UNIVERSITY OF CALIFORNIA, SAN DIEGO

Cortical Dynamics in Visual Processing

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Neurosciences with a Specialization in Computational Neuroscience

by

Kimberly Elise Reinhold

Committee in charge:

Professor Massimo Scanziani, Chair
Professor Edward Callaway
Professor Timothy Gentner
Professor Jeffry Isaacson
Professor John Reynolds

2015

The Dissertation of Kimberly Elise Reinhold is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____

Chair

University of California, San Diego

2015

TABLE OF CONTENTS

LIST OF ABBREVIATIONS

CDF	cortical decay function
LFP	local field potential
TRN	thalamic reticular nucleus
V1	primary visual cortex
dLGN	dorsal lateral geniculate nucleus
LP	lateral posterior nucleus
MU	multi-unit activity
SU	single-unit activity
EPSP	excitatory post-synaptic potential
fEPSP	field excitatory post-synaptic potential
EPSC	excitatory post-synaptic current
Norm.	normalized
Amp.	amplitude
L	cortical layer (e.g., L4 indicates cortical layer 4)
TC	thalamocortical
ITI	inter-trial interval
GABA	gamma aminobutyric acid
LED	light-emitting diode
ChR2	Channelrhodopsin 2, a light-activated cation channel
ArchT	Archaerhodopsin T, a light-activated proton pump
AAV	adeno-associated virus

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I acknowledge Massimo Scanziani for his invaluable scientific guidance and mentorship. I also thank Stephen Larson for his support. I thank Jamie Evora for help with genotyping and mouse husbandry. I would like to thank Jeffry Isaacson, Brenda Bloodgood, Charles Reinhold, Stephen Larson and members of the Scanziani and Isaacson labs for helpful discussions about the project and comments on manuscripts. This work was supported by the National Science Foundation Graduate Research Program Fellowship, the Gatsby charitable foundation and the Howard Hughes Medical Institute.

Chapters 1, 2, 3, 4 and 5, in part, have been submitted for publication of the material as it may appear in Reinhold, K., Lien, A.D. and Scanziani, M., 2015, Nature Neuroscience. The dissertation author was the primary investigator and author of this paper.

VITA

Summer 2004	Research Internship, MIT Media Laboratory
Summer 2005	Research Assistant, Joint Astronomy Centre Telescope
2005-2009	Bachelor of Science, Massachusetts Institute of Technology
Summer 2006	Research Assistant, MIT Computational Visual Cognition Laboratory
2007	Research Assistant, Harvard-MIT Health Sciences and Technology Program
Summer 2008	Research Internship, Columbia University
June 2007-May 2008, Sep. 2008-May 2009	Research Assistant, Picower Institute for Learning and Memory
2015	Doctor of Philosophy, University of California, San Diego

PUBLICATIONS

Osterweil, E., Krueger, D., Reinhold, K. & Bear, M. 2010. Hypersensitivity to, not hyperactivity of, mGluR5 in fragile X syndrome. Journal of Neuroscience.

Reinhold, K., Lien, A.D. & Scanziani, M. 2015. Distinct recurrent versus afferent dynamics in cortical visual processing. Nature Neuroscience.

FIELDS OF STUDY

Major Field: Neuroscience (Specialization or Focused Studies)

Studies in Systems Neuroscience
Professor Massimo Scanziani, University of California, San Diego
Professors Aude Oliva and Mark Bear, Massachusetts Institute of Technology

Studies in Artificial Intelligence
Professors Push Singh and Marco Ramoni, Massachusetts Institute of Technology

ABSTRACT OF THE DISSERTATION

Cortical Dynamics in Visual Processing

by

Kimberly Elise Reinhold

Doctor of Philosophy in Neurosciences with a Specialization in Computational Neuroscience

University of California, San Diego, 2015

Professor Massimo Scanziani, Chair

Recurrent circuits are a hallmark of mammalian sensory cortex. How they impact dynamics of sensory representation is not understood. Because recurrent circuits provide a majority of the synaptic excitation to cortical neurons in response to sensory stimulation, the intrinsic dynamics of these cortical recurrent circuits are expected to be a critical determinant of the timing of the sensory response in cortex. Previous methods could not isolate dynamics of these intra-cortical recurrent circuits from those of thalamic afferents during sensory processing. I now accomplish this by developing an approach to optogenetically silence thalamus in a model system: the mouse visual pathway. Silencing thalamus revealed the time course over which visually evoked activity in visual cortex was maintained by the intra-cortical recurrent circuits themselves, in isolation from thalamic input. I found that, at all time

points during the cortical sensory-evoked response, optogenetically silencing thalamus led to a fast decay of sensory-evoked activity in cortical recurrent circuits. This activity decay time course was fit by a 10 ms network time constant, similar to a neuron's integration time window. This decay time course was invariant across all tested visual stimulation conditions and behavioral states but depended on cortical inhibition. In awake mice, the dynamics of this time course predicted the time-locking of cortical activity to thalamic input at frequencies <15 Hz and the attenuation of the cortical response to higher frequencies. Under anesthesia, however, dynamics of depression at thalamocortical synapses disrupted the fidelity of sensory transmission. Thus, I determine sensory-evoked dynamics intrinsic to the intra-cortical recurrent circuits in isolation from thalamus and show how these dynamics transform afferent input in time.

Chapter 1:

Introduction and Background

1.1 Overview

A unique adaptation of the mammalian brain is the neocortex. Its hallmark is a highly complex and recurrent local circuit structure. It has been known for decades that these local recurrent circuits dominate the architecture of sensory neocortex, yet their unique role in perception is just beginning to be understood. While progress has been made in understanding how local intra-cortical recurrent circuits alter sensory representations of space [1-4], still very little is known about how recurrent circuits impact sensory-evoked activity in time.

The impact of cortical recurrent circuits on the temporal representation of neural activity is difficult to predict. If one population of cortical neurons drives a second population of cortical neurons, and that second population also drives the first, self-sustaining activity may result, prolonging the sensory-evoked response in cortex. However, different levels of cortical inhibition, which counteracts cortical recurrent excitation, may lead to different durations of the sensory response in cortex [5-8]. Thus, the following question must be addressed empirically: Does the recruitment of cortical recurrent circuits by a sensory stimulus impact the neural representation of that stimulus in time, and, if so, by what mechanism?

Answering this question has been technically challenging. The cortex receives its primary source of sensory input from the thalamus via thalamic afferents to cortex, and the cortex projects back to thalamus. This reciprocal interaction between thalamus and cortex could alter the cortical representation of sensory activity in time. Hence, to determine how intra-cortical recurrent circuits themselves impact sensory activity in time, we must factor out the thalamic contribution to the duration and dynamics of the sensory response. Previous attempts to determine how cortical circuits impact sensory activity in time compared the spiking response in thalamus with the spiking response in cortex [9-18]. However, the synapses of thalamic afferents connecting thalamus to cortex exhibit their own

dynamic properties, like frequency-dependent short-term depression, which may impact the sensory response of the cortex [19-23].

How can we determine the dynamics of intra-cortical recurrent circuits in response to sensory stimulation in isolation from the thalamic afferents, given that activity in thalamic afferents is needed to trigger a response in cortical recurrent circuits? Using the thalamocortical visual system of the mouse as a model, I have developed an optogenetic method to silence, with millisecond-precision, visual thalamus at any arbitrary delay following the presentation of a visual stimulus. This approach allows me to follow the duration of visually evoked activity continuing in intra-cortical recurrent circuits in the absence of thalamic input and thus to directly measure sensory-evoked dynamics intrinsic to recurrent circuits of cortex.

To test whether this measurement of cortical dynamics predicts the temporal transformation of sensory activity by cortex, I then compare the spiking response of cortex with the output of thalamic afferent synapses, detected in isolation from cortical recurrent circuits by silencing recurrent circuits using a second optogenetic technique. Therefore the new multi-focal and multi-functional optogenetic approach described here (Figure 1.1) identifies relative contributions of thalamic versus cortical connections to the temporal processing of visual information.

I discover that the intrinsic dynamics of visual cortical recurrent circuits in response to sensory stimulation are on the order of the integration time window of a single cortical neuron (Chapter 2). Consistent with this, I show that, in awake mice, these dynamics of cortical recurrent circuits predict how visual cortex transforms the amplitude of sensory input (Chapter 3). Thus, cortical recurrent circuit dynamics limit the rate of the sensory response in the awake state. I then show data indicating that cortical inhibition affects these recurrent circuit dynamics (Chapter 4).

Interestingly, my data also show that, under anesthesia, the dynamics of synaptic depression at thalamic afferents, rather than the dynamics of cortical recurrent circuits, are rate-limiting in sensory processing, leading to a pronounced disruption of the timing of the sensory response in the non-conscious state (Chapter 3). Therefore, this thesis work identifies two mechanisms underlying the

temporal transformation of visually evoked activity by cortex. The interplay of these mechanisms may be a general principle of thalamocortical interaction.

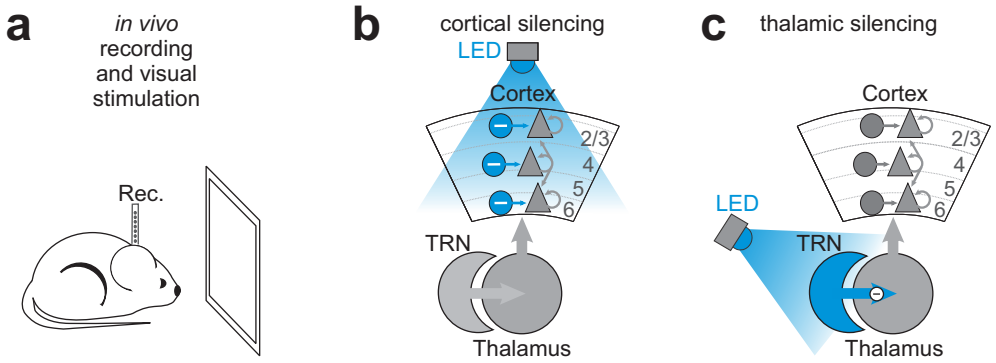

Figure 1.1: Multi-focal optogenetic approach to silence cortex or thalamus.

(a) Recording and visual stimulation configuration.

(b) Silencing cortical recurrent circuits by optogenetically activating local inhibitory interneurons expressing ChR2 to isolate thalamic afferent input.

(c) Silencing thalamus by optogenetically activating the thalamic reticular nucleus (TRN) expressing ChR2 to isolate cortical recurrent circuits.

1.2 Background

1.2.1 Strong recurrent circuits in cortex

Most synapses within neocortex are between cortical neurons. Far fewer synapses are afferents from other brain structures. Of these afferents to cortex from other brain structures, afferents from thalamus are a large and important fraction. However, even these thalamocortical afferents are small in quantity relative to the recurrent cortico-cortical connections within cortex. For example, in primary visual cortex (V1), even in the major thalamo-recipient layer, layer 4 (L4), only about 10% of the total excitatory synapses are feed-forward from thalamus [24-26]. The rest of the excitatory synapses are cortical recurrent connections. In other cortical layers, an even larger fraction of the total excitatory synapses are cortical recurrent connections [27,28]. Thus, recurrent circuits within cortex are anatomically dominant.

There is also evidence that these excitatory intra-cortical recurrent circuits are functionally dominant (i.e., stronger than the thalamic input). Experiments in cats [29] and mice [30-32] have assessed the strength of the thalamocortical synaptic excitation relative to the strength of the total recurrent synaptic excitation. To do this, these studies separated the thalamocortical synaptic excitation from recurrent synaptic excitation. In mice, researchers measured the thalamocortical synaptic excitation as the fraction of the sensory-evoked excitatory post-synaptic current (EPSC) to layer 4 neurons remaining when recurrent excitation in cortex was silenced. To silence recurrent excitatory circuits in cortex, these studies suppressed the spiking activity of cortical excitatory pyramidal cells by optogenetically photo-activating promiscuously connected inhibitory GABAergic interneurons expressing channelrhodopsin 2 (ChR2) that project to and inhibit pyramidal cells. Hence, during the optogenetic manipulation, the remaining synaptic excitation to cortical neurons in layer 4 was predominantly feedforward excitation from thalamic afferents. Silencing cortical excitatory recurrent circuits abolished more than two-thirds of the sensory-evoked EPSC in both primary auditory [32] and primary visual [30,31] cortices, suggesting that most of the cortical excitation depends on recurrent circuits and that direct afferent input from thalamus accounts for only a minor fraction of the total sensory-evoked EPSC, even in thalamo-recipient layer 4. Thus, these results suggest that cortical recurrent synaptic excitation may be larger in

magnitude than the thalamocortical synaptic excitation, at least during responses to simple sensory stimuli under anesthesia.

1.2.2 Impact of recurrent circuits on spatial structure of cortical response

A ubiquitous component of the cortical architecture, these strong cortical recurrent circuits may dramatically reshape the sensory-evoked response. We know something about how recurrent circuits reshape the spatial structure of the cortical response [1-4]. In cats [33] and primates [34-38], neurons within the same cortical column are more likely to be synaptically connected than are neurons across columns. As a specific example, in the visual cortex of carnivores, cortical columns called orientation columns [39] represent different visual stimulus orientations. Recurrent within-orientation column preferential connectivity is thought to facilitate the orientation column's coordinated response to a given visual stimulus orientation [40,41]. In this way recurrent circuits impact the spatial structure of the visual response in the visual cortex of carnivores.

In rodents, although visual cortex lacks spatially segregated orientation columns, neurons in the visual cortex are more likely to be synaptically connected if they prefer visual stimuli of the same orientation [1,4], consistent with findings in carnivores [42]. Thus, recurrent circuits in rodent visual cortex also influence both the spatial structure and stimulus feature-tuning of the visual response.

Cortical models propose a computational role for these intra-cortical recurrent circuits [5,7,8]. If neurons activated by similar sensory stimuli are preferentially connected by excitatory recurrent circuits, consistent with the literature (see discussion above), a sub-network in cortex excited by sensory input will feed back onto itself, specifically amplifying its own activity without activating other, perhaps competing, sub-networks. In these models [5,7,8], recurrent circuits rather than feed-forward circuits determine which features of sensory input are amplified. If connections of intra-cortical recurrent circuits can be learned or altered by top-down feedback, these recurrent circuits might adapt to amplify only features of the sensory environment that are relevant to sensory perception in the context of a given behavioral goal.

1.2.3 Impact of recurrent circuits on temporal structure of cortical response

Thus, intra-cortical recurrent circuits reshape the spatial structure of neural activity arriving from thalamus. Whether and how these recurrent circuits also impact the temporal structure of sensory activity arriving from thalamus is a matter of debate. Models addressing this question have arrived at, essentially, two opposed predictions – either intra-cortical recurrent circuits slow and prolong the sensory response, or intra-cortical recurrent circuits rapidly follow thalamic sensory activity. Therefore the impact of intra-cortical recurrent circuits on the temporal structure of the sensory response must be answered empirically. Answering this empirically requires a measurement of the dynamics of neural activity evoked by thalamic input and reverberating within intra-cortical recurrent circuits.

Measuring these dynamics is difficult. Because the cortex receives temporally modulated signals from thalamus and feeds back onto thalamus, dissociating the contribution of the cortical circuits from the contribution of thalamus to the temporal structure of the cortical response requires a method to isolate cortical circuits from their thalamic input during active sensory processing. This has not previously been possible. However, some studies have employed an indirect approach to estimate the impact of intra-cortical recurrent circuits on the timing of the sensory response. These studies compared the sensory-evoked spiking response in thalamus with the sensory-evoked spiking response in cortex [9-18,43], and some studies concluded that the difference between these responses represents dynamics intrinsic to cortical circuits [16,44-47]. This indirect approach produced widely varying estimates of cortical circuit dynamics, which appeared to be dependent on sensory stimulation conditions [46,48], model organism [9,49], and the animal's brain state [49-51]. In some experiments, cortex appeared to dramatically alter (low-pass filter) the temporal structure of thalamic input [12-17,40,44,46]. In other experiments, the temporal structure of the cortical response seemed to be time-locked to, and thus inherited from, thalamus [43,52-55].

Sensory cortex filters thalamic input

Some researchers have argued that sensory-evoked dynamics of cortical recurrent circuits are slow relative to dynamics of the input to cortex from thalamus. The evidence in support of this idea is

the observation that sensory cortex fails to follow rapid fluctuations of sensory input from thalamus in some experiments [9-16,44,46]. The inability of cortex to follow fluctuations of thalamic input would suggest that the timescale required for a response in cortex is slower than the frequency of input from the thalamus. A slow timescale for the evoked response in cortex, relative to the timescale of a rapidly fluctuating response in thalamus, implies that cortex attenuates high-frequency activity (Figure 1.2). This attenuation of the amplitude of the response to high frequencies is low-pass filtering, a transformation that has been observed in visual cortices of various species. In rodents [43,49], cats [10-13,56] and primates [15] under anesthesia, the visually evoked response to rapidly fluctuating stimuli is less reliable in V1 than in the dorsal lateral geniculate nucleus (dLGN), the primary thalamic relay between retina and V1, in large part, because V1 low-pass filters sensory input [9-12,14,15,56-58]. The low-pass filtering in some of these cases is dramatic. For instance, in cat V1, most neurons fail to respond to fluctuations of visual input at frequencies above 6 Hz [10], while a substantial fraction of the neurons in cat dLGN still have large-amplitude responses to visual frequencies up to 10-15 Hz [56]. This suggests a timescale for the response dynamics in visual cortex that may be slower than the period of a 6 Hz oscillation (~166 ms). This transformation of the temporal structure of sensory activity by visual cortex has been attributed by some researchers to cortical recurrent circuits [16,17,44,46,47,59,60]. However, a causal link between recurrent circuits and the observed low-pass filtering has not been established.

Sensory cortex follows thalamic input

On the other hand, in other experiments, sensory cortex follows fast fluctuations of sensory stimuli [3,43,54,61-65]. This result implies fast dynamics of recurrent circuits able to lock the timing of cortical activity to the timing of high-frequency sensory input [63,66]. Robust responses to high-frequency sensory stimuli have been observed in various sensory areas, including somatosensory cortex [61] and visual cortex [51,54,62,64,65,67]. The results in visual cortex seem to directly contradict the view that visual cortex strongly attenuates sensory input above 6 Hz, since in some studies visual cortex responds reliably to sensory input above 6 Hz. What explains the discrepancy in the literature?

Several explanations have been proposed to account for the differing reports. First, it has been proposed that details of the sensory stimulation affect cortical recurrent circuit dynamics. For example,

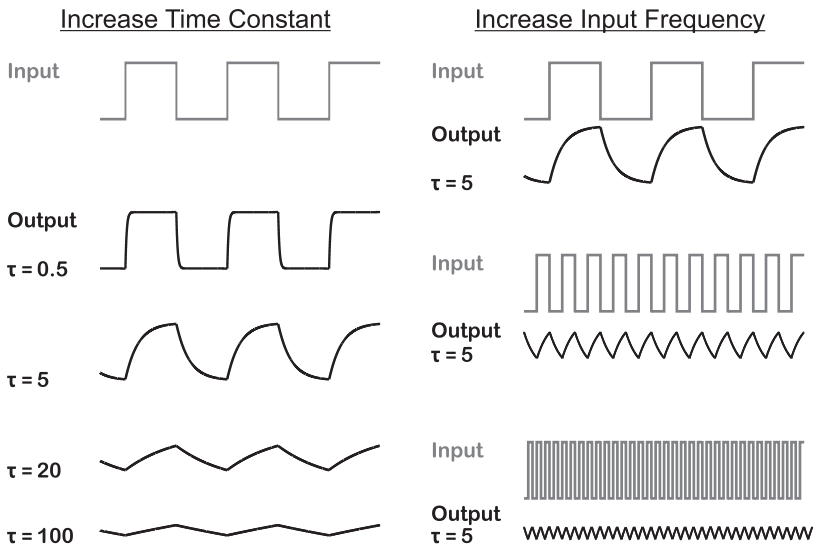

Figure 1.2: Predicted filtering as a function of the time constant and input frequency.

Left: A system's response rate in time is related to how it transforms the amplitude of an input signal. Slower rates of response (i.e., longer time constants) produce slower dynamics and greater attenuation of the amplitude of the response to the same input signal. If the time delay until the system reaches steady-state is long, then the system requires a long time to express its full-amplitude response to a change in input. If the input signal shuts off before the system reaches its maximal response, then the system discharges without achieving its maximal response.

Right: A system's response rate in time (i.e., time constant) is related to how it filters input as a function of temporal frequency. Given the same time constant, the system has a bigger response to lower input temporal frequencies, because slower input fluctuations provide sufficient time for the system to achieve its maximal response before the input shuts off. Thus a dynamic system with a time constant low-pass filters its input.

Nauhaus et al. [46] have reported slow, traveling waves of sensory-evoked activity in visual cortex in response to low- but not high-contrast visual stimuli. Second, it has been suggested that brain state might impact sensory-evoked cortical dynamics. For instance, in awake rabbits, alertness increases the response in visual cortex to visual stimuli of high temporal frequency [51]. Optogenetic activation of the basal forebrain desynchronizes cortical activity and improves the reliability of a cortical response to rapidly modulated sensory input [68]. Electrical stimulation of the mesencephalic reticular formation in anesthetized cats enhances multi-unit responses in V1 to high-frequency visual flicker [62]. Larger V1 responses to high stimulus temporal frequencies have been reported in awake [50], with respect to anesthetized [49], mice. These data suggest that brain state also affects cortical dynamics. However, recurrent circuits within cortex have not been established as the mechanism underlying these effects.

