THE UNIVERSITY OF CHICAGO

VISUALLY-EVOKED NORMALIZATION IN AWAKE MOUSE V1

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES

AND THE PRITZKER SCHOOL OF MEDICINE

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

INTERDISCIPLINARY SCIENTIST TRAINING PROGRAM:

COMPUTATIONAL NEUROSCIENCE

BY

ZAINA ADEL ZAYYAD

CHICAGO, ILLINOIS

AUGUST 2021

For my parents, for everything.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

I would like to thank my thesis advisors, Jason N. MacLean and John H.R. Maunsell, for their unwavering support, tireless efforts, and care in training me. I am constantly awed by their propensity to hit the mark; their clarity of thought has taught me how to think as a scientist. It has been my privilege and honor to learn from them. I would also like to thank the rest of my thesis committee, S. Murray Sherman and Nicho Hatsopoulos (Chair) for their insights and support in developing this project. Thank you to my lab mates, especially to Maayan Levy for her help with the partial pairwise correlation matrices. Thanks to Vaughn Spurrier, Joe Dechery, Jackson Cone, Julian Day-Cooney, Morgan Bade, and Supriya Ghosh for teaching me techniques required to collect this data. It has been a gift to live and learn with such outstanding colleagues and friends.

Additionally, I would like to thank the Medical Scientist Training Program at the University of Chicago and the Committee on Computational Neuroscience for providing a rich intellectual and collegial environment in which to train. These have been golden years.

I often reflect that I would not be on this path without the mentoring I received from my advisor at Yale, Maria Moreno, and from my teachers at Aqsa School. Thank you, thank you.

I would also like to thank my siblings Osama, Zaid, Bana, Shadin, and Mohammad, my grandparents and extended family, and my friends for their love, inspiration, and support.

Most of all, I would like to thank my parents, Fareeda Alaidi and Adel Zayyad. Truly, my accomplishments can always be traced back to their profound love and sacrifice. With the sheer strength of their love and against tremendous odds, they managed to give me and my siblings the world. They have always been my biggest support. Thank you Mama and Baba, for everything.

Alhamdullilah always.

ABSTRACT

Normalization is a canonical nonlinear computation that underlies multisensory integration, attention, and other sensory processing. In this computation, the response of one neuron to a set of stimuli is normalized by the weight of all the presented stimuli – not just its preferred stimulus. Normalization is ubiquitous and conserved across species, suggesting its fundamental role in brain processing. When central brain computations go awry, the consequences can be severe – and normalization has been implicated in brain disorders like autism. While normalization has been described in multiple modalities and species, its functional circuitry and mechanism remain unknown. This project aims to produce an experimental model to address that gap. Specifically, here, we investigate divisive normalization in awake mouse visual cortex. This mouse model will enable the use of powerful *in vivo* genetic, electrophysiology, imaging, optogenetic, and psychophysics techniques.

Using transgenic mice, excitatory neurons in mouse V1 were functionally labelled with GCaMP6s. Two-photon imaging was used to capture the activity of large V1 excitatory populations in awake mice presented with cross-inhibitory stimuli designed to evoke normalization. We recorded from hundreds of tuned and untuned neurons in an unbiased manner and observed tuned normalization similar to what has been observed in other species. These findings were cross-validated using electrophysiological recordings. Additionally, in our electrophysiology data, normalization strength was observed to have a depth dependence, suggesting the need for further study of laminar differences in this computation. Our data show that normalization can be visually evoked and measured in the V1 of awake, head-fixed mice, opening the door to further functional and mechanistic study.

Furthermore, we examined how normalization influences population pairwise noise correlations. We observed that normalization increased the correlation of similarly orientation-selective pairs with high normalization strength and decreased the correlation of oppositely orientation-selective pairs with high normalization strength. We also observed that in mice, in contrast with macaques, normalization decreased the correlation of similarly orientation-selective pairs with low normalization strength. Normalization mechanisms underlie the changes in pairwise correlations that can explain improved behavioral performance in attention, so further understanding of normalization will help better understand processes that modify pairwise correlations, like attention.

We anticipate that this work will provide groundwork for continuing to study normalization in the mouse and may ultimately lead to circuit dissection and behavioral assays, contributing to an understanding of the role of normalization in normal physiology and behavior, aberrant circuitry, and disease.

CHAPTER 1

INTRODUCTION

Normalization, a canonical brain computation

Over an evolutionary timescale, organisms evolve features that add survival value. Brain processes and computations are no exception. To interact with the world and meet the requirements for survival in a physical and competitive environment, the brain must keep pace with rapid stimulus bombardment from the physical world, extracting relevant stimuli to direct behaviors that are compatible with life. Failure to do so poses a risk of starving or becoming prey, among other problems: in short, managing rapid competing stimuli is a matter of life and death. To meet this need, the brain manages massive amounts of external and internal stimulus data every split second, processing that data into usable signals in a manner that appropriately prioritizes salience while discarding irrelevant information (H. Barlow 2001).

To better understand the ways that the brain can process such large quantities of information, neuroscience has turned to structural, descriptive, and functional examination of the brain at different scales (Marr 1982). In their seminal work, Hubel and Wiesel described a linear model for neuronal responses in cat V1 (D. H. Hubel and Wiesel 1959; 1962). Not long after, this linear model was supplemented with new information about cross-orientation suppression, contrast saturation, and other non-linearities in neuronal responses (Bonds 1989). The normalization model was proposed to better capture the nonlinearities present in V1 neuron responses, originally for when multiple stimuli were shown simultaneously in a cross-oriented plaid (Heeger 1992; M. Carandini and Heeger 1994; Matteo Carandini, Heeger, and Movshon 1997; D. G. Albrecht and Geisler 1991). Since then, normalization has been shown to underlie the

non-linearities present in fundamental processing in many brain regions and species (Matteo Carandini and Heeger 2011; Kouh and Poggio 2008). In normalization, the response of a neuron is normalized by the activity of its neighbors and ultimately weighted according to the relative strength of stimuli.

Fundamentally, normalization is a computation. The normalization model describes a transform of a set of inputs to produce an output. By creating models that capture data, falsifiable predictions can be further made and tested to inform further understanding of the brain and an iterative process can continue to refine the original model.

In a widely adopted model of normalization, the computation can be modelled as a divisive function, in which the neuron response is normalized by the total stimulus intensity (Matteo Carandini and Heeger 2011). Thus, for example, the linear sum of a neuron's responses to cross-oriented gratings is divided by the aggregate activity pooled across nearby neurons. In this form, normalization can be modeled by the following equation,

$$R = \frac{w_1 c_1 + w_2 c_2}{c_{50} + c_1 + c_2}$$

in which the numerator indicates linear summation corresponding to two stimuli with contrasts c_1 and c_2, and the denominator indicates the normalizing component according to the strength of the stimuli.

Thus, in the case of a stimulus that generates a strong response masked by one that generates no response, both at equal contrasts, the expected result would approximate a simple average of the responses to each stimulus. The normalization model also makes a prediction about the response to a pair of stimuli that each evoke a response, where one response is weak and the other is much stronger. In this case, the normalization model would predict a response heavily

weighted to the stronger stimulus, and so the response of the neuron would be expected to closely match the response to the stronger stimulus (Busse, Wade, and Carandini 2009; Matteo Carandini and Heeger 2011).

Normalization occurs on a population level, and importantly, the output from normalization becomes the input in downstream areas, affecting downstream processing, perception, cognition, and behavior (Busse, Wade, and Carandini 2009; Coen-Cagli and Schwartz 2013). Locally, the computation can also explain changes in functional network connectivity in the form of pairwise correlations and can even account for the changes in pairwise correlations observed in attention (Verhoef and Maunsell 2017). Because normalization manifests in many brain areas and modalities, the impact of normalization on early stimulus representation is magnified over the cortical processing hierarchy, profoundly influencing stimulus representation in the brain.

While normalization models were first developed to explain the nonlinear summation of visual responses, the computation is also observed in a variety of modalities and species ranging from olfaction in fruit flies to light adaptation in retina to decision-making in macaques (Laughlin 1981; Olsen, Bhandawat, and Wilson 2010; Louie, Grattan, and Glimcher 2011). Normalization underlies multisensory integration, influencing the relative representation of stimuli in the brain by emphasizing some aspects of the stimulus and de-emphasizing others (Ohshiro, Angelaki, and DeAngelis 2011). In this way, normalization is one way that the brain handles competing inputs, be they visual (i.e. cross-oriented gratings), olfactory (i.e. simultaneous presentation of different odors), or cognitive (i.e. simultaneous consideration of different reward options). Normalization plays a role in diverse processes, species, and modalities, suggesting it is a canonical brain computation (Matteo Carandini and Heeger 2011).

Along with its role in sensory processing, normalization also plays an important role in attention. Evidence from macaque V4 and MT shows that changes in neuronal firing rates due to attention can be explained by normalization. Furthermore, the changes in pairwise correlations observed in attention can also be explained by a normalization model (Ni, Ray, and Maunsell 2012; Ni and Maunsell 2017; Verhoef and Maunsell 2017). These correlations changes have previously been shown to be central to improving behavior in attention (Marlene R. Cohen and Maunsell 2009). Together, these findings highlight a central role of normalization in population representation of sensory information and indicate a need to better understand normalization.

Although normalization has been described in many species, modalities, and contexts, its functional properties and circuit mechanisms remain elusive. Here, we develop a mouse model of normalization to pave the way for further examination of this computation. In particular, we describe visually-evoked normalization in awake mice using two-photon population imaging and electrophysiology. In doing so, we characterize the properties of normalization in mice, comparing and contrasting with findings in other species. At the same time, we take advantage of population imaging to better understand changes in pairwise correlations and network structures engaged by this computation. It is our hope that a mouse model of visually evoked normalization will provide fertile ground for continued study of normalization, including its statistical properties, circuit underpinnings, and behavioral relevance.

Characterization of normalization in neuronal response properties

While some aspects of visual processing appear locally linear, this linearity breaks down and gives way to nonlinearities in multiple different visual processes. Some examples include contrast saturation, surround suppression, and cross-orientation inhibition (Matteo Carandini and

Heeger 2011). While different mechanisms may give rise to each of these three phenomena, they can all be united with one computation: normalization.

Contrast modulation of firing rate and many other visual neuronal response properties across different brain regions and species can be well-described by the normalization model (Heeger 1992; Tolhurst and Heeger 1997; Heuer and Britten 2002; 2007; Keller and Martin 2015; Sawada and Petrov 2017; Matteo Carandini, Heeger, and Movshon 1997; Simoncelli and Heeger 1998; Rust et al. 2006; Zoccolan, Cox, and DiCarlo 2005). The model manages to capture response properties of neurons across species and modalities, but normalization has been perhaps most extensively studied using a cross-species-validated model of cross-orientation suppression in primary visual cortex.

In addition to modeling contrast saturation and other phenomena with normalization, studies of normalization have been conducted using cross-orientated stimuli to elicit normalization in cat and macaque (Matteo Carandini, Heeger, and Movshon 1997; Busse, Wade, and Carandini 2009; Ruff, Alberts, and Cohen 2016). In the early 1990s, Heeger proposed the normalization model to explain non-linearities in the response properties of cat V1 neurons (Heeger 1992). Subsequently, Carandini and Heeger studied the nonlinear summation by V1 cells of cross-oriented stimuli presented to anesthetized macaques while recording from visual cortex. They found that rather than sum the expected responses to each stimulus, individual neurons weighted the stimuli relative to their contrasts, and spiked according to a weighted average of their responses to each individual stimulus. While cortical normalization was first characterized in early studies in the cat and macaque visual cortex, it has also been described in other species like the tree shrew, suggesting that it is a common and conserved computation (MacEvoy, Tucker, and Fitzpatrick 2009).

Further study emphasized that beyond normalization at the single neuron level, the computation occurs on a population-wide level (Busse, Wade, and Carandini 2009). In that study, Busse and colleagues showed cross-oriented plaid stimuli to anesthetized cat while using a Utah array to record from layer 2/3 of anesthetized cats. Because the array covered multiple orientation selective columns, the authors recorded from units with a broad range of stimulus preferences. When fitting the data from this population of neurons to the normalization model and to other models, the normalization model captured population responses best, highlighting that the computation occurs on a population level. In addition, in further experiments in the same study, recording from V1 and MT, showed that the output from normalizing populations becomes the input into downstream brain regions, as measured in MT (Busse, Wade, and Carandini 2009).

In addition to this work in cat, work in macaque has continued to characterize normalization in populations of V1, MT, and V4 neurons. In a study of normalization in passively viewing macaques, Ruff and his colleagues presented multiple sets of cross-oriented plaids at various orientations while recording from populations of neurons. Neuron normalization strength between conditions was correlated, suggesting that the computation does not depend on stimulus orientation, but rather on a neuron's propensity to normalize (Ruff 2016).

In addition to visually-evoked normalization, recent experiments have explored optogenetic activation as a method for evoking increased activity in neurons, meant to evoke normalization (Nassi et al. 2015; Histed 2018; Sato et al. 2016a). Nassi and his co-workers used optogenetic activation to activate distal neurons in macaques while recording from V1; they observed summation at low contrasts and division at high contrast (Nassi et al. 2015). Histed similarly optogenetically activated distal neurons in mouse V1 binocular zone while recording from the contralateral hemisphere during the presentation of a white noise stimulus, but reported

mostly linear summation in mouse V1 (Histed 2018). Most of the neurons identified in this study had low firing rates, raising the question of whether the stimuli were effective in generating responses from which normalization could be reliably investigated. Nevertheless, these studies point to the possibility of modulating normalization with optogenetics, an intriguing possibility for future study.

Normalization and attention

Normalization has been heavily implicated in attention. Spatial attention increases the emphasis placed on a particular section of the visual field and can dramatically influence the responses of visual cortical neurons (Desimone and Duncan 1995; Reynolds and Chelazzi 2004; J. H. R. Maunsell and Treue 2006). Recently, a model of normalization for attention has emerged that accounts both for the firing rates and correlation changes observed in attention (Boynton 2009; Reynolds and Heeger 2009; Lee and Maunsell 2009; 2010; Verhoef and Maunsell 2017). In both single cell and population studies, normalization has emerged as the key mechanism by which normalization influences neuronal changes due to attention: and a normalization model of attention accounts for many of the changes in firing rate depending on the attended stimulus and population changes due to attention (Lee and Maunsell 2009; Ni, Ray, and Maunsell 2012; Ni and Maunsell 2017; 2019; Verhoef and Maunsell 2017).

To different extents depending on the individual neuron, attention can increase the firing rate of neurons that respond to the receptive field to which attention is directed. Studies in macaque show that different neurons exhibit heterogeneous "tuned" levels of divisive normalization – some neurons exhibit standard weighted averaging, while others super-normalize or show no effects of normalizing (Ni, Ray, and Maunsell 2012). In studies of attention, this tuning held functional

relevance (Lee and Maunsell 2009; Ni and Maunsell 2017; Ni, Ray, and Maunsell 2012; Verhoef and Maunsell 2017). In addition to its fundamental roles in multisensory integration and information processing, modulation of population activity by attention depends on normalization mechanisms, and neurons that do not normalize largely do not show attention modulation (Lee and Maunsell 2009; 2010; Ni, Ray, and Maunsell 2012).

In other experiments, attention has been shown to decrease the correlations within areas of visual cortex, thought to improve population sensitivity (Marlene R. Cohen and Maunsell 2009). These changes in pairwise correlations are influenced by different factors, like choice selectivity, and contexts like difficulty: for example, pairwise noise correlations decrease to a greater extent in more difficult tasks (Ruff and Cohen 2014b; 2014a). Importantly, improvements in behavior due to attention correlate with changes in pairwise correlations. How these changes are related to computation and stimulus representation is an area of active investigation. Recent work has shown that these changes in correlation arise from normalization mechanisms (Verhoef and Maunsell 2017).

In this work, Verhoef et. al. trained macaques to perform a visual spatial attention task, then recorded from V4 while the monkeys completed the task. When attention was not directed to a pair of orthogonal stimuli in the receptive field, pairwise correlations between neurons with similar orientation selectivity were increased and correlations between pairs of neurons with opposite orientation selectivity were decreased, consistent with a normalization model of mutually suppressive neuron populations. Verhoef and Maunsell found that when the animal directed attention to stimuli in the receptive field, correlations between pairs of neurons that shared similar orientation selectivity decreased in correlation relative to stimuli the macaque was not attending to. Importantly, these findings can be explained by a model for populations of neurons that

mutually suppress each other according to a normalization model. In this model, the population that responding to an attended, preferred stimulus receives additional excitatory drive, suppressing the other subpopulation of neurons and decreasing the resultant suppression from that subpopulation, decreasing shared suppression and decorrelating the first population's pairwise correlations (Verhoef and Maunsell 2017).

Altogether, these studies have shown that the changes in neuronal representation of stimuli due to attention are largely influenced by local circuitry and the different extents to which neurons normalize. Globally, these findings strongly corroborate the "normalization model" of attention, in which the mechanism of attention relies on normalization circuitry (Lee and Maunsell 2009; Reynolds and Heeger 2009). This points to a need to better understand normalization, as it underlies important cognitive processes, including attention.

Biophysical, circuit, and network mechanisms for normalization

From a computational perspective, one computation, such as normalization, can be produced through multiple different processes (Matteo Carandini and Heeger 2011). Because normalization is an abstract computation that takes place in many different modalities and species, it can be effected through different biophysical mechanisms.

In some cases, including fruit fly olfaction, the mechanisms that produce normalization are known to be driven by inhibition (Olsen, Bhandawat, and Wilson 2010). However, there remains much debate as to the origin of normalization measured in the cortex, with some favoring a subcortical contrast saturation-based model (Freeman et al. 2002; Priebe and Ferster 2006), others positing cellular conductance-based mechanisms (M. Carandini and Heeger 1994), and still others favoring a local, cortical basis for normalization in which balanced excitation and inhibition result

in divisive gain modulation of neuronal response (Frances S. Chance, Abbott, and Reyes 2002). The field has not settled on a mechanism, and the participation of multiple biophysical mechanisms (possibly with distinct functional roles) cannot be excluded. Circuit level studies will be important for elucidating this open question.

Some studies point to a subcortical origin for normalization in the cortex. Data from cat shows that at high contrasts, there is contrast saturation in LGN neurons (Priebe and Ferster 2006). It has been posited that at overall contrasts past saturation, responses to a cross-inhibitory stimulus could result in decreased firing in V1 neurons compared to a linear sum of responses to individual stimulus components. This model would also predict that at low contrast, neurons should primarily sum activity. Thus, this model on its own cannot explain normalization at low contrast cross-inhibitory stimulus pairs, at which LGN neurons exhibit linear response properties (Busse, Wade, and Carandini 2009). Additionally, since normalization can be observed in populations, this model would require that a vast majority of neurons providing input from the dLGN to V1 exhibit contrast saturation. However, that is not the case: the dominant inputs to layer 2/3 from dLGN are remarkably linear in mice and other species (Seabrook et al. 2017). In the mouse for example, LGN neurons in mouse (that form the main input to V1) show linear responses, with an average contrast at half-maximal response of approximately 20% (Niell and Stryker 2008). Furthermore, evidence of normalization in higher visual cortex areas like V4 and MT and modulation by cognitive processes like attention and receptive field properties suggest an important cortical basis for cortical normalization (Heuer and Britten 2002; Ni, Ray, and Maunsell 2012; Verhoef and Maunsell 2017). While there may be a place for contrast saturation in the responses of some neurons, this model seems unlikely to be the primary explanatory mechanism for normalization.

One way to further study this question would be to examine normalization across cortical layers to observe if normalization emerges or is modified at any particular stage of cortical processing. The thalamic contrast saturation model would also predict that robust normalization should be observed as early in the cortical layer processing hierarchy as layer 4, which receives primary inputs from the dLGN. It is possible to record from different layers to gain a better understanding of the stages of cortical processing in which normalization takes place: for example, if normalization does not occur in layer 4, normalization cannot be taken to be simply inherited from the LGN. Surround suppression has been examined in this way in macaque visual cortex and suggests that cortical layers that receive receptive field input directly from thalamus reflect surround suppression latest (Bijanzadeh et al. 2018). Thus, there may be a layer-specificity to computations involving normalization that can be tested by recording across cortical layers.

Synaptic depression could also contribute to normalization (Matteo Carandini and Heeger 2011; Rosenbaum, Rubin, and Doiron 2013). In a representation of this model, as pre-synaptic current increases to a neuron, post-synaptic current ultimately plateaus, resulting in a diminished neuronal response. Thus, when multiple stimuli are presented simultaneously, neuronal responses are decreased relative to when they are presented alone.

On a more local, cortical level, shunting inhibition was an early theory offered to explain the mechanism that normalizes neuron responses. This model depends on the RC model of the cell membrane, in which the cell membrane is modelled as a resistor and capacitor in parallel. In this case, the neurons in a normalization pool effectively inhibit each other by increasing each other's conductance. Thus, increasing the stimulus energy would mediate shunting inhibition. Indeed, shunting inhibition was originally hypothesized to mediate normalization in early studies of normalization (M. Carandini and Heeger 1994; Matteo Carandini and Heeger 2011). Here, the

cellular processes of neurons would be sufficient to mediate normalization and would depend primarily on recurrent excitation (R. J. Douglas et al. 1995; Shushruth et al. 2012). Cortical neurons in awake animals are constantly in a high conductance state (Destexhe and Paré 1999), but in previous theoretical work this was paradoxically shown to result in increased excitability (Hô and Destexhe 2000). Experiments are required to better understand how recurrent excitation and conductance changes can influence neuronal activity.

In addition to classically tuned excitatory populations, the mouse neocortex also includes a rich diversity of inhibitory neurons, with different morphologies, connectivity, tuning properties, electrical properties, neurotransmission profiles, and gene expression (Isaacson 2010; Fino and Yuste 2011; Packer and Yuste 2011; Fino, Packer, and Yuste 2013; Kerlin et al. 2010; Runyan et al. 2010; Liu et al. 2009; Ma et al. 2010; Znamenskiy et al. 2018). In the mouse visual cortex, approximately 20% of neurons are inhibitory (Niell and Stryker 2008), among which this diversity of classes and subclasses can be observed even between layers. Different interneurons classes are present in varying proportions and numbers across the neocortex, suggesting distinct roles in signal processing (Vogels, Rajan, and Abbott 2005; Isaacson and Scanziani 2011; Adesnik et al. 2012; Nienborg et al. 2013). In one study, shared excitatory and inhibitory input was shown to be a mechanism for decorrelating neuronal activity, emphasizing the possibilities of inhibition to generate complex computations (Renart et al. 2010). While various mechanisms have been proposed to account for divisive normalization (Wilson et al. 2012; Keller and Martin 2015; Sato et al. 2016a), inhibition is currently the most widely considered theory (Wilson et al. 2012; Hauesler and Maass 2017).

It has been posited that separate inhibitory populations tuned to different orientations could provide inhibition to populations of excitatory neurons with orthogonal preferences in cross-

inhibition; however this would suggest that the normalization of a neuron varies depending upon stimulus preferences. This was shown not to be the case (Ruff, Alberts, and Cohen 2016). Thus, normalization is thought to be a property of the neuron itself or its embedded position in a homogenously pooling normalization network. This was hypothesized in a 1982 paper by Morrone and colleagues, in which simple cells in the cat striate cortex were shown to be evenly suppressed by cross-inhibitory gratings irrespective of orientation and spatial frequency (Morrone, Burr, and Maffei 1982). Nevertheless, there is some functional specificity of synaptic connectivity among excitatory neurons based on similarities in receptive field properties. In one study in mouse visual cortex, neurons imaged and shown to have similar orientation preferences in a two-photon microscopy experiment were twice as likely to form synaptic connections when later patched in vitro, in comparison with pairs that had orthogonal preferences. This suggests that there are local subnetworks of excitatory neurons that process similar stimulus information and that recurrent connectivity is related to shared receptive field properties and visual responses (Ko et al. 2011). Consequently, it is possible that these subnetworks can act in concert to suppress other subnetworks indirectly through pooled inhibition or other suppressive mechanisms.

