1987. doi: 10.1016/0006-8993(87)90921-8.

Ho Ko, Sonja B. Hofer, Bruno Pichler, Katherine A. Buchanan, P. Jesper Sjöström, and Thomas D. Mrsic-Flogel. Functional specificity of local synaptic connections in neocortical networks. *Nature*, 473(7345):87–91, 2011. doi: 10.1038/nature09880.

Nataliya Kraynyukova and Tatjana Tchumatchenko. Stabilized supralinear network can give rise to bistable, oscillatory, and persistent activity. *Proceedings of the National Academy of Sciences*, 115(13):3464–3469, 2018. doi: 10.1073/pnas.1700080115.

Stephen W. Kuffler. Discharge Patterns and Functional Organization of Mammalian Retina. *Journal of Neurophysiology*, 16(1):37–68, 1953. doi: 10.1152/jn.1953.16.1.37.

Sandra J. Kuhlman and Z. Josh Huang. High-resolution labeling and functional manipulation of specific neuron types in mouse brain by Cre-activated viral gene expression. *PLoS ONE*, 3(4), 2008. ISSN 19326203. doi: 10.1371/journal.pone.0002005.

Josef Ladenbauer, Sam McKenzie, Daniel Fine English, Olivier Hagens, and Srdjan Ostojic. Inferring and validating mechanistic models of neural microcircuits based on spike-train data. *Nature Communications*, 10(1):1–17, 2019. doi: 10.1038/s41467-019-12572-0.

Jung H. Lee, Christof Koch, and Stefan Mihalas. A Computational Analysis of the Function of Three Inhibitory Cell Types in Contextual Visual Processing. *Frontiers in Computational Neuroscience*, 11(28):1–15, 2017. doi: 10.3389/fncom.2017.00028.

Sangkyun Lee, Jiyoung Park, and Stelios M. Smirnakis. Internal gain modulations, but not changes in stimulus contrast, preserve the neural code. *The Journal of Neuroscience*, 39(9):2012–18, 2019. doi: 10.1523/jneurosci.2012-18.2019.

Johannes J. Letzkus, Björn M Kampa, and Greg J. Stuart. Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *Journal of Neuroscience*, 26(41):10420–10429, 2006. doi: 10.1523/JNEUROSCI.2650-06.2006.

William R. Levick. Receptive fields and trigger features of ganglion cells in the visual streak of the rabbit's retina. *The Journal of Physiology*, 188(3):285–307, 1967. doi: 10.1113/jphysiol.1967.sp008140.

Ya-tang Li, Weng-pei Ma, Ling-yun Li, Leena A. Ibrahim, Sheng-zhi Wang, and Huizhong Whit Tao. Broadening of Inhibitory Tuning Underlies Contrast-Dependent Sharpening of Orientation Selectivity in Mouse Visual Cortex. *Journal of Neuroscience*, 32(46):16466–16477, 2012. doi: 10.1523/JNEUROSCI.3221-12.2012.

Anthony D. Lien and Massimo Scanziani. Tuned thalamic excitation is amplified by visual cortical circuits. *Nature Neuroscience*, 16(9):1315–1323, 2013. doi: 10.1038/nn.3488.

Ashok Litwin-Kumar, Robert Rosenbaum, and Brent Doiron. Inhibitory stabilization and visual coding in cortical circuits with multiple interneuron subtypes. *Journal of Neurophysiology*, 115(3):1399–1409, 2016. doi: 10.1152/jn.00732.2015.

Liu D. Liu, Kenneth D. Miller, and Christopher C. Pack. A Unifying Motif for Spatial and Directional Surround Suppression. *The Journal of Neuroscience*, 38 (4):989–999, 2018. doi: 10.1523/JNEUROSCI.2386-17.2017.

Rodolfo R. Llinás. The contribution of Santiago Ramon y Cajal to functional neuroscience. *Nature Reviews Neuroscience*, 4(1):77–80, 2003. doi: 10.1038/nrn1011.

Wolfgang Maass. Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, 10(9):1659–1671, 1997. doi: https://doi.org/10.1016/S0893-6080(97)00011-7.

Ildefons Magrans de Abril, Junichiro Yoshimoto, and Kenji Doya. Connectivity inference from neural recording data: Challenges, mathematical bases and research directions. *Neural Networks*, 102:120–137, 2018. doi: 10.1016/j.neunet.2018.02.016.