1.2.4 Long-lasting responses in sensory cortex

Dynamics describe the rate of change of activity in a system in response to a sudden change in input. If recurrent circuits in cortex have slow dynamics, sensory-evoked activity in these recurrent circuits should decay slowly after sensory input from the thalamus ends. Responses that persist beyond the end of sensory stimulation have been observed in sensory cortex [69,70]. In visual cortex, various durations of post-stimulus sensory activity have been reported, ranging from 30 ms of stimulus-specific local field potential (LFP) activity that persists after the offset of a drifting grating [70] to many seconds of post-stimulus spiking activity in primate V1 after a flash of light [69]. In humans, electroencephalographic (EEG) signals related to features of the visual stimulus are observed in visual cortex over delay periods of several seconds [71]. Some researchers have attributed such long-lasting sensory responses to cortical recurrent circuits [6,16,17,44,46,59,60,72]. However, experiments have not excluded several alternate sources of the long-lasting, post-stimulus neural activity in cortex.

First, peripheral sensory structures have not been excluded as a driver of the long-lasting sensory responses in cortex. In the eye, photoreceptors respond over hundreds of milliseconds. The recovery time constant of cone photoreceptors (in mice) is 70 ms [73], and the rod photoreceptor time constant is 200 ms [73]. Even transient retinal ganglion cell (RGC) responses last for more than 50 ms [74].

Some recordings directly comparing post-stimulus sensory-evoked activity in retina and cortex suggest that the duration of the post-stimulus response in the retina may fully account for the duration of the post-stimulus cortical response [69].

Second, the cortico-thalamo-cortical loop is a potential mechanism sustaining long-lasting sensory responses observed in cortex. This loop sustains persistent spontaneous oscillations in certain brain states [75,76], but whether it can sustain sensory-evoked activity is not known. If the cortico-thalamo-cortical loop does sustain sensory-evoked activity in the absence of peripheral sensory input, the net effect of corticothalamic feedback must be excitation of thalamus, because thalamus excites cortex. Indeed, V1 sends a massive excitatory feedback projection to dLGN. However, V1 also sends a disynaptic inhibitory projection to dLGN via the thalamic reticular nucleus (TRN). Thus, the net effect of corticothalamic feedback onto dLGN is not obvious from the anatomy. Recordings from anesthetized mice indicate that corticothalamic feedback subtly suppresses visually evoked responses in dLGN [30,31,77]. Yet the effect of corticothalamic feedback onto dLGN in awake mice is not known. In awake primates, corticothalamic feedback has been proposed to serve a wide range of functions, including altering receptive fields of thalamic relay neurons, enhancing the synchronous activity of thalamic neurons, and modulating gain or reliability of the thalamic response, but these effects tend to be subtle [78,79]. *In vitro* slice experiments identify specific conditions in which the net effect of corticothalamic feedback switches from inhibition to excitation [80,81], although these states are not well-established *in vivo*. While theories abound, it is not yet known whether corticothalamic feedback through dLGN in awake animals sustains long-lasting sensory-evoked activity in visual cortex.

It is also possible that corticothalamic feedback through higher-order thalamic nuclei sustains long-lasting sensory-evoked activity in cortex in the absence of peripheral sensory input. For example, in rodents, V1 projects to the second-order visual thalamic nucleus, the lateral posterior nucleus (LP) [82], and LP projects back to V1 [83,84]. Projections to LP from V1 originate mostly from layer 5 [82,85], whereas projections to dLGN from V1 originate mostly from upper layer 6 [86,87], suggesting that cortico-thalamo-cortical loops via second-order thalamic nuclei (e.g., V1-LP-V1) may have a different function than loops via thalamic nuclei providing primary input to a cortical area (e.g., V1-dLGN-V1). Whether any

type of cortico-thalamo-cortical loop can exhibit a positive feedback mode capable of sustaining sensory-evoked activity in cortex after sensory stimulation ends is not yet known.

Thus, while long-lasting sensory-evoked activity is clearly observed in cortex, whether intra-cortical recurrent circuits are themselves sufficient to maintain this activity remains to be seen.

1.2.5 Evidence for prolonged activity sustained by cortical recurrent circuits

***In vitro* evidence**

The major difficulty in assessing dynamics intrinsic to intra-cortical recurrent circuits has been separating dynamics of the thalamic input from the dynamics of the intra-cortical circuits themselves. Studies *in vitro* accomplish this by cutting thalamocortical axons or by artificially manipulating thalamic input to cortex. These studies *in vitro* cannot study a true sensory response in cortex, but they can provide insight into activity patterns intrinsic to cortex.

In vitro studies demonstrate that cortical recurrent circuits can generate and sustain their own active states. In certain conditions (i.e., low Ca^{2+}, low Mg^{2+} and high K^{+} artificial cerebrospinal fluid), spontaneous bouts of prolonged spiking activity are generated and sustained in acute slices of cortex without input from thalamic afferents [59,88]. These active states can also be triggered by brief electrical stimulation [59]. Because thalamic afferents are cut in the slice, these active states are sustained by the cortical recurrent circuits. These thalamus-independent active states have been observed both in associative [59,88] and sensory [88] neocortices *in vitro*.

***In vivo* evidence**

In vivo, in anesthetized animals, spontaneously occurring active states appear in associative [89] and sensory [88-91] neocortical areas in the absence of sensory stimulation. These spontaneously occurring active states, called “Up” states, recur at frequencies of roughly 0.1-3 Hz. Each Up state lasts for several hundreds of milliseconds. Because Up states persist in cortex even after the thalamus has been disconnected [92] or pharmacologically lesioned [91], Up states are thalamus-independent. Moreover, models of cortex propose mechanisms to explain how cortical recurrent circuits generate and sustain Up states [93]. Thus, spontaneously occurring active states in cortex are sustained by the recurrent circuits.

These Up states are observed only when the animal is anesthetized [76]. In awake animals, spontaneous fluctuations of cortical activity do occur at frequencies between 1 and 4 Hz [76,94,95]. However, whether these slow oscillations in the awake state are independent of thalamic input is not yet known [95,96].

1.2.6 Comparing spontaneously occurring and input-evoked active states in cortex

Like Up states, are sensory-evoked active states in cortex sustained by recurrent circuits? There are some well-described differences between Up states and active states evoked by thalamic input.

For example, electrical stimulation of thalamic afferents, which may be analogous to the feed-forward recruitment of thalamic afferents by a sensory response [90], recruits in the cortex both strong excitation and strong disynaptic inhibition mediated by interneurons[*]. As a result, an initial bout of excitation is followed, within a few milliseconds, by a strong suppressive bout of inhibition [19,20,53,98]. The thalamocortical input is strongest to cortical layer 4 [42], and, consistent with this, strong thalamocortical synaptic excitation as well as strong feed-forward disynaptic inhibition is observed in this layer [19,20,53,98]. Moreover, sensory-evoked activity initiates in layer 4 [18]. Therefore thalamic afferents recruit strong disynaptic inhibition in layer 4, which may delimit the duration of cortical activity triggered by thalamic input.

In contrast to sensory-evoked responses that initiate in layer 4, spontaneous Up states *in vivo* [99-101] and *in vitro* [88] initiate in deep layers of cortex. Consistent with this, *in vitro*, Up states are triggered by electrical stimulation of layer 5 [59]. Hence, Up states and sensory-evoked active states may have different profiles of activity across cortical layers.

Furthermore, Up states are synchronous over large areas of cortex (over distances of >1 mm [102]), whereas sensory input canonically recruits focal cortical microcolumns [103] (diameter: <500 microns

[*] Strong disynaptic inhibition is observed *in vitro* [19,20,53,97,98] and *in vivo* [29] following the electrical [19,53,98] or optogenetic [20,97] stimulation of thalamic afferents. However, stimulation of a random population of thalamic afferents could have a dramatically different effect on cortical activity than stimulation of thalamic afferents in accord with the statistics of real sensory stimulation. Therefore it will be necessary to measure the dynamics intrinsic to intra-cortical recurrent circuits in response to real sensory stimulation.

[104-107]). Activating multiple cortical columns at the same time may drive competition between them [108]. Thus, Up states and sensory-evoked active states may have distinct profiles of activation across cortical space.

Based on these data, it seems likely that recruiting cortical activity via thalamic input triggers a type of cortical activity qualitatively distinct from Up states.

However, some researchers have argued the opposite point. These researchers have argued that sensory-evoked and spontaneously occurring cortical active states share an underlying mechanism – maintenance via intra-cortical recurrent circuits [6,16,17,44,46,59,60,109-112]. In this view, thalamic input only triggers the cortical response. After the first few tens of ms or so, recurrent circuits alone are responsible for determining the pattern of activity in cortex.

In support of this idea, some studies have suggested that the pattern of activity in cortex evoked by sensory input resembles thalamus-independent Up states. Some of these studies describe similarity between sensory-evoked active states and Up states in terms of the statistics of activation of neuronal ensembles over time [17,109,112]. Other studies juxtapose findings showing that Up states propagate across cortex and, similarly, sensory-evoked activity in superficial layers of cortex propagates across cortex in traveling waves [46]. Based on these similarities, these researchers have proposed that sensory-evoked active states might share with Up states a mechanism for maintenance by intra-cortical recurrent circuits. However, these studies failed to distinguish the component of cortical activity dependent on thalamocortical interactions from the component of cortical activity dependent on intra-cortical recurrent circuits.

What is needed is a direct, empirical approach to separate the cortical activity dependent on thalamic input from the cortical activity sustained by intra-cortical circuits.

1.2.7 Dynamics of thalamocortical synapses

Previous attempts have been made to separate the thalamic contribution to the dynamics of a sensory response from the contribution of intra-cortical recurrent circuits to these dynamics. Studies attempting this compared the spiking response in thalamus with the spiking response in cortex.

However, the synapses of thalamic afferents connecting thalamus to cortex may exhibit their own dynamic properties, which could impact the sensory response of cortex [20,21,23,113,114]. If thalamocortical synapses do filter the spiking activity from thalamus, we cannot infer dynamics of cortical recurrent circuits by a direct comparison between thalamus and cortex. Dynamics of thalamocortical synapses will confound the result.

Extensive literature describes dynamics at thalamocortical synapses. The classic dynamic property of these synapses is depression, a physiologic characteristic of the synapse that may arise from different underlying cellular mechanisms [115]. Depression is a decrease in the strength of the synapse both over time and as a function of pre-synaptic spike frequency. The synapse's strength decreases over the course of repetitive stimulation eliciting spiking in the pre-synaptic neuron, but, once spiking in the pre-synaptic neuron stops, the synapse gradually recovers to its initial higher strength. Thus the strength of the synapse decreases more rapidly as the frequency of spiking (or stimulation) increases, because the synapse has less time to recover between spikes. Hence, higher frequencies of pre-synaptic activity produce stronger reductions in the strength of the synapse [116,117]. However, above a certain frequency of pre-synaptic activity, the strength of the synapse does not further decrease. At this steady-state level of depression, sensory-evoked neurotransmission via the synapse at all sensory input frequencies is strongly attenuated. In this manner, synaptic depression powerfully filters neural activity as a function of pre-synaptic spike frequency.

The presence or absence of depression at a synapse may depend on specific conditions. For example, it has been shown that some synapses exhibiting strong depression *in vitro* do not depress *in vivo* [22]. The explanation often proposed for this finding is that depression is not observed *in vivo*, because the cellular mechanisms underlying depression are saturated *in vivo*. In these cases, spontaneous activity of the pre-synaptic cell is often higher *in vivo* than *in vitro* [21-23,113,114]. It is thought that this spontaneous activity fully depresses, or "pre-depresses", the synapse, even before stimulation of the pre-synaptic neuron begins to measure depression. Because depression is quantified with respect to the size of the response to the first stimulus in a train of stimuli, no depression is observed *in vivo*, as the synapses does not further depress over the train of stimuli.

Consistent with this profile, thalamocortical synapses show strong depression *in vitro* in response to electrical or optogenetic stimulation of thalamic afferents [19,20,53,97,98]. *In vivo*, depression at these synapses may or may not be observed [21,23,113,114]. In support of the idea that depression at thalamocortical synapses is affected by the spontaneous activity of the pre-synaptic thalamic relay cell, thalamocortical synapses in the visual system are less depressed when the inter-spike interval (ISI) is long preceding the spike in a thalamic relay cell, occurring, for example, during the first spike in a burst of relay cell spikes preceded by a long ISI, presumably because the long ISI enables the thalamocortical synapses to recover from preceding activation [113]. However, thalamocortical synapses are continually depressed when thalamic relay cells are tonically active, presumably because there is no opportunity for the synapses to recover from activation. Interestingly, tonic and long-ISI/bursting modes of activity in thalamus correspond to alert and non-alert behavioral states, respectively, suggesting that thalamocortical depression is modulated by brain state [113].

Thus, as a result of thalamocortical synaptic depression, measuring spiking activity in the thalamus does not reveal the magnitude of the synaptic input to cortical neurons. In order to study how intra-cortical recurrent circuits transform their sensory input from thalamus, we need an approach to measure the actual strength of transmission via thalamocortical synapses separately from the dynamics of activity maintained within cortical recurrent circuits. In the chapters that follow, I describe the development of this approach and what it reveals about dynamics in visual cortex during sensory processing as a function of brain state, helping to illuminate how intra-cortical recurrent circuits impact sensory-evoked activity in time.

Chapter 1, in part, has been submitted for publication of the material as it may appear in Reinhold, K., Lien, A.D. and Scanziani, M., 2015, Distinct Recurrent Versus Afferent Dynamics in Cortical Visual Processing, Nature Neuroscience. The dissertation author was the primary investigator and author of this paper.

Chapter 2:

Intrinsic Dynamics of Cortical Recurrent Circuits

2.1 Introduction

Does the recruitment of intra-cortical recurrent circuits by a sensory stimulus impact the representation of that stimulus in time, and, if so, by what mechanism? Answering this question was not previously possible, because it was not possible to dissociate cortical recurrent circuits from their thalamic input during sensory processing. Without this dissociation, we cannot determine which dynamics of the response observed in cortex are dependent on the intra-cortical recurrent circuits themselves, rather than inherited from thalamocortical afferents.

Thus, studying sensory-evoked dynamics of intra-cortical recurrent circuits will require a method to dissociate cortical recurrent circuits from their afferent input while sensory processing is ongoing. To factor out any contribution of the thalamic input to the rate at which sensory-evoked cortical activity changes in time, we need to shut off thalamic input and measure the resulting dynamics of the sensory response in cortical recurrent circuits in the absence of thalamic input. These cortical dynamics will then describe how intra-cortical recurrent circuits respond to changes in thalamic input and therefore how these recurrent circuits filter the temporal structure of sensory input from thalamus. Methods that lesion or pharmacologically inactivate thalamus are not appropriate, because they prevent the transmission of sensory activity via thalamus to cortex and therefore prevent the initiation of a sensory response in cortex. What we need is a way to rapidly inactivate thalamic input to cortex only after the sensory response in cortex has been initiated. Moreover, this manipulation must be reversible to ensure that the manipulation itself does not damage cortical recurrent circuits.

Toward these two goals, I have developed an optogenetic approach to rapidly (i.e., with millisecond-precision) and reversibly silence visual thalamus, *in vivo*, without directly affecting the sensory-evoked activity reverberating in cortical recurrent circuits. With this approach, I measure the timescale of sensory-evoked activity triggered in cortex and maintained by the recurrent circuits in the absence of thalamic input.

2.2 Methods: Approach to optogenetically silence thalamus

To silence the thalamus, I exploited an inhibitory projection* that powerfully suppresses the spiking activity of thalamocortical relay cells [118]. This is the projection from the thalamic reticular nucleus (TRN) onto sensory thalamic relay nuclei (Figure 1.1c). (The TRN does not itself project to cortex.) All neurons within the TRN are GABAergic and express Glutamate Decarboxylase 2 (Gad2) [119]. Therefore these TRN neurons express Cre recombinase under control of the Gad2 promoter in the Gad2-Cre transgenic mouse line [119]. Stereotactic injections of Cre-dependent adeno-associated virus (AAV) carrying the construct for Channelrhodopsin 2 (ChR2) specifically into the TRN of this mouse line† targeted expression of ChR2 to TRN. Therefore, driving GABAergic neurons in the TRN by photo-activation of ChR2 should drive powerful inhibition onto thalamocortical relay cells‡.

An additional advantage of this approach is that TRN projects to and inhibits both first-order and higher-order sensory thalamic relay nuclei. For example, in the rodent visual system, TRN projects to and inhibits both the dorsal lateral geniculate nucleus (dLGN), the first-order relay between retina and cortex, and the lateral posterior nucleus (LP), which relays activity from retina and superior colliculus to cortex. Thus, photo-activating TRN should suppress all thalamic sources of visual input to

* I chose this method over the method of expressing inhibitory opsins in thalamocortical relay cells to directly suppress their activity for two reasons. First, the suppression of spiking by the currently available inhibitory opsins is often incomplete. Second, the inhibitory projection from TRN is divergent onto thalamocortical relay cells; thus, expression of ChR2 in a subset of TRN cells should accomplish inhibition onto most or all thalamocortical relay cells, whereas inhibition of thalamocortical relay cells by expression of an inhibitory opsin is limited to the fraction of relay cells that express the opsin. It was critical for my study that suppression of thalamic input to cortex be nearly complete.

† Consistent with another study in the literature [120], I found that expressing ChR2 in the TRN under the control of the parvalbumin (PV) promoter produced suppression followed by a rebound in thalamocortical relay cells. However, it is not appropriate to use the PV-Cre transgenic mouse line to restrict the expression of ChR2 to GABAergic neurons in the thalamus, because excitatory thalamocortical relay cells themselves express PV [121]. The TRN is a thin structure. It would be difficult to achieve expression of ChR2 throughout the TRN without infecting neighboring thalamocortical relay cells. Hence, it is probable that, in past work using PV-Cre mice [120], thalamocortical relay cells as well as the TRN neurons expressed ChR2. Here I report only results from experiments in Gad2-Cre transgenic mice, which do *not* express Cre in thalamocortical relay cells [119].

‡ Spread of the virus (AAV Flexed-ChR2) led to expression of ChR2 in both the TRN and neighboring local GABAergic interneurons of visual relay thalamus in a subset of mice. Photo-activating these local GABergic interneurons of visual relay thalamus should further suppress the activity of excitatory thalamocortical relay cells. Therefore I did not take measures to avoid ChR2 expression in local GABAergic interneurons of thalamus.

primary visual cortex (V1), effectively isolating the intra-cortical recurrent circuits from their thalamic input[*].

[*] Moreover, visual stimulation was restricted to the monocular visual field providing input to (contralateral to) the hemisphere of the thalamus suppressed by TRN photo-activation. Consequently, visual stimulus information was not available to the visual cortex contralateral to this hemisphere, i.e., contralateral to the site of the optogenetic manipulation (thalamus) and the site of electrophysiology (visual cortex). Thus, unilateral, rather than bilateral, photo-activation of TRN should be sufficient to suppress all sources of feed-forward visual input to V1.

2.3 Results

2.3.1 Time course of cortical recurrent excitation versus thalamic afferent excitation in visual cortex

First I needed to establish the time course over which intra-cortical recurrent circuits contribute to the sensory response. To do this, in collaboration with Dr. Tony Lien (see Author Contributions, pg. 141), I considered data from recordings of visually evoked synaptic excitation in V1 collected by Dr. Lien during his Ph.D. work. Dr. Lien performed whole-cell voltage-clamp recordings from layer 4 neurons of primary visual cortex (V1) in response to the appearance of a static visual stimulus (e.g., oriented grating) in anesthetized mice. To quantify the contribution of visually evoked excitation mediated by cortical recurrent circuits (i.e., the cortical fraction), he compared the visually evoked excitatory post-synaptic current (EPSC) observed under control conditions with the EPSC observed on interleaved trials while silencing recurrent excitatory circuits (Figure 2.1a). Dr. Lien silenced recurrent excitatory circuits by optogenetically activating cortical inhibitory interneurons, as described in his previous work [30]. Hence, during this manipulation, remaining excitation to cortical neurons was direct, feed-forward input from thalamus [30].

We found that silencing recurrent circuits abolished the vast majority (79±9%, mean±s.e., as a fraction of the total integrated response; n=8 cells; 1.7 s-long visual stimuli; Figure 2.1a) of visually evoked synaptic excitation to cortical neurons in layer 4, the main thalamic afferent recipient layer [24], consistent with previous work in rodents [30,31] and carnivores [29]. The total synaptic excitatory current, including both cortical recurrent and direct thalamic components, peaked 150 ms after the onset of visually evoked excitation and decayed slowly over the next 1.35 s. The direct thalamic excitation to layer 4 was dominant over the first several tens of milliseconds (ms) of the response (time until thalamic fraction was less than 50% of total instantaneous EPSC amplitude: 43±6.5 ms, mean±s.e., n=25 cells; over the first 10 ms, 82±4% of excitation was of direct thalamic origin, as a fraction of the total integrated response; Figure 2.1a). The fraction of excitation dependent on cortical recurrent circuits (cortical fraction) grew progressively from 0% to 72±6% (as a fraction of the instantaneous

Figure 2.1: Time course of cortical recurrent excitation versus thalamic afferent excitation in visual cortex.