There exists strong feedforward inhibition in thalamocortical circuits, in which fast-spiking inhibitory neurons were shown to respond to feedforward excitation faster than their excitatory counterparts (Cruikshank, Lewis, and Connors 2007). Among the many different types of neocortical interneurons, parvalbumin (PV) expressing neurons tend to be fast-spiking and are good candidates for participation in divisive normalization: they synapse directly onto pyramidal soma and peri-somatic regions, spike at relatively high rates, have low input resistance synapses, and form electrical synapses, all of which allow these neurons to rapidly modulate the activity of neighboring neurons (Meyer et al. 2002; Kawaguchi and Kubota 1997). Additionally, PV neurons

can be tuned, albeit broadly, to orientation or direction, similar to their excitatory counterparts (Cardin, Palmer, and Contreras 2007; Runyan et al. 2010). Because they are tuned and form specific, differential synapses with distinctly tuned excitatory neurons but also a dense network of electrical synapses with other PV neurons, PV neurons could be positioned to pool activity related to a stimulus to normalize the activity of neighboring excitatory neurons. To evaluate the possibility that PV neurons contribute to normalization, excitatory and PV inhibitory population activity could be recorded simultaneously during evoked divisive normalization to investigate time-locked functional relationships between these populations.

Recent work has shown that inhibitory neurons in mouse receive input from excitatory cells with a wide range of orientation tuning, offering another insight into how inhibition functions in cortex (Bock et al. 2011; Hofer et al. 2011). There is minimal orientation and phase selectivity of inhibitory neurons in mouse visual cortex and presence of untuned inhibitory neurons in the visual cortex of other species (Cardin, Palmer, and Contreras 2007; Niell and Stryker 2008). This indicates that inhibitory neurons pooled local activity that could be leveraged for computation (Kerlin et al. 2010). Indeed, this feature has been offered as a mechanism for normalization, balanced excitation and inhibition (F. S. Chance and Abbott 2000; Frances S. Chance, Abbott, and Reyes 2002), coupling metabolism to brain activity (Buzsáki, Kaila, and Raichle 2007), and contrast-invariant tuning (Lauritzen and Miller 2003).

Relatedly, in the inhibitory stabilized network (ISN), blanket inhibition plays a critical role in maintaining many of the network's properties including normalization (Rubin, Van Hooser, and Miller 2015). There have been efforts to construct in silico models that produce normalization. Using a rate-based recurrent neural network designed with inhibition, Miller and colleagues found that an inhibition-stabilized network could produce response properties similar to those observed

in normalization when a contrast mask was superimposed onto a visual stimulus (Ahmadian, Rubin, and Miller 2013; Rubin, Van Hooser, and Miller 2015; Hennequin et al. 2018). This work advances a theory of inhibition-stabilized networks, in which balanced inhibition stabilizes the firing activity of a network of neurons. Importantly, the model also makes a prediction about paradoxical inhibition, in which both inhibition and excitation decrease when inhibitory units receive additional drive. Models like this can be used to make predictions about neural dynamics, like paradoxical inhibition, that can be tested experimentally. Cortical dynamics generally include a strong set of recurrent and balanced excitation and inhibition (Wolf et al. 2014), which would be in line with this model. Finally, this type of balanced excitation and inhibition could contribute to decorrelating activity in cortex, which makes it an intriguing candidate for a computation implicated in correlation changes in cortex (Renart et al. 2010).

While inhibition may play a network role in normalization, its direct synaptic effects on normalizing neurons has recently challenged by electrophysiologic studies. In recent work using patch-clamp electrophysiology, Sato and colleagues optogenetically activated binocular zone neurons in mice while recording from the contralateral monocular zone. When they showed a visual stimulus in conjunction with the optogenetic activation designed to evoke normalization, they found a decrease in excitation in comparison with activating the distal network without visual stimulus, in which neurons showed summation (Sato et al. 2016a). They did not find any increase in inhibitory post-synaptic potentials in any of the neurons they patched (n=10). This finding is consistent with a model of normalization involving strong local inhibitory networks (Rubin, Van Hooser, and Miller 2015; Ozeki et al. 2009), in which the excitatory/inhibitory ratio suppresses overall network activity during periods of excitation, but does not directly synaptically suppress excitatory neurons (Haider et al. 2006; Haider, Häusser, and Carandini 2013).

Katnzner and colleagues also found that normalization is immune to blockage of $GABA_A$ receptors in cat visual cortex, decreasing the likelihood that normalization occurs exclusively through inhibition (Katzner, Busse, and Carandini 2011). While it is possible that inhibition does not contribute to normalization, this experiment still leaves open the possibility of electrical synapses or other inhibitory neurotransmitters playing a role in the computation. It also leaves open the possibility that normalization exists endogenously in recurrent networks, but is sharpened by inhibitory networks, similar to other brain phenomena like tuning (Seabrook et al. 2017).

It is possible that one, some, or all of these potential mechanisms participate in mediating normalization. Distinguishing between these multiple mechanisms will be important to further understanding the computation, its functional implications, and pathologies in which normalization may be compromised. With a mouse model of normalization, it will become easier to elucidate the cell types, dynamics, and networks involved in normalization.

Possible computational roles for normalization in neuron responses

As normalization is so ubiquitous and conserved, it may improve an organism's processing of information. In terms of the representation of stimuli, marginalization and statistical independence have been mathematically shown to be effected by a normalization model (O. Schwartz and Simoncelli 2001; Simoncelli and Olshausen 2001), consistent with the efficient coding hypothesis (Horace Barlow 1961; Attneave 1954; H. Barlow 2001). This theoretical work suggests that normalization can improve the encoding of a stimulus by reducing or eliminating the statistical dependencies between linear filters, allowing for statistical independence across natural signals (O. Schwartz and Simoncelli 2001; Coen-Cagli and Schwartz 2013; Odelia Schwartz and Coen-Cagli 2013; Sanchez-Giraldo, Laskar, and Schwartz 2019). Furthermore, normalization has

been shown to factorize the probability distribution function of natural images (Malo and Laparra 2010). Artificial networks also point to lateral inhibition producing adversarial effects that improve population coding by increasing the distance between stimuli and sparsifying the code, making it less prone to perturbations (Paiton et al. 2020). Additionally, in improving the statistical independence of linear inputs to a system, normalization has been shown to decrease mutual information between neurons (Valerio and Navarro 2003).

Additional work shows that normalization could have an influence on individual neuronal variability. Fluctuations in variability of individual neurons as well as shared variability directly influence the capacity of a system to represent and transmit information (Tripp 2012). This could contribute to improved sensory transmission.

In all cases, these theoretical studies can be used to generate testable hypotheses. At the same time, more complete genetically, spatially, and functionally labelled data from experiments can be used to further constrain models to continue to advance the field's understanding of cortical computation.

Shared neuronal variability and normalization: pairwise noise correlations

On average, neurons in a local cortical circuit exhibit small, positive noise correlations, in which the activity of pairs of neurons covaries (Smith and Kohn 2008; M.R. Cohen and Kohn 2011). These noise correlations can be shaped by external stimuli as well as by synaptic and cellular properties and network properties like balanced synaptic input, shared excitatory and inhibitory inputs, and feedback inhibition (Kohn and Smith 2005; Renart et al. 2010; Litwin-Kumar et al. 2011; Tetzlaff et al. 2012; Rosenbaum, Rubin, and Doiron 2013; Wolf et al. 2014). Information encoding depends on the activity of individual neurons as well as the relationship

between neurons, including pairwise correlations (Shadlen and Newsome 1994; Kohn et al. 2016). Theoretically, shared noise can prevent a population from averaging out noise; even weak noise correlations can considerably diminish the signaling of a population (Shadlen and Newsome 1994). Noise correlations have been theorized and shown to be information-limiting in some cases (Moreno-Bote et al. 2014; Kohn et al. 2016; Kafashan et al. 2021). However, noise correlations do not always necessarily reduce information, and different readout mechanisms can influence this relationship between correlations and information (Abbott and Dayan 1999; Seung and Sompolinsky 1993; H. Barlow 2001). In behavioral studies, reductions in noise correlations can largely explain performance improvements in attention (Marlene R. Cohen and Maunsell 2009). Altogether, considerable work has shown the correlations influence the scope of computation that is possible in a way that can be both calculated theoretically and measured experimentally.

In addition to the changes in sensory stimulus representation in individual neurons that are discussed above, there have also been efforts to model the effects that normalization would impose on population pairwise correlations. Tripp proposed that on the whole, normalization should decorrelate neuronal activity and reduce variability, improving information transfer (Tripp 2012). Indeed, in a study of normalization in awake, passively viewing macaques using cross-inhibitory stimuli, Ruff and Cohen found that overall, correlations decreased as pairwise normalization increased (Ruff, Alberts, and Cohen 2016). Furthermore, experimental studies relating normalization to neuronal response variability on a single trial basis found that in trials where normalization was inferred to be strong, pairwise correlations decreased in pairs of neurons in units with similar tuning preferences (Coen-Cagli and Solomon 2019). In this study, macaques were either anesthetized or passively viewing drifting gratings (Coen-Cagli and Solomon 2019). Neither of these studies distinguished between pairs of units with similar versus different orientation

selectivity as in Verhoef and Maunsell's work (Verhoef and Maunsell 2017), but they both show that normalization may have a powerful effect on shared variability in a population.

In addition to influencing noise correlations in passively viewing animals, normalization has also been implicated in behavioral studies, particularly in the changes of noise correlations observed in attention. Indeed, in Verhoef and Maunsell's 2017 study in macaques trained on an attention task, changes in pairwise noise correlations (shared neuronal variability) were shown to change in a more structured way according to pairwise relational orientation selectivity model, and in agreement with a stochastic normalization model (Verhoef and Maunsell 2017). This was discussed here previously, but briefly, it was shown that a normalization model can explain changes in pairwise correlation due to attention (Verhoef and Maunsell 2017).

In recent theoretical work focused on attention and shared variability, attentional modulation of neuronal variability was modelled and hypothesized to result from the top-down modulation of inhibition (Kanashiro et al. 2017). Indeed, subsequent work from the same group showed that low-dimensional shared variability could be recapitulated by a spiking neural network matching the temporal and spatial scales of inhibitory coupling found in cortex (Huang et al. 2019). In this model, modulating inhibitory neurons via a top-down signal could capture the changes in pairwise correlations observed in experimental studies of attention.

System with modifiable interconnections improves capacity to hold information, shown in synaptic studies (Gardner-Medwin 1976; H. Barlow 2001). To extrapolate, this concept can also be applied in the case of functional networks and interconnections that are modified in a stimulus- and context-dependent manner (Dechery and MacLean 2018; Verhoef and Maunsell 2017; Marlene R. Cohen and Maunsell 2009; Ruff, Alberts, and Cohen 2016). Functional networks that can be modified may also improve information capacity. In humans, attention can increase the

mutual information between V1 and MT, two distinct areas of visual cortex, indicating a potential avenue by which attention may influence cognition and a way to quantify the effects of cortical processes like attention (Saproo and Serences 2014). In each case, these improvements in information encoding can be explained with divisive normalization, a canonical brain computation (Matteo Carandini and Heeger 2011). Understanding the effects of normalization on population activity and the cell types that contribute to it may elucidate key steps of processing in the brain, and may provide insight into how the brain cascades normalization and other computations to generate complex behaviors.

While further work must be done to test the nuances and other hypotheses presented, theoretical modelling in combination with careful experimentation can offer guidance into the computational role of phenomena like normalization.

Vision as a model for studying cortical computation

Visual cortical processing has been extensively studied both to better understand vision proper and to better understand cortical processing. This dates back to early single-unit studies of visual receptive fields in cats (D. H. Hubel and Wiesel 1959; 1962). These early studies provided a rich foundation for characterizing the stimuli that evoke visual responses, the architecture that processes these responses across the cortex, and the mechanisms that underlie this processing in carnivores and primates (D. H. Hubel and Wiesel 1968; J. H. Maunsell and Essen 1983; D. G. Albrecht and Geisler 1991; Ferster and Miller 2000, 200; Fitzpatrick 2000; Duane G. Albrecht et al. 2002; Ringach, Shapley, and Hawken 2002; Ringach 2004). From this work, different models have emerged for visual responses and computation (David H. Hubel, Wiesel, and Stryker 1977; Sompolinsky and Shapley 1997; Duane G. Albrecht et al. 2002; Matteo Carandini et al. 2005;

Lennie and Movshon 2005; Priebe and Ferster 2008; 2012). Through a rich tradition of psychophysics and electrophysiology recording techniques, the careful manipulation of visual stimuli and behavioral tasks have related neuronal responses to external stimuli, cognition, and behavior (Britten et al. 1992; Marlene R. Cohen and Maunsell 2009; Lennie and Movshon 2005).

In recent decades, the mouse visual cortex has emerged as an important model for the study of vision and cortical computation, as mouse vision shares many important similarities with other mammalian visual systems (Dräger 1975; Grubb and Thompson 2003; Mangini and Pearlman 1980; Niell and Stryker 2008; 2010; Liu et al. 2009; Gao, DeAngelis, and Burkhalter 2010; Kerlin et al. 2010; Ma et al. 2010; S. L. Smith and Häusser 2010; Andermann et al. 2011; Bonin et al. 2011; Piscopo et al. 2013; Glickfeld, Histed, and Maunsell 2013; Glickfeld et al. 2013; Glickfeld, Reid, and Andermann 2014; Roth et al. 2016; Van Hooser 2007; Seabrook et al. 2017, 201; Huberman and Niell 2011). Vision in mice has a basic organization similar to other mammals (Seabrook et al. 2017; Huberman and Niell 2011). Mouse eyes are positioned laterally on a mouse's head, so most representations of the visual field in mouse V1 are monocular; however, there is a 40° binocular zone with inputs from both eyes (Seabrook et al. 2017). Photons are first detected by the retina; while mice lack a fovea, resulting in an even spatial frequency preference at a range of eccentricities (Niell and Stryker 2008) and a uniform cortical magnification (Kalatsky and Stryker 2003), the retina does include some ethologically relevant organization similar to carnivores and primates (Seabrook et al. 2017). Retinal ganglion cells project to the LGN in a retinotopic organization and also project to other visual areas like superior colliculus (SC). While the dLGN lacks the layers present in primate dLGN, there is a spatial organization of retinal ganglion cell input types in the mouse dLGN. In the dLGN, the visual input is spatially linearly filtered by center-surround cells with different spatial frequency preferences, with large center

diameters relative to other species (Grubb and Thompson 2003). This is analogous to filters observed in other species: mouse dLGN has similar visual receptive field and center-surround properties that are qualitatively comparable to other species studied (Tang et al. 2016).

The dLGN sends inputs to layer 4 of the visual cortex, although in addition to this classical pathway, there are direction selective RGCs that bypass layer 4 to innervate superficial layers of the cortex, a finding that is also observed in koniocellular neurons in primate LGN (Cruz-Martín et al. 2014; Cheong et al. 2013). There are at least eleven distinct visual processing areas indicated by distinct retinotopic maps in mouse cortex, suggesting different processing role for each visual region of cortex analogous to other species (Wang and Burkhalter 2007; Garrett et al. 2014). In the mouse visual cortex, the average contrast at half-maximum response (c50) is approximately 20%, the average preferred spatial frequency preference is approximately 0.04 cpd, and average preferred temporal frequency is approximately 2 Hz (Niell and Stryker 2008). Layer 4 sends projections to layers 2/3 and layer 5, where further processing takes place and units develop orientation and direction preferences. Inter-layer connections in addition to lateral and cortico-cortical processing come together to process stimuli in behavioral and context-dependent manners (Niell and Stryker 2010).

Neurons in mouse V1 share many core characteristics, including receptive field, contrast response functions, spatial and temporal frequency preferences, orientation and direction tuning, contrast invariant tuning, and pairwise correlations, with other mammalian species. Despite the poor visual acuity (Prusky, West, and Douglas 2000), behavioral studies have shown that mice can detect contrast changes and use vision in behavioral decisions, such as in hunting (Hoy et al. 2016; Busse et al. 2011; Glickfeld, Histed, and Maunsell 2013). Altogether, it is likely that much can be

learned about vision and cortical processing from using the mouse as a model organism with many properties that are conserved across species.

In addition to these important similarities, the mouse offers several other advantages and differences for studying cortical visual processing. Notably, mice are small and less expensive to work with compared to their larger mammalian counterparts. Aside from practical matters of accessibility and maintenance, the mouse's anatomy offers specific advantages that lend themselves to better understanding cortical processing. Namely, large populations of neurons can be easily accessed through both electrophysiology and optical imaging methods. The lack of curvature across mouse V1 and the lissencephalic nature of the mouse cortex enables the imaging of a coplanar population of cells in the superficial layers, in fields of view on the order of several hundreds of microns. This allows for optical access to large fields of view with hundreds of neurons, which in turn can inform the understanding of cortical processing on scales ranging from the single cell level to network models of cortical micro-circuitry (Andermann, Kerlin, and Reid 2010; Dechery and MacLean 2018). Additionally, the salt-and-pepper spatial distribution of tuning preferences in mouse V1 permits access to large populations of identifiable, diversely tuned neurons in one optically accessible field of view (Niell 2015). Furthermore, the mouse is genetically accessible and manipulatable, making extensive dissection of network, circuit, and molecular properties possible (Luo, Callaway, and Svoboda 2008; O'Connor, Huber, and Svoboda 2009; Niell 2015) (Luo et al 2008, O'Connor et al, 2009, Niell 2015). Finally, the visual and behavioral repertoire and capabilities of mice are under active development on an international scale, further improving the robust toolset available in mice for studying visual cortical processing (Cone et al. 2019; 2020; International Brain Laboratory 2017). Activity in mouse V1 is modulated by locomotor activity as well as auditory cues, further enriching the field of questions that can be

addressed in mouse V1 to cross-modality modulation (Niell and Stryker 2010; Iurilli et al. 2012; Ibrahim et al. 2016; Dadarlat and Stryker 2017).

In addition to offering genetic opportunities, mouse V1 offers a window into the network structures and properties of cortical activity. Cortical connectivity, both excitatory and inhibitory, directly influences sensory coding (Rodney J. Douglas and Martin 2004, 200; Ozeki et al. 2009; Harris and Mrsic-Flogel 2013). This extends beyond highly responsive and tuned neurons that can be reliably measured in extracellular preparations. For example, recent work has shown that both tuned and untuned neurons are present in networks, and the functional relationships between and within these subpopulations are important for predicting neural activity and decoding visual stimuli (Chambers and MacLean 2016; Dechery and MacLean 2018; Levy, Sporns, and MacLean 2020). These so-called functional networks built from pairwise functional dependencies can change according to context: for example, in slice preparations, neuromodulation by acetylcholine has been shown to reduce pairwise interactions between unreliable cells, pruning functional circuitry (Runfeldt, Sadovsky, and MacLean 2014). These studies emphasize the position of neurons relative to each other within a network. To glean the information within such complex networks, it is important to sample broadly in an unbiased manner from labelled neurons to best understand the structure and statistics of neural activity and propagation.

The mouse makes such sampling and study accessible. Two-photon calcium imaging using calcium indicators can be used to record to the activity of large populations of neurons at single-cell resolution near-simultaneously (Yuste and Katz 1991; Stosiek et al. 2003; MacLean and Yuste 2009). The advent of genetically-encoded calcium indicators (GECIs) has made transgenic lines that express indicator in specific classes of neurons widely accessible (Dana et al. 2014). These indicators report on the presence of free calcium in the cell, such as during an action potential.

Calcium ions bind to the indicator, allowing the molecule to undergo a conformational change that allows the attached fluorophore to fluoresce when stimulated with the appropriate wavelength (Yuste and Katz 1991).

Even with sensitive indicators, one-photon imaging faces the problem of a lack of spatial resolution: exciting a field rich in fluorophore could result in extensive fluorescence. While some experiments still use one-photon excitation with algorithms constantly developed to improve disambiguation and assignment of activity, two-photon imaging largely avoids this problem. With two-photon microscopy, first theorized by Maria Goeppert-Mayer in her 1931 doctoral thesis, targeted lasing of a focused point in three planes could be accomplished by requiring two photons of lower energy to coincide in space and time to excite the fluorophore (Mayer and Jensen 1955; Masters and So 2004). This theory has since been realized with the advent of femtosecond laser pulsing technology, in which packets of low energy photons are packaged densely into a high power pulse of exceedingly short (femtosecond) duration – increasing the probability of two photons coinciding in time and space on a sample (Denk, Strickler, and Webb 1990).

This technological advance enables recording from hundreds of neurons in a single field of view in vivo, providing a powerful window into network and population activity (Chen et al. 2013; Dechery and MacLean 2018; Levy, Sporns, and MacLean 2020). While raster imaging offers slow frame rates, recent advances such as path scanning and resonant galvanometers allow for faster frame rates. For example, heuristically optimized path scanning (HOPS) allows for frame rates of ~20-30 Hz in populations of several hundred neurons (Sadovsky et al. 2011). Additional recent work has produced additional approaches to achieve high-speed scanning using two-photon imaging (Lu et al. 2020). Altogether, these technological advances provide improved temporal

resolution for capturing network structures closer to the order of synaptic integration (Chambers and MacLean 2016).

In addition to excellent optical imaging access via two-photon microscopy, mouse offers optogenetic access, opening the door to the manipulation of spatially and genetically selected neurons (Boyden et al. 2005; Deisseroth 2015). Thus, hypotheses can be tested and generated via the principled perturbation of populations, circuits, and networks (Cone et al. 2019; 2020). Combined with two-photon calcium imaging and extensive advances in electrophysiology (Jun et al. 2017; Steinmetz et al. 2018), there is now unprecedented access to interrogate cortical processing in the mouse.

With the toolkit available in mouse in combination with the toolkit bequeathed by visual neuroscience, it is now possible to investigate the circuit underpinnings, the cell types, and functional networks of populations of neurons that enable cortical processing by studying mouse visual cortex (Niell and Stryker 2008; Huberman and Niell 2011; Dechery and MacLean 2018; Levy, Sporns, and MacLean 2020). In turn, these findings can guide a more complete understanding of vision and other cortical processes across species, states, modalities, and pathologies.

Evidence pointing to the possibility of normalization in the mouse

Because normalization is ubiquitous across species and modalities, it can be anticipated that the mouse will exhibit normalization in its neuronal responses (Matteo Carandini and Heeger 2011). While most mammalian studies of normalization in the visual system have involved cat or primate as a model system, there remains no visually evoked study of normalization in populations in mouse. However, as the mouse's visual cortex varies in organization compared to other

mammalian species that have been studied, it is important to characterize whether normalization in mouse differs in any meaningful way from species whose cortices are organized differently. Such work will guide the continued study the cortical circuitry of normalization and cortical computation more broadly and is an area of opportunity.

No purely visually-evoked studies of normalization have been conducted in the mouse to the best of our knowledge. Nevertheless, there are some conflicting studies that point to a mixed likelihood of normalization in the mouse.

Recent work has posited that mouse visual responses are largely summation-based (Histed 2018). Those results are based on a non-specific, non-visual set of stimuli, including broad optogenetic activation, that may have influenced local circuit dynamics. Furthermore, the very low firing rates of a large proportion of units could have made it difficult to detect normalization, as the normalization equation's numerator is comprised of the driven response, so very low firing rates will yield unreliable measures for normalization. Finally, the measurements were taken from multiple layers of cortex and were not separated based on layer, raising the question of whether the variability in summation and normalization were related to this.

In another recent study, Sato and colleagues used optogenetic perturbations to activate neurons in the binocular zone contralateral to the hemisphere in which they patch-clamped neurons responding to a set of visual stimuli in an anesthetized mouse (Sato et al. 2016b). The neurons appeared to decrease activity during concurrent optogenetic stimulation, suggesting that normalization may play a role in mouse visual processing. However, it is possible that artifacts from optogenetic stimulation could have been measured, necessitating a purely visually-evoked study of normalization in mouse. While this study supports the possibility of normalization in mouse visual cortex, it is necessary to study cortical responses to a visual stimulus.

Since the early descriptions of normalization, there has been some effort to identify the cellular or circuit-level mechanisms that underlie the computation. In Wilson and Sur's paper, they attempt to activate either parvalbumin or somatostatin neurons to measure their respective effects on neighboring neurons during the presentation of a drifting grating in mice. Here, the lab worked to suggest that interneurons subtypes are involved in division and subtraction, suggesting that parvalbumin interneurons are involved in dividing the activity of a neuron. While this is an intriguing possibility, it does not inform the understanding of whether under physiologic conditions, mouse visual cortex shows visually-evoked normalization. Nevertheless, such studies highlight the promise of exploring normalization circuitry using a visually-evoked paradigm of normalization in mouse (Wilson et al. 2012).