Emily M. Martersteck, Karla E. Hirokawa, Mariah Evarts, Amy Bernard, Xin Duan, Yang Li, Lydia Ng, Seung W. Oh, Benjamin Ouellette, Joshua J. Royall, Michelle Stoecklin, Quanxin Wang, Hongkui Zeng, Joshua R. Sanes, and Julie A. Harris. Diverse Central Projection Patterns of Retinal Ganglion Cells. *Cell Reports*, 18 (8):2058–2072, 2017. doi: 10.1016/j.celrep.2017.01.075.

James A. Mazer, William E. Vinje, Josh McDermott, Peter H. Schiller, and Jack L. Gallant. Spatial frequency and orientation tuning dynamics in area V1. *Proceedings of the National Academy of Sciences of the United States of America*, 99(3):1645–1650, 2002. doi: 10.1073/pnas.022638499.

David A. McCormick and Diego Contreras. On The Cellular and Network Bases of Epileptic Seizures. *Annual Review of Physiology*, 63(1):815–846, 2001. doi: 10.1146/annurev.physiol.63.1.815.

David A. McCormick, Dennis B. Nestvogel, and Biyu J. He. Neuromodulation of Brain State and Behavior. *Annual Review of Neuroscience*, 43:391–415, 2020. doi: 10.1146/annurev-neuro-100219-105424.

Richard Miles and Jean-Christophe Poncer. Paired recordings from neurones. *Current Opinion in Neurobiology*, 6(3):387–394, 1996. doi: 10.1016/S0959-4388(96)80124-3.

Kenneth D. Miller and Todd W. Troyer. Neural noise can explain expansive, power-law nonlinearities in neural response functions. *Journal of Neurophysiology*, 87 (2):653–659, 2002. doi: 10.1152/jn.00425.2001.

Daniel J. Millman, Gabriel Koch Ocker, Shiella Caldejon, India Kato, Josh D. Larkin, Eric Kenji Lee, Jennifer Luviano, Chelsea Nayan, Thuyanh V. Nguyen, Kat North, Sam Seid, Cassandra White, Jerome Lecoq, Clay Reid, Michael A. Buice, and Saskia E.J. de Vries. VIP interneurons in mouse primary visual cortex selectively enhance responses to weak but specific stimuli. *eLife*, 9:1–22, 2020. doi: 10.7554/eLife.55130.

Gianluigi Mongillo, Simon Rumpel, and Yonatan Loewenstein. Inhibitory connectivity defines the realm of excitatory plasticity. *Nature Neuroscience*, 21(10): 1463–1470, 2018. doi: 10.1038/s41593-018-0226-x.

Alexandra K. Moore, Aldis P. Weible, Timothy S. Balmer, Laurence O. Trussell, and Michael Wehr. Rapid Rebalancing of Excitation and Inhibition by Cortical Circuitry. *Neuron*, 97(6):1341–1355.e6, 2018. doi: 10.1016/j.neuron.2018.01.045.

Joshua L. Morgan and Jeff W. Lichtman. Why not connectomics? *Nature Methods*, 10(6):494–500, 2013. doi: 10.1038/nmeth.2480.

Nicolás A. Morgenstern, Jacques Bourg, and Leopoldo Petreanu. Multilaminar networks of cortical neurons integrate common inputs from sensory thalamus. *Nature Neuroscience*, 19(8):1034–1040, 2016. doi: 10.1038/nn.4339.

Dennis B. Nestvogel and David A. McCormick. Visual thalamocortical mechanisms of waking state-dependent activity and alpha oscillations. *Neuron*, 2021. doi: https://doi.org/10.1016/j.neuron.2021.10.005.

Jilda S. Nettleton and William J. Spain. Linear to supralinear summation of AMPA-mediated EPSPs in neocortical pyramidal neurons. *Journal of Neurophysiology*, 83(6):3310–3322, 2000. doi: 10.1152/jn.2000.83.6.3310.

Cristopher M. Niell and Michael P. Stryker. Highly Selective Receptive Fields in Mouse Visual Cortex. *Journal of Neuroscience*, 28(30):7520–7536, 2008. doi: 10.1523/jneurosci.0623-08.2008.

Cristopher M. Niell and Michael P. Stryker. Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. *Neuron*, 65(4):472–479, 2010. doi: 10.1016/j.neuron.2010.01.033.