(a) Comparison of cortical recurrent versus thalamic afferent excitation to V1 layer 4.
Left: Experimental configuration: whole-cell voltage-clamp of layer 4 (L4) neurons to record visually evoked excitatory post-synaptic currents (EPSCs) with or without optogenetic silencing of cortical excitatory recurrent circuits.
Center: Time course of synaptic excitation. Vis. stim. is static for 1.7 s. See detailed list of visual stimuli in all figures in Appendix B. Top: Example EPSC in response to the appearance of a static grating (arrow), comparing control (black) versus cortical silencing with LED to isolate the thalamic EPSC (blue). Larger downward deflection is larger current. Middle: Mean EPSC and s.e. of 8 similar experiments (8 cells). Bottom: Fraction of EPSC mediated by cortical excitation (cortical fraction, i.e., difference between black and blue, divided by black).
Right Column: Same as middle but on expanded timescale showing immediate dynamics following stimulus onset. Average also includes experiments using shorter stimulus presentations (n=17 250 ms-long vis. stim., plus 8 at left). Note the slow time course of the EPSC and progressive build-up of the cortical fraction.

(b) Spiking activity in V1 cortical circuits.
Left: Expt. configuration: extracellular multi-channel recordings spanning all V1 layers.
Center: Multi-unit (MU) response to appearance of static grating. Vis. stim. is static for 3 s. Top: Mean and s.e. over 1.5 s window. Amplitude in spikes per second (Hz). Bottom: MU mean response subdivided according to cortical layer. The time course of the spiking response is comparable to the time course of the EPSC in (b). (See Figure 2.8 for dLGN response.)
Right: Single-unit (SU) response to appearance of static grating. Vis. stim. as left. Top: Mean and s.e. over 1.5 s window. Bottom: Example SUs (left) and scatter plot (right) of response duration of SU subdivided according to cortical layer and visual stimulus type. Each point is a unit (n=207). Fast-spiking units are outlined (n=5). Each color is a different visual stimulus. Cyan: 3 s-long static grating (mean±s.d. of unit response duration: 0.96±0.28 s). Dark blue: 10-100 ms luminance step (0.38±0.25 s). Red: 3 s-long luminance step (0.46±0.24 s). Purple: checkerboard reversal every 3 s (0.62±0.18 s). Note that all static stimuli elicit transient visually evoked responses lasting for several hundred ms.

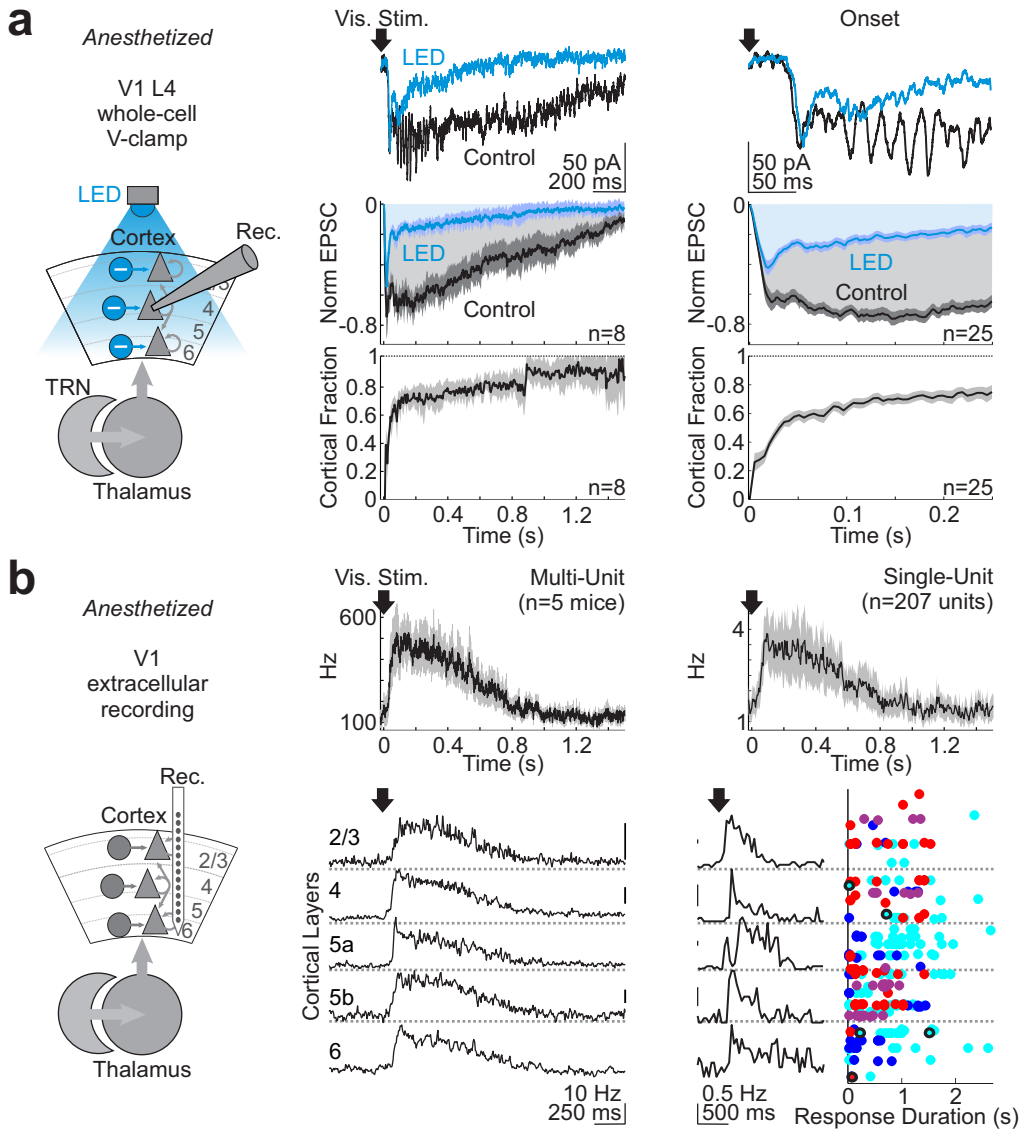
a
Anesthetized
V1 L4
whole-cell
V-clamp
LED
Cortex
Rec.
TRN
Thalamus
Vis. Stim.
LED
Control
50 pA
200 ms
Norm EPSC
Cortical Fraction
n=8
Time (s)
Onset
50 pA
50 ms
n=25
b
Anesthetized
V1
extracellular
recording
Rec.
Cortex
Thalamus
Multi-Unit
(n=5 mice)
Hz
Single-Unit
(n=207 units)
Cortical Layers
2/3
4
5a
5b
6
10 Hz
250 ms
0.5 Hz
500 ms
Response Duration (s)

EPSC amplitude, n=25, including 1.7 s- and 250 ms-long stimuli) over the first 250 ms of the response, overwhelming the direct thalamic component. Thus, by 250 ms, cortical circuits amplified the thalamic input more than three-fold. Therefore in response to visual stimulation, the contribution of cortical recurrent excitation progressively builds up to overwhelm the contribution of direct thalamic excitation to visual cortex.

If the cortical neurons in recurrent circuits are the source of the cortical component of the excitatory synaptic current, the spiking of neurons in visual cortex should follow a time course similar to that of cortical synaptic excitation. To test this, I recorded multi-unit and single-unit activity across all cortical layers detected with extracellular linear probes (Appendix A: Supplementary Methods, "Electrophysiology", pg. 113) during responses to the same visual stimuli in anesthetized mice. Over the first 800 ms of the visual response, multi-unit and single-unit activity in V1 had a time course comparable to that of recurrent synaptic excitation (Figure 2.1b). Thus, powerful recurrent excitation from other spiking cortical neurons is larger in magnitude than the direct thalamic input after the first several tens of ms of the evoked response in V1.

Once recurrent excitation builds up in V1, over what timescale do cortical recurrent circuits alone sustain the sensory response in the absence of continued thalamic input?

2.3.2 Silencing thalamic input to visual cortex

To answer this, I needed to precisely gate thalamic input to cortex. I developed an optogenetic approach to rapidly silence thalamic input at any arbitrary time following the onset of a sensory response in visual cortex. Critically, unlike electrically [122] or optogenetically [123] stimulating thalamic afferents, which recruits a random population of afferents, this method allowed me to measure the duration of activity reverberating through cortical recurrent circuits in response to actual visual stimuli.

To gate thalamic input, I exploited the inhibitory projection from the thalamic reticular nucleus (TRN, which does not itself project to cortex) onto relay nuclei of thalamus (e.g., dLGN), which transmit sensory activity to cortex. I conditionally expressed Channelrhodopsin 2 [124], a light-activated cation channel (ChR2, conditional on Cre), in the TRN [125] using Gad2-Cre mice [119] (expressing Cre in

GABAergic cells; Appendix A: Supplementary Methods, "Optogenetic manipulations", pg. 116; Figures 2.2 and 2.3a-b). Photo-activating the TRN (step pulse, 470 nm) rapidly and powerfully suppressed visually evoked (>96%) and spontaneous (>88%) activity in dLGN of anesthetized mice (Figures 2.3c, 2.4 and 2.5; 3.7 ms time constant of dLGN silencing; for duration of silencing, see Appendix A: Supplementary Methods, "Optogenetic manipulations", pg. 116). TRN photo-activation also suppressed activity in the lateral posterior (LP) nucleus, the second-order visual thalamic nucleus in rodents (Figure 2.6).

2.3.3 Time course of shut-off of visually evoked cortical activity after silencing thalamus

I measured multi-unit activity in visual cortex of anesthetized mice after rapidly silencing thalamus. Silencing thalamus in the absence of visual stimulation had no effect on cortical spontaneous active states, called "Up" states (Figures 2.3d-e and 2.7), bouts of high activity occurring spontaneously in the anesthetized cortex[*], indicating that cortical recurrent circuits can generate and sustain spiking activity for several hundreds of milliseconds in the absence of thalamic input, consistent with previous work [88,91,92]. I triggered sensory-evoked active states in cortex using the static visual stimuli described above (as in Figure 2.1). These sensory-evoked active states decayed slowly after removal of the visual stimulus (i.e., continued for several hundreds of milliseconds, Figures 2.1 and 2.8). If cortical recurrent circuits sustain this slowly decaying activity, then silencing thalamus should have little effect on the time course of this decay. Strikingly, silencing thalamus 250 ms into the sensory-evoked response led to a decay of visually evoked activity in V1 that was more than two orders of magnitude faster than the decay observed following simply the removal of the visual stimulus (Figure 2.3f, c.f. Figures 2.1 and

[*] In most experiments, mice were anesthetized with isoflurane. While this anesthetic does not produce Up states that are as reliably periodic and stereotyped as are the Up states observed under ketamine or urethane anesthesia, I found that isoflurane, like ketamine and urethane [89], produced a bimodal distribution of V1 activity (activity as either multi-unit firing rate or local field potential ratio [68]). This bimodal distribution of cortical activity is one indication of Up states. Moreover, this bimodal distribution of spontaneous activity was unaffected by thalamic silencing, consistent with the thalamus-independent nature of Up states in cortex. Hence, I refer to spontaneous activity in cortex under isoflurane anesthesia as Up states. To be sure that the results reported here are not specific to isoflurane anesthesia, I replicated the results shown in Figure 2.3e-f in a subset of mice anesthetized with urethane rather than isoflurane.

2.8). I call the time course of this decay of cortical activity after silencing thalamus the cortical decay function (CDF). The CDF was fit by an exponential time constant of 9±3 ms (mean±std. dev., n=26 mice, Figure 2.3f; 3 ms delay before start of exponential to account for time it takes thalamus to shut off, see Figure 2.9 and Appendix A: Supplementary Methods, "Accounting for shut-off delay in LGN", pg. 127). Thus, even after the build-up of cortical recurrent excitation in response to sensory stimulation, cortical recurrent circuits sustain sensory-evoked activity for only a few tens of milliseconds without thalamic input[*].

I verified that this fast, ~10 ms CDF was not influenced by off-target expression of ChR2 in visual cortex interneurons and that ChR2 expression in TRN axons to dLGN was the main histologic correlate of the shut-off of visually evoked activity in V1 (Figure 2.12 and Appendix A: Supplementary Methods, "Excluding animals with ChR2 expression in V1", pg. 114). The CDF was fast across all V1 layers, even in the layers that do not receive the major thalamocortical input, although I did observe subtle but significant differences across layers (mean fit to CDF ± s.e. subdivided by layer – L23: 9.8±1.7 ms, L4: 9.0±2.2 ms, L5a: 8.9±1.3 ms, L5b: 15.7±2.5 ms, L6: 7.6±1.5 ms, p-values in Table 2.1; Figure 2.13). Furthermore, although it has been suggested that slow cortical sensory-evoked dynamics emerge at low contrast [46], I found that the CDF was independent of stimulus contrast (Figure 2.14a-b). The CDF was also the same at later time points in the visual response (300, 600 and 900 ms after stimulus onset, Figure 2.14a-b), demonstrating that cortical sensory-evoked activity locks to thalamic input over the full time course of the sensory response. Moreover, the CDF was independent of strength of the thalamic silencing (Figure 2.14b-c), indicating that cortex exhibits the same fast sensory-evoked dynamics even when partial thalamic drive remains.

Consistent with the multi-unit data, the average CDF of visually responsive, isolated single units in V1 was fit by a 12±1 ms time constant (mean±s.e., n=297 units from all cortical layers; Figure 2.15). Furthermore, the vast majority (>90%) of single units considered individually had fast CDFs (time constant <20 ms, Figures 2.3g-h and 2.15), and only a small minority in layer 5 were a few tens of

[*] Local field potential (LFP) activity in V1 was also abolished by silencing thalamic input to the cortex (Figures 2.10 and 2.11).

ms slower. Are the neurons that exhibit spontaneous Up states part of a different sub-network than neurons responsive to visual stimulation? No, many of the visually responsive units also participated in Up states, indicating that thalamus-locked sensory activity and thalamus-independent Up states engage many of the same neurons in the cortical circuits (Figure 2.3i), despite dramatic differences in the dynamics of these two types of cortical activity and their dependence on thalamic input[*] (Figures 2.16 and 2.17).

Taken together, these results show that, even at time points when cortical recurrent circuits generate the majority of visually evoked synaptic excitation (i.e., more than 43 ms after the onset of the visually evoked response, Figure 2.1a), activity in cortical recurrent circuits time-locks to thalamic input. Hence, cortical recurrent circuits have fast dynamics.

Previous work suggests that dynamics in recurrent circuits might be very different between the anesthetized and awake brain states [62,68,76,91] and between primary and higher-order cortical areas [126]. I found that the CDF was similar in awake mice (Figure 2.18a-b; awake time constant fit to CDF: 10±3 ms, mean±s.d., n=9 mice, anesthesia: 9±3 ms, n=26 mice, unpaired p=0.15 – all CDF comparisons in text are two-tailed t-tests, see Figure 2.19 for CDF distribution), even when mice were running [127], a behavior that increases the gain of V1's response [57] (running: 9±6 ms, stationary: 10±6 ms, n=4 mice, paired p=0.16; Figures 2.19 and 2.20). Furthermore, in contrast to Up states under anesthesia, spontaneous cortical activity in awake mice decreased after silencing thalamus (73±7% decrease in single-unit baseline as mean±s.e., Figures 2.18c and 2.21) at a rapid rate consistent with the CDF (Figure 2.18c; awake spontaneous CDF: 14±4 ms, n=7 mice), indicating that spontaneous cortical

[*] I found that the mean visually evoked firing rate of neurons in visual cortex returns within a few tens of ms to the pre-stimulus baseline after silencing thalamus. However, the mean firing rate is a gross summary of cortical activity. If, during thalamic silencing in the middle of the visual response, some neurons in visual cortex maintained persistent higher levels of activity while other neurons maintained persistent lower levels of activity, the mean firing rate across the population might return to the pre-stimulus baseline, yet relative firing rates within the population might retain information about the visual stimulus, even in the absence of thalamic input. To test this, I considered whether each single unit returns to its own pre-stimulus spontaneous firing rate after silencing the thalamus. The answer was yes (Figure 2.17b, right). I found a strong correlation between each unit's pre-stimulus spontaneous firing rate and that unit's spontaneous firing rate when thalamic silencing was preceded by the presentation of a visual stimulus. Thus, the presentation of a visual stimulus does not systematically impact the subsequent pattern of spontaneous activity in visual cortex under anesthesia. Hence there is not information retained about the visual stimulus by the thalamus-independent spontaneous activity.

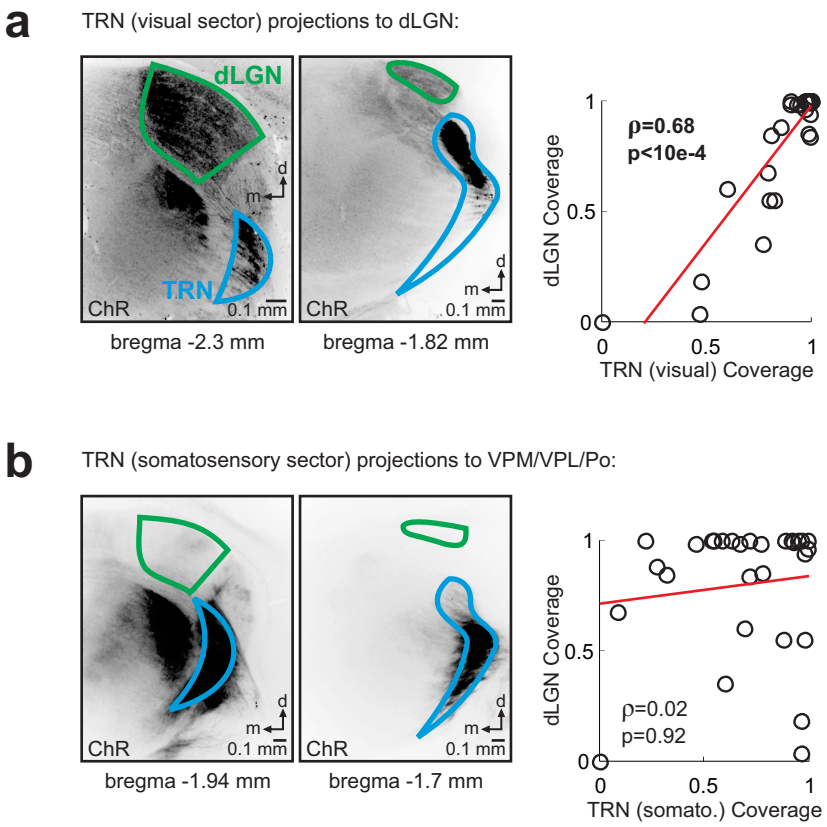

Figure 2.2: Identification of visual sector of TRN that projects to dLGN.

Expression of fluorescent reporter TdTomato fused to Flexed-ChR2 in Gad2-Cre mice in two sets of coronal sections from two mice (one mouse in (a) and the other in (b)). Coverage is the fraction of pixels expressing TdTomato within region-of-interest (ROI) outlining dLGN (green outline) or TRN (blue outline). Scatter plots compare reporter coverage of TRN region with reporter coverage of dLGN. For brain structures with >1 mm anterior-posterior extent, coverage is average of coverage across three coronal sections spanning extent of structure. Coverage of TRN's caudo-dorsal sector (putative visual sector) but not TRN's somatosensory sector correlates with coverage of dLGN. ρ is correlation coefficient; p is p-value of correlation.

(a) Expression of TdTomato-ChR2 in visual sector of TRN (caudo-dorsal sector) and in TRN axons projecting to dLGN.

(b) Expression of TdTomato-ChR2 in somatosensory sector of TRN (ventral sector) and in TRN axons projecting to ventral posterior medial and lateral (VPM/VPL) somatosensory thalamic relay nuclei. Note lack of correlation between coverage of TRN's somatosensory sector and coverage of dLGN.

Figure 2.3: Time course of shut-off of visually evoked activity during thalamic silencing.

(a) Experimental configuration. Silencing thalamus by optogenetically activating TRN while recording in dLGN with an extracellular multi-channel probe. Refers to (c).

(b) Expression of TdTomato fused to ChR2 in coronal sections through TRN (outlined in blue) at -1.58 mm A-P, dLGN (green) at -2.46 mm A-P, and V1 at -3.08 mm A-P. Insets: close-up of areas within outlined structures. Inset scale bars: 5 μm.

(c) Activity in dLGN is suppressed by TRN photo-activation. **Left:** Single-unit (SU) firing rates (pre-stimulus baseline-subtracted, n=157) during visual stimulation in control (LED Off) or plus TRN photo-activation (LED On). Vis. stim. is 3 s long (Supplementary Fig. 2a legend for other stimuli). Broken line is unity; points below line are suppressed. Inset: effect on spontaneous activity (not baseline-subtracted; X and Y axes: 0 to 7 Hz; n=38). **Center:** Mean amplitude of dLGN response (summed SU, normalized, see "PSTH normalization", pg. 126) to visual stimulus (arrow) followed by TRN photo-activation (blue) or in control (black, without TRN photo-activation). Blue bar: LED illumination. Error bars are s.d. across mice. **Right:** Expanded timescale of suppression of dLGN activity at LED onset. Green: single exponential fit (τ is time constant of fit).