Aims of this thesis: Critical gaps in knowledge that will be addressed

As an ostensibly ubiquitous and conserved computation, normalization can teach us about more abstract principles relevant for brain processing, not least of all the mechanisms that drive changes in pairwise noise correlations. To better understand attention and associated disorders, it is imperative to better understand normalization, including its circuit underpinnings. Yet no mouse model of visually-evoked normalization exists to dissect and manipulate these circuits. Although divisive normalization may have implications for understanding brain processing and brain disease, little is known about its dynamics, cell types, and how the properties of each vary in different behavioral states such as locomotion and attention.

This project studies the awake mouse V1 as a model for divisive normalization. Here, we will probe divisive normalization in a mouse model of visually-evoked divisive normalization using a synergistic set of genetic, optical, electrophysiology, optogenetic, and computational

techniques. We anticipate that this research will open the door to further genetic dissection, manipulation, and behavioral experiments to guide a mechanistic understanding of divisive normalization and its potential links to disease.

Normalization is a canonical computation that processes information for further processing downstream. When core brain processes go awry, pathologies may ensue. There is some evidence suggesting that deficits in normalization are observed in various disease states ranging from epilepsy and major depression to schizophrenia (Matteo Carandini and Heeger 2011). Much work remains to better understand normalization and its role in normal brain processing as well as brain pathologies.

Better understanding normalization will elucidate the cortical information processing that undergirding complex behaviors. The mouse offers a powerful opportunity for studying normalization, including better understanding the single-cell, pairwise correlation-level, and network structure-level characteristics of normalization. With such a model and the myriad genetic tools available in mice, it may also be possible to further dissect the circuitry that produces normalization, a ubiquitous and conserved brain computation.

CHAPTER 2

MATERIALS AND METHODS

Mice

C57BL/6J (B6) mice and mice expressing GCaMP6s under the Thy-1 promotor (Dana et al. 2014; Chen et al. 2013) were purchased from the Jackson Laboratory, then bred and maintained in the Animal Research Center at the University of Chicago. All mice were bred and maintained in accordance with the animal care and use regulations of the University of Chicago Institutional Animal Care and Use Committee. Male and female mice age P60-P200 were used across experiments.

Craniotomy and window implantation

Anesthetized adult mice were surgically implanted with a head bar and 3mm or 3.5mm chronic imaging window centered over V1, as described in elsewhere. Briefly, mice were anesthetized using isoflurane (1.5-2.0% in 50% O_2) with (electrophysiology experiments) or without (two-photon experiments) ketamine (40 mg/kg, i.p.), and xylazine (2 mg/kg i.p.). Mice were provided buprenorphine for pain management as well as dexamethasone and meloxicam to reduce inflammation in advance of surgery. Upon reaching a state of deep anesthesia as evidenced by lack of response to a toe pinch, each mouse's head was shaved and an incision made to reveal the cranium. The fascia was removed and the surface of the skull lightly scored with a scalpel blade to improve adhesion of dental cement, then a titanium head post was cemented in place with dental cement (C&B Metabond, Parkell). Next, a 3.0-3.5 mm craniotomy was drilled, centered over left visual cortex at 4.0 mm lateral from the midline and 0.5 mm anterior of lambda.

Throughout the surgery, the surface of the cortex was kept irrigated. After the craniotomy, the dura was also removed to minimize the likelihood of bone regrowth. An optical window (0.8 mm thick; Tower Optical) and coverslip were coated with silicone grease and placed over the craniotomy, then cemented in place with dental cement. Stitches were placed in the surrounding skin to improve healing and minimize the risk of infection. The mice were allowed to recover before further experimentation.

Intrinsic imaging to identify primary visual cortex

In all experiments, V1 was identified by intrinsic imaging (Kalatsky and Stryker 2003; Andermann, Kerlin, and Reid 2010). At least two weeks after the craniotomy, intrinsic imaging was conducted to identify V1. In the two-photon experimental animals, a repeating drifting bar was presented to anesthetized mice (1.25% isoflurane, 50% O_2) approximately 23 cm from the right eye during intrinsic imaging to generate a retinotopic map of responses in V1. In the electrophysiology experiments, mice were head-fixed and awake during the presentation of a stimulus set shown on a screen anterior to the mouse. Drifting Gabor stimuli were presented in each of five locations tiling binocular and monocular zones in the right hemifield. For both groups of experiments, a retinotopic map of visual cortex was produced to guide imaging or electrode placement.

Stimulus presentation

After recovery, mice were head-fixed for stimulus presentation (Dechery and MacLean 2018). Stimuli were displayed with gamma correction on a monitor positioned approximately 23 cm from the mouse's right eye, and covering ~60° x ~70° of the visual field. Two sets of stimuli

were shown: full-contrast gratings drifting in eight or twelve directions to measure direction tuning and preferences, similar to previous experiments, and a cross-inhibitory stimulus to elicit divisive normalization.

In the two-photon experiments, drifting grating stimuli in eight equally spaced directions were shown in a pseudorandom order to determine the responsivity and tuning of each neuron. The drifting grating stimuli had a square-wave spatial profile (0.04-0.06 cpd, 2 Hz) and were shown for 4 s per direction, interleaved with 3 s of a luminance matched gray screen to allow the calcium indicator to decay back to baseline. Each stimulus direction was shown at least 20 times.

In the electrophysiology experiments, drifting grating stimuli in twelve equally spaced directions were shown in a pseudorandom order to determine the responsivity and tuning of each neuron. The drifting grating stimuli again had a square-wave spatial profile (0.04 cpd, 2 Hz) and were shown for 1 s per direction, interleaved with 1 s of a luminance matched gray screen. Each stimulus direction was shown at least 20 times.

To assess divisive normalization, a cross-inhibitory stimulus composed of orthogonal gratings similar to previous studies in cat and macaque were presented. Cross-orientation inhibition was selected for studying normalization because it is the most robust and well-established stimulus for exploring divisive normalization in V1 and has been validated in multiple studies of divisive normalization (M. Carandini and Heeger 1994; Busse, Wade, and Carandini 2009). Spatially offset stimuli are not appropriate for V1 receptive fields, because they push the limits of visual acuity (Niell 2015). In each experiment type, the cross-orientation stimuli consisted of orthogonal pairs of square-wave gratings superimposed to create a plaid, presented with varying contrasts (each at 0%, 6%, 12%, 25%, or 50% contrast). In six out of the eight calcium imaging datasets, stimuli were counter-phased to mitigate the low possibility of eliciting pattern motion

direction responses. In two of the eight calcium imaging datasets, the stimuli were drifting. We did not detect differences in the results from these two stimulus types so we combined the data from both groups.

Twenty-five contrast combinations of the two orientations were formed and were presented in a pseudorandom order that varied by recording day but remained constant during a given recording session. The sequence of stimuli was fixed each session to limit the effects of preceding stimuli on noise correlations. Three dummy stimuli were presented at the beginning of each recording session and later excluded to mitigate the effects of adaptation. In the two-photon imaging experiments, each stimulus was shown for 5 s, interleaved with a 3 s luminance-matched gray epoch. The onset and offset of stimuli on the video display were tracked by a photodiode and a small black (off) or white (on) square in the bottom corner of the monitor. Each stimulus was shown 20 times. In the electrophysiology experiments, the stimuli were shown for 1 s, interleaved with a 1 s luminance-matched gray epoch. Each stimulus condition was presented 100 times.

Two-photon imaging

Layer 2/3 of V1 was imaged using a custom-built two-photon microscope. The awake mouse was secured in a head-fixed position on a custom-made linear treadmill. Treadmill velocity was measured with a rotary encoder (Niell and Stryker 2010; Dechery and MacLean 2018). The mouse's right eye was oriented towards a monitor approximately 23 cm away. A light shield between the objective and the imaging window was applied to block any light from the imaging field. Then, a field of view (FOV) was identified in layer 2/3 of V1 using raster scanning. After a suitable FOV was identified, a 4x4 grid of raster scans was acquired of the FOV to allow for automatic neuronal detection. Neurons were identified by an automated algorithm (Dechery and

MacLean 2018); a typical FOV consisted of 200-400 identified excitatory neurons. Subsequently, an optimal path for scanning was determined for the population in the field of view and the HOPS system was used to image neurons via path scan at ~30Hz (Dechery and MacLean 2018; Sadovsky et al. 2011). V1 is advantageous for these specific experiments because of its surface accessibility, minimal curvature, and well-documented association with reportable visually-evoked activity (Niell 2015; Dechery and MacLean 2018). Optical recording is advantageous because it permits simultaneous recording from large numbers of individual neurons with known genetic and spatial identities and relationships.

Electrophysiology

Extracellular electrophysiological recordings were conducted to cross-validate findings from optical recordings. Silicone multi-channel electrodes (NeuroNexus Technologies; 32-site model 4x8-100-200-177) that could span multiple layers of cortex were used. The electrode contacts were plated to impedances between 200 and 500 kΩ at 1000 Hz using a mixture of gold solution and carbon nanotubes (Cone et al. 2020). At the start of the recording experiments, mice were head-fixed and anesthetized (1.25-1.5% isoflurane, 50% O_2) and the eyes irrigated with 0.9% saline. The optical window was removed and the recording site identified. The electrode was inserted to a depth of 800-1000 μm in the V1 representation of the lateral visual field. The craniotomy was covered with 2.5% agarose (MilliporeSigma) dissolved in aCSF (Tocris Bioscience) and the resultant setup was covered in silicone oil to prevent desiccation (MilliporeSigma). After the electrode was placed, the anesthesia was removed and the mouse allowed to awaken. After an hour of recovery time in which the electrode settled into the cortex, the recording commenced. To minimize movement artifacts, mice were awake but stationary in

custom plastic sleds during recording sessions. Signals from the recordings were amplified, filtered, and sampled at threshold crossings (BlackRock). Spikes were be sorted offline (Offline Sorter, Plexon).

Normalization metric

The resultant fluorescence traces were analyzed for characterization of normalization in mouse V1. We used a normalization index defined as:

$$NI = \frac{(R[p] + R[n]) - R[p+n]}{(R[p] + R[n]) + R[p+n]}$$

Here, R[p] is the response to the 50% contrast in the preferred direction, R[n] is the response to the 50% contrast in the non-preferred direction, and R[p+n] is the response to the 50%/50% plaid composed of both the preferred and non-preferred directions. Behavioral state, as encoded by running velocity, was considered in the calculation of this metric (Dipoppa et al. 2018; Niell and Stryker 2010; McBride, Lee, and Callaway 2019). In our two-photon imaging data, mice ran for ~5% of total stimulus presentations. In some experiments, mice did not run at all. While there was a slight increase in the normalization index when including running trials versus excluding running, there was no difference in the distribution of our measure of normalization between these cases (Supplemental Figure 3-1), so all data were pooled in further analyses. Furthermore, because running trials only composed ~5% of total trials, these data were not analyzed separately.

Spike sorting

Units were identified as clusters of waveforms that were separated from the noise and from other clusters throughout the duration of the recording. Additionally, units were inspected for inter-spike interval histograms consistent with absolute and relative refractory periods and only units consistent with single or small multi-unit clusters were included in our analyses.

Pairwise noise correlations

In the two-photon experiment analysis, for each pair of imaged neurons, we computed pairwise noise correlations as the mean partial correlation between the first two seconds of neurons' fluorescence changes across individual stimuli, while accounting for three control variables: mean response of each of the two neurons in all other presentations of the stimulus (to account for the signal correlation) and within-stimulus mean fluorescence of all other neurons (to account for population-wide co-variability like behavioral state and running speed effects). The first two second segments of the response a neuron to the same stimulus were concatenated in groups of five and smoothed with a running average window of 10 frames in order to remove discontinuities due to concatenation. The Matlab (MathWorks) function *parcor* was used to compute the pairwise partial correlation between the response vectors. Correlations were calculated separately in this way for each stimulus condition.

Noise correlations were calculated for the electrophysiology data as previously described (Verhoef and Maunsell 2017). Specifically, Pearson's correlation coefficient was used to compute spike-count correlations between spike counts from pairs of units. The spike count window was 250 ms starting 50 ms after stimulus onset to stimulus conclusion.

Analysis

Data were analyzed using custom-written Matlab (MathWorks) scripts. In the case of the two-photon imaging experiments, raw data was preprocessed to calculate changes in fluorescence over baseline as previously described (Dechery and MacLean 2018). For both two-photon and electrophysiology data, direction and orientation tuning using statistical criteria and tuning fit was assessed by identifying visually responsive neurons and fitting an asymmetric Gaussian to visually evoked activity as described previously (Mazurek, Kager, and Van Hooser 2014; Dechery and MacLean 2018; Levy, Sporns, and MacLean 2020). The non-preferred suppression index was calculated from the fitted non-preferred alpha parameter of the tuned normalization model (Ni, Ray, and Maunsell 2012; Verhoef and Maunsell 2017). Statistical analyses were completed and plots were generated using custom scripts written in Matlab (MathWorks).

CHAPTER 3

MOUSE VISUAL CORTEX EXHIBITS VISUALLY-EVOKED NORMALIZATION

Introduction

Normalization can be considered a canonical brain computation; in normalization, the activity of a neuron is normalized relative to the activity of its neighbors and to the relative strength of its preferred stimuli to other stimuli (Matteo Carandini and Heeger 2011). The computation is conserved across species and modalities, suggesting that it plays an fundamental role in signal processing (Matteo Carandini and Heeger 2011). Furthermore, normalization underlies multi-sensory integration, and has also been shown to underlie decision-making in macaques (Ohshiro, Angelaki, and DeAngelis 2011; Louie, Grattan, and Glimcher 2011).

Recent work has shown that different neurons can exhibit different normalization strengths. This has been described as tuned normalization, in which different units and neurons show different extents of normalization ranging from linear summation to simple averaging of responses to super-normalization (Ni, Ray, and Maunsell 2012; Ni and Maunsell 2017). This tuning property has been shown to have a close relationship with attention, as neurons that have stronger normalization are also more likely to show changes in firing rate due to attention (Ni, Ray, and Maunsell 2012). Recent work has shown that indeed, normalization underlies the influences of attention on changes in both firing rate and pairwise noise correlations (Lee and Maunsell 2009; Ni, Ray, and Maunsell 2012; Ni and Maunsell 2017; Verhoef and Maunsell 2017) that largely explain attention's effects on behavioral improvements (Marlene R. Cohen and Maunsell 2009) Specifically, attention influences pairwise noise correlations by biasing normalization mechanisms

(Verhoef and Maunsell 2017). It is likely that understanding normalization will lead to improved understanding of attention – among other important brain phenomena.

To better understand normalization, it will be important to understand the circuit mechanisms and cell types that are involved in the computation. Mouse offers distinct advantages for such a circuit and cell type level of analysis. For example, the mouse visual cortex offers many advantages for genetic and optical access (Luo, Callaway, and Svoboda 2018; Niell 2015). Prior to this study, it was not known if mouse V1 exhibits visually-evoked normalization, and if so, whether it exhibits similar properties, like tuned normalization, that have been observed in other species like the macaque (Ni, Ray, and Maunsell 2012). Studying normalization in the mouse visual cortex would offer many advantages for dissecting the circuit underpinnings of normalization. Additionally, establishing a mouse model of normalization can pave the way to mechanistic study of how pairwise correlations are generated in cortex.

Here, we study normalization in mouse visual cortex using both two-photon imaging and electrophysiology. Specifically, we investigate the visual response properties of awake mice to cross-inhibitory stimuli that classically evoke normalization (M. Carandini and Heeger 1994; Matteo Carandini, Heeger, and Movshon 1997; Busse, Wade, and Carandini 2009). We demonstrate that mouse visual cortex exhibits tuned normalization similar in quality to that observed in macaques (Ni, Ray, and Maunsell 2012). However, we also note that there are some differences, including the somewhat weaker strength of normalization in mouse. Finally, we report on recording-depth differences observed in our electrophysiology studies, highlighting the availability of the mouse to uncover new information about normalization.

Results

Cross-inhibitory stimuli elicit normalization in mouse visual cortex

Despite a burgeoning role for the mouse in the study of visual cortical processing (Niell 2015), no previous studies have examined visually-evoked normalization in the mouse. Here, we recorded from awake, head-fixed mice. We recorded in the mouse visual cortex using both two-photon imaging and electrophysiology and observed normalization in visually responsive units.

In the two-photon imaging experiments, we recorded from a total of 2,120 excitatory neurons from layer 2/3 in eight imaging sites and five mice (Figure 3-1). Drifting grating stimuli were presented in eight full-contrast drifting grating stimuli equally spaced around a circle to determine visual responsiveness and tuning preferences. Of the neurons that were imaged, 1,689 were responsive to drifting grating stimuli, and of those, 1,211 were tuned to one of the eight drifting grating directions (Figure 3-2). As mouse V1 shows a salt-and-pepper distribution of direction and orientation tuning (Niell 2015), our imaging fields of view captured diversely tuned neurons. The proportions of responsive and tuned neurons are in agreement with previously reported measurements from mouse visual cortex (Niell and Stryker 2008).

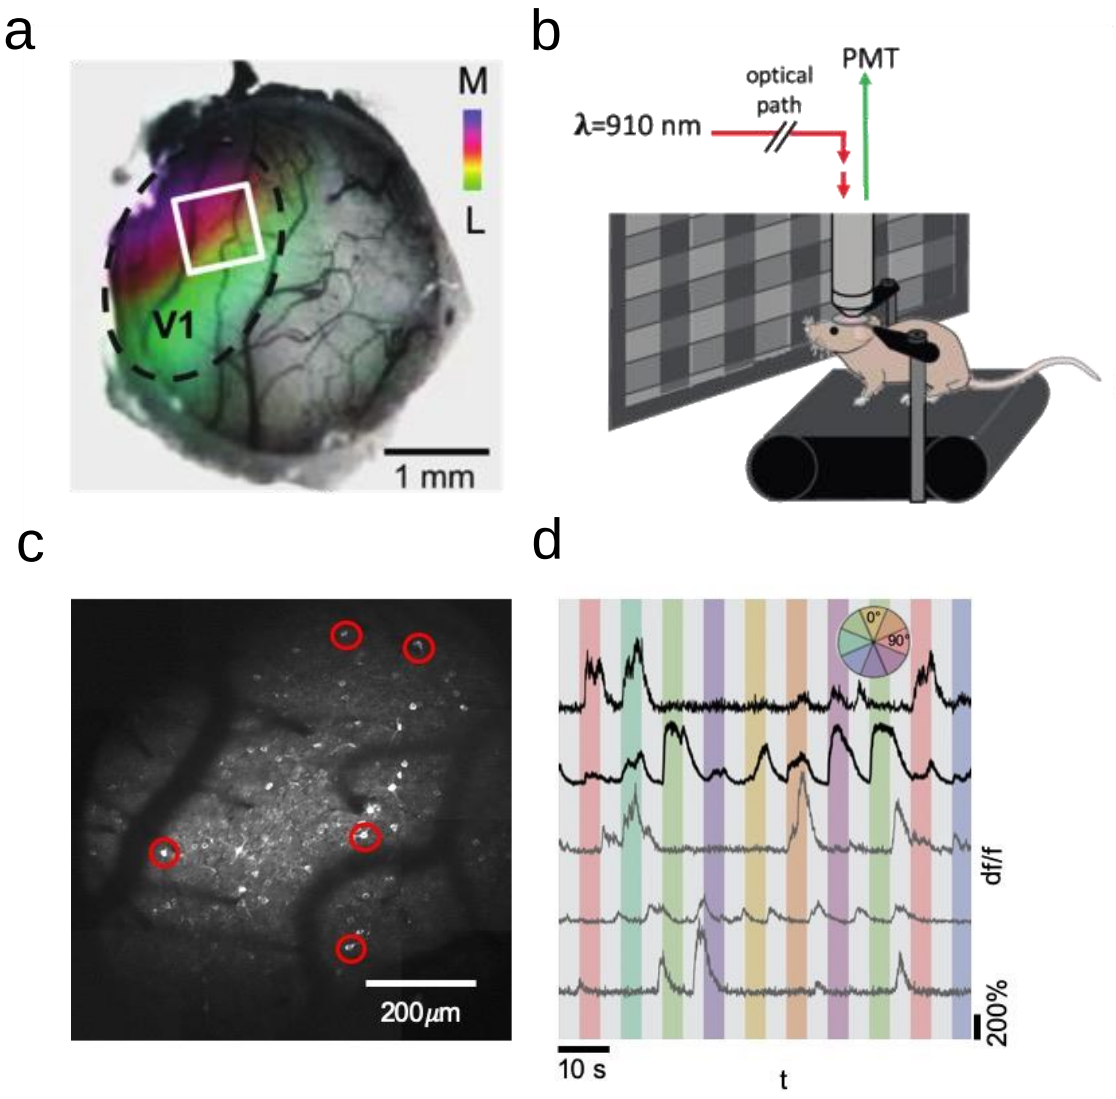

Figure 3-1. Two-photon imaging in awake mouse V1

a, Intrinsic imaging of cortex identifies visual cortex. Dashed oval represents V1; white square represents the approximate field of view for two-photon imaging shown in Figure 3-1c. **b,** Schematic of two-photon imaging. The mouse was head-fixed for duration of recording from the left visual cortex while stimuli were presented to the contralateral (right) eye. During recording, mice were able to run on a treadmill. A photodiode was used to temporally match stimulus presentation and recording. **c,** An example field of view from mouse V1 (identified in Figure 3-1a). Red circles label five example neurons, from which example traces are shown in Figure 3-1d. **d,** Example fluorescence traces from neurons identified in Figure 3-1c. Different colors represent different direction stimuli; the gray epochs represent interleaved gray screens. The black traces represent the responses of example tuned neurons and the gray traces represent the responses of example untuned neurons.

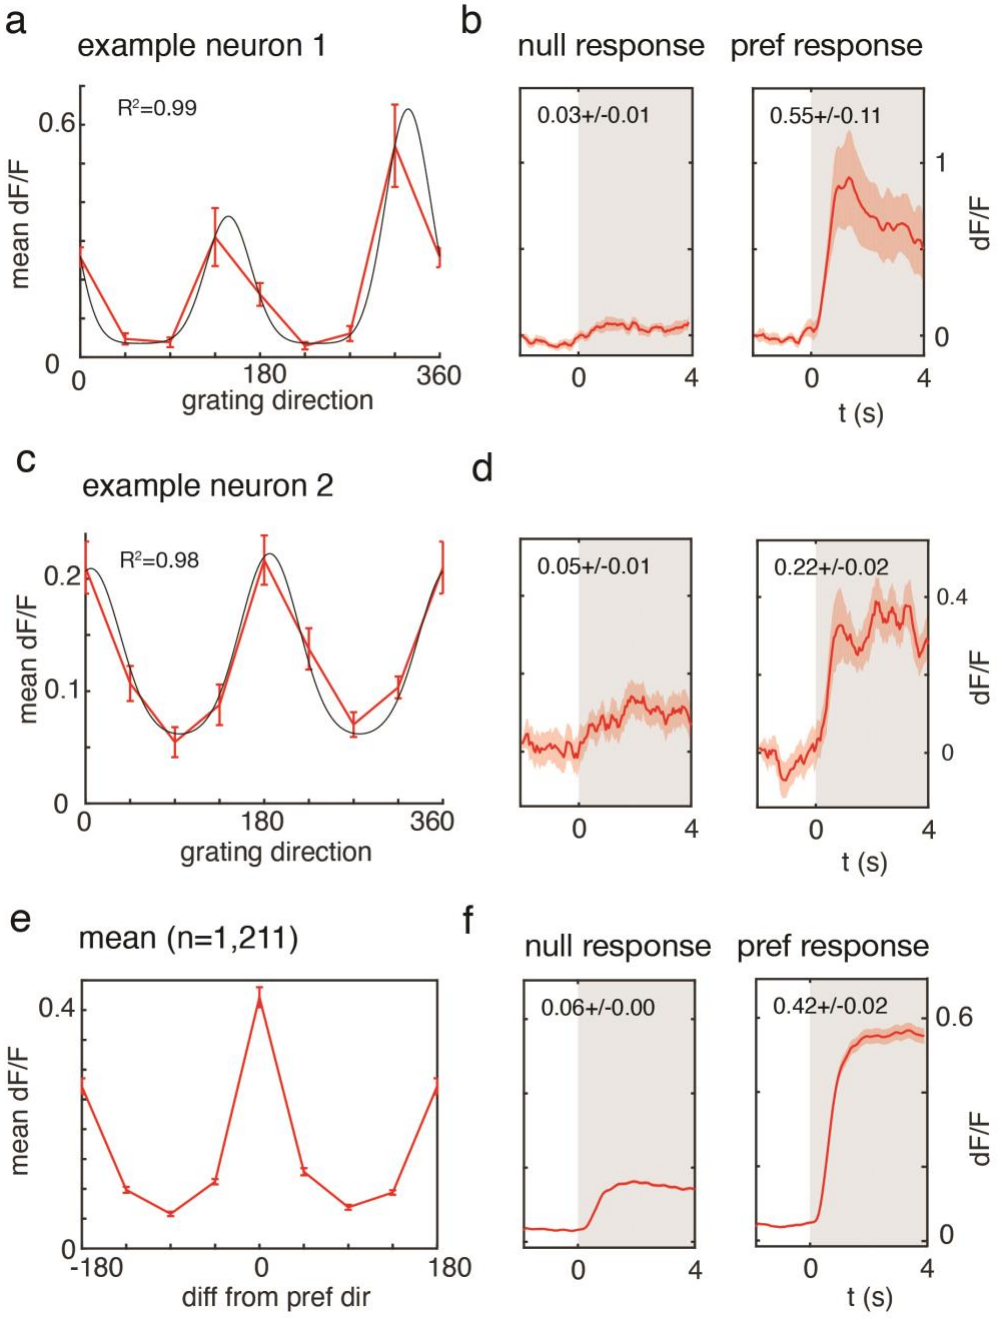

Figure 3-2. Orientation and direction tuned neurons imaged in two-photon experiments

a, An example tuning curve from a direction-selective neuron in mouse V1. The red line represents mean df/f of responses to drifting grating stimuli and the black line represents the tuning curve fit (r^2=0.99). **b,** Mean response of the neuron shown in Figure 3-2a to the drifting grating stimulus in its null and preferred stimuli. Gray represents the period of stimulus on. **c-d,** Same as Figure 3-2a and Figure 3-2b for an orientation-selective neuron (r^2=0.98).