Lionel G. Nowak and Pascal Barone. Contrast adaptation contributes to contrast-invariance of orientation tuning of primate V1 cells. *PLoS ONE*, 4(3), 2009. doi: 10.1371/journal.pone.0004781.

Dina Obeid and Kenneth D Miller. Supralinear Stabilized Network : Model of layer 2/3 of the Primary Visual Cortex. *bioRxiv*, 2019. doi: 10.1101/2020.12.30.424892.

Bruno A. Olshausen and David J. Field. How close are we to understanding V1? *Neural Computation*, 17(8):1665–1699, 2005. doi: 10.1162/0899766054026639.

Adam M. Packer and Rafael Yuste. Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: A canonical microcircuit for inhibition? *Journal of Neuroscience*, 31(37):13260–13271, 2011. doi: 10.1523/JNEUROSCI.3131-11.2011.

Agostina Palmigiano, Francesco Fumarola, Daniel P. Mossing, Nataliya Kraynyukova, Hillel Adesnik, and Kenneth D. Miller. Structure and variability of optogenetic responses identify the operating regime of cortex. *bioRxiv*, 2020. doi: 10.1101/2020.11.11.378729.

Jonathan W. Peirce. The potential importance of saturating and supersaturating contrast response functions in visual cortex. *Journal of Vision*, 7(6):1–10, 2007. doi: 10.1167/7.6.13.

Erez Persi, David Hansel, Lionel Nowak, Pascal Barone, and Carl van Vreeswijk. Power-law input-output transfer functions explain the contrast-response and tuning properties of neurons in visual cortex. *PLoS Computational Biology*, 7(2), 2011. doi: 10.1371/journal.pcbi.1001078.

Carsten K. Pfeffer, Mingshan Xue, Miao He, Z. Josh Huang, and Massimo Scanziani. Inhibition of inhibition in visual cortex: The logic of connections between molecularly distinct interneurons. *Nature Neuroscience*, 16(8):1068–1076, 2013. doi: 10.1038/nn.3446.

Denise M. Piscopo, Rana N. El-Danaf, Andrew D. Huberman, and Cristopher M. Niell. Diverse Visual Features Encoded in Mouse Lateral Geniculate Nucleus. *Journal of Neuroscience*, 33(11):4642–4656, 2013. doi: 10.1523/JNEUROSCI.5187-12.2013.

Pierre-Olivier Polack, Jonathan Friedman, and Peyman Golshani. Cellular mechanisms of brain state-dependent gain modulation in visual cortex. *Nature Neuroscience*, 16(9):1331–1339, 2013. doi: 10.1038/nn.3464.

Nicholas J. Priebe, Ferenc Mechler, Matteo Carandini, and David Ferster. The contribution of spike threshold to the dichotomy of cortical simple and complex cells. *Nature Neuroscience*, 7(10):1113–1122, 2004. doi: 10.1038/nn1310.

Clay R. Reid and Jose Manuel Alonso. Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378(6554):281–284, 1995. doi: 10.1038/378281a0.

Alfonso Renart, Nicolas Brunel, and Xiao-Jing Wang. Mean-Field Theory of Irregularly Spiking Neuronal Populations and Working Memory in Recurrent Cortical Networks. In Jianfeng Feng, editor, *Computational neuroscience: A comprehensive approach*, chapter 15, pages 431–490. Chapman & Hall/CRC, London, 1st edition, 2003. doi: 10.1201/9780203494462.ch15.

Dario L. Ringach, Michael J. Hawken, and Robert Shapley. Dynamics of orientation tuning in macaque primary visual cortex. *Nature*, 387(6630):281–284, 1997. doi: 10.1038/387281a0.

Robert W. Rodieck. Receptive Fields in the Cat Retina: A New Type. *Science*, 157 (3784):90–92, 1967. doi: 10.1126/science.157.3784.90.

Juan L. Romero-Sosa, Helen Motanis, and Dean V. Buonomano. Differential excitability of PV and SST neurons results in distinct functional roles in inhibition stabilization of Up-states. *The Journal of Neuroscience*, 41(34):JN–RM–2830–20, 2021. doi: 10.1523/jneurosci.2830-20.2021.