(d) Experimental configuration. Silencing thalamus on interleaved trials while recording across V1 layers with an extracellular multi-channel probe. Refers to (e-h).

(e) Spontaneous Up states in V1. **Left:** Example raw data from single trials as local field potential (LFP), multi-unit (MU), and summed isolated SU with thalamic silencing (blue shaded area) or without (unshaded area) thalamic silencing. **Center:** Mean MU activity in V1 during Up states in control (black, without thalamic silencing) or during Up states beginning within 300 ms of thalamic silencing (blue). **Right:** As Center but on expanded timescale at LED onset. MU was normalized at 1, not baseline-subtracted. Green: fit to dLGN shut-off for comparison (from (c)).

(f) Sensory-evoked activity in V1. Data are presented as in (e) but here in response to appearance of 3 s-long static vis. stimulus (arrow). Onset of LED to silence thalamus 250 ms after stimulus onset. Right: Dark blue line is single exponential fit to time course of shut-off of visually evoked activity (Cortical Decay Function, in light blue) upon LED onset. Green: fit to dLGN shut-off for comparison (from (c)). Note that visually evoked active states shut off rapidly upon LED onset.

(g) Example response of single unit to appearance of static visual stimulus (arrow). **Top:** raster of spikes. **Bottom:** peri-stimulus time histogram (PSTH). Blue shaded area: thalamic silencing.

(h) Time constants of single exponential fits to shut-offs of SU in V1 subdivided by cortical layers. Each point is a time constant fit to one unit's trial-averaged time course at LED onset (n=297). Filled circles are fast-spiking (n=9). No significant difference between fast- and regular-spiking CDFs (p=0.16, Wilcoxon rank-sum). Analogous to scatter plot in Fig. 1c but here at moment of thalamic silencing instead of at stimulus offset. Inset: Example PSTH of SU. Blue shaded area: thalamic silencing. Blue: fit to shut-off. Gray: unit mean baseline.

(i) Spontaneous versus visually evoked firing rates of V1 SU (n=208). Note lack of correlation.

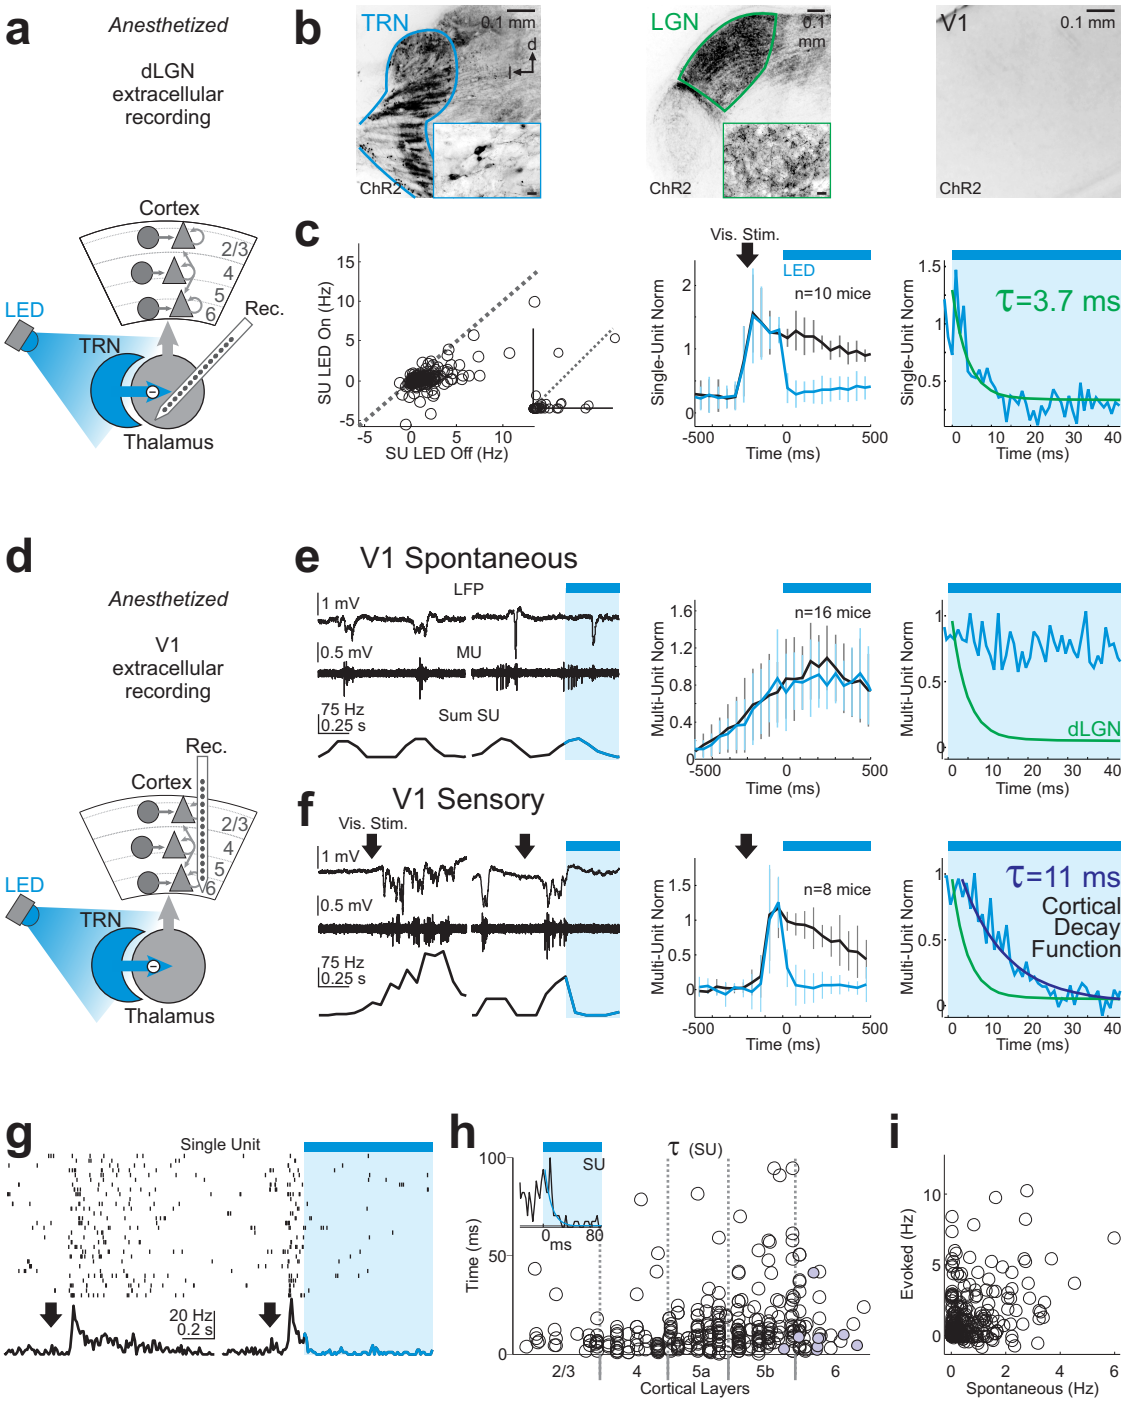
a
Anesthetized
dLGN extracellular recording
Cortex
2/3
4
5
6
Rec.
LED
TRN
Thalamus
b
TRN
0.1 mm
ChR2
LGN
0.1 mm
ChR2
V1
0.1 mm
ChR2
c
SU LED On (Hz)
SU LED Off (Hz)
Vis. Stim.
LED
n=10 mice
Single-Unit Norm
Time (ms)
τ=3.7 ms
d
Anesthetized
V1 extracellular recording
Rec.
Cortex
2/3
4
5
6
LED
TRN
Thalamus
e
V1 Spontaneous
LFP
1 mV
0.5 mV
MU
75 Hz
0.25 s
Sum SU
n=16 mice
Multi-Unit Norm
Time (ms)
dLGN
f
V1 Sensory
Vis. Stim.
1 mV
0.5 mV
75 Hz
0.25 s
n=8 mice
τ=11 ms
Cortical Decay Function
g
Single Unit
20 Hz
0.2 s
h
SU
τ (SU)
Time (ms)
2/3
4
5a
5b
6
Cortical Layers
i
Evoked (Hz)
Spontaneous (Hz)

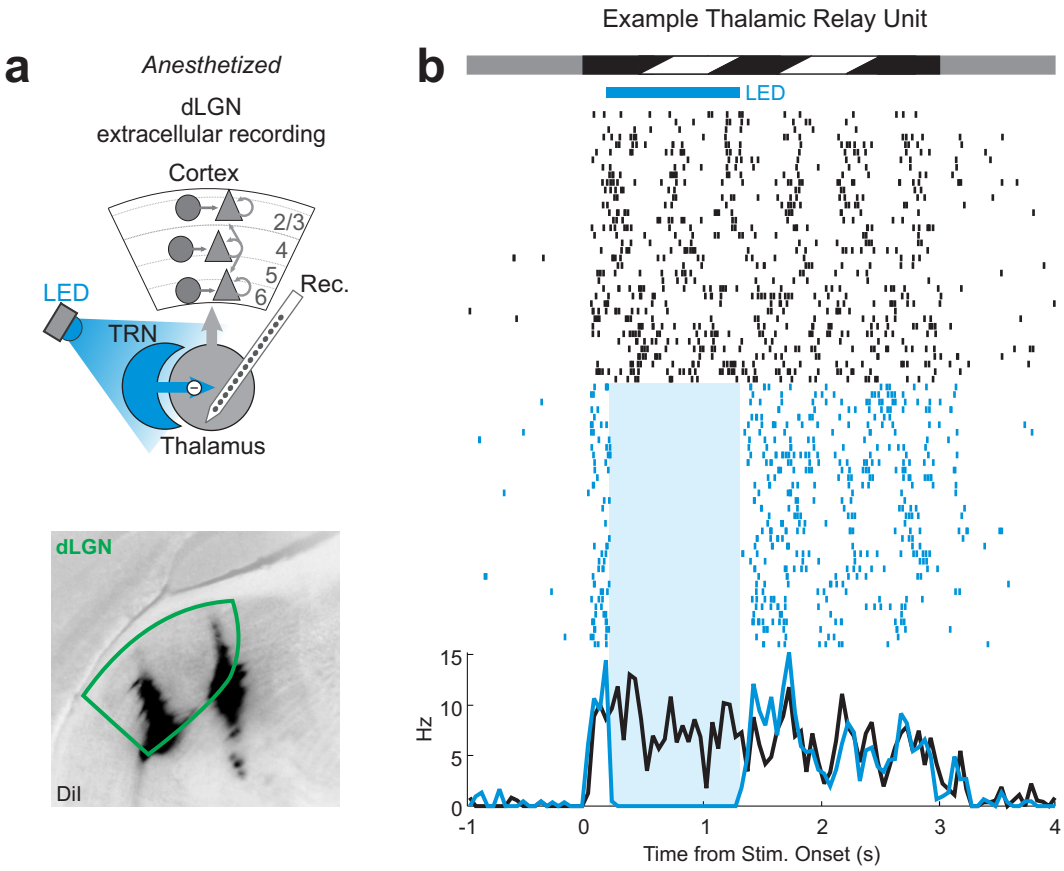

Figure 2.4: TRN photo-activation silences spiking in example thalamic relay cell in dLGN under anesthesia.

(a) Top: Schematic of thalamic silencing by photo-activation of TRN and recording in dLGN. **Bottom:** Recording tracks through dLGN marked by DiI at time of recording and examined by post-mortem histology. Black is dye (DiI).

(b) Activity in example thalamic relay cell. Blue bar and shading indicates LED illumination of TRN. Black is control (no illumination of TRN). Black and white angled bars indicate moving grating stimulus (3 s duration). In raster plot, each line is a spike. Peri-stimulus time histogram (PSTH) at bottom is average across trials.

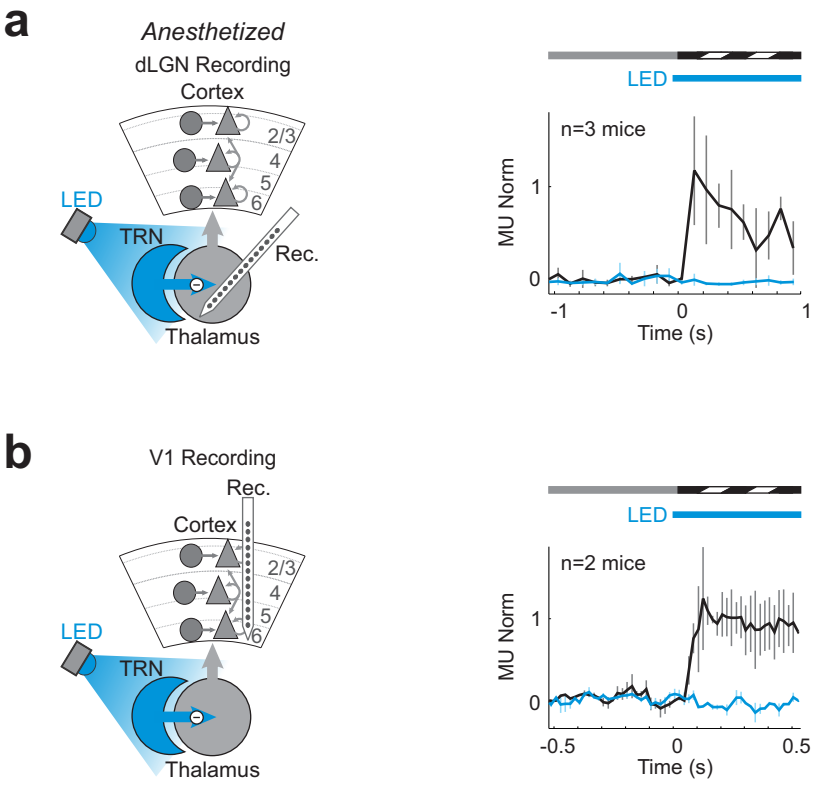

Figure 2.5: TRN photo-activation prevents the onset of a visually evoked response in dLGN and V1.

(a) Left: Schematic of thalamic silencing by photo-activation of TRN and recording in dLGN.
Right: TRN photo-activation prior to onset of a visual stimulus (moving grating) prevents the onset of a visual response in dLGN. Shown here: mean and std. dev. of multi-unit (MU) PSTH across mice (normalized to evoked response amplitude, see "PSTH normalization", pg. 126).

(**b) Left:** As in (a), but here recording in V1.
Right: As in (a), but here V1 MU data.

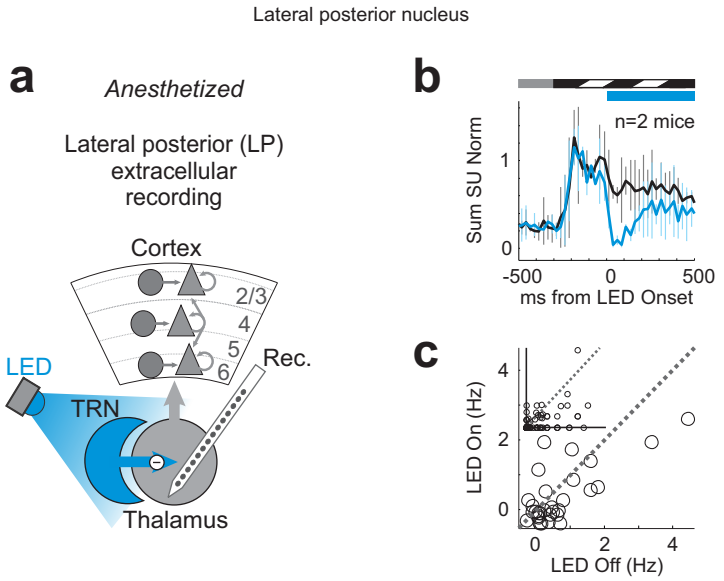

Figure 2.6: TRN photo-activation suppresses activity in the lateral posterior nucleus (LP) of anesthetized mice.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in LP.

(b) TRN photo-activation suppresses ongoing visually evoked response in LP measured as normalized ("PSTH normalization", pg. 126) single units (SU) summed within each mouse. Control: black trace. With TRN photo-activation: blue trace. Error bars are std. dev. across mice. Blue bar indicates LED illumination of TRN. Black and white angled bars indicate moving grating stimulus (3 s duration).

(c) Activity of relay cells in LP. Firing rate of SU (n=30, baseline-subtracted) during visual stimulation (LED Off) versus visual stimulation with TRN photo-activation (LED On). Broken line is unity. Dots below line are suppressed. Inset: effect on SU spontaneous activity during blank screen (X and Y axes: 0 to 12 Hz, n=87 SU).

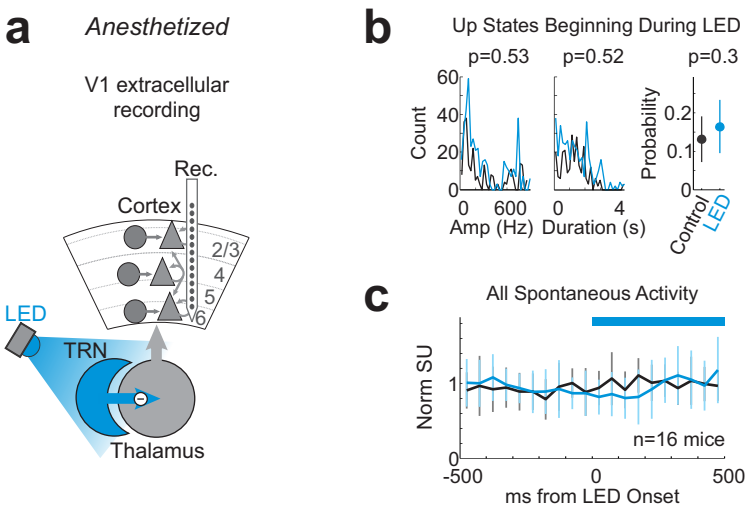

Figure 2.7: Effect of thalamic silencing on the initiation of spontaneously occurring Up states in V1 under anesthesia.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1 under anesthesia.

(b) Left and Center: Histograms of Up state multi-unit amplitude (left) and duration (center) when Up states begin during period of thalamic silencing. Up states detected in n=14 mice; all p-values are from Wilcoxon rank-sum tests.
Right: Probability of Up state initiation over 1 s window during thalamic silencing (blue) or control (black).

(c) Single-unit (SU) spontaneous activity from Up and Down states is summed within each recording session and normalized to 1. SU activity verifies that all spikes are from V1 neurons. Lines are mean, and error bars are std. dev. of this spontaneous activity across mice. Blue bar shows LED illumination of TRN. Spontaneous activity in V1 is not affected by silencing thalamus.

Table 2.1: Network time constant (NTC) fit to CDF across conditions.

This table lists values obtained for time constants of single exponential fits to the CDF as well as a non-parametric measurement of the CDF, the time delay until the CDF reaches half of its max, across experimental conditions. Columns 1 and 2 show the experimental condition being tested. Column 3 lists the network time constants (NTC) fit to the CDF. Column 4 gives the time delay until the CDF reaches its half-max. Column 5 gives the number of mice in which independent observations were made. Below each condition is a statistical test of whether the NTC and/or time-to-half-max changes as a function of that experimental condition.

Condition		Mean NTC (ms) and s.e.	Mean ms to half-max	# mice
Anesthetized, Across Layers	L2/3	9.8±1.7	4.8±1.1	26
	L4	9.0±2.2	5.4±0.9	26
	L5a	8.9±1.3	6.8±0.9	26
	L5b	15.7±2.5	9.4±0.9	26
	L6	7.6±1.5	6.8±0.9	26
		1D bal. ANOVA p=0.01	1D bal. ANOVA p=0.01	
		L5a faster than L6	L2/3 and L4 faster than L5b	
Anesthetized, LED Onset Delay	200 ms	11.8±2.6		6
	300 ms	9.0±0.6		26
	600 ms	13.8±2.1		3
	900 ms	8.1±4.5		2
		1D unbal. ANOVA p=0.14		
Anesthetized, Stimulus Contrast	0.1	10.4±3.5		3
	0.4	11.3±3.1		4
	0.7	10.0±1.8		3
	1	9.0±0.6		26
		1D unbal. ANOVA p=0.63		
Anesthetized, Stimulus Type	Drifting grating	9.0±0.6		26
	Persistent, all	11.0±0.6		8
		two-tailed unpaired t-test p=0.03		
	Persistent, full-field flash	10.1		1
	Persistent, reversing checkerboard	12.0		1
	Persistent, static grating	11.0±2.7		5
	Persistent, 100 ms full-field flash	7.5		1
Anesthetized, Fractional Suppression	30-50%	10.5±1.0		4
	50-70%	14.0±8.8		2
	70-90%	12.0±2.1		5
	>90%	17.4±3.1		5
		1D unbal. ANOVA p=0.65		
Brain State, Anesthetized vs. Awake	Anesthetized, drifting grating	9.0±0.6		26
	Awake, drifting grating	10.5±0.7		9
		two-tailed unpaired t-test p=0.15		
Awake, Across Layers	L2/3	6.1±0.5	4.1±1.2	11
	L4	7.1±2.3	3.2±0.6	11
	L5a	9.7±2.5	5.6±1.8	11
	L5b	14.1±1.9	14±2.6	11
	L6	14.3±2.1	6.1±0.8	11
		1D bal. ANOVA p=0.03	1D bal. ANOVA p=9.2e-5	
		5b slower than 2/3,4	L5b slower than all others	
Awake, Stimulus Contrast	0.1	15.0±3.4		5
	0.4	14.4±2.5		7
	0.7	11.6±2.9		5
	1	10.5±0.7		9
		1D unbal. ANOVA p=0.24		4

Table 2.1: Network time constant (NTC) fit to CDF across conditions (continued).