Figure 3-2, continued

e-f, Same as Figure 3-2a and Figure 3-2b, except averaged across all tuned neurons recorded (n=1,211). Error bars and bands represent ±1 SEM.

Next, we presented a set of counter-phasing cross-inhibitory stimuli while imaging from the same field of view. The stimuli were designed to elicit normalization, as previously shown in cats and macaques (M. Carandini and Heeger 1994; Matteo Carandini, Heeger, and Movshon 1997; Busse, Wade, and Carandini 2009). Specifically, the stimulus set consisted of 25 combinations of cross-oriented gratings with component contrasts ranging from 0% to 50%. In classic normalization, a response to a cross-oriented stimulus composed of two component gratings of equal stimulus intensity would be an equally weighted average of the response to the two individual component gratings. In contrast, in linear summation, the response to the cross-oriented stimulus would be a simple linear sum of the responses to the two individual component gratings. In tuned normalization, an unevenly weighted average of the responses to the individual component gratings would be expected (Ni, Ray, and Maunsell 2012).

To quantify normalization strength in our experiments, we used a normalization index (NI), defined as the difference between the true response to the cross-oriented (plaid) stimulus and the linear sum of responses to the component stimuli (gratings), all over the sum of these two quantities. In this quantification, a normalization index of zero represents linear summation and a normalization index of 0.33 represents classic normalization. Values between 0 and 0.33 represent tuned normalization; values greater than 0.33 represent super-normalization, a differently weighted example of tuned normalization. Values less than zero represent supra-linear summation.

After the normalization stimulus set was presented, responsive neurons were identified. Our stimulus set only included two orientations, which were not the preferred stimuli for most neurons. Thus, on average, the mean change in fluorescence among these neurons was low. From among tuned and visually-responsive neurons identified using the drifting grating stimulus, 1,187 were deemed responsive to either the vertical, horizontal, or plaid counterphase stimuli at 50%

contrast according to a Kruskal-Wallis test with alpha set to 0.1. Furthermore, to manage irreducible noise and its systematic effect on the normalization index metric, we set a threshold for minimum difference in the df/f compared to the gray. We measured at which level of change in fluorescence over baseline we could stably measure normalization, setting the threshold to a minimum df/f change of 0.1 over baseline in either the plaid or preferred component response (Supplemental Figure 3-2). Using this criterion, we ultimately included 341 neurons for further normalization analysis.

Among the units included, we observed clear responses to the stimuli presented (Figure 3-3). The responses measured during the presentation of the component stimuli were weaker those measured during the presentation of the drifting gratings used for the direction tuning measurements, but this was expected because the stimuli presented in the normalization stimulus set were composed of components at a maximal contrast of 50%, whereas the stimuli presented in the direction tuning measurements were set to 100% contrast. Nevertheless, we observed responses that increased with increasing contrast in both the preferred and non-preferred orientation, and a normalized response to the plaid stimuli on average. Some neurons showed approximately classic normalization (Figure 3-3a), while others showed tuned normalization (Figure 3-3b).

Overall, we observed a range of normalization strengths. On the whole, we found the median normalization index to be 0.15 (IQR=0.03-0.25) among tuned, responsive neurons in mouse visual cortex (Figure 3-4). This range represents a departure from classic normalization, in which the expected measurement for normalization is 0.33. The range of normalization strengths observed indicates tuned normalization, in which neurons exhibit different levels of non-preferred stimulus suppression. This phenomenon was first described in the macaque, in which units were shown to have normalization strengths ranging from linear summation (normalization index of

zero) to classic normalization (normalization index of 0.33) (Ni, Ray, and Maunsell 2012) (Ni 2012). The data in mouse are qualitatively in agreement with previous studies of tuned normalization in macaques, indicating tuned normalization in the mouse. However, the overall normalization strength in mouse is somewhat weaker than that reported in macaque area MT, in which population normalization medians were closer to ~0.2 (Ni, Ray, and Maunsell 2012).

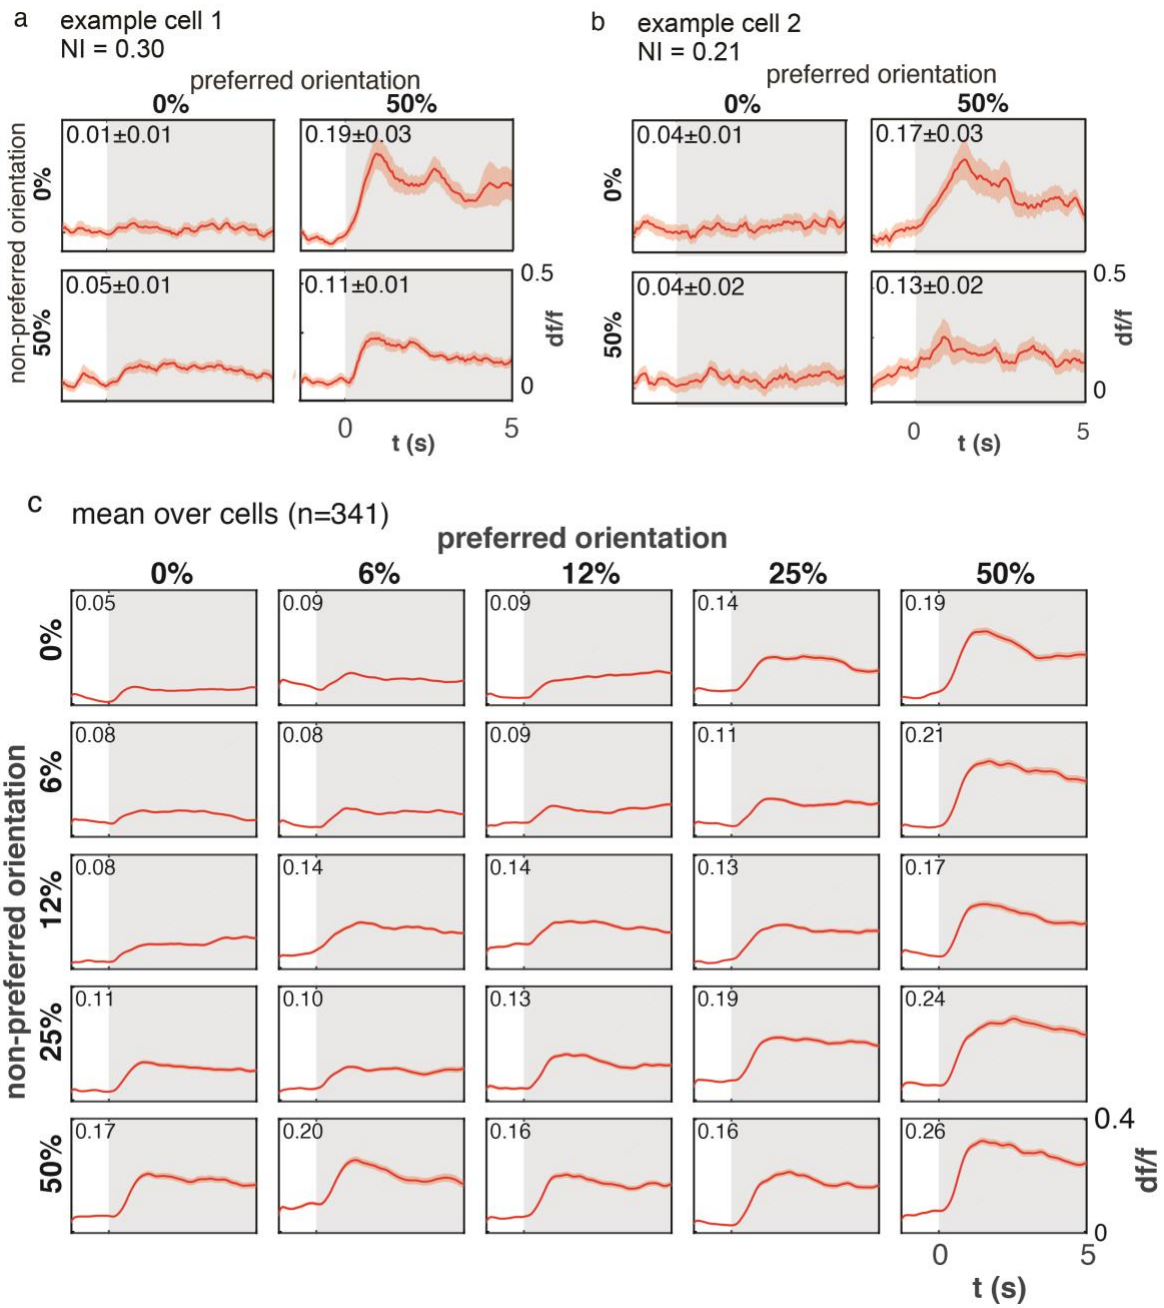

Figure 3-3. Visually-evoked normalization in mouse V1 measured using two-photon imaging

a-b, Responses of example tuned neurons to normalization stimuli (same neurons as in Figure 3-2). On average, the units show robust responses to the preferred stimulus orientation relative to baseline (across horizontal axis) and some response to the non-preferred stimulus orientation (vertical axis). When the stimuli are superimposed to create a cross-inhibitory stimulus, the driven response to the resulting stimulus is sublinear. Each stimulus was presented 20 times. Error bands represent ±1 SEM. **b**, Mean normalization stimulus set responses across neurons. All mean response SEMs were less than or equal to 0.01. The SEM bands are not visible at this scale.

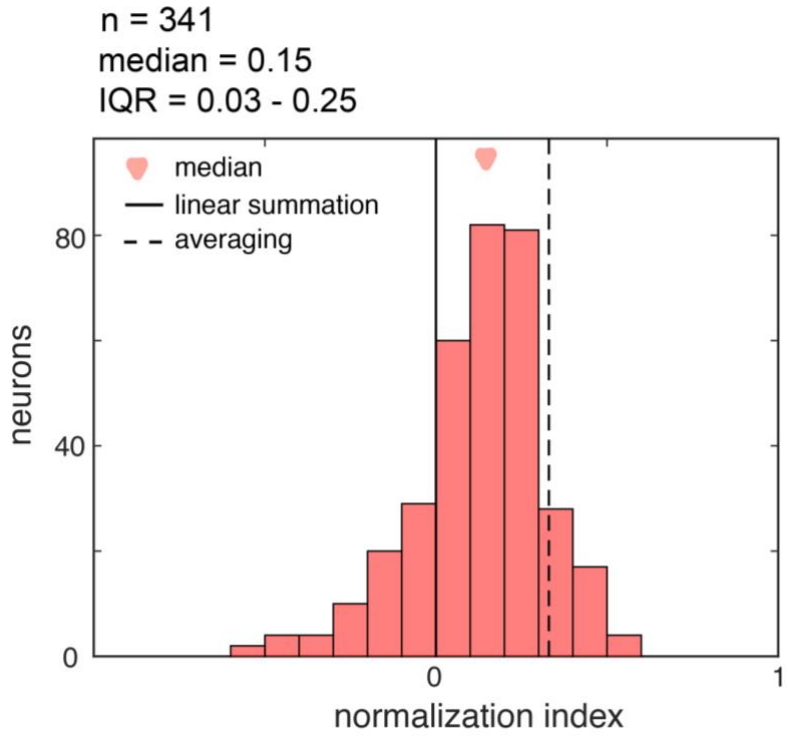

Figure 3-4. Population histogram of normalization strength in mouse V1 measured using two-photon imaging

Normalization index distribution across neurons (n = 341) as measured using two-photon imaging (median = 0.15, IQR = 0.03 - 0.25). An index of zero represents linear summation of responses; the dashed line at 0.33 represents the value for averaging of responses. The median is indicated by the pink triangle.

Next, electrophysiology experiments were conducted to provide improved temporal resolution, cross-validation, and recording at different depths. We used multi-channel silicone probes with electrode sites spanning 700 μm of V1 (Figure 3-5). We recorded from awake, head-fixed mice while presenting visual stimuli. First, we presented a set of drifting grating stimuli to determine visual responsiveness and tuning preferences, as in the two-photon imaging experiments. Overall, units were recorded from 13 sites from seven mice. We recorded stably from 114 units and small multi-unit clusters that had an evoked firing rate at least 0.3 spikes/s. Of those recorded, 72 units were responsive according to a Kruskal-Wallis test with alpha set to 0.1. Of those, 63 were tuned, showing a preference to either orientation or direction (Figure 3-6). Those 63 units were included in further analysis.

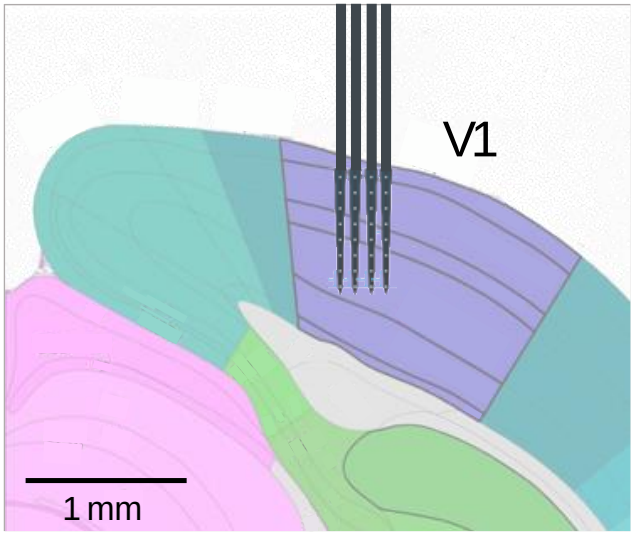

Figure 3-5. Electrophysiology recordings in awake mouse visual cortex

Schematic of electrode placement in mouse visual cortex (purple). Electrodes with four linear probes (NeuroNexus), including eight channels spaced at 100 microns on each probe, were inserted in visual cortex to depths of 800-1000 microns using a micromanipulator and according to intrinsic imaging results. This image was adapted from the Mouse Brain Atlas (Allen Institute).

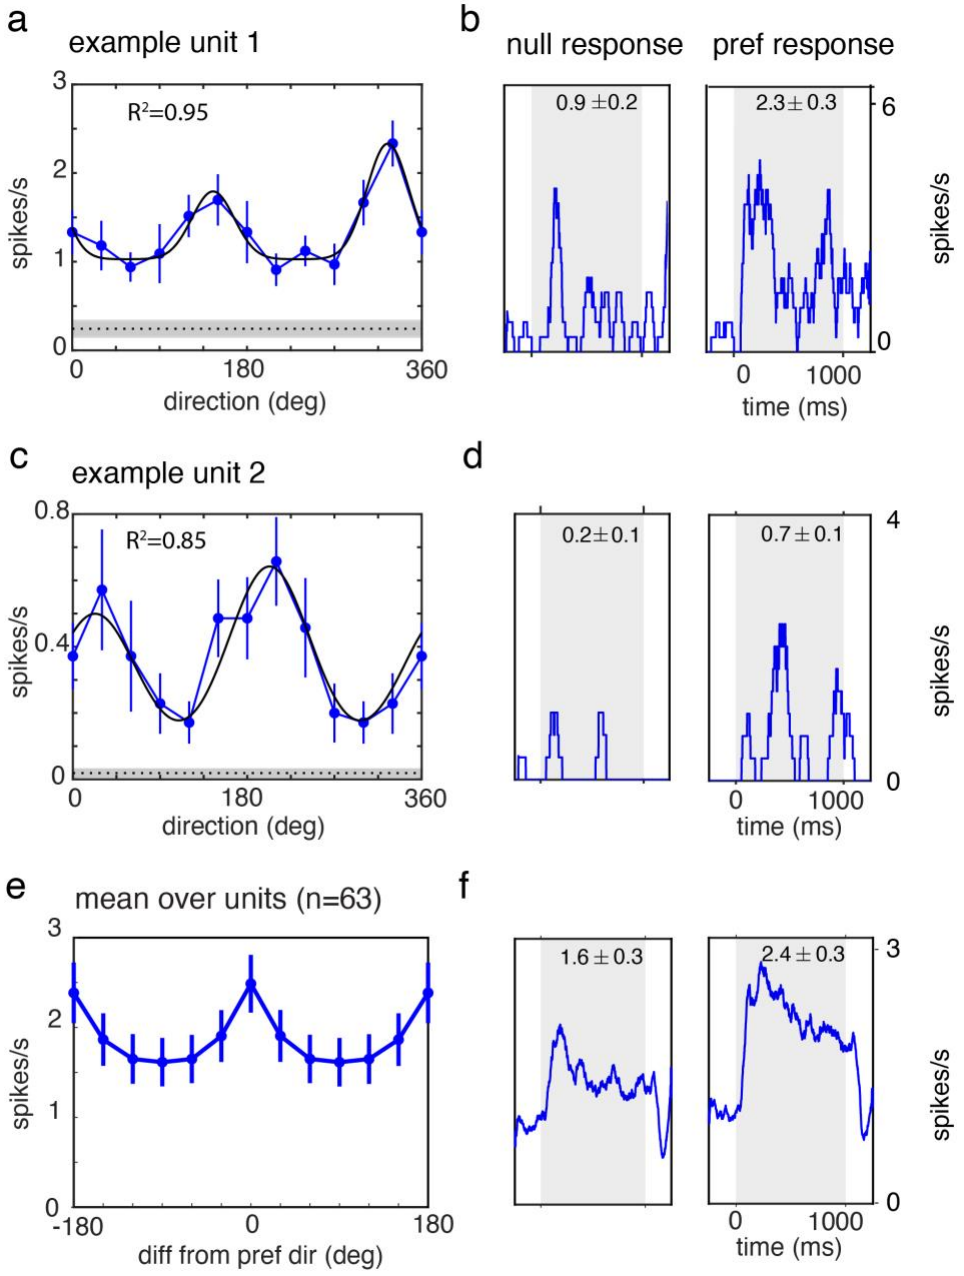

Figure 3-6. Orientation and direction tuned units recorded in electrophysiology experiments

a, An example tuning curve from a direction-selective unit in mouse V1. The blue line represents mean responses to drifting grating stimuli and the black line represents the tuning curve fit (r^2=0.95). The dashed line represents mean spontaneous activity. **b,** Mean response of the neuron shown in panel a to the drifting grating stimulus in the orthogonal orientation and preferred direction stimuli. Gray represents the stimulus on period.

Figure 3-6, continued
c-d, Same as Figure 3-6a and Figure 3-6b for an orientation-selective neuron (r2=0.85). **e-f,** Same as Figure 3-6a and Figure 3-6b, except across all tuned neurons recorded (n=63). All error bars and bands represent ±1 SEM.

Next, we presented the normalization stimulus set of 25 pseudo-randomly ordered plaids, again consisting of combinations of component gratings with contrasts ranging from 0% to 50%, as in the two-photon imaging experiments. We observed visually-evoked responses to the component stimuli that increased as contrast increased, but which showed normalization when the stimulus components were superimposed to form a plaid (Figure 3-7). Again, we observed a range of normalization strengths. Some units showed response properties consistent with classic normalization (Figure 3-7a), in which the response to the plaid stimulus in the 50%/50% condition was an evenly weighted average of the unit's responses to the component 50% stimuli. Others showed tuned normalization (Figure 3-7b).

Overall, the units we recorded in mouse V1 using electrophysiology showed tuned normalization (Figure 3-8). A range of normalization indices was observed, ranging from linear summation (NI=0) to simple weighted averaging (NI=0.33), with the bulk of units falling somewhere between these two paradigms in the tuned normalization range. We found the median NI to be 0.13 (IQR=-0.01-0.26), which was close to the value measured using two-photon imaging. Again, the finding of tuned normalization in mouse V1 is qualitatively similar to that observed in macaques, but quantitatively, the strength of normalization as measured by the normalization index suggests weaker normalization in the mouse compared to macaque (Ni, Ray, and Maunsell 2012)

The populations we recorded using both two-photon imaging and electrophysiology showed a range of visually evoked normalization values, indicating tuned normalization.