Paolo M. Rossini, Riccardo Di Iorio, Marina Bentivoglio, Giuseppe Bertini, Florinda Ferreri, Christian Gerloff, Risto J. Ilmoniemi, Francesca Miraglia,

Michael A. Nitsche, Franco Pestilli, Mario Rosanova, Yuichiro Shirota, Chiara Tesoriero, Yoshikazu Ugawa, Fabrizio Vecchio, Ulf Ziemann, and Mark Hallett. Methods for analysis of brain connectivity: An IFCN-sponsored review. *Clinical Neurophysiology*, 130(10):1833–1858, 2019. doi: 10.1016/j.clinph.2019.06.006.

Daniel B. Rubin, Stephen D. Van Hooser, and Kenneth D. Miller. The Stabilized Supralinear Network: A Unifying Circuit Motif Underlying Multi-Input Integration in Sensory Cortex. *Neuron*, 85(2):402–417, 2015. doi: 10.1016/j.neuron.2014.12.026.

Bernardo Rudy, Gordon Fishell, SooHyun Lee, and Jens Hjerling-Leffler. Three groups of interneurons account for nearly 100% of neocortical GABAergic neurons. *Developmental Neurobiology*, 71(1):45–61, 2011. doi: 10.1002/dneu.20853.

Sadra Sadeh and Claudia Clopath. Patterned perturbation of inhibition can reveal the dynamical structure of neural processing. *eLife*, 9:1–29, 2020a. doi: 10.7554/eLife.52757.

Sadra Sadeh and Claudia Clopath. Inhibitory stabilization and cortical computation. *Nature Reviews Neuroscience*, 2020b. doi: 10.1038/s41583-020-00390-z.

Sadra Sadeh and Stefan Rotter. Orientation Selectivity in Inhibition-Dominated Networks of Spiking Neurons: Effect of Single Neuron Properties and Network Dynamics. *PLoS Computational Biology*, 11(1), 2015. doi: 10.1371/journal.pcbi.1004045.

Alessandro Sanzeni, Bradley Akitake, Hannah C. Goldbach, Caitlin E. Leedy, Nicolas Brunel, and Mark H. Histed. Inhibition stabilization is a widespread property of cortical networks. *eLife*, 9:1–39, 2020. doi: 10.7554/eLife.54875.

Shreya Saxena and John P. Cunningham. Towards the neural population doctrine. *Current Opinion in Neurobiology*, 55:103–111, 2019. doi: 10.1016/j.conb.2019.02.002.

William R. Schafer. The Worm Connectome: Back to the Future. *Trends in Neurosciences*, 41(11):763–765, 2018. doi: 10.1016/j.tins.2018.09.002.

Benjamin Scholl, Kenneth W. Latimer, and Nicholas J. Priebe. A retinal source of spatial contrast gain control. *Journal of Neuroscience*, 32(29):9824–9830, 2012. doi: 10.1523/JNEUROSCI.0207-12.2012.

Gary Sclar and Ralph D. Freeman. Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental Brain Research*, 46(3):457–461, 1982. doi: 10.1007/BF00238641.

Stephanie C. Seeman, Luke Campagnola, Pasha A. Davoudian, Alex Hoggarth, Travis A. Hage, Alice Bosma-Moody, Christopher A. Baker, Jung Hoon Lee, Stefan Mihalas, Corinne Teeter, Andrew L. Ko, Jeffrey G. Ojemann, Ryder P. Gwinn, Daniel L. Silbergeld, Charles Cobbs, John Phillips, Ed Lein, Gabe J. Murphy, Christof Koch, Hongkui Zeng, and Tim Jarsky. Sparse recurrent excitatory connectivity in the microcircuit of the adult mouse and human cortex. *bioRxiv*, pages 1–27, 2018. doi: 10.1101/292706.

Carla J. Shatz. Emergence of order in visual system development. *Journal of Physiology-Paris*, 90(3-4):141–150, 1996. doi: 10.1016/S0928-4257(97)81413-1.

Adam M. Sillito, John A. Kemp, James A. Milson, and Nicoletta Berardi. A re-evaluation of the mechanisms underlying simple cell orientation selectivity. *Brain Research*, 194(2):517–520, 1980. doi: 10.1016/0006-8993(80)91234-2.

Wolf Singer. Synchronization of Cortical Activity and Learning. *Brain*, 1993.

Bernt C. Skottun, Arthur Bradley, Gary Sclar, Izumi Ohzawa, and Ralph D. Freeman. The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *Journal of Neurophysiology*, 57(3):773–786, 1987.