Condition		Mean NTC (ms) and s.e.	Mean ms to half-max	# mice
Awake, Stimulus Size	5 degrees	14.0±4.3		4
	7 degrees	14.6±2.9		3
	15 degrees	13.7±1.6		6
	Full-field	10.5±0.7		9
		1D unbal. ANOVA p=0.21		
Awake, Stimulus Type	Drifting grating	10.5±0.7		9
	Persistent, all	13.1±1.0		11
		two-tailed unpaired t-test p=0.04		
	Persistent, full-field flash	13.0±1.6		9
	Persistent, reversing checkerboard	13.5±1.9		5
	Persistent, static grating	12.5±2.7		3
	Persistent, 100 ms full-field flash	14.0		1
Awake, Behavioral State	Running	9.4±3.1		4
	Non-Running	10.3±3.0		4
		two-tailed paired t-test p=0.16		
Awake, Evoked vs. Spontaneous Baseline	Drifting grating	10.5±0.7		9
	Spontaneous baseline	13.7±1.5		7
		two-tailed paired t-test p=0.18		
		two-tailed unpaired t-test p=0.08		
Awake, V1 vs. V2	V1	10.5±0.7		9
	V2	15.8±2.4		3
		two-tailed unpaired t-test p=0.003		

activity in awake mice relies on ongoing thalamic input[*]. Also, the CDF was only a few milliseconds slower in a higher-order visual cortex (Figure 2.18d; awake V2 lateral [128], Appendix A: Supplementary Methods, "Verification of V2 recording sites", pg. 116; V2 CDF: 16±5 ms, n=3 mice, unpaired vs. V1 p=0.003). Thus, visually evoked activity in cortical recurrent circuits locks to the timing of the thalamic input even in awake conditions and secondary visual areas (see Table 2.1 for a complete list of tested conditions, Figures 2.9-2.26).

2.3.4 Features of the cortical decay function in V1

Over the 50 ms initially following thalamic silencing, the trajectory of the trial-averaged V1 multi-unit and single-unit activity is well described by a single exponential decay function with a time constant of ~10 ms (Figures 2.3 and 2.18). However, over the approximately 250 ms it takes for V1 activity in anesthetized mice to return to the pre-stimulus baseline following thalamic silencing, a slow, low-amplitude recovery component of the CDF is observed (see Figure 2.9c). Because this component is quite small in amplitude relative to the fast decay captured by the ~10 ms time constant, and because this low-amplitude recovery component is variable across mice, I focused on describing the initial fast shut-off of V1 sensory-evoked activity, which dominates the CDF[†]. However, the low-amplitude recovery component may have implications for how cortex filters thalamic input. I will return to this in Chapter 3.

[*] The dependence of spontaneous activity in V1 on thalamic input in awake mice contrasts with the thalamus-independence of spontaneous activity in V1 under anesthesia. This change was related to anesthesia, because spontaneous activity in V1 became dependent on thalamic input as each mouse awoke from anesthesia as I continued to record from the same V1 site across this brain state transition (Figure 2.22).

[†] Interestingly, the ~10 ms CDF of cortex is also consistent with the dynamics of cortical gamma-frequency activity (30-80 Hz), a temporal pattern of neural activity believed to be generated by the intra-cortical recurrent circuits [64,129]. The period of a 60 Hz gamma cycle is 17 ms, and the CDF is ~10 ms (approximately half of a gamma cycle). In support of the possibility that the CDF is related to the mechanism underlying gamma activity in cortex, I found that the time required for cortex to shut off after silencing thalamus deviated by a few ms across trials as a function of the phase of gamma activity in cortex (Figure 2.27). Silencing thalamus just after the trough of cortical gamma activity in the LFP led to a slightly longer time delay (~5-10 ms) before visually evoked activity in cortex disappeared. However, because cortical gamma frequency activity does not reliably phase-lock to the visual stimulus, averaging across trials (and thus across gamma phases at the moment of silencing thalamus) always produced the ~10 ms CDF.

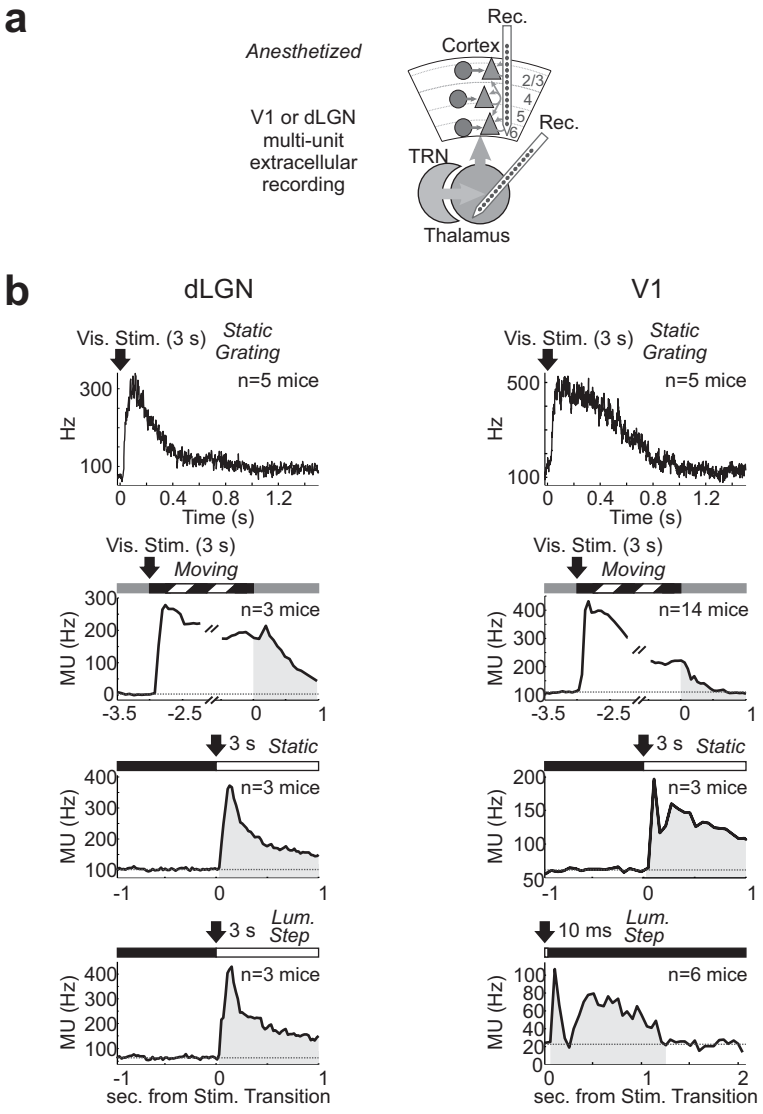

Figure 2.8: Time course of thalamic and cortical responses to visual stimuli.
(a) Schematic of extracellular recordings in dLGN or V1 of anesthetized mice.
(b) Time course of multi-unit activity (MU) in response to visual stimuli. **Left column:** MU in dLGN. **Right column:** MU in V1 (different animals than dLGN at left). **Top row:** Average peri-stimulus time histogram (PSTH) of response to appearance of static grating (static for 3 s). V1 PSTH (right) is from Figure 2.1. **Second row:** PSTH response to onset and offset of moving grating (moving for 3 s). **Third row:** PSTH response to appearance of static stimulus (both 3 s-long static gratings and 3 s-long luminance step). **Bottom row:** PSTH response to luminance step only (3 s duration, left) in dLGN or 10 ms-long luminance step (10 ms duration, occurring at 0 s) in V1 (right). Response to even this very brief (10 ms-long) stimulus is long-lasting. Gray shading is integral of response after vis. stim. transition (stops moving).

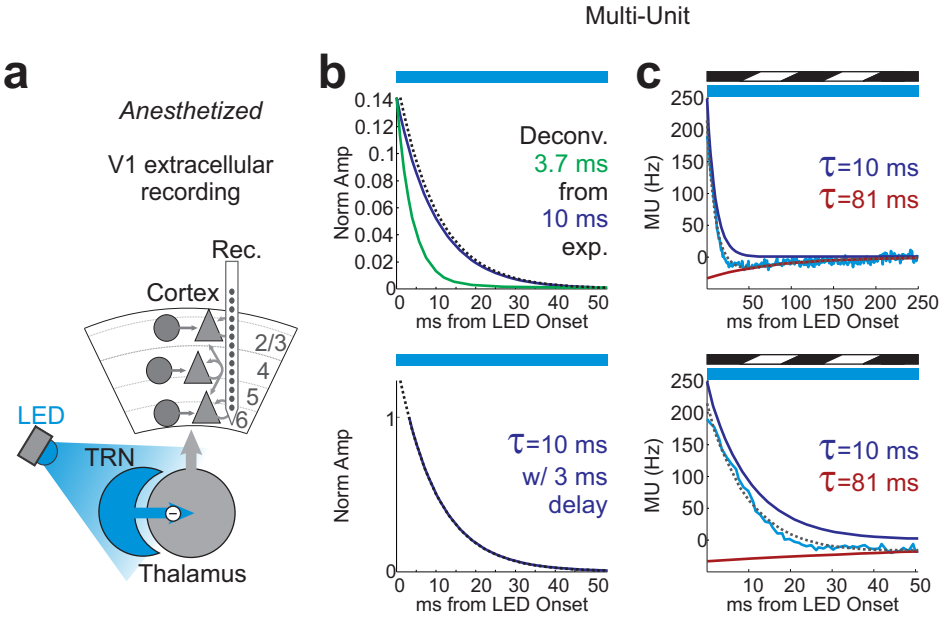

Figure 2.9: Fitting the CDF in V1 of anesthetized mice.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) To determine the exact time course of activity decay in V1 following thalamic silencing, i.e., the time course of the cortical decay function (CDF), we need to remove (deconvolve) the time course of shut-off in the thalamus from the time course of shut-off in the cortex. Blue bar indicates LED illumination of TRN.
Top: Deconvolution (dotted black) over 2-50 ms (using fits to dLGN=green and V1=blue decays, see "Accounting for shut-off delay in dLGN", pg. 127).
Bottom: The deconvolution (dotted black) closely matches a 10 ms single exponential decay (blue) that begins at 3 ms after onset of the LED to silence the thalamus. Because the deconvolution introduces error in real, noisy data, we use a single exponential function, beginning at 3 ms after LED onset, to fit the CDF.

(c) Fitting the CDF time course with two exponentials over longer timescales. Blue bar indicates LED illumination of TRN. Black and white angled bars indicate moving grating stimulus (3 s duration).
Top: Time course of average baseline-subtracted multi-unit (MU) activity in V1 after silencing thalamus with LED. Dark blue: exponential fit component 1, a fast decay. Red: exponential fit component 2, a slow recovery component (red amplitude is 13% of dark blue). Gray dotted: sum of dark blue and red fits.
Bottom: As left panel but on an expanded timescale. Elsewhere in the text we report a single exponential fit over the initial 50-80 ms of the CDF, which is dominated by the fast component.

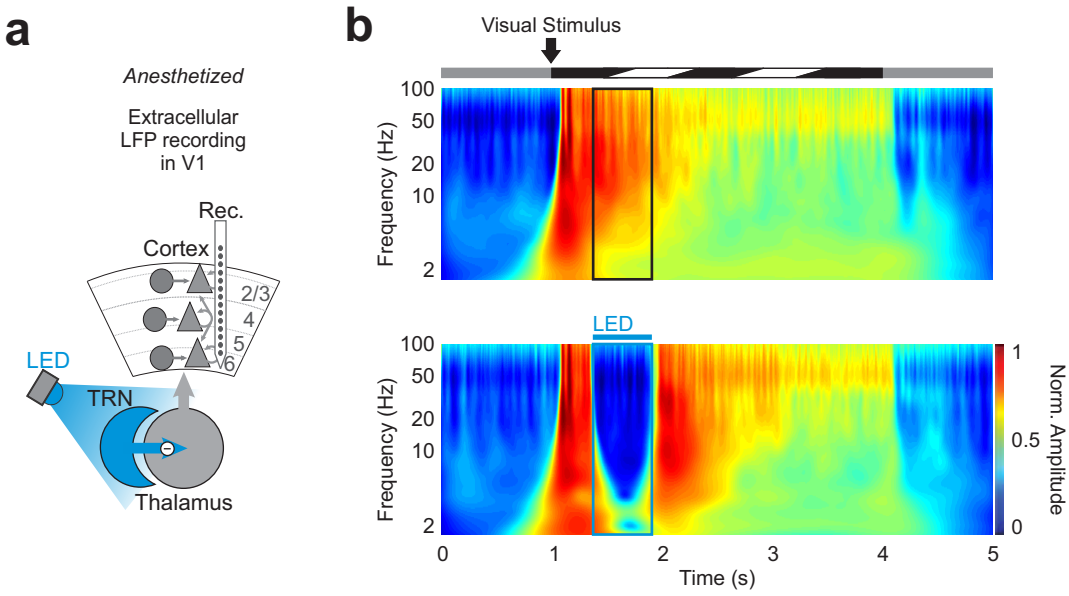

Figure 2.10: Visually evoked local field potential (LFP) in V1 shuts off upon silencing thalamus.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) Trial-averaged spectrograms of LFP (normalized amplitude, whitened) in representative example mouse in response to onset of moving grating (3 s duration) visual stimulus. Temporal frequency of moving grating is 2 Hz. Angled black and white bars (at top) represent moving grating.
Top: LFP spectrogram response to visual stimulus.
Bottom: LFP spectrogram response to visual stimulus with thalamic silencing. Blue bar indicates duration of LED illumination of TRN to silence thalamus.

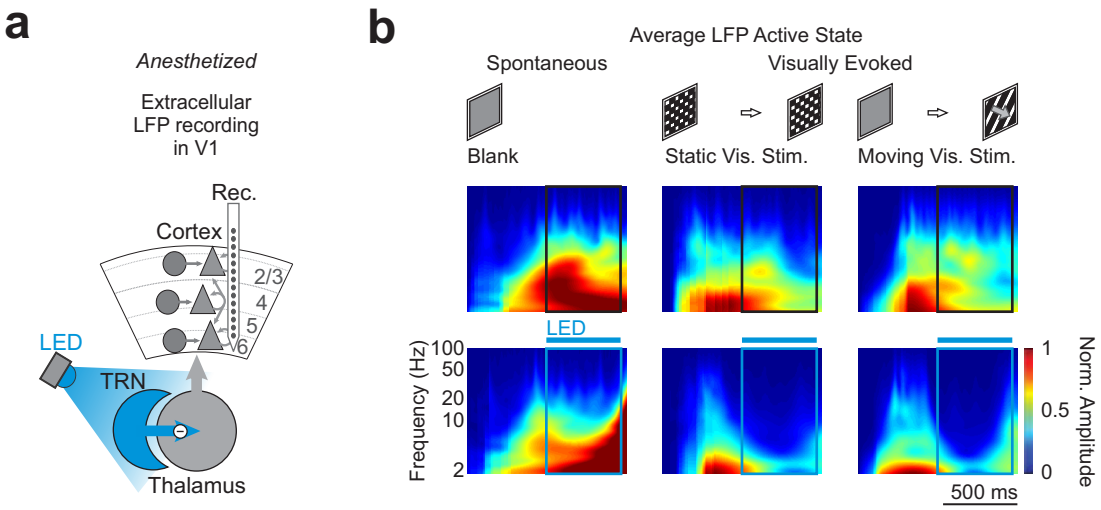

Figure 2.11: Visually evoked but not spontaneous activity in the local field potential (LFP) in V1 shuts off upon silencing thalamus.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1 under anesthesia.

(b) Average LFP spectrograms (amplitude across frequency bands as a function of time) triggered at the onset of active state in V1 (active states from single representative example mouse). Active states separated according to whether they were spontaneously occurring (**left,** Up states) in the absence of visual stimulation (i.e., blank gray screen) or visually evoked (middle and right) by reversal of a static checkerboard pattern (**middle**) or the onset of a moving grating (**right**, 3 s duration of moving grating). Visual stimuli schematized at top. Amplitude in normalized, arbitrary units.
Top row: Active states in V1 with intact input from thalamus (control).
Bottom row: Active states in V1 beginning just prior to thalamic silencing. Blue bar indicates LED illumination of TRN to silence thalamus. Note that spontaneously occurring active states (Up states) do not shut off, but visually evoked active states shut off.

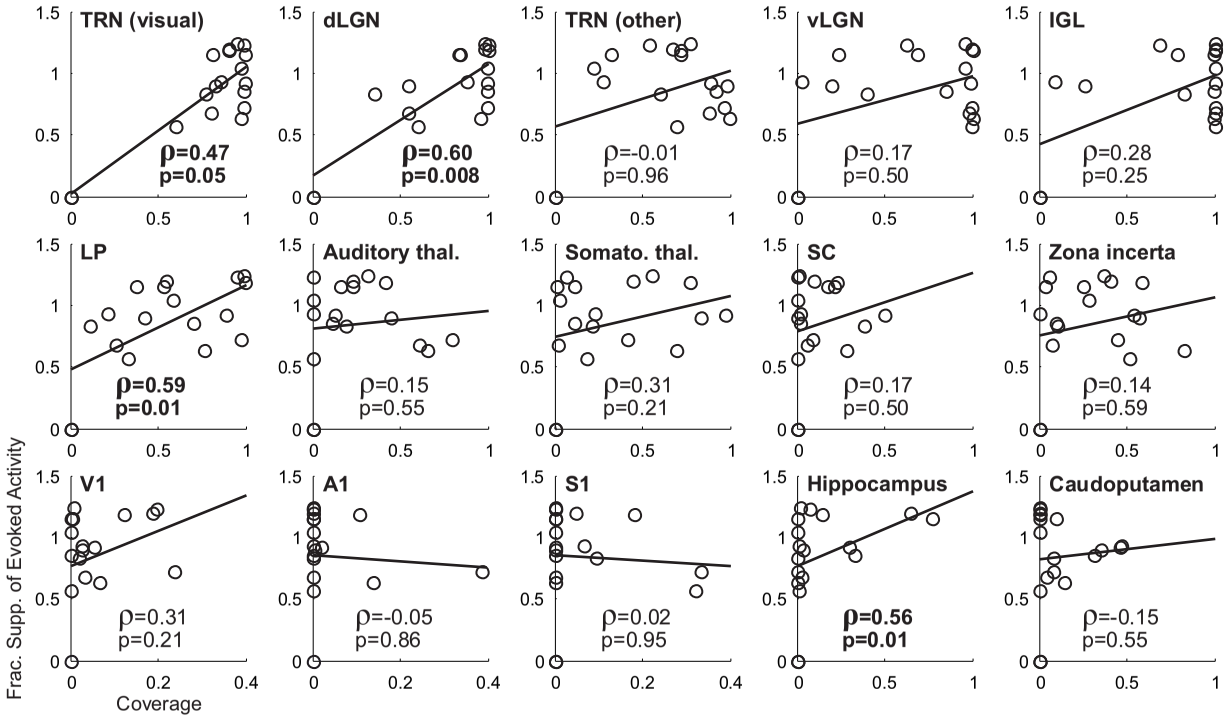

Figure 2.12: dLGN inhibition correlates with shut-off of visually evoked response in V1.

Coverage is fraction of pixels expressing TdTomato within region-of-interest (ROI) containing brain area. For brain structures with >1 mm anterior-posterior extent, coverage is average of coverage across three coronal sections spanning extent of structure. Scatter plots show coverage of different brain areas plotted against fractional suppression of V1 evoked response (frac. supp.) following photo-activation of TRN in anesthetized animals. Each point is one mouse. ρ is correlation coefficient; p is p-value of correlation. Coverage of TRN (visual sector) and dLGN are the only expression patterns necessary to observe shut-off of V1 evoked response (i.e., linear fit passes near the origin, has a positive slope). Note that expression of ChR2 in TRN axons to LP is not necessary to observe shut-off of V1 response; even when ChR2 expression in axons to LP is weak, we may observe a strong shut-off of V1 evoked response. Abbreviations: A1=primary auditory cortex, vLGN=ventral lateral geniculate nucleus, IGL=intergeniculate leaflet, LP=lateral posterior nucleus, SC=superior colliculus, Auditory thal.=medial geniculate nucleus, S1=primary somatosensory cortex, Somato. thal.=ventral posterior medial/lateral nuclei.