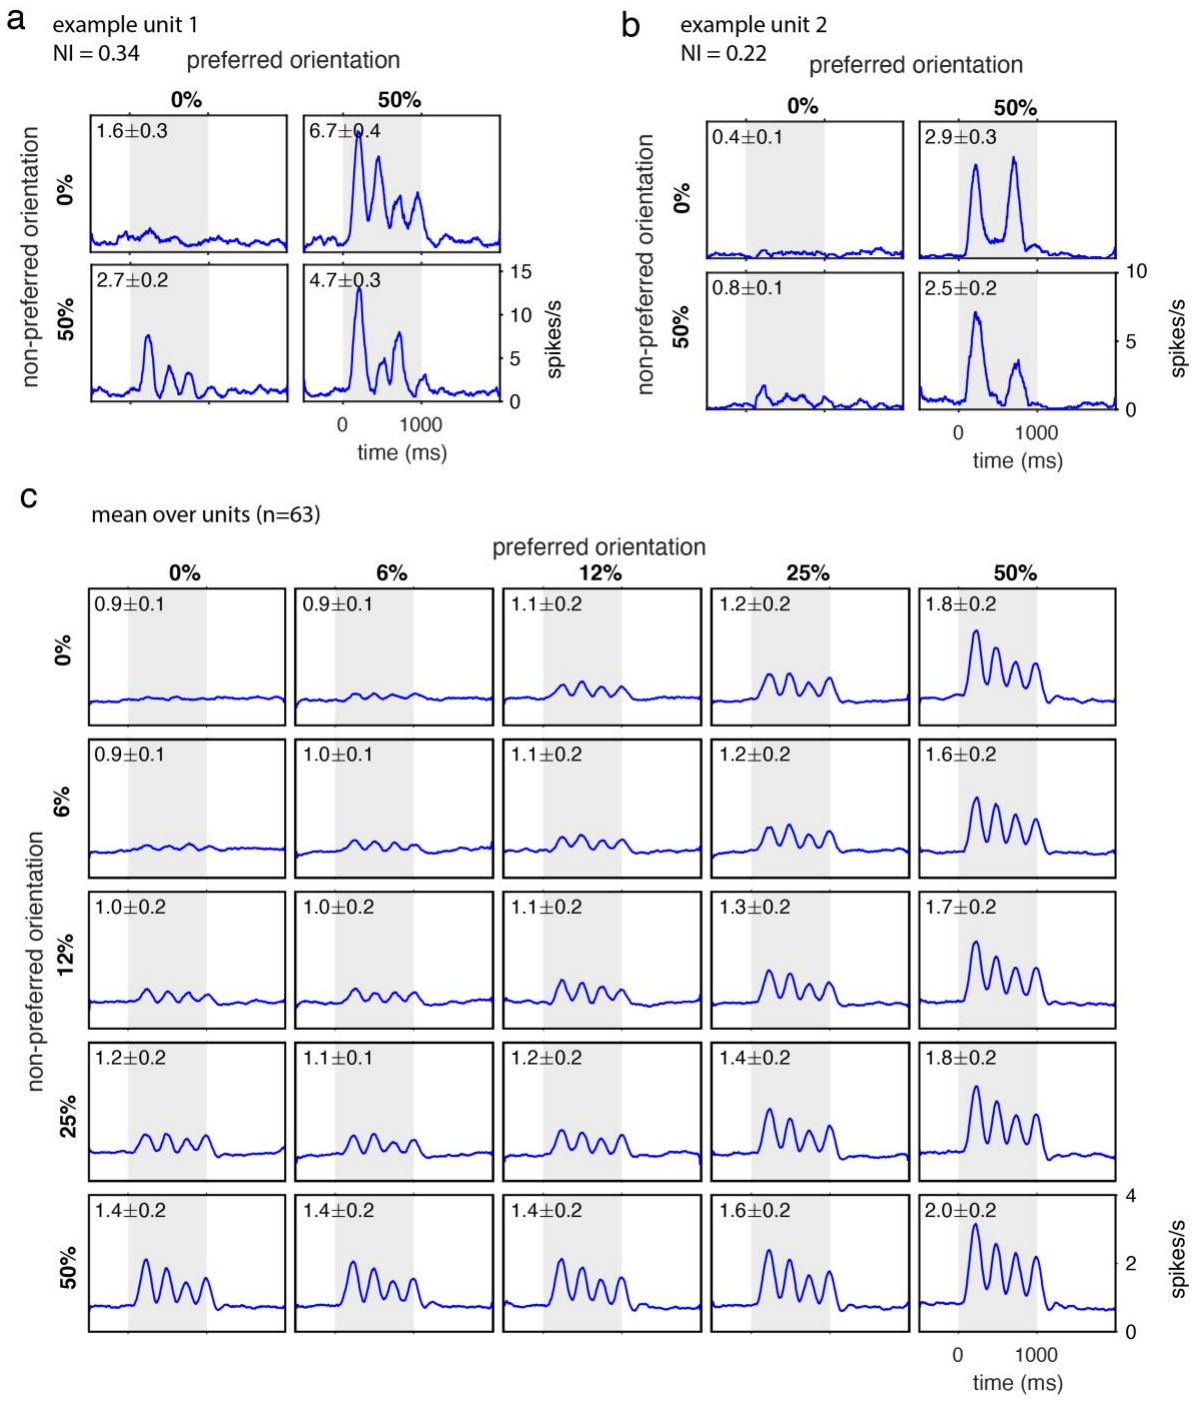

Figure 3-7. Visually-evoked normalization in mouse V1 measured using electrophysiology

a-b, Responses of example neurons to normalization stimuli (same neurons as in Figure 3-6). On average, the units responded well to the preferred stimulus orientation (across horizontal axis) and less to the orthogonal stimulus orientation (vertical axis). When the stimuli are superimposed to create a cross-inhibitory stimulus, the responses were less than those for the preferred component grating, as expected for normalization. Each stimulus was presented 100 times.

Figure 3-7, continued
c, Mean normalization stimulus set responses across neurons. Neurons with a minimum driven firing rate > 0.3 spikes/s were included. Overall, 63 tuned, responsive neurons from 13 recording sessions from 7 mice were included. The mean PSTHs across units shows clear responses to preferred stimulus orientation that increases monotonically with contrast (across top row) and less response to the non-preferred stimulus orientation that also increases with contrast (left column). When the stimuli are superimposed to create a cross-inhibitory stimulus, the driven response to the resulting stimulus is sublinear on average, but on average exceeded the response to the preferred stimulus. Labels correspond to mean rate ±1 SEM.

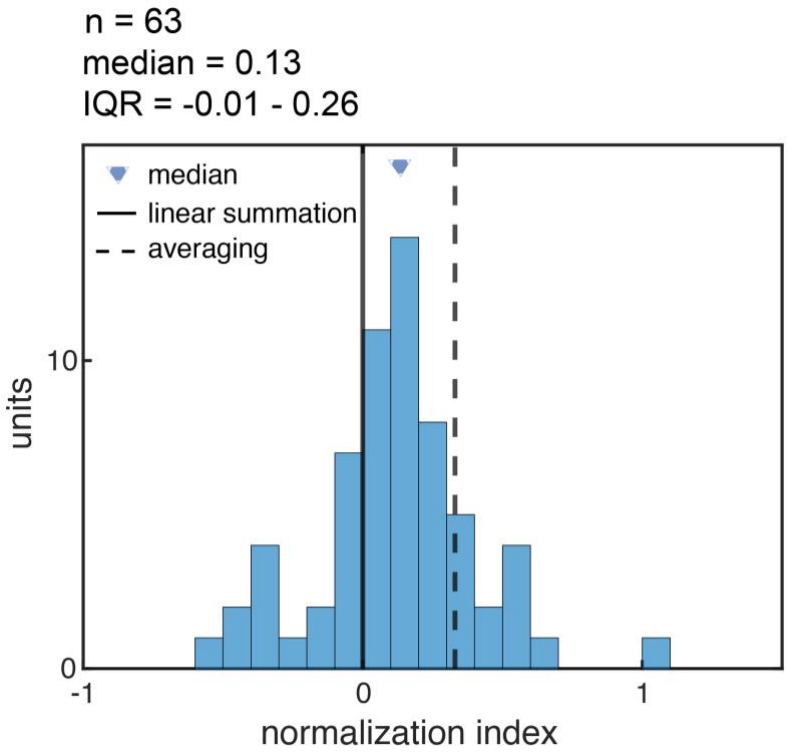

Figure 3-8. Population histogram of normalization strength in mouse V1 measured using electrophysiology

Normalization index distribution across neurons (n = 63) as measured using electrophysiology (median = 0.13, IQR = -0.01 – 0.26). An index of zero represents linear summation of responses; the dashed line at 0.33 represents the value for simple averaging of responses. The median is indicated by the blue triangle.

Normalization strength as a function of recording depth

We examined whether the strength of normalization in the units varied with recording depth. Previous studies examining cortical normalization have primarily considered neurons in layer 2/3 (Busse, Wade, and Carandini 2009; Verhoef and Maunsell 2017). This can be partly attributed to the greater thickness of cortex in larger mammals like cats and macaques. It can also be attributed in part to the preference for the use of Utah arrays over linear arrays, as Utah arrays enable the recording of many units simultaneously though spanning only the partial thickness of the cortex. As mouse V1 is only ~1 mm thick, we took advantage of the opportunity to record at multiple depths in mouse and relate recording depth to normalization strength.

In our electrophysiology recordings, the electrode was inserted at a medial posterior region of visual cortex to cover the lateral visual field at approximately zero elevation (Figure 3-5). The electrode spanned a 700 μm recording span. Channel contacts were spaced at even 100 μm intervals on four shanks for a total of 32 contacts. With knowledge of the electrode insertion depth and electrode contact depths, a recording depth was determined for each unit. We compared the normalization strength across recording depths and found significant differences across depths (Kruskal-Wallis test, $p < 0.01$) (Figure 3-9). On average, units at a depth of 400-500 μm showed significantly weaker normalization compared with units at depths of 100-200 μm ($p < 0.01$) and 500-600 μm ($p < 0.05$). Indeed, units at a recording depth of 400-500 μm more closely showed linear summation than normalization. Firing rates across depths were not significantly different (Kruskal-Wallis test, n.s.). This analysis is limited in that the depth assignments were not corrected for any angular offsets from normal. However, it points to differences in normalization across recording depths and to a need for further study of this phenomenon.

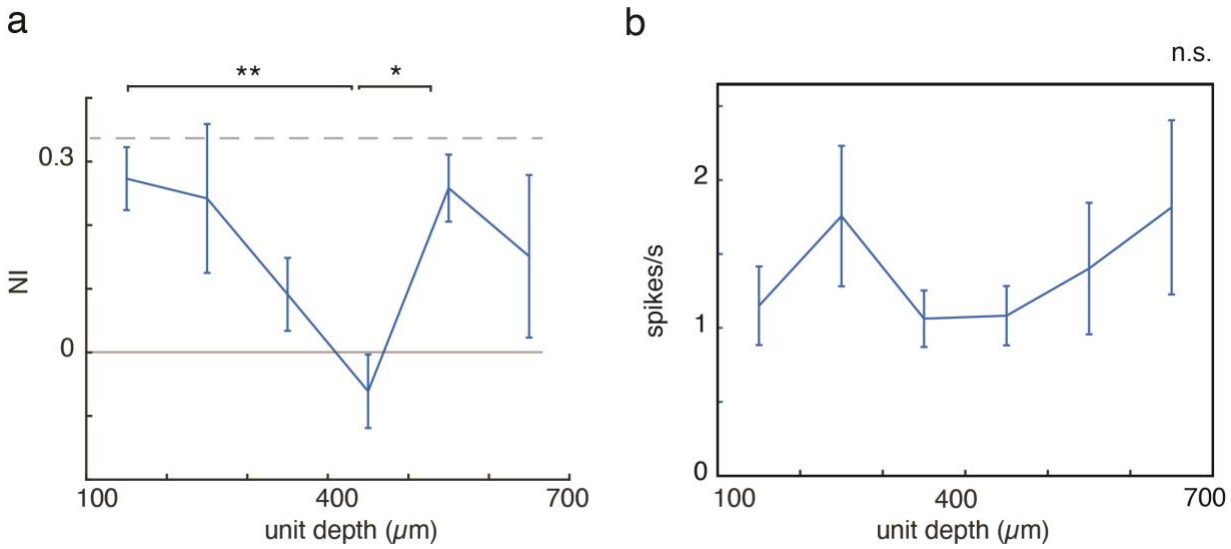

Figure 3-9. Normalization and recording depth indicate depth-specificity

a, Normalization index according to unit depth. The median normalization index of the units recorded is indicated by the dashed blue line (median=0.13, IQR= -0.01 - 0.26). The expected normalization index for linear summation is indicated by the solid horizontal gray line and the expected normalization index for a simple average of responses is indicated by the horizontal dashed gray line. Significance was determined with Tukey-Cramer-corrected Kruskal-Wallis test ($p < 0.01$). * indicates $p < 0.05$. ** indicates $p < 0.01$. **b**, The driven firing rate of the included units does not vary significantly by recording depth (Kruskal-Wallis test, n.s.). The bars show the number of units included at each depth. Units were binned according to depth (100-199 μm; 200-299 μm, and so on). We did not consider bins with fewer than 3 units. All error bars represent ±1 SEM.

Discussion

Our cumulative findings reveal visually-evoked normalization in awake mouse V1 as measured by both two-photon imaging and electrophysiology. Neurons and units that we recorded showed a sublinear response cross-oriented plaid normalization stimuli in comparison to the summed responses to individual component stimuli. On a scale in which linear summation was represented by a normalization index of zero and classic normalization or a simple weighted average was indicated by a normalization index of 0.33, the populations we recorded from had median normalization index of 0.13 (electrophysiology, IQR = -0.01 – 0.26) and 0.15 (two-photon imaging, IQR = 0.03 – 0.25). This range of normalization strengths represents tuned normalization. This is similar to findings in macaques, in which different units and neurons showed different extents of normalization (Rust et al. 2006; Ni, Ray, and Maunsell 2012). However, the overall strength of normalization in the neurons and units we recorded in mouse V1 was weaker than that observed in macaques, which had a median normalization index of ~0.2 (Ni, Ray, and Maunsell 2012).

While the medians observed in the two-photon imaging data and the electrophysiology data recorded in mouse V1 are in approximate agreement, we did find differences in normalization strength between recording depths in our electrophysiology data. This may reflect distinct roles of cortical processing at different cortical depths, and warrants further investigation into possible laminar dependency of normalization. While some studies have used linear probes to study normalization in cortex, most studies on cortical normalization have primarily examined layer 2/3 due to electrode array depth constraints (Busse, et al. 2009; Verhoef and Maunsell 2017). While further work must be done to characterize any potential differences in layer-specific normalization, it is important to consider the possibility that normalization does not necessarily have to operate

uniformly across layers. If such differences exist, it would be important to characterize whether and how this influences population dynamics and visual processing. Differences in the cell types and network structures observed at different cortical depths can also constrain models of biophysical mechanisms that underlie normalization and by extension, the rich cognitive processes (such as attention and multi-sensory integration) that normalization undergirds (Verhoef and Maunsell 2017; Ohshiro, Angelaki, and DeAngelis 2011).

There were some limitations in our study. Our normalization stimulus set consisted of cross-oriented plaids composed of only two orientations (vertical and horizontal) due to limits on recording time and the need to have a suitable number of stimulus repetitions on which to base conclusions. Mouse V1 shows a salt-and-pepper distribution of direction and orientation tuning preferences, so in any field of view or electrophysiology recording session, we recorded from a diversely tuned set of neurons and units. These neurons and units were strongly driven by the full-contrast drifting grating stimuli that were used to determine visual responsiveness and direction tuning preference, but because many neurons did not prefer the component orientations of the plaid stimuli presented, responses were generally weaker in the (lower-contrast) plaid stimulus set compared with the tuning stimulus set. This led to a problem of irreducible noise, requiring us to set inclusion criteria. Thus, many units and neurons were excluded from analysis. This decreased the number of units included for analysis. We were still able to conduct analyses on these neurons and units, but our analysis of the units that did meet responsivity and tuning criteria and the prevalence of normalization in those units indicate that the presentation of a broader range of stimuli would uncover normalization in the majority of visually responsive neurons in mouse visual cortex.

On the whole, this work shows that tuned normalization can be visually-evoked in the mouse visual cortex, allowing for further study of the circuit, cell-type, and network underpinnings of this computation. The mouse offers a ready model for such study and characterization.

Supplemental Information

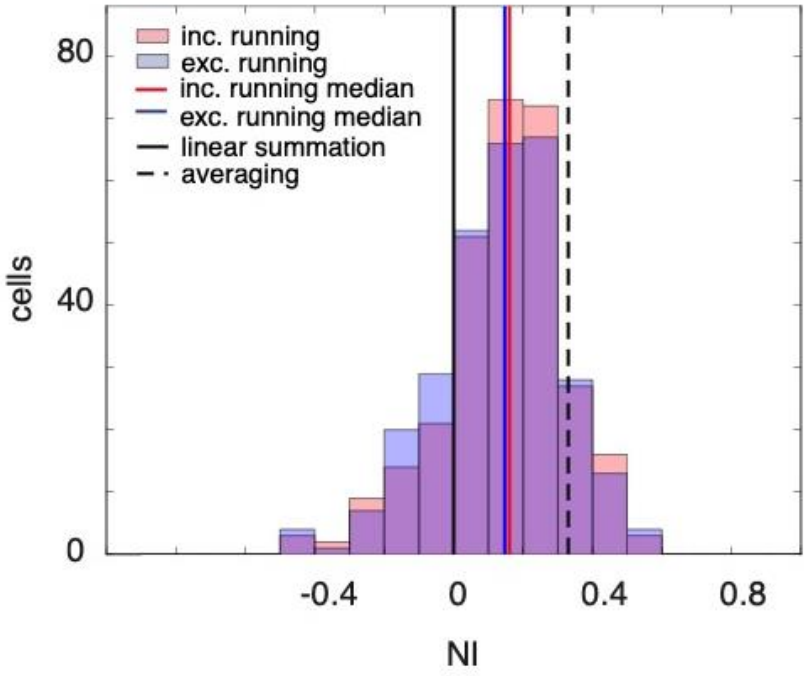

Supplemental Figure 3-1. Mice ran during a small fraction of trials that influenced the median but not the overall distribution of normalization index across neurons

Normalization index measured with and without running trials in cells that passed inclusion criteria in both conditions. The red line represents the median of the included cells for the "include running" condition (median=0.16, IQR=0.06-0.27), and the blue line represents the median for the "exclude running" condition (median=0.15, IQR=0.03-0.26). There was a significant difference in the median normalization index between the include versus exclude running conditions (Mann-Whitney U test, $p < 0.001$). No difference in the overall distribution of normalization indices was statistically detectable (two-sample Kolmogorov-Smirnov test, n.s.). The solid line represents the expected normalization index for linear summation and the dashed line represents the expected normalization index for classic normalization or simple averaging.

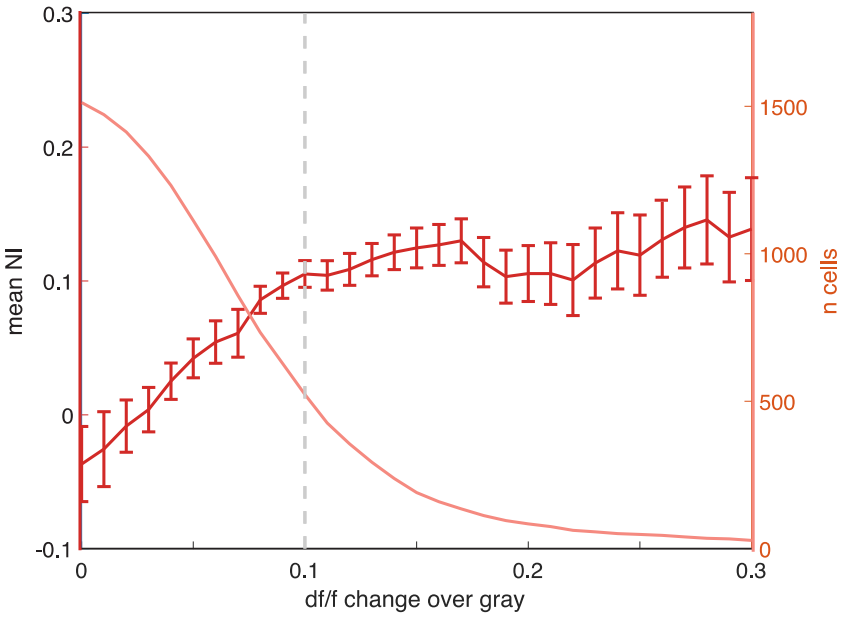

Supplemental Figure 3-2. Visually-evoked normalization varies systematically with calcium signal strength

The red line shows the normalization index as a function of minimum absolute df/f change of the response to the preferred and plaid stimulus over the gray condition. As the df/f change over baseline approaches zero, the normalization index also approaches zero, which would be predicted due to irreducible noise. The pink line shows the number neurons that are included at each threshold level for df/f change over gray. The number of neurons included at each threshold level decreases as the minimum df/f change over gray increases. The dashed gray line represents the threshold selected, at which the exponential function plateaus and at which the largest number of neurons could be included in analysis of normalization. This provides a comparison to early studies in which normalization was measured with electrophysiology, in which only active neurons could be recorded. Mean ± SEM is indicated.

CHAPTER 4

PAIRWISE NOISE CORRELATIONS IN MOUSE V1 COVARY WITH NORMALIZATION

Introduction

Normalization has previously been shown to influence noise correlation patterns in macaques (Ruff, Alberts, and Cohen 2016; Verhoef and Maunsell 2017). Importantly, information encoding depends on activity of individual neurons as well as the relationship between neurons, including pairwise correlations; theoretically, shared noise can prevent a population from averaging out noise and even weak noise correlations can considerably diminish the signaling of a population (Shadlen and Newsome 1994). While not all correlations necessarily limit information (Abbott and Dayan 1999), there are instances of noise correlations - in particular those between units with similar preferences - that have been theorized to be information-limiting (Moreno-Bote et al. 2014). Recent work corroborates these theories (Kafashan et al. 2021).

In behavioral studies, reductions in noise correlations can largely explain performance improvements in attention (Marlene R. Cohen and Maunsell 2009), and these correlations have been shown to occur between pairs of units with similar orientation selectivity and strong normalization (Verhoef and Maunsell 2017). Indeed, normalization has been theorized to decrease correlations and its circuitry provides an infrastructure for decorrelation, seen in passively viewing macaques (Tripp 2012; Ruff, Alberts, and Cohen 2016; Coen-Cagli and Solomon 2019). In behavioral studies, correlations changes due to attention can be entirely explained by attention's biasing of normalization mechanisms (Verhoef and Maunsell 2017). These findings point to a central role for normalization in modulating the neuronal and population representation of stimuli.

One of the main advantages for developing a mouse model of normalization is the opportunity to better understand the circuit mechanisms and cell types that give rise to pairwise correlation structures in the cortex. Thus it is important to observe how normalization influences pairwise correlations in the mouse, and how these patterns compare and contrast with findings in primates. Here, we examined population-level effects of normalization by examining pairwise noise correlations in both two-photon imaging data and electrophysiology collected from awake mice presented with normalization stimuli.

Results

Noise correlations were measured in mouse V1 populations recorded during the presentation of normalization stimuli

Noise correlations were examined in awake mouse V1 using both two-photon imaging and electrophysiology data during the presentation of a stimulus set designed to evoke normalization (M. Carandini and Heeger 1994; Matteo Carandini, Heeger, and Movshon 1997; Busse, Wade, and Carandini 2009). Namely, the set of stimuli included 25 stimuli composed of cross-oriented vertical and horizontal component gratings ranging from 0% to 50% contrast each. A different order of stimuli was generated for each recording day, but on any given recording day, the order of stimuli was maintained from trial to trial in order to limit the effects of preceding stimuli on noise correlations.

In classic normalization, the “plaid” normalization stimulus is expected to produce an equally weighted average of a neuron’s responses to the two “grating” component stimuli. This is in contrast with linear summation, in which the plaid stimulus would be expected to produce a

simple linear sum of a neuron's responses to the two grating components. In tuned normalization, previously described in macaques (Rust et al. 2006; Ni, Ray, and Maunsell 2012) and here in mouse (Figure 3-4 and Figure 3-8), normalization strength can vary such that an unequally weighted average of the responses to the component stimuli is produced when the plaid stimulus is presented. To quantify normalization strength in Chapter 3, we used a normalization index (NI) defined as the difference between the true response to the cross-oriented (plaid) stimulus and the linear sum of responses to the component stimuli (gratings), all over the sum of these two quantities. In this quantification, a normalization index of zero represents linear summation and a normalization index of 0.33 represents classic normalization. Values between 0 and 0.33 represent tuned normalization; values greater than 0.33 represent super-normalization, a differently weighted example of tuned normalization. Values less than zero represent supra-linear summation.

We examined data collected from excitatory neurons in V1 layer 2/3 using two-photon imaging as well as data collected from different cortical depths using electrophysiology experiments. The majority of neurons and units that we recorded from had a normalization index ranging between zero and 0.33 (two-photon imaging: median = 0.15, IQR = 0.03 - 0.25; electrophysiology: median = 0.13, IQR = -0.01 – 0.26) (Figure 3-4 and Figure 3-8). Noise correlations were calculated between direction tuned neurons within each dataset using partial pairwise correlations (two-photon data; n = 17,273 pairs from 341 neurons) and by taking the spike count correlation (electrophysiology; n = 206 pairs from 63 units). The electrophysiology data was examined according to depth, but the limited number of pairs available within depths and within experiments precluded reliable comparison. Thus, all pairs from the electrophysiology experiments were combined for further analyses.

Noise correlation distributions are distinct in plaid versus grating stimulus presentation

Distributions of noise correlations in the grating and plaid conditions were compared (Figure 4-1). Overall, correlations were small and positive as observed in other species, such as macaque (Verhoef and Maunsell 2017). The medians of the distributions were not significantly different in both the two-photon datasets (grating median = 0.03, IQR = -0.08 – 0.15; plaid median = 0.03, IQR = -0.09 – 0.15; Mann-Whitney U test, n.s.) and the electrophysiology data (grating median = 0.04, IQR = -0.04 – 0.10; plaid median = 0.03; IQR=-0.04 – 0.10; Mann-Whitney U test, n.s.). However, in the two-photon datasets, the distributions of noise correlations were found to be significantly different between the grating and plaid conditions (Kolmogorov-Smirnov test, $p < 10^{-4}$). In both data types, higher correlations were observed in pairs of neurons with similar direction tuning preferences compared with pairs that had opposite or oblique direction tuning preferences (Figure 4-1b and Figure 4-1d). Furthermore, correlations were higher in the grating (component) condition compared with the plaid (normalization) condition for pairs of neurons that had orthogonal or oblique preferences relative to one another in the two-photon imaging data.