David C. Somers, Sacha B. Nelson, and Mriganka Sur. An emergent model of orientation selectivity in cat visual cortical simple cells. *The Journal of Neuroscience*, 15(8):5448–5465, 1995. doi: 10.1523/JNEUROSCI.15-08-05448.1995.

Sen Song, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B. Chklovskii. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology*, 3(3):0507–0519, 2005. doi: 10.1371/journal.pbio.0030068.

Daniel Soudry, Suraj Keshri, Patrick Stinson, Min-hwan Oh, Garud Iyengar, and Liam Paninski. A shotgun sampling solution for the common input problem in neural connectivity inference. *arXiv*, 2013.

Klaus M. Stiefel, Bernhard Englitz, and Terrence J. Sejnowski. Origin of intrinsic irregular firing in cortical interneurons. *Proceedings of the National Academy*

of Sciences of the United States of America, 110(19):7886–7891, 2013. doi: 10.1073/pnas.1305219110.

Chris Tailby, Samuel G. Solomon, Neel T. Dhruv, Najib J. Majaj, Sach H. Sokol, and Peter Lennie. A new code for contrast in the primate visual pathway. *Journal of Neuroscience*, 27(14):3904–3909, 2007. doi: 10.1523/JNEUROSCI.5343-06.2007.

Andrew Y.Y. Tan, Brandon D. Brown, Benjamin Scholl, Deepankar Mohanty, and Nicholas J. Priebe. Orientation selectivity of synaptic input to neurons in mouse and cat primary visual cortex. *Journal of Neuroscience*, 31(34):12339–12350, 2011. doi: 10.1523/JNEUROSCI.2039-11.2011.

Zhenjun Tan, Hang Hu, Z. Josh Huang, and Ariel Agmon. Robust but delayed thalamocortical activation of dendritic-targeting inhibitory interneurons. *Proceedings of the National Academy of Sciences of the United States of America*, 105(6):2187–2192, 2008. doi: 10.1073/pnas.0710628105.

Arthur W. Toga, Kristi A. Clark, Paul M. Thompson, David W. Shattuck, and John Darrell Van Horn. Mapping the human connectome. *Neurosurgery*, 71(1):1–5, 2012. doi: 10.1227/NEU.0b013e318258e9ff.

Elaine Tring and Dario L. Ringach. On the Subspace Invariance of Population Responses. *arXiv*, 2018.

Misha V. Tsodyks, William E. Skaggs, Terrence J. Sejnowski, and Bruce L. McNaughton. Paradoxical Effects of External Modulation of Inhibitory Interneurons. *The Journal of Neuroscience*, 17(11):4382–4388, 1997. doi: 10.1523/JNEUROSCI.17-11-04382.1997.

Stephen D. Van Hooser. Similarity and diversity in visual cortex: is there a unifying theory of cortical computation? *The Neuroscientist*, 13(6):639–656, 2007. doi: 10.1177/1073858407306597.

Carl van Vreeswijk and Haim Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726, 1996. doi: 10.1126/science.274.5293.1724.

Julia Veit, Richard Hakim, Monika P. Jadi, Terrence J. Sejnowski, and Hillel Adesnik. Cortical gamma band synchronization through somatostatin interneurons. *Nature Neuroscience*, 20(7):951–959, 2017. doi: 10.1038/nn.4562.

Saurabh Vyas, Matthew D. Golub, David Sussillo, and Krishna V. Shenoy. Computation through Neural Population Dynamics. *Annual Review of Neuroscience*, 43:249–275, 2020. doi: 10.1146/annurev-neuro-092619-094115.

Edgar Y. Walker, Fabian H. Sinz, Erick Cobos, Taliah Muhammad, Emmanouil Froudarakis, Paul G. Fahey, Alexander S. Ecker, Jacob Reimer, Xaq Pitkow, and Andreas S. Tolias. Inception loops discover what excites neurons most using deep predictive models. *Nature Neuroscience*, 22(12):2060–2065, 2019. doi: 10.1038/s41593-019-0517-x.

John G. White, Erica Southgate, J. N. Thomson, and S. Brenner. The structure of the nervous system of the nematode Caenorhabditis elegans. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 314(1165): 1–340, 1986. doi: 10.1098/rstb.1986.0056.

Hugh R. Wilson and Jack D. Cowan. Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons. *Biophysical Journal*, 12(1):1–24, 1972. doi: 10.1016/S0006-3495(72)86068-5.