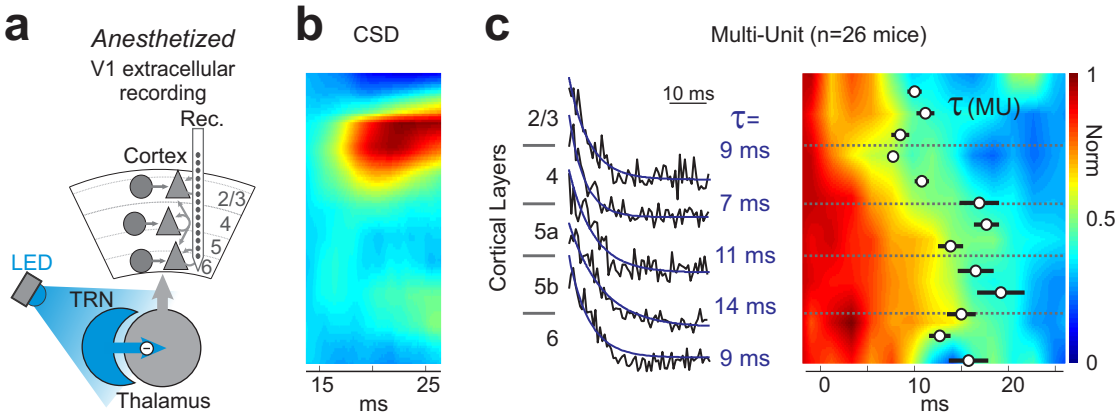

Figure 2.13: CDF across V1 cortical layers in anesthetized mice.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) Current source density (CSD) plot across cortical depths (average of 26 mice) in response to the appearance of a moving grating (gray screen to grating transition). Hotter colors (i.e., red) are current sinks. Colder colors (i.e., blue) are current sources (inverse of standard color scheme so that red indicates greater activity in CSD and in multi-unit heatmap at right). Normalized color scheme at right of (c). Cortical depths are aligned to cortical depths in (c). Time is in ms from appearance of moving grating.

(c) Left: Multi-unit (MU) CDF across putative cortical layers (n=26 mice) and fits (dark blue). **Right:** Heat-map of MU firing rates normalized across recording depths. 1 (red) is MU rate before LED onset; 0 (blue) is baseline rate. LED onset at 0 ms. Dotted lines show putative layer boundaries based on CSD. Black overlay: Mean±s.e. of time constant of single exponential fit to CDF. Note slightly slower shut-off in deeper layers.

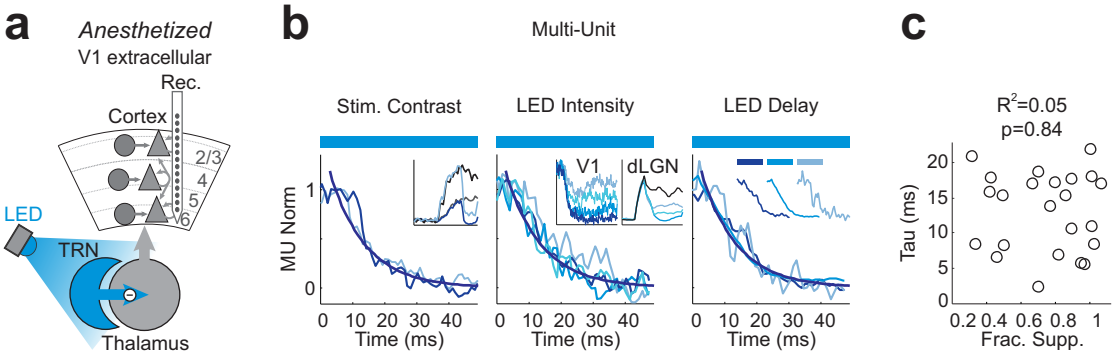

Figure 2.14: CDF in anesthetized mice is constant across experimental conditions.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) The CDF (mean normalized multi-unit activity, MU) is not affected by stimulus contrast, LED intensity or LED onset delay. Traces are scaled and superimposed. Dark blue is 11 ms fit.
Left: Contrast 1 is light blue. Contrast 0.25 is dark blue (n=5 mice). Inset: Zoomed out MU evoked response in example mouse (thalamic silencing not complete in this example) to low- (gray/dark blue) or high-contrast (light blue/black) moving grating with (light blue/dark blue) or without (black/gray) thalamic silencing; y axis is 200 Hz, x axis is 800 ms.
Middle: As left but varying LED intensity to achieve 30-50%, 50-70%, 70-90%, or >90% suppression of V1 evoked response (n=2 mice). Inset: (left inset) V1 CDFs un-normalized; (right inset) different levels of silencing of dLGN during LED intensities 0.5, 3.6 and 6.5 mW; y axis is 150 Hz, x axis is 1.5 s, PSTHs smoothed with 250 ms window.
Right: As left but for LED onset delays of 200 ms, 300 ms or 900 ms relative to stimulus onset (n=3 mice). Inset: CDFs separated by onset delay, left to right: 200, 300, 900 ms. Table 1 for more conditions.

(c) Single exponential fit (Tau) to CDF as a function of fractional suppression (Frac. Supp.) of visually evoked response in V1 across recordings and varying LED intensity. Note lack of correlation.

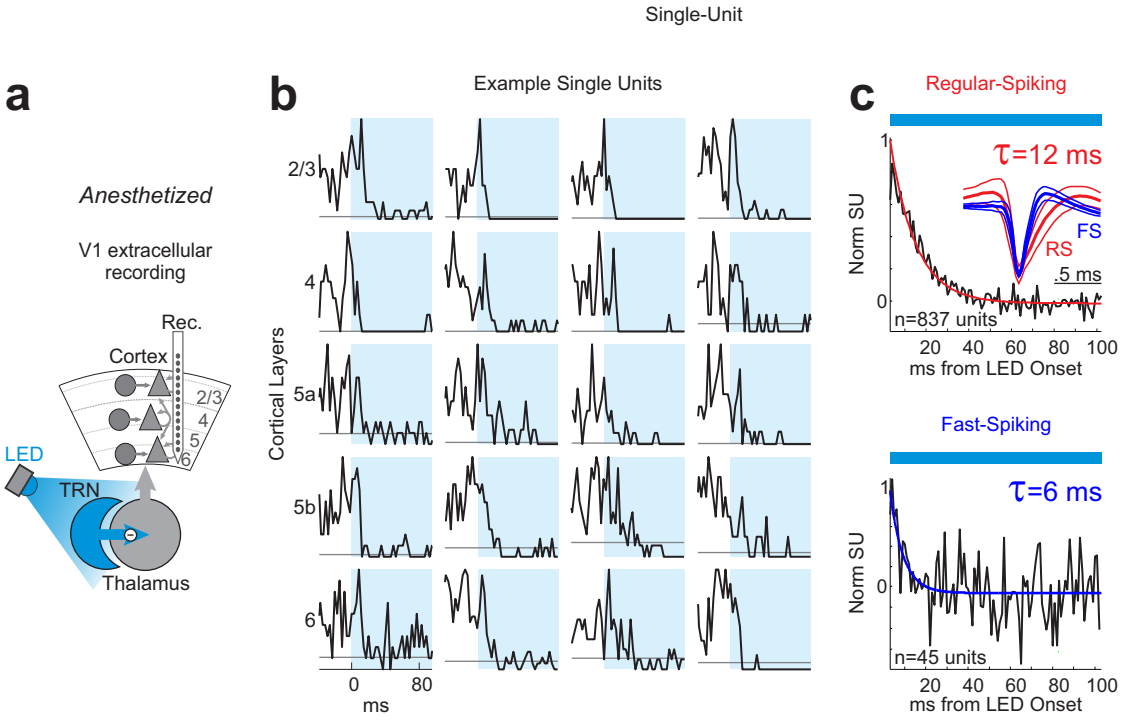

Figure 2.15: Single-unit CDF in anesthetized mice.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) Example single units (SU) during visual stimulation at onset of LED illumination of TRN. Light blue shading: LED on. Gray line: mean pre-stimulus baseline of each unit. Unit examples arranged by putative layer (see labels at left). All traces normalized to peak.

(c) Top: Mean normalized shut-off of regular-spiking (RS) units across cortical layers (red is fit). Inset shows mean and std. dev. of waveforms for RS and fast-spiking (FS) units (see "Separating RS and FS units in the cortex", pg. 121, for separation of these unit types).
Right: As at (top) but for fast-spiking (FS) units (blue is fit).

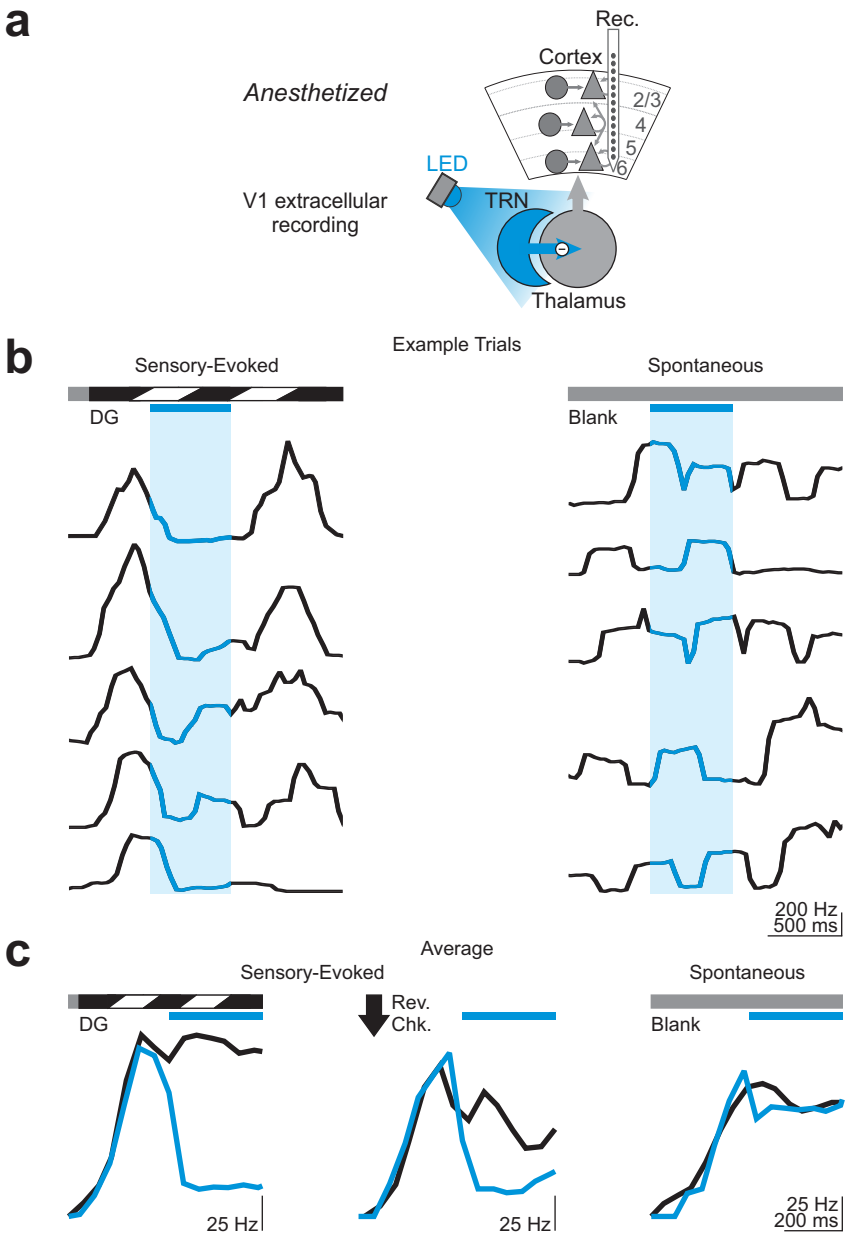

Figure 2.16: Sensory-evoked but not spontaneous active states in V1 shut off upon silencing the thalamus under anesthesia.
Blue bar/shading indicates LED illumination of TRN to silence thalamus.
(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1 under anesthesia.
(b) Example single trials of V1 multi-unit activity in response to visual stimulus (**left**, vis. stim. is moving grating, 3 s duration, indicated by angled black and white bars) or during spontaneous Up states (**right**). Multi-unit activity is binned into 50 ms bins and smoothed with 250 ms window.
(c) Average V1 active states during (blue) or without (black) thalamic silencing. Active states detected by LFP ratio ("V1 active states...", pg. 121) and displayed as summed single-unit activity in V1 in response to visual stimuli (moving, **left**; static, **middle**) or during spontaneous Up states (blank screen, **right**).

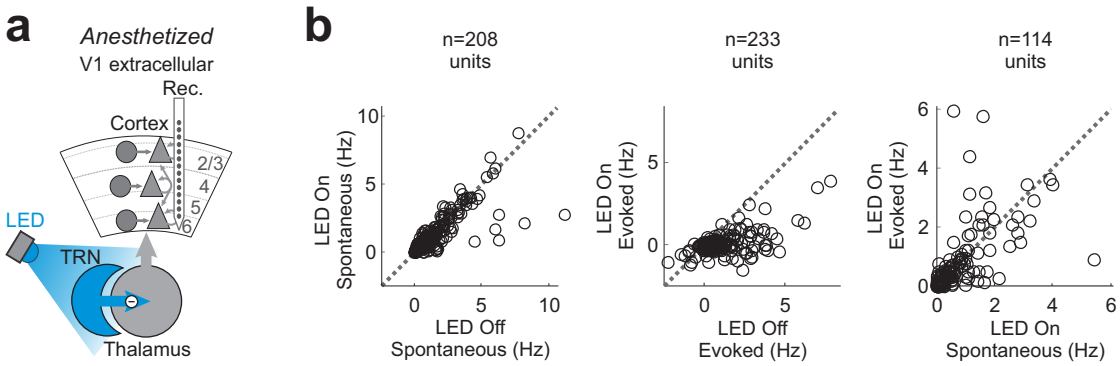

Figure 2.17: Effect of thalamic silencing on single-unit activity in V1 of anesthetized mice.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) Single-unit activity at moment of silencing thalamus. Each point is a unit. Dotted line is unity. Points below line are suppressed.
Left: Spontaneous activity in V1 is unaffected.
Middle: Visually evoked activity is suppressed (3 s-long moving visual stimulus).
Right: Analysis of whether each single unit in V1 returns to its own spontaneous baseline after the thalamus is silenced during a sensory response. X axis is each unit's spontaneous rate (no visual stimulus, blank screen) during the LED, and Y axis is that unit's evoked response (moving grating present) during the LED. No systematic change from spontaneous activity when the thalamus is off during the visual stimulus.

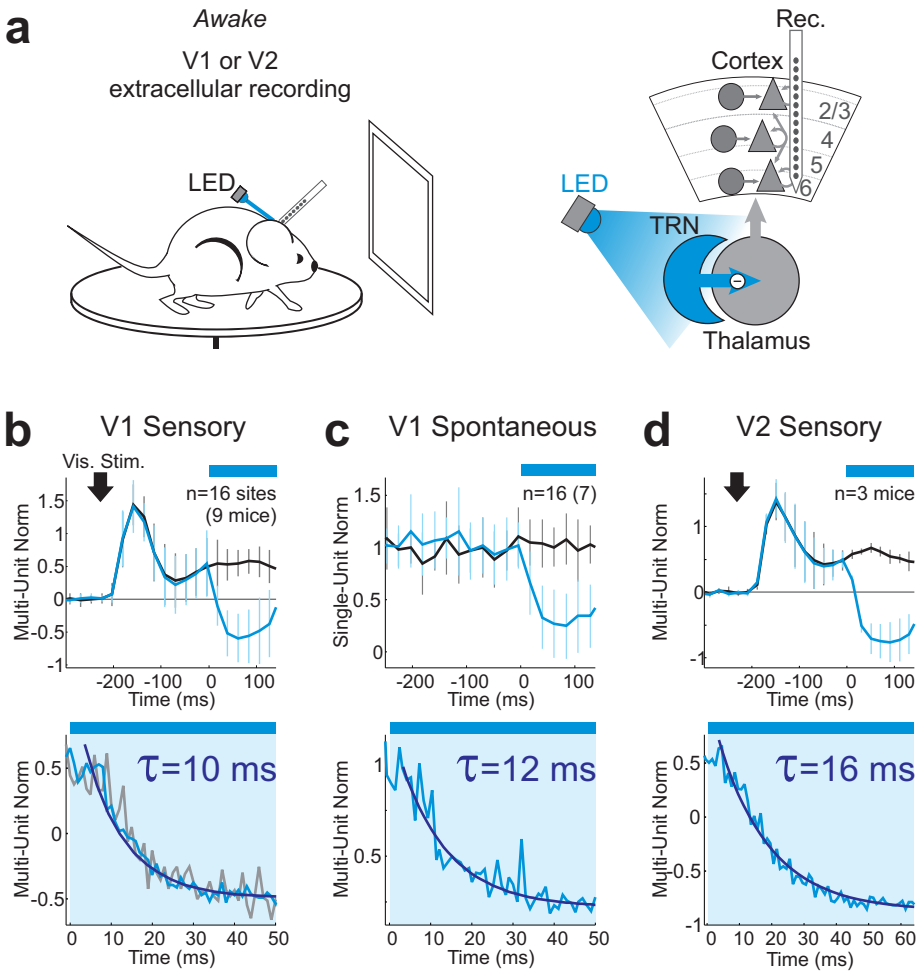

Figure 2.18: Time course of shut-off of visually evoked activity in V1 and V2 of awake mice during thalamic silencing.

(a) Experimental configuration.
Left: Awake mice, head-fixed on circular treadmill.
Right: Silencing thalamus while recording across cortical layers in V1 or V2 with an extracellular multi-channel probe.

(b) Sensory-evoked activity in V1 of awake mice.
Top: Mean multi-unit (MU) activity in V1 in response to appearance of visual stimulus (arrow) of duration 3 s followed by thalamic silencing (blue) or no optogenetic manipulation (control, black). MU is baseline-subtracted and normalized to first 150 ms of control evoked response (Appendix B for stimuli). Gray line: pre-stimulus baseline.
Bottom: Expanded timescale showing the decay of cortical activity at LED onset (cortical decay function, CDF) during visual stimuli, either static (gray) or moving (light blue) gratings. Dark blue: single exponential fit to CDF. Data include both running and non-running mice (Figure 2.19 for further break-down).

(c) As in (b) but here spontaneous activity in V1 (not baseline-subtracted).

(d) As in (b) but here sensory-evoked activity recorded in V2.

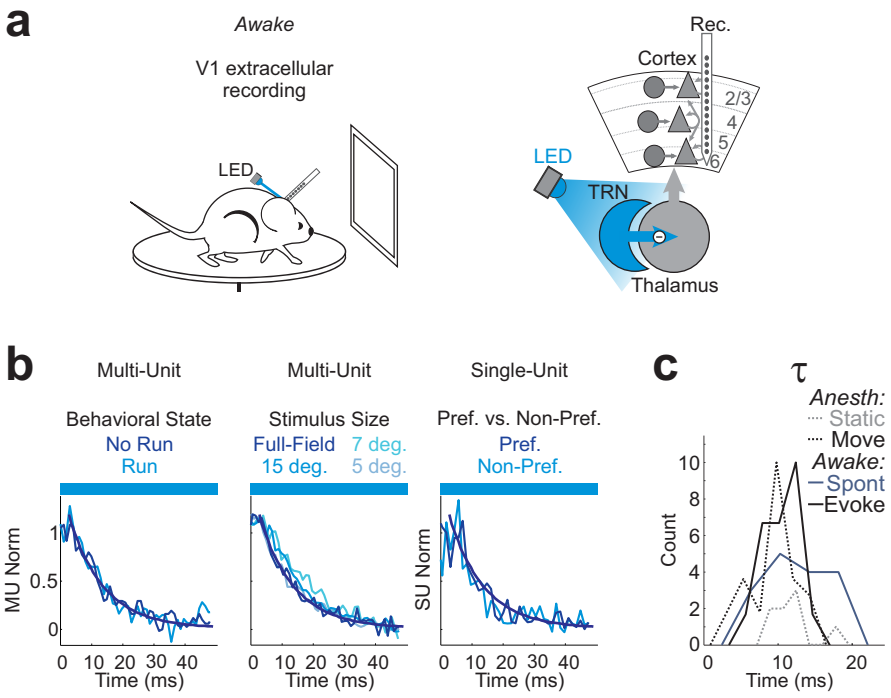

Figure 2.19: CDF in awake mice is constant across experimental conditions.

(a) Awake extracellular recordings in V1 sand optogenetic configuration to silence thalamus by photo-activating TRN.

(b) Cortical decay function (CDF) is independent of animal's behavioral state (**left**; superimposed multi-unit activity, MU, shut-off for running and non-running trials; n=4 mice), stimulus diameter (**middle**; as MU; 5, 7, 15 degrees or full-field; n=3 mice), and single unit's stimulus preference (**right**; i.e., grating orientation; mean single-unit activity, SU, during stimulation with preferred versus non-preferred orientation of moving grating; n=127 SU). Dark blue fit is 11 ms. See Table 2.1 for more conditions.

(c) Histogram of time constants of single exponential fits to CDF (MU) across conditions. Count is mice. Static=CDF for static stimuli eliciting visually evoked activity in anesthetized mice. Move=CDF for moving gratings eliciting visually evoked activity in anesthetized mice. Spont=CDF for spontaneous activity in awake mice. Evoke=CDF for visually evoked activity in awake mice. See Table 2.1 for further break-down.