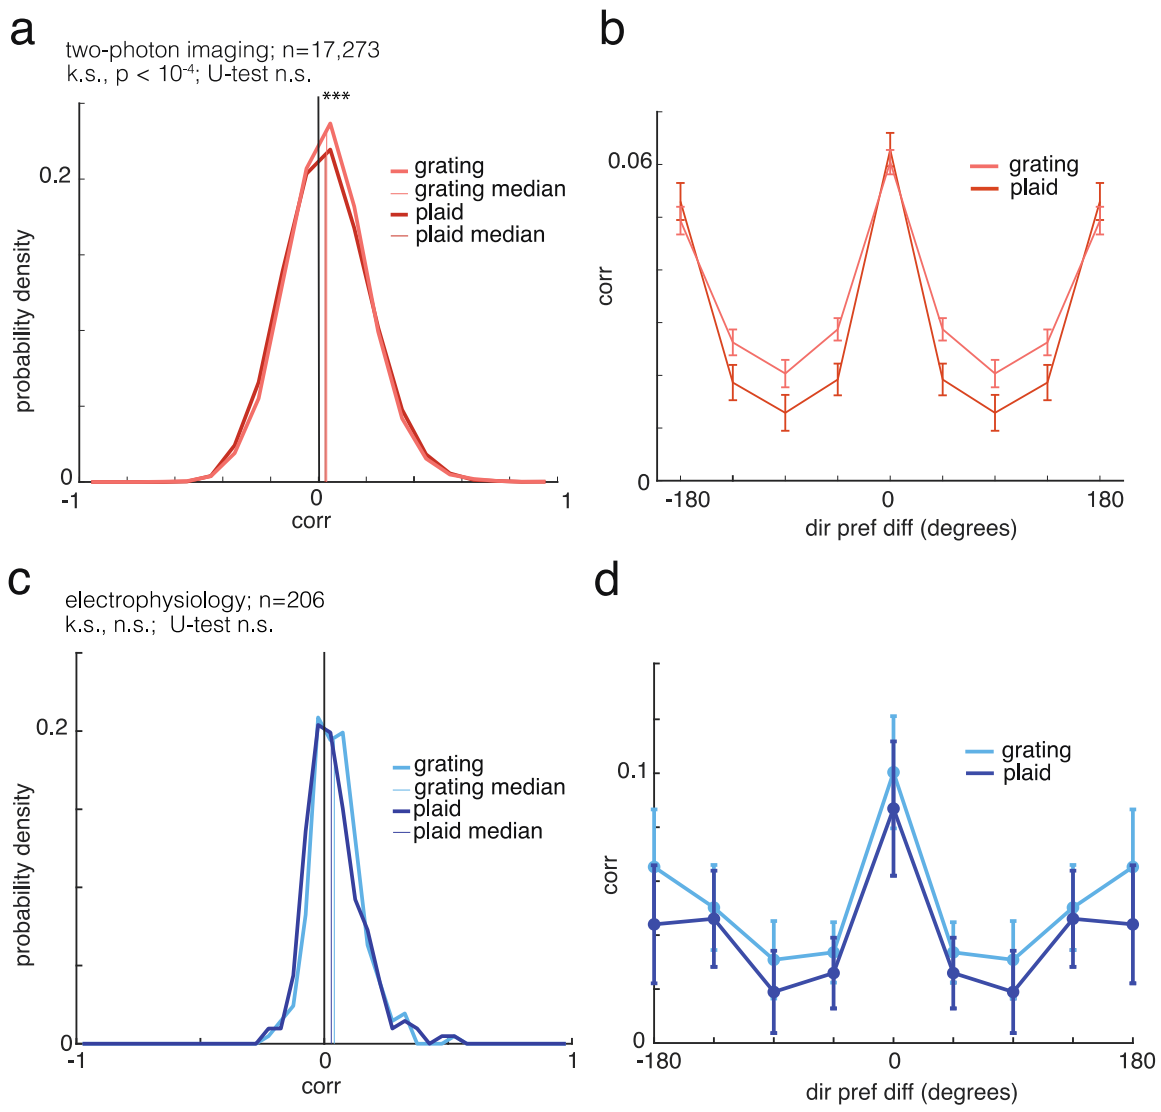

Figure 4-1. Distributions and pairwise tuning-dependence of noise correlations in mouse V1

a, The distributions of pairwise noise correlations measured using two photon-imaging during the presentation of grating (component) versus plaid (normalization) stimuli (n = 17,273). The distributions are significantly different (Kolmogorov-Smirnov test, $p < 10^{-4}$) between the two conditions. However, the medians are overlapping and are not significantly different (grating median = 0.03, IQR = -0.08-0.15; plaid median = 0.03, IQR = -0.09 – 0.15; Mann-Whitney U test, n.s.). **b,** Noise correlations are plotted against pairwise tuning direction preference difference, showing higher correlations on average for pairs with similar direction tuning preferences. Correlations were higher in the grating (component) condition compared with the plaid (normalization) condition for pairs of neurons that had orthogonal or oblique preferences relative to one another. **c,** Same as Figure 4-1a, except the pairs examined are from electrophysiology experiments (n = 206 pairs). In this data, no statistical difference was detected

Figure 4-1, continued
in the distributions or medians of the noise correlation distributions between the grating (component) and plaid (normalization) stimuli (Kolmogorov-Smirnov test, n.s.; Mann-Whitney U test, n.s.). The medians were comparable with the two-photon imaging data, but the inter-quartile range was narrower (grating median = 0.04, IQR = -0.04 – 0.10; plaid median = 0.03; IQR=-0.04 – 0.10). **d,** Same as Figure 4-1b, except the pairs examined are from electrophysiology experiments. Pairwise noise correlations are higher for pairs with similar direction tuning preferences.

Pairwise noise correlations are influenced by normalization strength and orientation selectivity

Next, we examined whether pairwise noise correlations were influenced by pairwise orientation preferences and normalization strengths, as previously shown in macaque (Verhoef and Maunsell 2017).

For the measure of orientation selectivity, we defined the orientation selectivity index (OSI), which is calculated as the difference in responses to the preferred and null component stimuli at 50% contrast divided by the sum of those responses. An OSI of zero indicates no orientation selectivity, and an OSI of one represents perfect selectivity to one of the two component stimuli. The pairwise OSI was calculated as the geometric mean OSI of the individual neurons or units in a pair, signed according to whether the pair preferred the same or opposite orientations.

For the measure of normalization, we defined the pairwise non-preferred suppression index (NSI) similar to previous work (Verhoef and Maunsell 2017). In that work, data from all 25 plaid stimuli presented were fitted to identify normalization parameters. The non-preferred suppression index for each unit was taken to represent the strength of normalization for each unit (Verhoef and Maunsell 2017). As in previous work, the non-preferred suppression index for each neuron or unit in a pair was here taken as the fitted alpha parameter of the tuned normalization model (Ni, Ray, and Maunsell 2012; Verhoef and Maunsell 2017). The pairwise NSI was calculated as the mean of fitted alpha parameters for each neuron or unit pair. It should be noted that this is a different normalization index than that presented in Chapter 3; in that chapter, a normalization index (NI) was calculated based on responses to preferred and non-preferred stimuli and the combined plaid from those components, again for close comparison with previous studies (Ni, Ray, and Maunsell 2012). However, it should be noted that here and in previous work in macaques, these two

measures (NI and NSI) are correlated, suggesting that these capture normalization strength similarly (Ni, Ray, and Maunsell 2012).

Data were binned along the medians for both pairwise OSI and pairwise NSI, resulting in evenly populated bins. Correlation strength was compared between the plaid and component stimulus conditions for each quadrant defined by OSI and NSI median-split bins. In the two-photon imaging data, noise correlations increased in pairs with similar orientation selectivity and strong normalization (Mann-Whitney U test, $p < 0.05$), decreased in pairs with opposite selectivity and strong normalization (Mann-Whitney U test, $p < 0.01$), and decreased in pairs with similar orientation selectivity and strong normalization (Mann-Whitney U test, $p < 0.05$) when plaid (normalization) stimuli were presented compared to when grating (component) stimuli were shown (Figure 4-2a). No significant differences in correlations were detected in the electrophysiology data, though this may indicate an under-powering of the comparison in the electrophysiology data (Mann-Whitney U tests, n.s. for all four quadrants). Nevertheless, the data are reported here for completeness (Figure 4-2b). For the electrophysiology data, it was difficult to draw reliable conclusions from the small number of pairs. This highlights the utility of simultaneous recordings from many neurons, as in the two-photon imaging experiments, and points to an application for high-density electrophysiology electrodes that have recently become commercially available (Steinmetz et al. 2018).

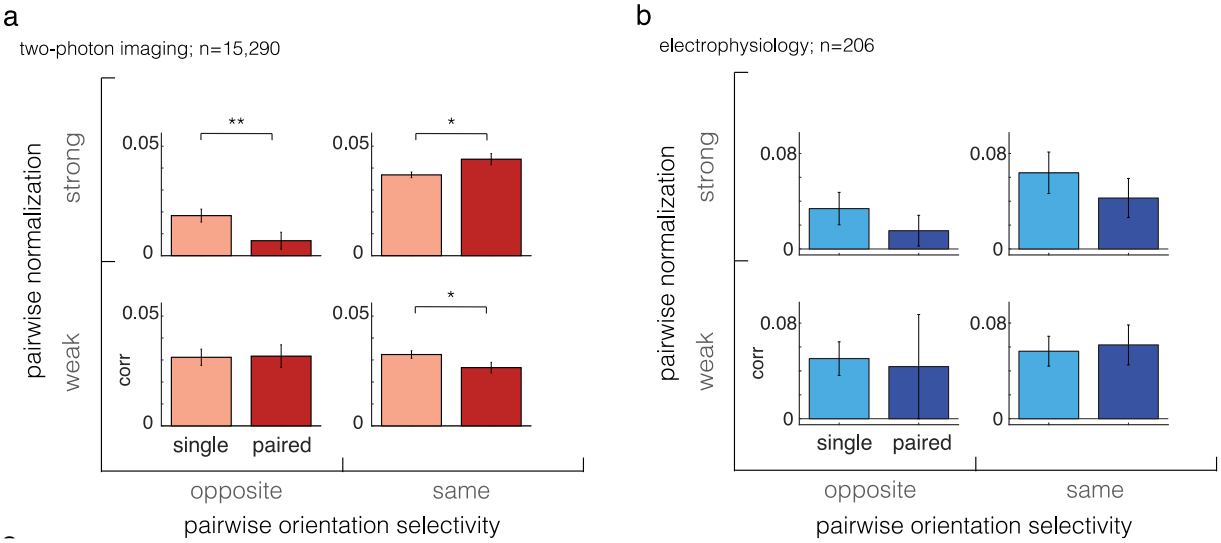

Figure 4-2. Pairwise noise correlations are influenced by normalization strength and orientation selectivity

a, Noise correlations measured in two-photon data (n=15,290) were divided according to median splits of the "pairwise orientation selectivity" index and "pairwise normalization" non-preferred suppression index such that all four quadrants were evenly populated. Noise correlations increased in pairs with similar orientation selectivity and strong normalization (Mann-Whitney U test, $p < 0.05$), decreased in pairs with opposite selectivity and strong normalization (Mann-Whitney U test, $p < 0.01$), and decreased in pairs with similar orientation selectivity and weak normalization (Mann-Whitney U test, $p < 0.05$) when plaid (normalization) stimuli were presented compared to when grating (component) stimuli were shown. **b,** Same as Figure 4-2a, except showing noise correlations measured in the electrophysiology data (n=206). No significant differences in correlations were detected (Mann-Whitney U tests, n.s. for all four quadrants).

Discussion

Here, we examine noise correlations in awake mouse V1 during the presentation of stimuli that elicit normalization using both two-photon imaging and electrophysiology. We observed small and positive correlations in the data in both recording modalities. Additionally, we observed a change in the distributions of correlations between grating (component) and plaid (normalization) stimuli. Furthermore, we report decreased correlations in pairs of neurons with opposite or oblique direction preferences during the presentation of plaid normalization stimuli as compared to component stimuli, suggesting that normalization systematically influences correlation structures in the mouse cortex as observed in macaque. Changes in cortical correlations in the context of normalization have been theorized and described previously (Tripp 2012; Ruff, Alberts, and Cohen 2016; Verhoef and Maunsell 2017), highlighting the role that normalization plays in shared cortical variability and emphasizing the opportunity that mouse presents for better understanding the fundamental mechanisms underlying pairwise correlation changes in cortex.

We next studied whether changes in pairwise correlations could be related to normalization strength and orientation selectivity in mouse, as previously observed in macaque (Verhoef and Maunsell 2017). We noted similarities and differences between the normalization-mediated changes in pairwise noise correlations in mouse and macaque. In particular, we found that normalization increases noise correlations between similarly orientation selective pairs that showed strong normalization. This is consistent with results from macaque, in which pairs of neurons with similar orientation selectivity and strong normalization had increased correlations in the presence of normalization stimuli compared with just component stimuli (Verhoef and Maunsell 2017). Importantly, correlations between similarly and strongly selective neurons are the same that are modified (decreased) by increased attention, highlighting the role of normalization

circuitry in creating the effects of attention that are important in behavior (Verhoef and Maunsell 2017; Marlene R. Cohen and Maunsell 2009). We also observed other similarities when comparing the mouse data with reports from macaque: namely, in macaques, pairs with opposite orientation selectivity and strong normalization show decreased correlations. We also observe this in mouse V1.

Additionally, we observed that pairs with strong normalization and similar orientation selectivity showed decreased correlations when normalization stimuli were presented compared with when component stimuli were presented. This has not been previously observed in macaque. These differences could represent true differences between the mouse and macaques, as we use the normalization index of fitted normalization parameters as studied in previous work. Differences could be related to differences in the organization of visual cortex in mouse and macaque; for example, macaque visual cortex is organized in columns, whereas mouse visual cortex exhibits a salt-and-pepper organization (Niell 2015). Further work is needed to better understand the differences described here.

The conclusions drawn here were largely derived from the correlations measured in the two-photon data, as the sample size in the electrophysiology was relatively much smaller. This was a limitation in this study, but also pointed to an opportunity to use dense sampling electrodes in future electrophysiology studies examining correlations in mouse (Steinmetz et al. 2018).

On a fundamental level, in addition to their important role in attention, noise correlations have been theorized and shown to limit information in populations of neurons (Shadlen and Newsome 1994; Moreno-Bote et al. 2014; Kafashan et al. 2021). Thus, accessing the study of noise correlations in a species like mouse now allows for a deeper understanding of the way that noise correlations are produced and modified. On the whole, these findings in mouse open the door

to further exploration of the circuit mechanism and network properties underlying normalization, pairwise correlations, and by extension, cognitive functions like attention and multi-sensory integration (Marlene R. Cohen and Maunsell 2009; Verhoef and Maunsell 2017; Ohshiro, Angelaki, and DeAngelis 2011).

CHAPTER 5

GENERAL DISCUSSION

Overview

In this work, we report visually-evoked normalization in mouse visual cortex, opening the door to further study of the circuit mechanism, cell types, and networks that underlie this canonical computation (Matteo Carandini and Heeger 2011). The data presented here was measured using both electrophysiology and two-photon imaging in awake, head-fixed mice; both recording modalities showed similar distributions of normalization strength, with differences in normalization strength noted across layers in the electrophysiology data. Normalization strength was shown to be tuned, as seen in primates, but systematically weaker than previous reports in macaques, in which normalization strength is closer to classic normalization (Ni, Ray, and Maunsell 2012). Finally, we also investigated pairwise noise correlations and observed that normalization affects patterns of noise correlations in relation to orientation selectivity and normalization strength. These findings highlight the mouse as a model for further elucidating the underlying circuit mechanisms and networks that underlie normalization, a canonical computation that profoundly influences stimulus representations in the brain.

Visually-evoked and tuned normalization measured in awake mouse V1

We recorded from the V1 of awake mice during the presentation of normalization stimuli using both two-photon imaging and electrophysiology. The majority of responsive and tuned neurons in the mouse visual cortex show sublinear responses to plaid stimuli when compared with the responses to the component gratings. This is consistent with normalization, and was observed

in both imaging and electrophysiology experiments. We show that in both sets of experiments, neurons and units exhibit a range of normalization strengths. We quantified this difference using a normalization index in which linear summation, in which the response to a plaid stimulus is a sum of the response to component stimuli, was set to zero. In our normalization index, a simple weighted average consistent with classic normalization, was represented by a normalization index of 0.33. We observed a median normalization index among neurons of 0.15 (two-photon imaging; IQR = 0.03 – 0.25) and among units of 0.13 (electrophysiology; IQR = -0.01 – 0.26). Among neurons and units, we observed a distribution of normalization strengths, with ranges largely overlapping between the two-photon and electrophysiology data. This indicates that normalization can be studied in mouse using either or a combination of these recording modalities.

Our finds are consistent with the phenomenon of tuned normalization observed in macaques previously in which units normalized to different extents (Rust et al. 2006; Ni, Ray, and Maunsell 2012; Ni and Maunsell 2017). However, there are notable quantitative differences in normalization observed in the mouse; namely, the data from macaque are closer to classic normalization, with an approximate median normalization index ~0.2 in macaques (Ni, Ray, and Maunsell 2012). In comparison, the data in mouse are systematically offset towards weaker normalization. Normalization tuning is related to attentional modulation of neuronal firing rates; neurons that show stronger normalization also show stronger attentional modulation. This highlights that the normalization strength of neurons is computationally relevant. Indeed, modulations in firing rate due to attention and related to tuned normalization bias normalization mechanisms to modulate correlations in attention (Verhoef and Maunsell 2017). Tuned normalization can be further studied in the mouse through circuit, cell-type, and network studies, opening the door to further insights into this property of neuronal responses.

Depth-dependence of normalization

In addition to measuring distributions of normalization strengths in the neurons and units we recorded, we also took advantage of our electrophysiology recordings (spanning the depth of V1) to compare normalization across different cortical recording depths. Much of the literature regarding normalization reports on samples collected from layers 2/3, reflecting greater cortical thickness of larger mammals like cats and primates as well as the accessibility of multi-site recording technologies (specifically Utah arrays) that primarily enable access to layer 2/3 (Busse, Wade, and Carandini 2009; Verhoef and Maunsell 2017). We observed that normalization was significantly stronger at superficial recording depths and decreased as recording depth increased, up to a depth of 500 μm, at which point normalization abruptly and significantly strengthened. We did not control for angular deviations from normal in our recordings. However, this the previously unreported and notable differences across recording depths can provide valuable clues as to the origin of normalization and the processes like attention and multi-sensory integration that it underlies (Ohshiro, Angelaki, and DeAngelis 2011; Verhoef and Maunsell 2017). Along with other work in normalization, this direct measurement from across the cortical layers highlights the likely cortical origin of normalization and emphasizes the possibility of flexibility gating the computation at a cortical level, enabling its manipulation by higher order cognitive functions like attention. Our finding warrants further investigation into the laminar organization of normalization mechanisms.

The role of normalization in cortical noise correlations

Beyond the single-cell level of analysis, we also examined pairwise noise correlations of populations of neurons during normalization. In this work, we demonstrate that mouse visual cortical pairwise correlations are influenced by normalization and normalization stimuli.

The median noise correlations were small and positive in our data across both plaid (normalization) and grating (component) stimuli. However, we noted a difference in correlation distributions between the plaid and the grating stimuli, as measured using two-photon imaging. When examining the data according to orientation selectivity and normalization strength, we observed increased correlations in similarly selective pairs that strongly normalize. Furthermore, we observed a decrease in correlations among pairs with opposite orientation selectivity that exhibited strong normalization. These findings are consistent with previous data from macaque (Verhoef and Maunsell 2017). We also observed differences relative to macaque. For example, we observed significantly decreased correlations in similarly selective pairs that have weak normalization, which was not described in the macaque data. The differences observed here may represent differences between macaque and mouse that could be due to differences in the organization of visual cortex between monkeys and mice (Verhoef and Maunsell 2017).

On the whole, these findings can help better understand noise correlations in the cortex. Correlations can limit the information carried by a population of neurons (Moreno-Bote et al. 2014). However, not all correlations are information-limiting (Abbott and Dayan 1999). Recent work has hypothesized the existence of "information-limiting correlations," and showed that reducing correlations between pairs of similarly tuned neurons specifically could limit information capacity of a network (Moreno-Bote et al. 2014). Experimental work estimating the Fisher information in a population of simultaneously imaged mouse V1 neurons showed that indeed, reductions in these correlations increased the asymptotic information limit in the neurons' populations (Kafashan et al. 2021).

Behaviorally, reductions in correlations during attention tasks largely explained improvements in macaques' performance in a visual attention task (Marlene R. Cohen and

Maunsell 2009). It is interesting to note that the correlations that are reduced in attention are specifically those in pairs that are similarly orientation selective that also have strong normalization (Verhoef and Maunsell 2017). Further work in this direction could elucidate how normalization, attention, and information processing are related in the brain.

Although previous studies have demonstrated summation and division can occur in mouse visual cortex using non-visual stimuli (including optogenetic stimulation and current injection), this is the first work to our knowledge that describes visually-evoked normalization in mouse (Ko et al. 2011), opening the door to further elucidating the mechanisms and structures that influence pairwise correlations in the cortex.

Opportunities in mouse

The finding that mouse visual cortex exhibits normalization adds this computation to a growing list of similarities in mouse vision and vision studied in other species (Seabrook et al. 2017). Combined with the advantages inherent to the mouse as a model species and to vision as a model cortical modality, this opens doors to better understanding cortical processing as a whole.

Because mice can be imaged while head-fixed, it is possible to test the influence of various behaviors on normalization and the population changes it effects. In these studies, mice were head-fixed and awake, similar to previous experiments in other species (Ni, Ray, and Maunsell 2012; Ruff, Alberts, and Cohen 2016; Verhoef and Maunsell 2017). However, the mice were able to locomote freely. While in this study only ~5% of all trials included running, it is possible to design experiments in which a constitutively revolving treadmill requires mice to locomote according to a specified trial structure. Should changes be observed in normalization, in this case, they could be related to mechanisms discovered in running. For example, Ayaz and colleagues found that

decreased surround suppression contributes to increased firing in mouse visual cortex during locomotion, and this decrease in suppression is related to changes in the activity of somatostatin neurons (Ayaz et al. 2013). Somatostatin neurons are disinhibited by VIP neurons in layer 1, and are the only currently known inhibitory cell type that do not observe a surround-suppression model, which positions them to compute surround suppression (Fino, Packer, and Yuste 2013). Studies investigating the relationship between running, surround suppression, and normalization are within reach using the mouse as a model for normalization.

Furthermore, although our work demonstrates a role for normalization in visual cortical sensory representations, our work does not directly implicate a biophysical or cellular mechanism. Much work remains to further elucidate these mechanisms. However, by leveraging visually-evoked activity in the mouse visual cortex and the many other experimental tools available in mice, many avenues for further study are made available. A powerful combination of visual manipulations, genetic tools, optogenetic techniques, and improved recording technologies all enable the detailed dissection of the cell types and circuits that are involved in cortical computation (Luo, Callaway, and Svoboda 2018; Deisseroth 2015; Steinmetz et al. 2018; Seabrook et al. 2017).

Finally, further study can investigate the underlying networks that give rise to normalization. Specifically, it would be interesting to study functional network structures change with normalization. Because of the widespread changes in correlations during normalization, it is implied that the functional network itself may change, influencing the information carrying capacity of a neural population. Further work using static network analyses and temporal graph analysis can be used to further elucidate how decreasing such correlations can influence visual perception (Chambers and MacLean 2016; Dechery and MacLean 2017; Levy, Sporns, and MacLean 2020). For example, subnetworks of cell types can be studied to better understand how

to the network structures of similarly or oppositely tuned neurons change with normalization. Indeed, it will be useful to better understand the feedforward and recurrent networks of neurons involved in normalization. With the imaging access that mouse provides, such studies involving near-simultaneously imaged and spatially identifiable neurons are now possible (Sadovsky et al. 2011; Dechery and MacLean 2017; Lu et al. 2020). These studies can provide insight into the networks involved in normalization, a canonical computation (Matteo Carandini and Heeger 2011).

Conclusion

Prior to this work, visually-evoked normalization in mouse visual cortex had not been described. Here, we put forth the mouse as a model for further elucidating the circuitry and functional networks that underlie normalization, a ubiquitous and consequential canonical neural computation that underlies many sensory and cognitive processes, including multi-sensory integration and attention. Additionally, further understanding normalization could improve our understanding of pairwise correlations that are linked to important cognitive processes like attention. By understanding the circuit and network basis of normalization, we can gain insight into the fundamental underpinnings of how the brain processes sensory signals. Indeed, it may one day become possible to link basic science and clinical pathologies in which normalization processes are interrupted.