Florentin Wörgötter and Christof Koch. A detailed model of the primary visual pathway in the cat: Comparison of afferent excitatory and intracortical inhibitory connection schemes for orientation selectivity. *Journal of Neuroscience*, 11(7): 1959–1979, 1991. doi: 10.1523/jneurosci.11-07-01959.1991.

Hongkui Zeng and Joshua R. Sanes. Neuronal cell-type classification: Challenges, opportunities and the path forward. *Nature Reviews Neuroscience*, 18(9):530–546, 2017. doi: 10.1038/nrn.2017.85.

List of publications

Nataliya Kraynyukova*, **S Renner***, Gregory Born, Yannik Bauer, Martin A. Spacek, Georgi Tushev, Laura Busse, and Tatjana Tchumatchenko. In vivo extracellular recordings of thalamic and cortical visual responses reveal v1 connectivity rules. *Proceedings of the National Academy of Sciences*, 119(41): e2207032119, 2022. doi: 10.1073/pnas.2207032119.

S Renner, N Kraynyukova, AH Kotkat, Y Bauer, G Born, M Spacek, T Tchumatchenko, and L Busse. Contrast-invariant orientation tuning in mouse primary visual cortex probed via briefly flashed gratings. In *Bernstein Conference 2019*, 2019. doi: 10.12751/nncn.bc2019.0080.

S Renner, N Kraynyukova, Y Bauer, G Born, AH Kotkat, X Liu, M Spacek, G Tushev, T Tchumatchenko, and L Busse. Inference of network connectivity from responses to briefly flashed gratings in mouse V1 using a stabilized supralinear network model (SSN). In *Bernstein Conference 2020*, 2020. doi: 10.12751/nncn.bc2020.0111.

C Stockinger, B Thürer, **S Renner**, A Focke, and T Stein. Intermanuelle Transfermechanismen beim motorischen Lernen. In J. Hermsdörfer, W. Stadler, and L. Johannsen, editors, *The Athlete's Brain – Neuronale Aspekte motorischer Kontrolle im Sport*, volume 241, Schriften der Deutschen Vereinigung für Sportwissenschaft, page 125f. Czwalina, Hamburg, 2015.

Author Contributions

Manuscript 1

Nataliya Kraynyukova (NK), Laura Busse (LB), and Tatjana Tchumatchenko (TT) conceived the study. NK developed the mathematical proof of how contrast-invariance constrains connectivity, and designed the connectivity inference method. Simon Renner (SR) recorded and curated V1 data in one mouse, developed reverse correlation analysis techniques for flashed gratings stimuli, optogenetic tagging analysis techniques, waves shape classification techniques, contrast-invariance analysis techniques, programmed the data analysis software, and showed contrast-invariance of the recorded V1 data. SR performed all analyses in Figures 1, S1, S2, S3, S4 and made these Figures. NK and SR pre-processed data for model fits, where SR developed and tested two-dimensional response functions and constructed population responses. SR and NK reviewed and analyzed V1 connectivity measures available in the experimental literature in relation to the inferred connectivity with input by TT and LB. Here, SR selected appropriate studies from the literature based on methodological considerations and extracted connectivity parameters from publicly available data sets. Georgi Tushev contributed to the connectivity inference algorithm. Gregory Born, Yannik Bauer and Martin Spacek recorded and curated dLGN data and additional V1 data, and developed data analysis software. NK, SR, LB, and TT contributed conceptual ideas, discussed and coordinated the project at all stages, wrote and edited the manuscript. All authors contributed to the final version of the manuscript. A later version of this manuscript has been published as Kraynyukova* et al. (2022).

Place, date	Nataliya Kraynyukova
Place, date	Simon Renner
Place, date	Laura Busse

Manuscript 2

Laura Busse (LB) and Simon Renner (SR) conceived the study. SR recorded and curated data V1 data in one mouse and developed the data analysis software. Gregory Born, Yannik Bauer and Martin Spacek recorded and curated dLGN data and additional V1 data. Emma Müller-Seydlitz developed analysis techniques for spontaneous activity data related to Figure 2. SR expanded and finalized analysis techniques for variability and oscillations of spontaneous activity. SR performed data analysis for all figures, including identification of SbC neurons, wave shape classification, layer location analysis, and stimulus tuning analysis. SR made the figures. SR and LB contributed conceptual ideas, discussed and coordinated the project at all stages, wrote and edited the manuscript.

Place, date	Simon Renner
Place, date	Laura Busse