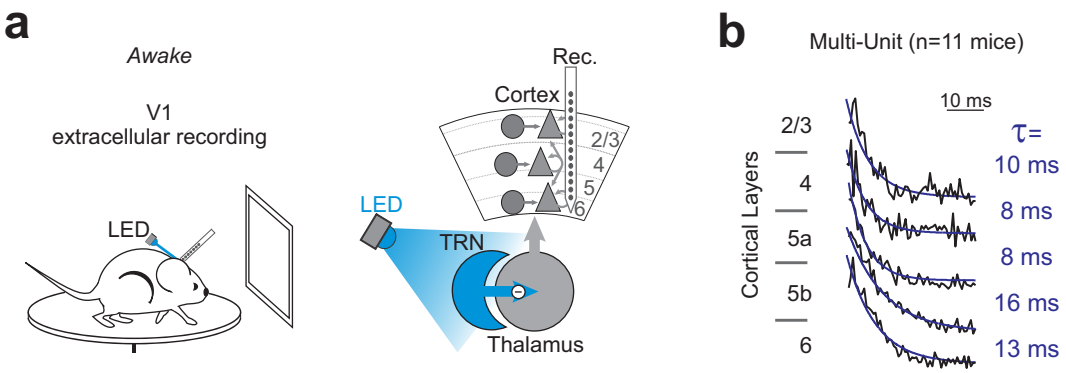

Figure 2.20: CDF across V1 cortical layers in awake mice.

(a) Awake extracellular recordings in V1 and optogenetic configuration to silence thalamus by photo-activating TRN.

(b) Single exponential fit (blue) to cortical decay function (CDF, black, multi-unit activity) across putative cortical layers (n=11 awake mice).

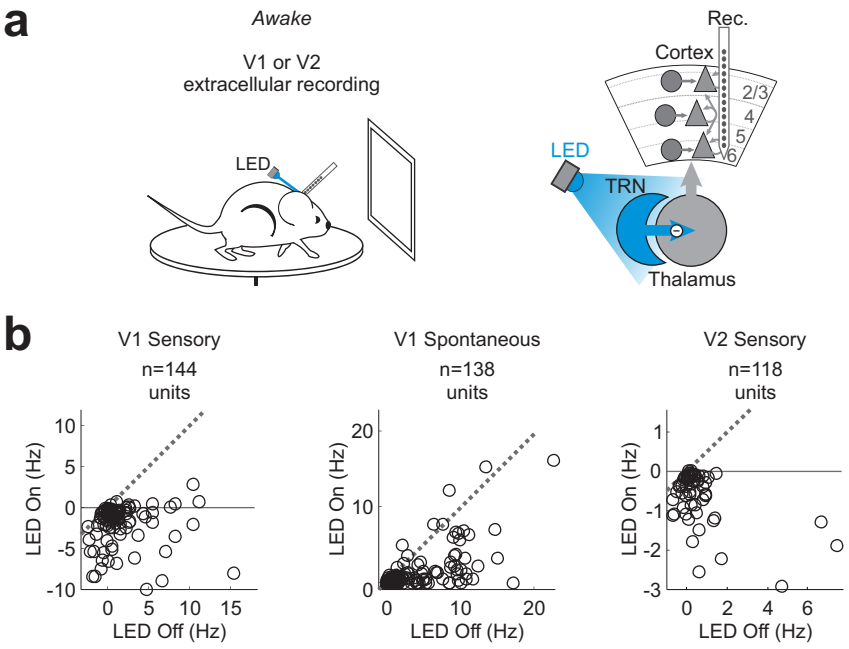

Figure 2.21: Effect of thalamic silencing on single-unit activity in V1 and V2 of awake mice.

(a) Awake extracellular recordings in V1 or V2 and optogenetic configuration to silence thalamus by photo-activating TRN.

(b) Single-unit activity at moment of silencing thalamus. Each point is a unit. Dotted lines are unity (i.e., no effect). Points below unity line are suppressed. **Left:** Visually evoked activity (pre-stimulus baseline-subtracted) in V1 of awake mice is suppressed. Points below zero (solid horizontal gray line) are suppressed below baseline.
Middle: Spontaneous activity in V1 of awake mice is also suppressed by silencing thalamus. Not baseline-subtracted.
Right: Visually evoked activity (baseline-subtracted) in V2 of awake mice is suppressed. Points below zero (solid horizontal gray line) indicate suppression below baseline.

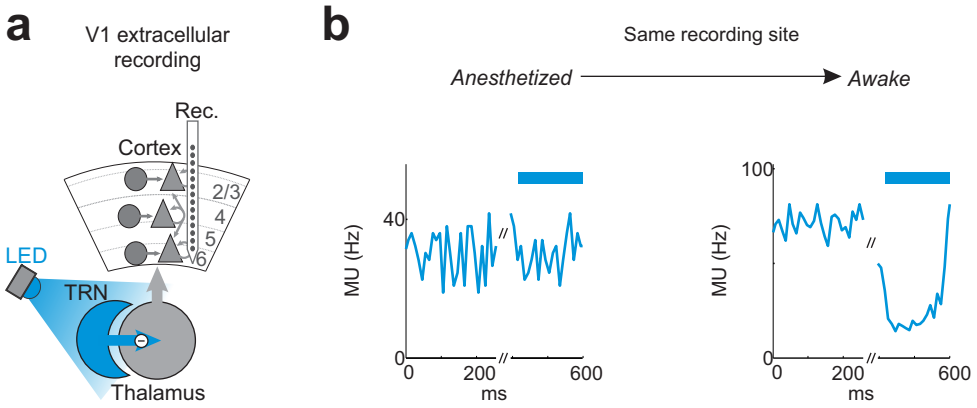

Figure 2.22: Waking up from anesthesia leads to the appearance of an effect of silencing thalamus on V1 spontaneous activity recorded within the same mouse.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) Recording within an example mouse during the transition from isoflurane anesthesia to the awake state. Blue bar indicates LED illumination of TRN. Note that as animal awakes from anesthesia, spontaneous activity increases and becomes thalamus-dependent.
Left: Effects of TRN photo-activation on baseline multi-unit (MU) activity under anesthesia.
Right: As left but animal is awake.

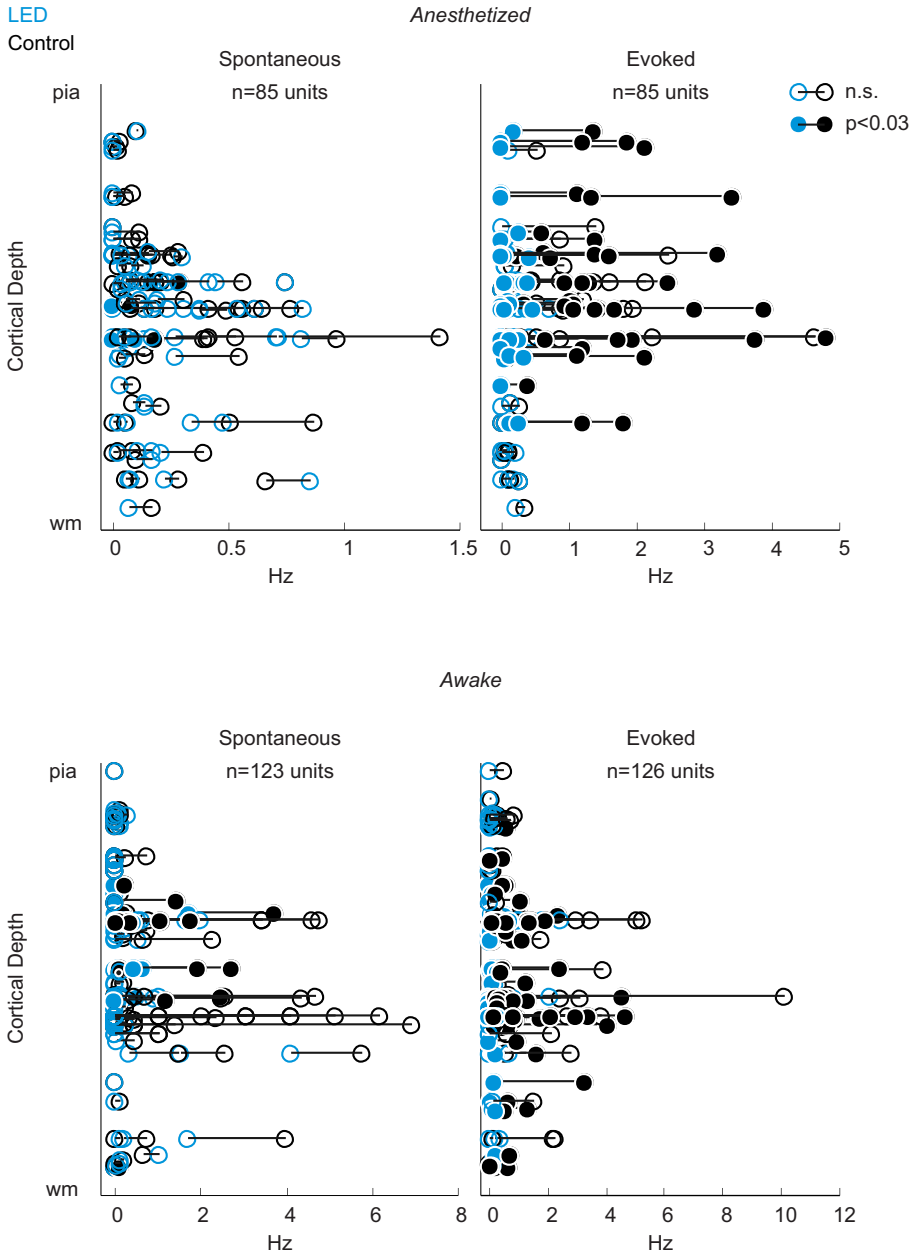

Figure 2.23: Effect of silencing thalamus on unit activity in V1 across cortical depths.

Each pair of points joined by a line represents one V1 single unit's activity, comparing its activity in control conditions (black, thalamic activity intact) with its activity during thalamic silencing by TRN photo-activation (blue). Units are arranged by cortical depth. Filled circles represent a statistically significant ($p<0.03$) change in the unit's activity, using a two-tailed paired t-test on the unit's firing rate across trials. Unfilled circles represent lack of a statistically significant change.

Top: Anesthetized. **Bottom:** Awake. **Left:** Spontaneous activity. **Right:** Visually evoked activity (not baseline-subtracted).

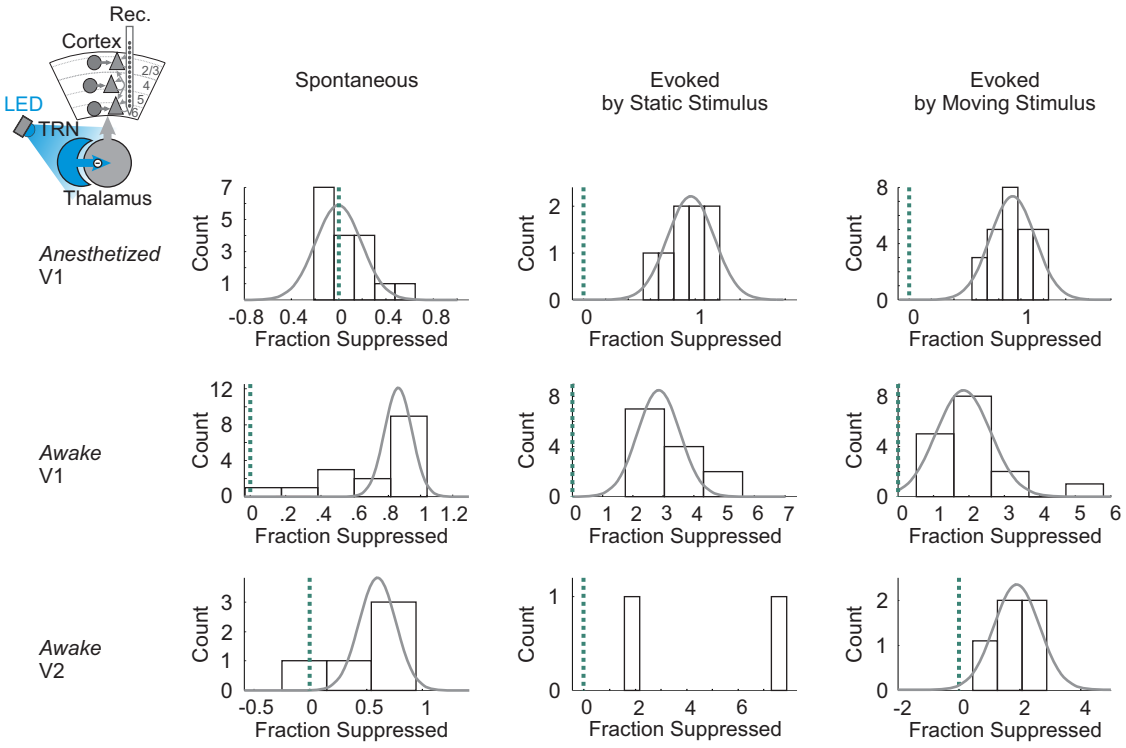

Figure 2.24: Histograms across mice of fraction of multi-unit activity in visual cortex suppressed by silencing thalamus.

Inset schematic at upper left: silencing thalamus by photo-activation of TRN. For each histogram, gray line shows gaussian fit; dotted green line shows no effect of silencing thalamus (fraction suppressed=0). Complete suppression of activity indicated by fraction suppressed=1.

Top Row: Anesthetized V1.
Middle Row: Awake V1.
Bottom Row: Awake V2 (lateral).

Left Column: Spontaneous activity (no visual stimulation).
Middle Column: Activity evoked by static visual stimulus (all static visual stimulus types combined).
Right Column: Activity evoked by moving visual stimulus (moving grating).

Note that all types of activity in visual cortex are abolished as a result of silencing thalamus, except spontaneously occurring Up states under anesthesia.

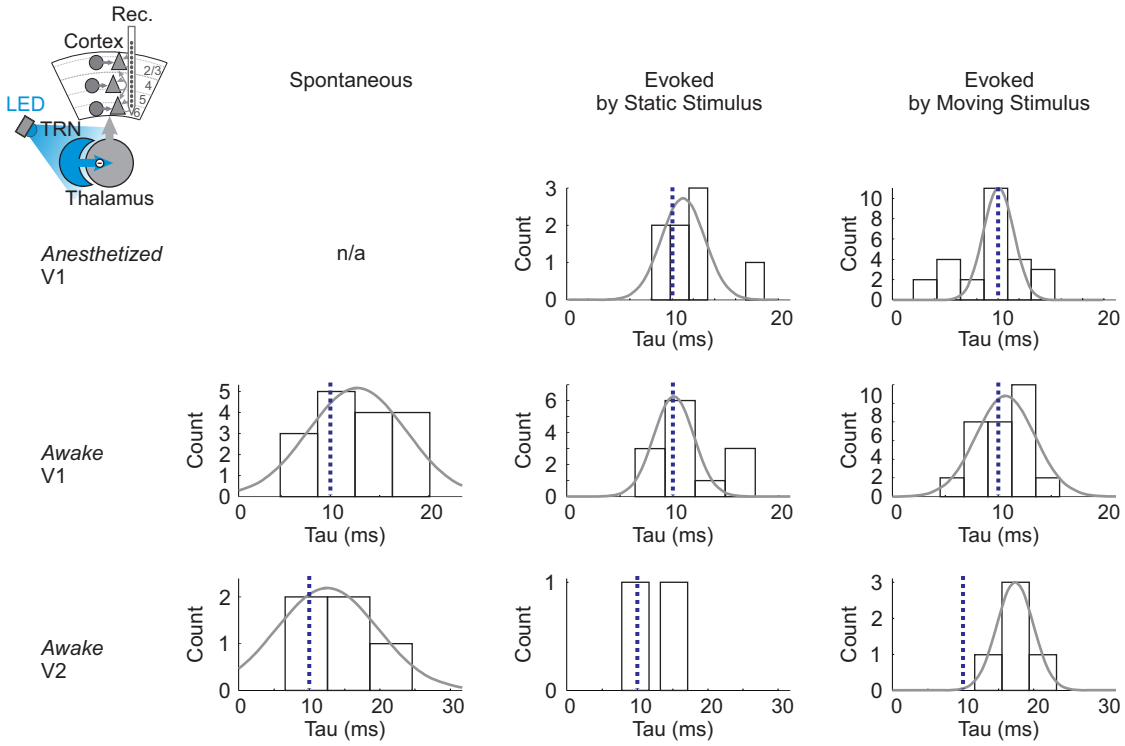

Figure 2.25: Histograms across mice of time constant (tau) of single exponential decay fit to shut-off of multi-unit activity in visual cortex (CDF) upon silencing thalamus.

Inset schematic at upper left: silencing thalamus by photo-activation of TRN. For each histogram, gray line shows gaussian fit; dotted blue line shows tau=10 ms.

Top Row: Anesthetized V1. Note that spontaneous activity under anesthesia does not shut off upon silencing thalamus; therefore, no CDF to fit in this case.
Middle Row: Awake V1.
Bottom Row: Awake V2 (lateral).

Left Column: Spontaneous activity (no visual stimulation).
Middle Column: Activity evoked by static visual stimulus (all static visual stimulus types combined).
Right Column: Activity evoked by moving visual stimulus (moving grating).

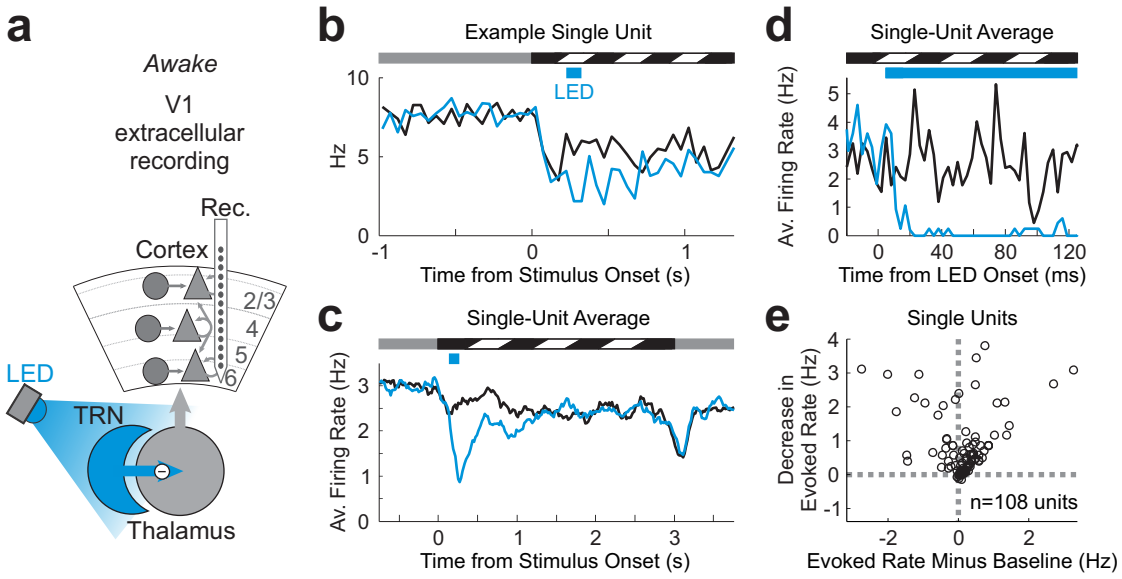

Figure 2.26: Units suppressed by visual stimulus are further suppressed by silencing thalamus in V1 of awake mice.

In all parts, blue bar indicates LED illumination of TRN.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) Peri-stimulus time histogram (PSTH) of example single-unit (SU) in response to the appearance of a moving grating visual stimulus (3 s duration, black and white angled bars) followed by 250 ms-long TRN illumination.

(c) Average across 10 similar units showing suppression of activity in response to visual stimulation. Average is smoothed with 250 ms time window.

(d) Average single-unit activity of 10 units (as in (c)) at moment of silencing thalamus. Not smoothed.

(e) Firing rate of single-unit population. Each point is a unit. X axis is the mean visually evoked firing rate minus the pre-stimulus baseline. Y axis is the decrease in firing rate as a result of silencing thalamus during visual stimulation (i.e., control evoked rate minus rate during thalamic silencing). Points to left of vertical dotted line are suppressed by visual stimulation. Points above horizontal dotted line show a decrease in activity during thalamic silencing. Note that most units show a decrease in activity during thalamic silencing, even when these cells are suppressed by visual stimulation.

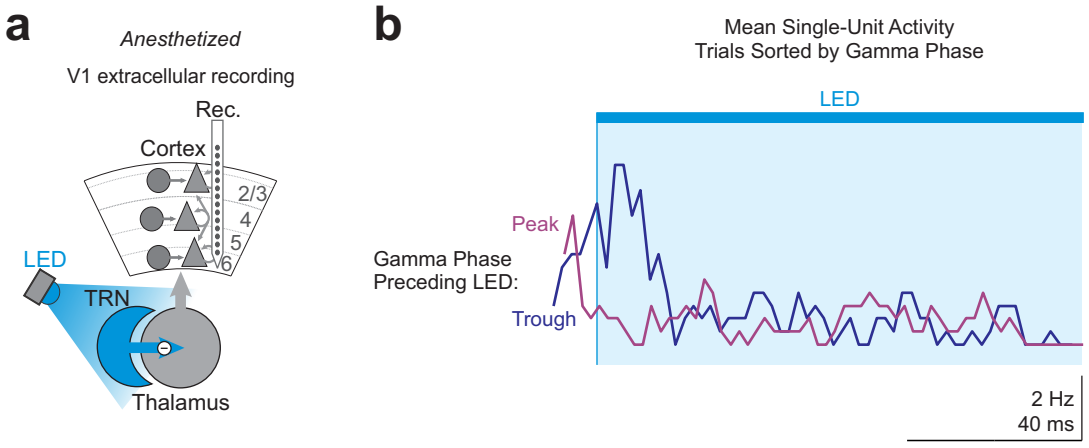

Figure 2.27: Single-trial time course of shut-off of V1 visually evoked activity upon silencing thalamus as a function of gamma phase.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) Mean single-unit (SU) activity in V1 of example mouse at moment of silencing thalamus sorted by phase of gamma frequency activity in local field potential (LFP). Blue bar indicates LED illumination of TRN to silence thalamus. Dark blue line: mean SU activity in V1 across trials where gamma trough precedes LED onset. Purple line: mean SU activity in V1 across trials where gamma peak precedes LED onset.