Future Directions

1. Are there layer-specific instantiations of normalization?
2. What are the within- and between-layer circuit mechanisms that give rise to normalization in the mouse visual cortex?
3. How does locomotion interact with normalization?
4. How does mutual information between pairs of neurons change in conditions of normalization and attention?
5. In which pairs or subnetworks of neurons is information carried most robustly?
6. What are the network properties of normalization in mouse visual cortex?
7. How are higher order cortical network relationships, such as triplet motif representation, influenced by normalization?
8. How do normalization networks enable the robust transfer of information and network information capacity?
9. How do different cell types (excitatory and inhibitory, including inhibitory subtypes) contribute to normalization?
10. How are these cortical circuits modulated by attention?

REFERENCES

Abbott, L. F., and P. Dayan. 1999. “The Effect of Correlated Variability on the Accuracy of a Population Code.” *Neural Computation* 11 (1): 91–101.

Adesnik, Hillel, William Bruns, Hiroki Taniguchi, Z. Josh Huang, and Massimo Scanziani. 2012. “A Neural Circuit for Spatial Summation in Visual Cortex.” *Nature* 490 (7419): 226–31. https://doi.org/10.1038/nature11526.

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

Albrecht, D. G., and W. S. Geisler. 1991. “Motion Selectivity and the Contrast-Response Function of Simple Cells in the Visual Cortex.” *Visual Neuroscience* 7 (6): 531–46. https://doi.org/10.1017/s0952523800010336.

Albrecht, Duane G., Wilson S. Geisler, Robert A. Frazor, and Alison M. Crane. 2002. “Visual Cortex Neurons of Monkeys and Cats: Temporal Dynamics of the Contrast Response Function.” *Journal of Neurophysiology* 88 (2): 888–913. https://doi.org/10.1152/jn.2002.88.2.888.

Andermann, Mark L., Aaron M. Kerlin, and Clay Reid. 2010. “Chronic Cellular Imaging of Mouse Visual Cortex during Operant Behavior and Passive Viewing.” *Frontiers in Cellular Neuroscience* 4. https://doi.org/10.3389/fncel.2010.00003.

Attneave, Fred. 1954. “Some Informational Aspects of Visual Perception.” *Psychological Review* 61 (3): 183–93. https://doi.org/10.1037/h0054663.

Ayaz, Aslı, Aman B. Saleem, Marieke L. Schölvinck, and Matteo Carandini. 2013. “Locomotion Controls Spatial Integration in Mouse Visual Cortex.” *Current Biology: CB* 23 (10): 890–94. https://doi.org/10.1016/j.cub.2013.04.012.

Barlow, H. 2001. “Redundancy Reduction Revisited.” *Network (Bristol, England)* 12 (3): 241–53.

Barlow, Horace. 1961. “Possible Principles Underlying the Transformation of Sensory Messages.” In *Sensory Communication*. MIT Press.

Bijanzadeh, Maryam, Lauri Nurminen, Sam Merlin, Andrew M. Clark, and Alessandra Angelucci. 2018. “Distinct Laminar Processing of Local and Global Context in Primate Primary Visual Cortex.” *Neuron* 100 (1): 259-274.e4. https://doi.org/10.1016/j.neuron.2018.08.020.

Billeh, Yazan N., and Michael T. Schaub. 2018. “Feedforward Architectures Driven by Inhibitory Interactions.” *Journal of Computational Neuroscience* 44 (1): 63–74. https://doi.org/10.1007/s10827-017-0669-1.

Bock, Davi D., 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. 2011. “Network Anatomy and in Vivo Physiology of Visual Cortical Neurons.” *Nature* 471 (7337): 177–82. https://doi.org/10.1038/nature09802.

Bonds, A. B. 1989. “Role of Inhibition in the Specification of Orientation Selectivity of Cells in the Cat Striate Cortex.” *Visual Neuroscience* 2 (1): 41–55. https://doi.org/10.1017/s0952523800004314.

Boyden, Edward S., Feng Zhang, Ernst Bamberg, Georg Nagel, and Karl Deisseroth. 2005. “Millisecond-Timescale, Genetically Targeted Optical Control of Neural Activity.” *Nature Neuroscience* 8 (9): 1263–68. https://doi.org/10.1038/nn1525.

Boynton, Geoffrey M. 2009. “A Framework for Describing the Effects of Attention on Visual Responses.” *Vision Research* 49 (10): 1129–43. https://doi.org/10.1016/j.visres.2008.11.001.

Britten, K. H., M. N. Shadlen, W. T. Newsome, and J. A. Movshon. 1992. “The Analysis of Visual Motion: A Comparison of Neuronal and Psychophysical Performance.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 12 (12): 4745–65.

Busse, Laura, Asli Ayaz, Neel T. Dhruv, Steffen Katzner, Aman B. Saleem, Marieke L. Schölvinck, Andrew D. Zaharia, and Matteo Carandini. 2011. “The Detection of Visual Contrast in the Behaving Mouse.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 31 (31): 11351–61. https://doi.org/10.1523/JNEUROSCI.6689-10.2011.

Busse, Laura, Alex R. Wade, and Matteo Carandini. 2009. “Representation of Concurrent Stimuli by Population Activity in Visual Cortex.” *Neuron* 64 (6): 931–42. https://doi.org/10.1016/j.neuron.2009.11.004.

Buzsáki, György, Kai Kaila, and Marcus Raichle. 2007. “Inhibition and Brain Work.” *Neuron* 56 (5): 771–83. https://doi.org/10.1016/j.neuron.2007.11.008.

Carandini, M., and D. J. Heeger. 1994. “Summation and Division by Neurons in Primate Visual Cortex.” *Science (New York, N.Y.)* 264 (5163): 1333–36.

Carandini, Matteo, Jonathan B. Demb, Valerio Mante, David J. Tolhurst, Yang Dan, Bruno A. Olshausen, Jack L. Gallant, and Nicole C. Rust. 2005. “Do We Know What the Early

Visual System Does?" *The Journal of Neuroscience* 25 (46): 10577–97. https://doi.org/10.1523/JNEUROSCI.3726-05.2005.

Carandini, Matteo, and David J. Heeger. 2011. "Normalization as a Canonical Neural Computation." *Nature Reviews. Neuroscience* 13 (1): 51–62. https://doi.org/10.1038/nrn3136.

Carandini, Matteo, David J. Heeger, and J. Anthony Movshon. 1997. "Linearity and Normalization in Simple Cells of the Macaque Primary Visual Cortex." *The Journal of Neuroscience* 17 (21): 8621–44. https://doi.org/10.1523/JNEUROSCI.17-21-08621.1997.

Cardin, Jessica A., Larry A. Palmer, and Diego Contreras. 2007. "Stimulus Feature Selectivity in Excitatory and Inhibitory Neurons in Primary Visual Cortex." *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience* 27 (39): 10333–44. https://doi.org/10.1523/JNEUROSCI.1692-07.2007.

Chambers, Brendan, and Jason N. MacLean. 2016. "Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks." *PLoS Computational Biology* 12 (8): e1005078. https://doi.org/10.1371/journal.pcbi.1005078.

Chance, F. S., and L. F. Abbott. 2000. "Divisive Inhibition in Recurrent Networks." *Network (Bristol, England)* 11 (2): 119–29.

Chance, Frances S., L. F. Abbott, and Alex D. Reyes. 2002. "Gain Modulation from Background Synaptic Input." *Neuron* 35 (4): 773–82. https://doi.org/10.1016/s0896-6273(02)00820-6.

Chen, Tsai-Wen, Trevor J. Wardill, Yi Sun, Stefan R. Pulver, Sabine L. Renninger, Amy Baohan, Eric R. Schreiter, et al. 2013. "Ultra-Sensitive Fluorescent Proteins for Imaging Neuronal Activity." *Nature* 499 (7458): 295–300. https://doi.org/10.1038/nature12354.

Cheong, Soon Keen, Chris Tailby, Samuel G. Solomon, and Paul R. Martin. 2013. "Cortical-like Receptive Fields in the Lateral Geniculate Nucleus of Marmoset Monkeys." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 33 (16): 6864–76. https://doi.org/10.1523/JNEUROSCI.5208-12.2013.

Coen-Cagli, Ruben, and Odelia Schwartz. 2013. "The Impact on Midlevel Vision of Statistically Optimal Divisive Normalization in V1." *Journal of Vision* 13 (8): 13–13. https://doi.org/10.1167/13.8.13.

Coen-Cagli, Ruben, and Selina S. Solomon. 2019. "Relating Divisive Normalization to Neuronal Response Variability." *The Journal of Neuroscience* 39 (37): 7344–56. https://doi.org/10.1523/JNEUROSCI.0126-19.2019.

Cohen, Marlene R., and John H. R. Maunsell. 2009. “Attention Improves Performance Primarily by Reducing Interneuronal Correlations.” *Nature Neuroscience* 12 (12): 1594–1600. https://doi.org/10.1038/nn.2439.

Cohen, M.R., and A. Kohn. 2011. “Measuring and Interpreting Neuronal Correlations.” *Nature Neuroscience* 14 (7): 811–19. https://doi.org/10.1038/nn.2842.

Cone, Jackson J., Morgan L. Bade, Nicolas Y. Masse, Elizabeth A. Page, David J. Freedman, and John H. R. Maunsell. 2020. “Mice Preferentially Use Increases in Cerebral Cortex Spiking to Detect Changes in Visual Stimuli.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 40 (41): 7902–20. https://doi.org/10.1523/JNEUROSCI.1124-20.2020.

Cone, Jackson J., Megan D. Scantlen, Mark H. Histed, and John H. R. Maunsell. 2019. “Different Inhibitory Interneuron Cell Classes Make Distinct Contributions to Visual Contrast Perception.” *ENeuro* 6 (1). https://doi.org/10.1523/ENEURO.0337-18.2019.

Cruikshank, Scott J., Timothy J. Lewis, and Barry W. Connors. 2007. “Synaptic Basis for Intense Thalamocortical Activation of Feedforward Inhibitory Cells in Neocortex.” *Nature Neuroscience* 10 (4): 462–68. https://doi.org/10.1038/nn1861.

Cruz-Martín, Alberto, Rana N. El-Danaf, Fumitaka Osakada, Balaji Sriram, Onkar S. Dhande, Phong L. Nguyen, Edward M. Callaway, Anirvan Ghosh, and Andrew D. Huberman. 2014. “A Dedicated Circuit Links Direction-Selective Retinal Ganglion Cells to the Primary Visual Cortex.” *Nature* 507 (7492): 358–61. https://doi.org/10.1038/nature12989.

Dadarlat, Maria C., and Michael P. Stryker. 2017. “Locomotion Enhances Neural Encoding of Visual Stimuli in Mouse V1.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 37 (14): 3764–75. https://doi.org/10.1523/JNEUROSCI.2728-16.2017.

Dana, Hod, Tsai-Wen Chen, Amy Hu, Brenda C. Shields, Caiying Guo, Loren L. Looger, Douglas S. Kim, and Karel Svoboda. 2014. “Thy1-GCaMP6 Transgenic Mice for Neuronal Population Imaging In Vivo.” *PLOS ONE* 9 (9): e108697. https://doi.org/10.1371/journal.pone.0108697.

Dechery, Joseph B., and Jason N. MacLean. 2017. “Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons.” *BioRxiv*, December, 238345. https://doi.org/10.1101/238345.

———. 2018. “Functional Triplet Motifs Underlie Accurate Predictions of Single-Trial Responses in Populations of Tuned and Untuned V1 Neurons.” *PLoS Computational Biology* 14 (5): e1006153. https://doi.org/10.1371/journal.pcbi.1006153.

Deisseroth, Karl. 2015. “Optogenetics: 10 Years of Microbial Opsins in Neuroscience.” *Nature Neuroscience* 18 (9): 1213–25. https://doi.org/10.1038/nn.4091.

Denk, W., J. H. Strickler, and W. W. Webb. 1990. “Two-Photon Laser Scanning Fluorescence Microscopy.” *Science (New York, N.Y.)* 248 (4951): 73–76. https://doi.org/10.1126/science.2321027.

Desimone, R., and J. Duncan. 1995. “Neural Mechanisms of Selective Visual Attention.” *Annual Review of Neuroscience* 18: 193–222. https://doi.org/10.1146/annurev.ne.18.030195.001205.

Destexhe, A., and D. Paré. 1999. “Impact of Network Activity on the Integrative Properties of Neocortical Pyramidal Neurons in Vivo.” *Journal of Neurophysiology* 81 (4): 1531–47. https://doi.org/10.1152/jn.1999.81.4.1531.

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

Douglas, R. J., C. Koch, M. Mahowald, K. A. Martin, and H. H. Suarez. 1995. “Recurrent Excitation in Neocortical Circuits.” *Science* 269 (5226): 981–85. https://doi.org/10.1126/science.7638624.

Douglas, Rodney J., and Kevan A. C. Martin. 2004. “Neuronal Circuits of the Neocortex.” *Annual Review of Neuroscience* 27: 419–51. https://doi.org/10.1146/annurev.neuro.27.070203.144152.

Ferster, D., and K. D. Miller. 2000. “Neural Mechanisms of Orientation Selectivity in the Visual Cortex.” *Annual Review of Neuroscience* 23: 441–71. https://doi.org/10.1146/annurev.neuro.23.1.441.

Fino, Elodie, Adam M. Packer, and Rafael Yuste. 2013. “The Logic of Inhibitory Connectivity in the Neocortex.” *The Neuroscientist: A Review Journal Bringing Neurobiology, Neurology and Psychiatry* 19 (3): 228–37. https://doi.org/10.1177/1073858412456743.

Fino, Elodie, and Rafael Yuste. 2011. “Dense Inhibitory Connectivity in Neocortex.” *Neuron* 69 (6): 1188–1203. https://doi.org/10.1016/j.neuron.2011.02.025.

Fitzpatrick, D. 2000. “Seeing beyond the Receptive Field in Primary Visual Cortex.” *Current Opinion in Neurobiology* 10 (4): 438–43. https://doi.org/10.1016/s0959-4388(00)00113-6.

Freeman, Tobe C. B., Séverine Durand, Daniel C. Kiper, and Matteo Carandini. 2002. "Suppression without Inhibition in Visual Cortex." *Neuron* 35 (4): 759–71. https://doi.org/10.1016/s0896-6273(02)00819-x.

Gardner-Medwin, A. R. 1976. "The Recall of Events through the Learning of Associations between Their Parts." *Proceedings of the Royal Society of London. Series B, Biological Sciences* 194 (1116): 375–402. https://doi.org/10.1098/rspb.1976.0084.

Garrett, Marina E., Ian Nauhaus, James H. Marshel, and Edward M. Callaway. 2014. "Topography and Areal Organization of Mouse Visual Cortex." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 34 (37): 12587–600. https://doi.org/10.1523/JNEUROSCI.1124-14.2014.

Glickfeld, Lindsey L., Mark H. Histed, and John H. R. Maunsell. 2013. "Mouse Primary Visual Cortex Is Used to Detect Both Orientation and Contrast Changes." *Journal of Neuroscience* 33 (50): 19416–22. https://doi.org/10.1523/JNEUROSCI.3560-13.2013.

Grubb, Matthew S., and Ian D. Thompson. 2003. "Quantitative Characterization of Visual Response Properties in the Mouse Dorsal Lateral Geniculate Nucleus." *Journal of Neurophysiology* 90 (6): 3594–3607. https://doi.org/10.1152/jn.00699.2003.

Haider, Bilal, Alvaro Duque, Andrea R. Hasenstaub, and David A. McCormick. 2006. "Neocortical Network Activity In Vivo Is Generated through a Dynamic Balance of Excitation and Inhibition." *Journal of Neuroscience* 26 (17): 4535–45. https://doi.org/10.1523/JNEUROSCI.5297-05.2006.

Haider, Bilal, Michael Häusser, and Matteo Carandini. 2013. "Inhibition Dominates Sensory Responses in Awake Cortex." *Nature* 493 (7430): 97–100. https://doi.org/10.1038/nature11665.

Harris, Kenneth D., and Thomas D. Mrsic-Flogel. 2013. "Cortical Connectivity and Sensory Coding." *Nature* 503 (7474): 51–58. https://doi.org/10.1038/nature12654.

Hauesler, Stefan, and Wolfgang Maass. 2017. "Inhibitory Networks Orchestrate the Self-Organization of Computational Function in Cortical Microcircuit Motifs through STDP." *BioRxiv*, December, 228759. https://doi.org/10.1101/228759.

Heeger, D. J. 1992. "Normalization of Cell Responses in Cat Striate Cortex." *Visual Neuroscience* 9 (2): 181–97.

Hennequin, Guillaume, Yashar Ahmadian, Daniel B. Rubin, Máté Lengyel, and Kenneth D. Miller. 2018. "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. https://doi.org/10.1016/j.neuron.2018.04.017.

Heuer, Hilary W., and Kenneth H. Britten. 2002. “Contrast Dependence of Response Normalization in Area MT of the Rhesus Macaque.” *Journal of Neurophysiology* 88 (6): 3398–3408. https://doi.org/10.1152/jn.00255.2002.

———. 2007. “Linear Responses to Stochastic Motion Signals in Area MST.” *Journal of Neurophysiology* 98 (3): 1115–24. https://doi.org/10.1152/jn.00083.2007.

Histed, Mark H. 2018. “Feedforward Inhibition Allows Input Summation to Vary in Recurrent Cortical Networks.” *ENeuro* 5 (1). https://doi.org/10.1523/ENEURO.0356-17.2018.

Hô, N., and A. Destexhe. 2000. “Synaptic Background Activity Enhances the Responsiveness of Neocortical Pyramidal Neurons.” *Journal of Neurophysiology* 84 (3): 1488–96. https://doi.org/10.1152/jn.2000.84.3.1488.

Hofer, Sonja B., Ho Ko, Bruno Pichler, Joshua Vogelstein, Hana Ros, Hongkui Zeng, Ed Lein, Nicholas A. Lesica, and Thomas D. Mrsic-Flogel. 2011. “Differential Connectivity and Response Dynamics of Excitatory and Inhibitory Neurons in Visual Cortex.” *Nature Neuroscience* 14 (8): 1045–52. https://doi.org/10.1038/nn.2876.

Hoy, Jennifer L., Iryna Yavorska, Michael Wehr, and Cristopher M. Niell. 2016. “Vision Drives Accurate Approach Behavior during Prey Capture in Laboratory Mice.” *Current Biology: CB* 26 (22): 3046–52. https://doi.org/10.1016/j.cub.2016.09.009.

Huang, Chengcheng, Douglas A. Ruff, Ryan Pyle, Robert Rosenbaum, Marlene R. Cohen, and Brent Doiron. 2019. “Circuit Models of Low-Dimensional Shared Variability in Cortical Networks.” *Neuron* 101 (2): 337-348.e4. https://doi.org/10.1016/j.neuron.2018.11.034.

Hubel, D. H., and T. N. Wiesel. 1959. “Receptive Fields of Single Neurones in the Cat’s Striate Cortex.” *The Journal of Physiology* 148 (3): 574–91. https://doi.org/10.1113/jphysiol.1959.sp006308.

———. 1962. “Receptive Fields, Binocular Interaction and Functional Architecture in the Cat’s Visual Cortex.” *The Journal of Physiology* 160 (1): 106-154.2.

———. 1968. “Receptive Fields and Functional Architecture of Monkey Striate Cortex.” *The Journal of Physiology* 195 (1): 215–43. https://doi.org/10.1113/jphysiol.1968.sp008455.

Hubel, David H., Torsten N. Wiesel, and Michael P. Stryker. 1977. “Orientation Columns in Macaque Monkey Visual Cortex Demonstrated by the 2-Deoxyglucose Autoradiographic Technique.” *Nature* 269 (5626): 328–30. https://doi.org/10.1038/269328a0.

Huberman, Andrew D., and Cristopher M. Niell. 2011. “What Can Mice Tell Us about How Vision Works?” *Trends in Neurosciences* 34 (9): 464–73. https://doi.org/10.1016/j.tins.2011.07.002.

Ibrahim, Leena A., Lukas Mesik, Xu-Ying Ji, Qi Fang, Hai-Fu Li, Ya-Tang Li, Brian Zingg, Li I. Zhang, and Huizhong Whit Tao. 2016. "Cross-Modality Sharpening of Visual Cortical Processing through Layer-1-Mediated Inhibition and Disinhibition." *Neuron* 89 (5): 1031–45. https://doi.org/10.1016/j.neuron.2016.01.027.

International Brain Laboratory. 2017. "An International Laboratory for Systems and Computational Neuroscience." *Neuron* 96 (6): 1213–18. https://doi.org/10.1016/j.neuron.2017.12.013.

Isaacson, Jeffry S. 2010. "Odor Representations in Mammalian Cortical Circuits." *Current Opinion in Neurobiology* 20 (3): 328–31. https://doi.org/10.1016/j.conb.2010.02.004.

Isaacson, Jeffry S., and Massimo Scanziani. 2011. "How Inhibition Shapes Cortical Activity." *Neuron* 72 (2): 231–43. https://doi.org/10.1016/j.neuron.2011.09.027.

Iurilli, Giuliano, Diego Ghezzi, Umberto Olcese, Glenda Lassi, Cristiano Nazzaro, Raffaella Tonini, Valter Tucci, Fabio Benfenati, and Paolo Medini. 2012. "Sound-Driven Synaptic Inhibition in Primary Visual Cortex." *Neuron* 73 (4): 814–28. https://doi.org/10.1016/j.neuron.2011.12.026.

Jun, James J., Nicholas A. Steinmetz, Joshua H. Siegle, Daniel J. Denman, Marius Bauza, Brian Barbarits, Albert K. Lee, et al. 2017. "Fully Integrated Silicon Probes for High-Density Recording of Neural Activity." *Nature* 551 (7679): 232–36. https://doi.org/10.1038/nature24636.

Kafashan, MohammadMehdi, Anna W. Jaffe, Selmaan N. Chettih, Ramon Nogueira, Iñigo Arandia-Romero, Christopher D. Harvey, Rubén Moreno-Bote, and Jan Drugowitsch. 2021. "Scaling of Sensory Information in Large Neural Populations Shows Signatures of Information-Limiting Correlations." *Nature Communications* 12 (1): 473. https://doi.org/10.1038/s41467-020-20722-y.

Kalatsky, Valery A., and Michael P. Stryker. 2003. "New Paradigm for Optical Imaging: Temporally Encoded Maps of Intrinsic Signal." *Neuron* 38 (4): 529–45.

Kanashiro, Tatjana, Gabriel Koch Ocker, Marlene R. Cohen, and Brent Doiron. 2017. "Attentional Modulation of Neuronal Variability in Circuit Models of Cortex." *ELife* 6 (June). https://doi.org/10.7554/eLife.23978.

Katzner, S., L. Busse, and M. Carandini. 2011. "GABAA Inhibition Controls Response Gain in Visual Cortex." *Journal of Neuroscience* 31 (16): 5931–41. https://doi.org/10.1523/JNEUROSCI.5753-10.2011.

Kawaguchi, Y., and Y. Kubota. 1997. "GABAergic Cell Subtypes and Their Synaptic Connections in Rat Frontal Cortex." *Cerebral Cortex (New York, N.Y.: 1991)* 7 (6): 476–86.

Keller, Andreas J., and Kevan A. C. Martin. 2015. “Local Circuits for Contrast Normalization and Adaptation Investigated with Two-Photon Imaging in Cat Primary Visual Cortex.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 35 (27): 10078–87. https://doi.org/10.1523/JNEUROSCI.0906-15.2015.

Kerlin, Aaron M., Mark L. Andermann, Vladimir K. Berezovskii, and R. Clay Reid. 2010. “Broadly Tuned Response Properties of Diverse Inhibitory Neuron Subtypes in Mouse Visual Cortex.” *Neuron* 67 (5): 858–71. https://doi.org/10.1016/j.neuron.2010.08.002.

Ko, Ho, Sonja B. Hofer, Bruno Pichler, Katherine A. Buchanan, P. Jesper Sjöström, and Thomas D. Mrsic-Flogel. 2011. “Functional Specificity of Local Synaptic Connections in Neocortical Networks.” *Nature* 473 (7345): 87–91. https://doi.org/10.1038/nature09880.

Kohn, Adam, Ruben Coen-Cagli, Ingmar Kanitscheider, and Alexandre Pouget. 2016. “Correlations and Neuronal Population Information.” *Annual Review of Neuroscience* 39 (July): 237–56. https://doi.org/10.1146/annurev-neuro-070815-013851.

Kohn, Adam, and Matthew A. Smith. 2005. “Stimulus Dependence of Neuronal Correlation in Primary Visual Cortex of the Macaque.” *The Journal of Neuroscience* 25 (14): 3661–73. https://doi.org/10.1523/JNEUROSCI.5106-04.2005.

Laughlin, S. 1981. “A Simple Coding Procedure Enhances a Neuron’s Information Capacity.” *Zeitschrift Fur Naturforschung. Section C, Biosciences* 36 (9–10): 910–12.

Lauritzen, Thomas Z., and Kenneth D. Miller. 2003. “Different Roles for Simple-Cell and Complex-Cell Inhibition in V1.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 23 (32): 10201–13.

Lee, Joonyeol, and John H. R. Maunsell. 2009. “A Normalization Model of Attentional Modulation of Single Unit Responses.” *PloS One* 4 (2): e4651. https://doi.org/10.1371/journal.pone.0004651.

———. 2010. “Attentional Modulation of MT Neurons with Single or Multiple Stimuli in Their Receptive Fields.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 30 (8): 3058–66. https://doi.org/10.1523/JNEUROSCI.3766-09.2010.

Lennie, Peter, and J. Anthony Movshon. 2005. “Coding of Color and Form in the Geniculostriate Visual Pathway (Invited Review).” *JOSA A* 22 (10): 2013–33. https://doi.org/10.1364/JOSAA.22.002013.

Levy, Maayan, Olaf Sporns, and Jason N. MacLean. 2020. “Network Analysis of Murine Cortical Dynamics Implicates Untuned Neurons in Visual Stimulus Coding.” *Cell Reports* 31 (2): 107483. https://doi.org/10.1016/j.celrep.2020.03.047.

Litwin-Kumar, Ashok, Anne-Marie M. Oswald, Nathaniel N. Urban, and Brent Doiron. 2011. "Balanced Synaptic Input Shapes the Correlation between Neural Spike Trains." *PLoS Computational Biology* 7 (12): e1002305. https://doi.org/10.1371/journal.pcbi.1002305.

Liu, Bao-hua, Pingyang Li, Ya-tang Li, Yujiao J. Sun, Yuchio Yanagawa, Kunihiko Obata, Li I. Zhang, and Huizhong W. Tao. 2009. "Visual Receptive Field Structure of Cortical Inhibitory Neurons Revealed by Two-Photon Imaging Guided Recording." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 29 (34): 10520–32. https://doi.org/10.1523/JNEUROSCI.1915-09.2009.

Louie, Kenway, Lauren E. Grattan, and Paul W. Glimcher. 2011. "Reward Value-Based Gain Control: Divisive Normalization in Parietal Cortex." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 31 (29): 10627–39. https://doi.org/10.1523/JNEUROSCI.1237-11.2011.

Lu, Rongwen, Yajie Liang, Guanghan Meng, Pengcheng Zhou, Karel Svoboda, Liam Paninski, and Na Ji. 2020. "Rapid Mesoscale Volumetric Imaging of Neural Activity with Synaptic Resolution." *Nature Methods* 17 (3): 291–94. https://doi.org/10.1038/s41592-020-0760-9.

Luo, Liqun, Edward M. Callaway, and Karel Svoboda. 2008. "Genetic Dissection of Neural Circuits." *Neuron* 57 (5): 634–60. https://doi.org/10.1016/j.neuron.2008.01.002.

———. 2018. "Genetic Dissection of Neural Circuits: A Decade of Progress." *Neuron* 98 (2): 256–81. https://doi.org/10.1016/j.neuron.2018.03.040.

Ma, Wen-pei, Bao-hua Liu, Ya-tang Li, Z. Josh Huang, Li I. Zhang, and Huizhong W. Tao. 2010. "Visual Representations by Cortical Somatostatin Inhibitory Neurons--Selective but with Weak and Delayed Responses." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 30 (43): 14371–79. https://doi.org/10.1523/JNEUROSCI.3248-10.2010.

MacEvoy, Sean P., Thomas R. Tucker, and David Fitzpatrick. 2009. "A Precise Form of Divisive Suppression Supports Population Coding in the Primary Visual Cortex." *Nature Neuroscience* 12 (5): 637–45. https://doi.org/10.1038/nn.2310.

MacLean, Jason N., and Rafael Yuste. 2009. "Imaging Action Potentials with Calcium Indicators." *Cold Spring Harbor Protocols* 2009 (11): pdb.prot5316. https://doi.org/10.1101/pdb.prot5316.

Malo, Jesús, and Valero Laparra. 2010. "Psychophysically Tuned Divisive Normalization Approximately Factorizes the PDF of Natural Images." *Neural Computation* 22 (12): 3179–3206. https://doi.org/10.1162/NECO_a_00046.

Marr, David. 1982. “Vision: A Computational Investigation into the Human Representation and Processing of Visual Information.” In *New York, NY: W.H. Freeman and Company*. CUMINCAD. http://papers.cumincad.org/cgi-bin/works/Show?fafa.

Masters, Barry R., and Peter T. C. So. 2004. “Antecedents of Two-Photon Excitation Laser Scanning Microscopy.” *Microscopy Research and Technique* 63 (1): 3–11. https://doi.org/10.1002/jemt.10418.

Maunsell, J. H., and DC van Essen. 1983. “The Connections of the Middle Temporal Visual Area (MT) and Their Relationship to a Cortical Hierarchy in the Macaque Monkey.” *Journal of Neuroscience* 3 (12): 2563–86. https://doi.org/10.1523/JNEUROSCI.03-12-02563.1983.

Maunsell, John H. R., and Stefan Treue. 2006. “Feature-Based Attention in Visual Cortex.” *Trends in Neurosciences* 29 (6): 317–22. https://doi.org/10.1016/j.tins.2006.04.001.

Mayer, Maria Goeppert, and J. Hans D. Jensen. 1955. *Elementary Theory of Nuclear Shell Structure*. Wiley.

Mazurek, Mark, Marisa Kager, and Stephen D. Van Hooser. 2014. “Robust Quantification of Orientation Selectivity and Direction Selectivity.” *Frontiers in Neural Circuits* 8: 92. https://doi.org/10.3389/fncir.2014.00092.

McBride, Ethan G., Su-Yee J. Lee, and Edward M. Callaway. 2019. “Local and Global Influences of Visual Spatial Selection and Locomotion in Mouse Primary Visual Cortex.” *Current Biology: CB* 29 (10): 1592-1605.e5. https://doi.org/10.1016/j.cub.2019.03.065.

Meyer, Axel H., István Katona, Maria Blatow, Andrei Rozov, and Hannah Monyer. 2002. “In Vivo Labeling of Parvalbumin-Positive Interneurons and Analysis of Electrical Coupling in Identified Neurons.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 22 (16): 7055–64. https://doi.org/20026742.

Moreno-Bote, Rubén, Jeffrey Beck, Ingmar Kanitscheider, Xaq Pitkow, Peter Latham, and Alexandre Pouget. 2014. “Information-Limiting Correlations.” *Nature Neuroscience* 17 (10): 1410–17. https://doi.org/10.1038/nn.3807.

Morrone, M. C., D. C. Burr, and L. Maffei. 1982. “Functional Implications of Cross-Orientation Inhibition of Cortical Visual Cells. I. Neurophysiological Evidence.” *Proceedings of the Royal Society of London. Series B, Biological Sciences* 216 (1204): 335–54. https://doi.org/10.1098/rspb.1982.0078.

Nassi, Jonathan J., Michael C. Avery, Ali H. Cetin, Anna W. Roe, and John H. Reynolds. 2015. “Optogenetic Activation of Normalization in Alert Macaque Visual Cortex.” *Neuron* 86 (6): 1504–17. https://doi.org/10.1016/j.neuron.2015.05.040.

Ni, Amy M., and John H. R. Maunsell. 2017. “Spatially Tuned Normalization Explains Attention Modulation Variance within Neurons.” *Journal of Neurophysiology* 118 (3): 1903–13. https://doi.org/10.1152/jn.00218.2017.

———. 2019. “Neuronal Effects of Spatial and Feature Attention Differ Due to Normalization.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 39 (28): 5493–5505. https://doi.org/10.1523/JNEUROSCI.2106-18.2019.

Ni, Amy M., Supratim Ray, and John H. R. Maunsell. 2012. “Tuned Normalization Explains the Size of Attention Modulations.” *Neuron* 73 (4): 803–13. https://doi.org/10.1016/j.neuron.2012.01.006.

Niell, Cristopher M. 2015. “Cell Types, Circuits, and Receptive Fields in the Mouse Visual Cortex.” *Annual Review of Neuroscience* 38 (July): 413–31. https://doi.org/10.1146/annurev-neuro-071714-033807.

Niell, Cristopher M., and Michael P. Stryker. 2008. “Highly Selective Receptive Fields in Mouse Visual Cortex.” *Journal of Neuroscience* 28 (30): 7520–36. https://doi.org/10.1523/JNEUROSCI.0623-08.2008.

———. 2010. “Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex.” *Neuron* 65 (4): 472–79. https://doi.org/10.1016/j.neuron.2010.01.033.

Nienborg, Hendrikje, Andrea Hasenstaub, Ian Nauhaus, Hiroki Taniguchi, Z. Josh Huang, and Edward M. Callaway. 2013. “Contrast Dependence and Differential Contributions from Somatostatin- and Parvalbumin-Expressing Neurons to Spatial Integration in Mouse V1.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 33 (27): 11145–54. https://doi.org/10.1523/JNEUROSCI.5320-12.2013.

O’Connor, Daniel H., Daniel Huber, and Karel Svoboda. 2009. “Reverse Engineering the Mouse Brain.” *Nature* 461 (7266): 923–29. https://doi.org/10.1038/nature08539.

Ohshiro, Tomokazu, Dora E. Angelaki, and Gregory C. DeAngelis. 2011. “A Normalization Model of Multisensory Integration.” *Nature Neuroscience* 14 (6): 775–82. https://doi.org/10.1038/nn.2815.

Olsen, Shawn R., Vikas Bhandawat, and Rachel I. Wilson. 2010. “Divisive Normalization in Olfactory Population Codes.” *Neuron* 66 (2): 287–99. https://doi.org/10.1016/j.neuron.2010.04.009.

Ozeki, Hirofumi, Ian M. Finn, Evan S. Schaffer, Kenneth D. Miller, and David Ferster. 2009. “Inhibitory Stabilization of the Cortical Network Underlies Visual Surround Suppression.” *Neuron* 62 (4): 578–92. https://doi.org/10.1016/j.neuron.2009.03.028.

Packer, Adam M., and Rafael Yuste. 2011. “Dense, Unspecific Connectivity of Neocortical Parvalbumin-Positive Interneurons: A Canonical Microcircuit for Inhibition?” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 31 (37): 13260–71. https://doi.org/10.1523/JNEUROSCI.3131-11.2011.

Paiton, Dylan M., Charles G. Frye, Sheng Y. Lundquist, Joel D. Bowen, Ryan Zarcone, and Bruno A. Olshausen. 2020. “Selectivity and Robustness of Sparse Coding Networks.” *Journal of Vision* 20 (12): 10. https://doi.org/10.1167/jov.20.12.10.

Priebe, Nicholas J, and David Ferster. 2006. “Mechanisms Underlying Cross-Orientation Suppression in Cat Visual Cortex.” *Nature Neuroscience* 9 (4): 552–61. https://doi.org/10.1038/nn1660.

Priebe, Nicholas J., and David Ferster. 2008. “Inhibition, Spike Threshold, and Stimulus Selectivity in Primary Visual Cortex.” *Neuron* 57 (4): 482–97. https://doi.org/10.1016/j.neuron.2008.02.005.

———. 2012. “Mechanisms of Neuronal Computation in Mammalian Visual Cortex.” *Neuron* 75 (2): 194–208. https://doi.org/10.1016/j.neuron.2012.06.011.

Prusky, G. T., P. W. West, and R. M. Douglas. 2000. “Behavioral Assessment of Visual Acuity in Mice and Rats.” *Vision Research* 40 (16): 2201–9. https://doi.org/10.1016/s0042-6989(00)00081-x.

Renart, Alfonso, Jaime de la Rocha, Peter Bartho, Liad Hollender, Néstor Parga, Alex Reyes, and Kenneth D. Harris. 2010. “The Asynchronous State in Cortical Circuits.” *Science (New York, N.Y.)* 327 (5965): 587–90. https://doi.org/10.1126/science.1179850.

Reynolds, John H., and Leonardo Chelazzi. 2004. “Attentional Modulation of Visual Processing.” *Annual Review of Neuroscience* 27: 611–47. https://doi.org/10.1146/annurev.neuro.26.041002.131039.

Reynolds, John H., and David J. Heeger. 2009. “The Normalization Model of Attention.” *Neuron* 61 (2): 168–85. https://doi.org/10.1016/j.neuron.2009.01.002.

Ringach, Dario L. 2004. “Mapping Receptive Fields in Primary Visual Cortex.” *The Journal of Physiology* 558 (Pt 3): 717–28. https://doi.org/10.1113/jphysiol.2004.065771.

Ringach, Dario L., Robert M. Shapley, and Michael J. Hawken. 2002. “Orientation Selectivity in Macaque V1: Diversity and Laminar Dependence.” *Journal of Neuroscience* 22 (13): 5639–51. https://doi.org/10.1523/JNEUROSCI.22-13-05639.2002.

Rosenbaum, Robert, Jonathan E. Rubin, and Brent Doiron. 2013. “Short-Term Synaptic Depression and Stochastic Vesicle Dynamics Reduce and Shape Neuronal Correlations.” *Journal of Neurophysiology* 109 (2): 475–84. https://doi.org/10.1152/jn.00733.2012.

Rubin, Daniel B., Stephen D. Van Hooser, and Kenneth D. Miller. 2015. “The Stabilized Supralinear Network: A Unifying Circuit Motif Underlying Multi-Input Integration in Sensory Cortex.” *Neuron* 85 (2): 402–17. https://doi.org/10.1016/j.neuron.2014.12.026.

Ruff, Douglas A., Joshua J. Alberts, and Marlene R. Cohen. 2016. “Relating Normalization to Neuronal Populations across Cortical Areas.” *Journal of Neurophysiology* 116 (3): 1375–86. https://doi.org/10.1152/jn.00017.2016.

Ruff, Douglas A., and Marlene R. Cohen. 2014a. “Attention Can Either Increase or Decrease Spike Count Correlations in Visual Cortex.” *Nature Neuroscience* 17 (11): 1591–97. https://doi.org/10.1038/nn.3835.

———. 2014b. “Global Cognitive Factors Modulate Correlated Response Variability between V4 Neurons.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 34 (49): 16408–16. https://doi.org/10.1523/JNEUROSCI.2750-14.2014.

Runfeldt, Melissa J., Alexander J. Sadovsky, and Jason N. MacLean. 2014. “Acetylcholine Functionally Reorganizes Neocortical Microcircuits.” *Journal of Neurophysiology* 112 (5): 1205–16. https://doi.org/10.1152/jn.00071.2014.

Runyan, Caroline A., James Schummers, Audra Van Wart, Sandra J. Kuhlman, Nathan R. Wilson, Z. Josh Huang, and Mriganka Sur. 2010. “Response Features of Parvalbumin-Expressing Interneurons Suggest Precise Roles for Subtypes of Inhibition in Visual Cortex.” *Neuron* 67 (5): 847–57. https://doi.org/10.1016/j.neuron.2010.08.006.

Rust, Nicole C., Valerio Mante, Eero P. Simoncelli, and J. Anthony Movshon. 2006. “How MT Cells Analyze the Motion of Visual Patterns.” *Nature Neuroscience* 9 (11): 1421–31. https://doi.org/10.1038/nn1786.

Sadovsky, Alexander J., Peter B. Kruskal, Joseph M. Kimmel, Jared Ostmeyer, Florian B. Neubauer, and Jason N. MacLean. 2011. “Heuristically Optimal Path Scanning for High-Speed Multiphoton Circuit Imaging.” *Journal of Neurophysiology* 106 (3): 1591–98. https://doi.org/10.1152/jn.00334.2011.

Sanchez-Giraldo, Luis G, Md Nasir Uddin Laskar, and Odelia Schwartz. 2019. “Normalization and Pooling in Hierarchical Models of Natural Images.” *Current Opinion in Neurobiology*, Machine Learning, Big Data, and Neuroscience, 55 (April): 65–72. https://doi.org/10.1016/j.conb.2019.01.008.

Saproo, Sameer, and John T. Serences. 2014. “Attention Improves Transfer of Motion Information between V1 and MT.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 34 (10): 3586–96. https://doi.org/10.1523/JNEUROSCI.3484-13.2014.

Sato, Tatsuo K, Bilal Haider, Michael Häusser, and Matteo Carandini. 2016a. “An Excitatory Basis for Divisive Normalization in Visual Cortex.” *Nature Neuroscience* 19 (4): 568–70. https://doi.org/10.1038/nn.4249.

Sato, Tatsuo K., Bilal Haider, Michael Häusser, and Matteo Carandini. 2016b. “An Excitatory Basis for Divisive Normalization in Visual Cortex.” *Nature Neuroscience* 19 (4): 568–70. https://doi.org/10.1038/nn.4249.

Sawada, Tadamasa, and Alexander A. Petrov. 2017. “The Divisive Normalization Model of V1 Neurons: A Comprehensive Comparison of Physiological Data and Model Predictions.” *Journal of Neurophysiology* 118 (6): 3051–91. https://doi.org/10.1152/jn.00821.2016.

Schwartz, O., and E. P. Simoncelli. 2001. “Natural Signal Statistics and Sensory Gain Control.” *Nature Neuroscience* 4 (8): 819–25. https://doi.org/10.1038/90526.

Schwartz, Odelia, and Ruben Coen-Cagli. 2013. “Visual Attention and Flexible Normalization Pools.” *Journal of Vision* 13 (1). https://doi.org/10.1167/13.1.25.

Seabrook, Tania A., Timothy J. Burbridge, Michael C. Crair, and Andrew D. Huberman. 2017. “Architecture, Function, and Assembly of the Mouse Visual System.” *Annual Review of Neuroscience* 40 (July): 499–538. https://doi.org/10.1146/annurev-neuro-071714-033842.

Seung, H. S., and H. Sompolinsky. 1993. “Simple Models for Reading Neuronal Population Codes.” *Proceedings of the National Academy of Sciences of the United States of America* 90 (22): 10749–53. https://doi.org/10.1073/pnas.90.22.10749.

Shadlen, Michael N., and William T. Newsome. 1994. “Noise, Neural Codes and Cortical Organization.” *Current Opinion in Neurobiology* 4 (4): 569–79. https://doi.org/10.1016/0959-4388(94)90059-0.

Shushruth, S., Pradeep Mangapathy, Jennifer M. Ichida, Paul C. Bressloff, Lars Schwabe, and Alessandra Angelucci. 2012. “Strong Recurrent Networks Compute the Orientation Tuning of Surround Modulation in the Primate Primary Visual Cortex.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 32 (1): 308–21. https://doi.org/10.1523/JNEUROSCI.3789-11.2012.

Simoncelli, E. P., and D. J. Heeger. 1998. “A Model of Neuronal Responses in Visual Area MT.” *Vision Research* 38 (5): 743–61. https://doi.org/10.1016/s0042-6989(97)00183-1.

Simoncelli, E. P., and B. A. Olshausen. 2001. “Natural Image Statistics and Neural Representation.” *Annual Review of Neuroscience* 24: 1193–1216. https://doi.org/10.1146/annurev.neuro.24.1.1193.

Smith, Matthew A., and Adam Kohn. 2008. “Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex.” *The Journal of Neuroscience : The Official*

Journal of the Society for Neuroscience, November. https://doi.org/10.1523/jneurosci.2929-08.2008.

Sompolinsky, H., and R. Shapley. 1997. "New Perspectives on the Mechanisms for Orientation Selectivity." *Current Opinion in Neurobiology* 7 (4): 514–22. https://doi.org/10.1016/s0959-4388(97)80031-1.

Steinmetz, Nicholas A, Christof Koch, Kenneth D Harris, and Matteo Carandini. 2018. "Challenges and Opportunities for Large-Scale Electrophysiology with Neuropixels Probes." *Current Opinion in Neurobiology*, Neurotechnologies, 50 (June): 92–100. https://doi.org/10.1016/j.conb.2018.01.009.

Stosiek, Christoph, Olga Garaschuk, Knut Holthoff, and Arthur Konnerth. 2003. "In Vivo Two-Photon Calcium Imaging of Neuronal Networks." *Proceedings of the National Academy of Sciences of the United States of America* 100 (12): 7319–24. https://doi.org/10.1073/pnas.1232232100.

Tang, Jiaying, Silvia C. Ardila Jimenez, Subhojit Chakraborty, and Simon R. Schultz. 2016. "Visual Receptive Field Properties of Neurons in the Mouse Lateral Geniculate Nucleus." *PLoS ONE* 11 (1). https://doi.org/10.1371/journal.pone.0146017.

Tetzlaff, Christian, Christoph Kolodziejski, Marc Timme, and Florentin Wörgötter. 2012. "Analysis of Synaptic Scaling in Combination with Hebbian Plasticity in Several Simple Networks." *Frontiers in Computational Neuroscience* 6: 36. https://doi.org/10.3389/fncom.2012.00036.

Tolhurst, D. J., and D. J. Heeger. 1997. "Contrast Normalization and a Linear Model for the Directional Selectivity of Simple Cells in Cat Striate Cortex." *Visual Neuroscience* 14 (1): 19–25. https://doi.org/10.1017/s0952523800008725.

Tripp, Bryan P. 2012. "Decorrelation of Spiking Variability and Improved Information Transfer through Feedforward Divisive Normalization." *Neural Computation* 24 (4): 867–94. https://doi.org/10.1162/NECO_a_00255.

Valerio, Roberto, and Rafael Navarro. 2003. "Optimal Coding through Divisive Normalization Models of V1 Neurons." *Network (Bristol, England)* 14 (3): 579–93.

Verhoef, Bram-Ernst, and John H. R. Maunsell. 2017. "Attention-Related Changes in Correlated Neuronal Activity Arise from Normalization Mechanisms." *Nature Neuroscience* 20 (7): 969–77. https://doi.org/10.1038/nn.4572.

Vogels, Tim P., Kanaka Rajan, and L. F. Abbott. 2005. "Neural Network Dynamics." *Annual Review of Neuroscience* 28: 357–76. https://doi.org/10.1146/annurev.neuro.28.061604.135637.

Wang, Quanxin, and Andreas Burkhalter. 2007. “Area Map of Mouse Visual Cortex.” *The Journal of Comparative Neurology* 502 (3): 339–57. https://doi.org/10.1002/cne.21286.

Wilson, Nathan R., Caroline A. Runyan, Forea L. Wang, and Mriganka Sur. 2012. “Division and Subtraction by Distinct Cortical Inhibitory Networks in Vivo.” *Nature* 488 (7411): 343–48. https://doi.org/10.1038/nature11347.

Wolf, Fred, Rainer Engelken, Maximilian Puelma-Touzel, Juan Daniel Flórez Weidinger, and Andreas Neef. 2014. “Dynamical Models of Cortical Circuits.” *Current Opinion in Neurobiology* 25 (April): 228–36. https://doi.org/10.1016/j.conb.2014.01.017.

Yuste, R., and L. C. Katz. 1991. “Control of Postsynaptic Ca2+ Influx in Developing Neocortex by Excitatory and Inhibitory Neurotransmitters.” *Neuron* 6 (3): 333–44. https://doi.org/10.1016/0896-6273(91)90243-s.

Znamenskiy, Petr, Mean-Hwan Kim, Dylan R. Muir, Maria Florencia Iacaruso, Sonja B. Hofer, and Thomas D. Mrsic-Flogel. 2018. “Functional Selectivity and Specific Connectivity of Inhibitory Neurons in Primary Visual Cortex.” *BioRxiv*, April, 294835. https://doi.org/10.1101/294835.

Zoccolan, Davide, David D. Cox, and James J. DiCarlo. 2005. “Multiple Object Response Normalization in Monkey Inferotemporal Cortex.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 25 (36): 8150–64. https://doi.org/10.1523/JNEUROSCI.2058-05.2005.