2.4 Discussion

The cortical decay function (CDF) describes a property of recurrent circuit dynamics, the timescale over which sensory activity within recurrent circuits of visual cortex can be sustained without thalamic input. I found that the CDF was on the order of the integration time window of a single cortical neuron (fit by single exponential decay with an approximately 10 ms time constant[*]), suggesting that recurrent dynamics in the cortex are fast enough to lock the cortical network response to the timing of thalamic input, consistent with models of cortex [5,7,8]. The CDF varied by no more than a few milliseconds across all experimental conditions tested (i.e., anesthetized versus awake brain states, time point in the visual response, stimulus contrast, stimulus pattern, extent of thalamic silencing, stimulus size, V1 versus V2, and running versus non-running behavioral states). Therefore the sensory-evoked dynamics intrinsic to recurrent circuits in visual cortex are fast and robust.

This was true even when intra-cortical recurrent circuits provided the majority of the sensory-evoked synaptic excitation to cortical neurons in layer 4, after the first ~40 ms of the response in cortex (Figure 2.1a). At these later time points beyond ~40 ms, thalamocortical input was only the minor source of synaptic excitation, yet, even at these later time points, silencing thalamus led to a fast decay of sensory-evoked activity in V1. Thus, even when thalamocortical input was the minor source of synaptic excitation, it determined the time course of the sensory response in cortex, which otherwise decayed slowly over hundreds of ms when thalamic input was intact (Figure 2.3f). How does this minor thalamocortical excitatory synaptic input control the timing of the cortical response?

This question is even more puzzling if we consider that, consistent with models of cortex [5,7], cortico-cortical recurrent synaptic excitation (measured as the cortical fraction of the EPSC), in response to a high-contrast visual stimulus, may be larger in magnitude than the direct thalamocortical synaptic excitation, in response to a low-contrast visual stimulus. If the thalamocortical synaptic

[*] The CDF also appeared consistent with the period of cortical gamma-frequency activity. Because cortical gamma-frequency activity has the same frequency across mammalian species and across cortical areas, a mechanistic relationship between the CDF and gamma-frequency cortical activity might imply that the CDF is the same in other mammalian species and in other cortical areas. Moreover, inhibition is thought to be the mechanism underlying gamma-frequency activity in cortex [129,130], suggesting cortical inhibition might also underlie the CDF.

excitation, in response to a low-contrast visual stimulus, is able to drive a response in cortex, which it clearly is, then why can't an even larger-magnitude synaptic excitation from cortico-cortical recurrent circuits sustain the response in cortex?

This reasoning suggests that thalamocortical input is somehow privileged in its ability to drive sensory-evoked activity in cortex. One possible explanation, proposed in models [5,7,8,55], is that thalamic input recruits more excitation than inhibition in cortex and is thus able to drive spiking activity in post-synaptic cortical neurons, but cortico-cortical transmission, which does not recruit relatively more excitation than inhibition, is not sufficient, on its own, to drive cortical spiking. In this view, the strong cortical recurrent excitation is counterbalanced by strong cortical recurrent inhibition. Hence the net reversal potential of cortico-cortical synaptic input, which represents the combined effects of both excitation and inhibition, may not even exceed the threshold for action potential firing. Thalamocortical input, on the other hand, is able to drive the membrane potential of the post-synaptic neuron above spike threshold. Once thalamocortical input is withdrawn, in this view, spiking activity in cortical recurrent circuits decays rapidly. In these models, therefore, it is the balanced effects of recurrent excitation and inhibition that provide fast dynamics of the cortical circuits. If this explanation is sufficient to account for the dynamics of cortex, then cortico-cortical connections should never drive spiking of cortical neurons in the absence of thalamic input.

However, clearly, spontaneous Up states under anesthesia drive spiking in cortical neurons and are sustained by cortico-cortical connections in the absence of thalamic input [91,92]. This indicates that the cortico-cortical connections recruited by Up states have different effects on the post-synaptic neuron than do cortico-cortical connections recruited by sensory input.

Are the cortical neurons in cortex recruited by Up states different from the cortical neurons recruited by visual stimulation? I found that Up states and visually evoked active states recruited largely overlapping populations of cortical neurons (Figure 2.3i), but it may still be the case that these populations differ sufficiently in composition to ensure that Up states persist in the absence of thalamic input but sensory-evoked active states do not. Alternatively, other differences may exist that explain

why sensory-evoked input fails to recruit the same cortico-cortical connections that sustain Up states without thalamic input.

Importantly, in the awake state, the thalamus-independent Up states are replaced by, or overwhelmed by, spontaneous activity in cortex that requires thalamic input (about 80% of the spontaneous activity in the awake state disappears during thalamic silencing, Figure 2.18c). Perhaps the remaining 20% of the spontaneous activity in visual cortex in awake mice shares a mechanism with Up states under anesthesia. The absolute magnitude, in average spike rate, of the possibly thalamus-independent spontaneous activity in the awake state (the remaining 20%) does appear similar to the magnitude of Up states under anesthesia (Figure 2.21b, middle, c.f. Figure 2.17b, left). Future studies will need to address the mechanism of this small, possibly thalamus-independent component of spontaneous activity in the awake state. What is obvious from my data is that a new, dominant component of spontaneous cortical activity (the 80% dependent on thalamic input) appears in the awake state. Consistent with this, I found that spontaneous activity in the thalamus increased more than 5-fold as mice woke up from isoflurane anesthesia (see Chapter 3, Figure 3.11). Thus, it seems likely that this higher level of spontaneous activity in the thalamus may drive the higher level of spontaneous activity in cortex in the awake state. Future work will be needed to understand the source of the increased spontaneous activity in the thalamus in the awake state.

Taking all these results together with the literature, we are led to the following view of cortical function. Sensory stimulation triggers activity in cortex through feed-forward thalamocortical connections capable of driving post-synaptic spiking, followed very rapidly by the recruitment of intra-cortical recurrent circuits in proportion to the feed-forward thalamic input (in proportion, as discussed in other work [5,30-32,63,131], because lower levels of thalamic input recruit less intra-cortical recurrent excitation, and higher levels of thalamic input recruit more intra-cortical recurrent excitation). The factor determining this proportionality between thalamocortical and intra-cortical synaptic excitation may evolve over the course of the sensory response (Figure 2.1a). At all time points, however, even when the intra-cortical recurrent circuits strongly amplify thalamic input, cortico-cortical recurrent connections driven by sensory input are not capable of sustaining cortical spiking activity for more than

a few tens of ms without thalamic input, consistent with the principle of amplification (no input implies no output). Up states under anesthesia activate an alternative pathway of connections in cortex, possibly by-passing some strong feed-forward inhibition that delimits the duration of sensory-evoked activity. In the awake state, however, most of the spontaneous cortical activity resembles feed-forward sensory-evoked activity, in that both depend on thalamic input. At least 80% of the spontaneous activity and 100% of the sensory-evoked activity in cortex in the awake state are therefore tightly coupled to the timing of thalamic input. Thus, cortex and thalamus are a tightly coupled functional unit in the conscious state, in terms of dynamics.

No unit in the visual cortex had a CDF longer than several tens of ms (Figure 2.3h). However, interestingly, a small number of cells at a cortical depth corresponding to layer 5b or possibly superficial layer 6 did sustain sensory-evoked spiking activity for ~50-100 ms after silencing thalamus (Figure 2.3h). Moreover, the multi-unit CDF at this cortical depth was longer-lasting than the CDF measured at other cortical depths (Figures 2.13 and 2.20). This suggests that some neurons in the deeper cortical layers may integrate thalamic input over timescales longer than ~10-30 ms, up to ~40-80 ms. If so, these neurons in the deeper cortical layers may show sustained, integrated activity in response to thalamic activity at frequencies above ~10-20 Hz. If these deeper-layer cortical neurons drive activity in thalamic nuclei projecting back to visual cortex, then activity above ~10-20 Hz could potentially be sustained by the interaction between visual cortex and visual thalamus. Alternatively, it may be that the long-lasting sensory-evoked response observed in thalamus is inherited from the retina, which sends both a sustained, low-frequency signal to dLGN as well as a transient, high-frequency signal, both carrying visual information [69,74].

Here I propose that the ~10 ms CDF, which summarizes the response across all layers of V1, enables cortical sensitivity to high-frequency sensory input. If this is true, V1 should precisely follow thalamic input up to ~15 Hz, in a manner consistent with temporal filtering by the CDF (Figure 2.28).

But might higher-order cortical areas discard an ability to follow high-frequency thalamic input in return for dynamics that integrate and therefore sustain sensory-evoked activity over timescales longer than a few tens of ms? I found that the CDF in V2 was 16 ms (Figure 2.18), a few ms slower

than the CDF in V1. However, it is not clear that this finding indicates slower dynamics intrinsic to V2, because this 16 ms CDF in V2 does not account for the decay time of inputs to V2, i.e., the decay time of sensory-evoked activity in V1 or in V1 projections to V2 via LP (the second-order visual thalamic nucleus in rodents). 16 ms does set an upper bound on the CDF of V2. Thus, the CDF in V2 is not as long as some previous estimates of intrinsic intra-cortical dynamics (hundreds of ms or more [6,16,17,44,46,59,60,72,109-112]).

It may be that in higher associative cortical areas, as in the visual cortex, long-lasting sensory-evoked, or delay-period, spiking activity in cortex depends on interactions with the thalamus. Consistent with this possibility, pharmacological silencing of the mediodorsal (MD) thalamic nucleus, the major thalamic input to prefrontal cortex, produces a profound reduction of delay period activity in prefrontal cortex [132].

Chapter 2, in part, has been submitted for publication of the material as it may appear in Reinhold, K., Lien, A.D. and Scanziani, M., 2015, Distinct Recurrent Versus Afferent Dynamics in Cortical Visual Processing, Nature Neuroscience. The dissertation author was the primary investigator and author of this paper.

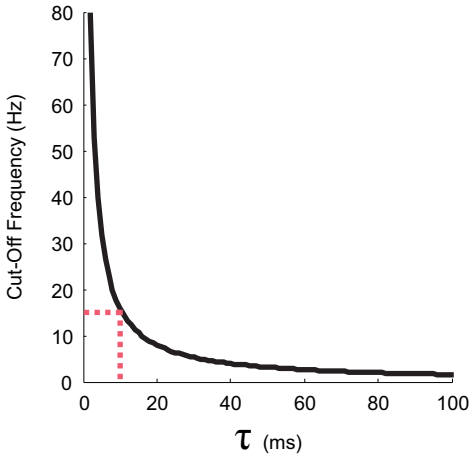

Figure 2.28: Time constant versus cut-off frequency.

Plot shows the time constant (τ) in ms of a linear system versus the predicted cut-off frequency in Hz of that system. The pink dotted lines indicate that the cut-off frequency predicted by a time constant of 10 ms is 15 Hz.

Chapter 3:

Rate-Limiting Processes in Visually Evoked Dynamics

3.1 Introduction

Do dynamics of intra-cortical recurrent circuits limit the rate of the visually evoked response in visual cortex? In Chapter 2, I measured dynamics of intra-cortical recurrent circuits as the decay of visually evoked activity in these intra-cortical recurrent circuits upon optogenetically silencing thalamus. This decay is described by the cortical decay function (CDF) and fit by single exponential decay with a time constant of ~10 ms.

To test whether this CDF is the primary constraint[*] on the cortical transformation of visually evoked activity in time, I needed to compare the temporal structure of sensory activity entering the cortex with the temporal structure of the cortical spiking response. The CDF predicts that visual cortex should follow fluctuations of sensory-evoked thalamic activity up to ~15 Hz but increasingly attenuate frequencies above ~15 Hz, filtering out more than two-thirds of the thalamic signal's amplitude at temporal frequencies above 30 Hz.

[*] A theoretical point: The CDF is a measure of the cortical response to a step change in thalamic input (i.e., near-instantaneous offset of thalamic input). Therefore we can consider the CDF a measurement of the impulse response function (IRF) of cortex. The IRF describes how a system responds to an instantaneous input. In linear systems theory, the IRF can be used to predict the system's response to any arbitrary input using the convolution. Therefore the IRF is an extremely powerful and fundamental property of a linear dynamic system. Whether the CDF provides similar insight into cortical dynamics is a question that must be addressed empirically.

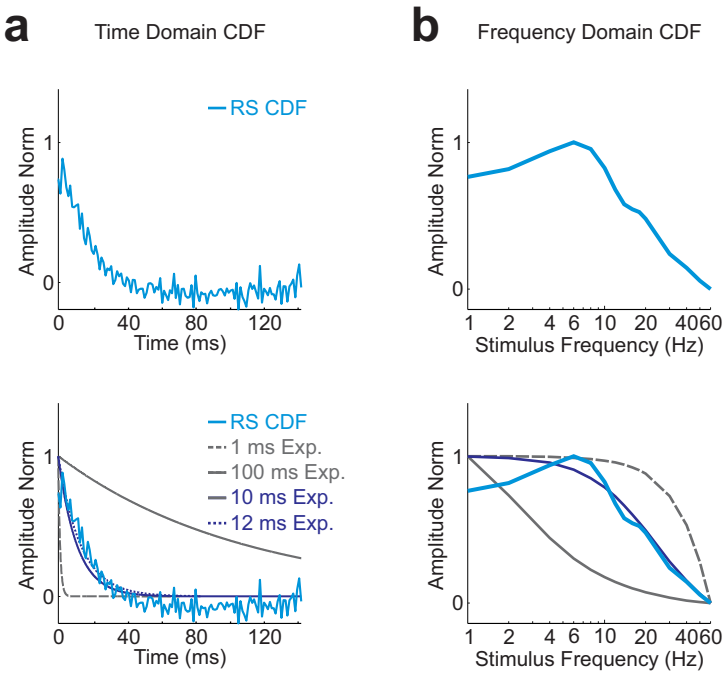

Figure 3.1: Filtering of input frequencies predicted by the cortical decay function (CDF).

(a) CDF in time domain.
Top: Cyan: CDF of regular-spiking (RS) units.
Bottom: CDF of RS units (cyan) with single exponential fits. Dark blue=10 ms time constant, dotted blue=12 ms time constant, solid gray=100 ms time constant, dotted gray=1 ms time constant. 1 and 100 ms time constants shown for comparison.

(b) CDF represented in the frequency domain.
Top: As a function of input frequency, amplitude spectrum (Fourier transform) of cyan RS CDF in (a). This cyan filter is used to predict V1's response.
Bottom: Colors as in (a, bottom) but curves are amplitude spectra in frequency domain (all scaled).

3.2 Methods: Measuring the transformation of sensory activity by cortical recurrent circuits

First, I compared spiking activity in the dLGN in response to visual stimuli with the spiking response in V1. Second, I measured the output of thalamocortical synapses in V1 and compared this signal with the spiking response in V1. These comparisons provided data to address whether the CDF can explain how cortex transforms its sensory input in time.

If the CDF captures the dynamics of the response in visual cortex to fluctuations of thalamic input, the CDF should precisely predict the spiking response of cortex, given the thalamic input[*]. To understand what the CDF predicts about how cortex filters sensory input from thalamus, I plotted the Fourier transform of the CDF, which indicates how the CDF is expected to transform the amplitude of sensory input at different temporal frequencies (Figure 3.1). As the CDF is a precise measurement in the time domain, the Fourier transform of the CDF (called the amplitude spectrum of the CDF) is a precise function in the temporal frequency domain. The amplitude spectrum of the CDF predicted that visual cortex should precisely follow rapid fluctuations of sensory-evoked thalamic activity up to ~15 Hz but increasingly attenuate frequencies above ~15 Hz.

To test this prediction, I used a visual stimulus set to map responses in dLGN and V1 to various temporal frequencies of sensory input (between 1 and 60 Hz). In order to compare data from recording sites in dLGN with data from non-retinotopically aligned recording sites in V1, the presented visual stimuli were full-field and spatially uniform. Full-field luminance was modulated as a sinusoidal function of time, called the visual flicker stimulus.

I measured the amplitude of the response in each brain structure at various temporal frequencies to determine that brain structure's frequency response. The amplitude of the response at the

[*] In a linear system, the decay function, often called the impulse response function (IRF), exactly predicts the system's output. This is the definition of a linear system. We cannot assume that the cortex is a linear system. However, evidence suggests that a variety of temporal features of the network-level response in early sensory cortical areas can be well-approximated using linear models [2,5,31,133,134 but see 135]. This may be because linearity preserves information [5], and early sensory cortical areas need to preserve information about the temporal structure of the visual stimulus. In any case, within this regime of approximately linear cortical operation, it may be that the CDF provides a useful first description of dynamics intrinsic to visual cortex. But this must be tested empirically.

temporal frequency of the stimulus represents how well the neural activity follows the stimulus. To separate the stimulus-driven response from non-stimulus-driven oscillations arising internally within the brain, I averaged the neural activity across trials. This averaging eliminates oscillations that are not phase-locked to the presentation of the visual stimulus.

As discussed in Chapter 1, thalamocortical synapses may exhibit their own dynamic properties, potentially transforming the frequency response measured as extracellular spiking at the cell bodies of thalamic relay neurons. Dynamics of intra-cortical recurrent circuits act downstream of the output of thalamocortical synapses. Thus, to test whether the CDF accounts for the effect of intra-cortical recurrent circuits on sensory input, the optimal approach is to compare the output of thalamocortical synapses, rather than spiking measured in dLGN, with the spiking response of visual cortex.

To determine the output of thalamocortical synapses, I recorded field potentials in layer 4 of visual cortex while silencing excitatory intra-cortical recurrent circuits (Figure 3.2), as in previous studies [30-32] and as in Chapter 2 (Figure 2.1a). Silencing excitatory intra-cortical connections leaves intact the thalamocortical projection as the source of fields in layer 4 of visual cortex. Hence, the visually evoked field measured in layer 4 is dominated by the depolarization of layer 4 neurons resulting from synaptic output of thalamocortical afferents.

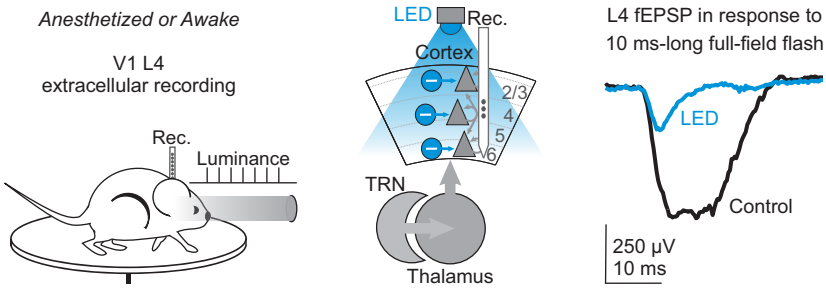

Figure 3.2: Isolating the output of thalamocortical synapses in V1 by silencing cortical excitatory recurrent circuits.

Left and Middle: Schematic of experiment to measure thalamocortical synaptic depression *in vivo* by recording in layer 4 of V1 while silencing cortical excitatory circuits (LED photo-activation of cortical inhibitory interneurons). Visual stimulus is a 4.5 Hz pulse train of 10 ms-long flashes of light.

Right: Example field potential response to flash of light in control conditions (black) and during silencing of cortical excitatory circuits (LED, blue) on interleaved trials. Note that field potential response during cortical silencing is smaller and transient.

3.3 Results

3.2.1 Depression of thalamic afferent synapses limits the rate of the visual response under anesthesia

What are the implications of such a consistent ~10 ms CDF for how cortical recurrent circuits impact the timing of sensory activity in visual cortex? If this fast CDF sets the rate of the cortical sensory response, then cortex should precisely follow fast fluctuations in sensory activity up to ~15 Hz but increasingly attenuate frequencies above ~15 Hz, the frequency cut-off predicted by the filtering properties of the ~10 ms CDF (Figure 3.1).

I determined V1's ability to follow the timing of sensory activity from thalamus. I measured the spike rate response to spatially uniform, full-field fluctuations of luminance (flicker) as a function of flicker temporal frequency (between 1 and 60 Hz), called the frequency response, in dLGN[*] and V1 (Figure 3.3). In contrast to the prediction by the CDF, however, under anesthesia V1 could not follow visually evoked thalamic activity at frequencies above ~6 Hz (Figure 3.3b-c, left). This low-pass filtering of the sensory response by V1 is consistent with previous measurements of V1's frequency response [49] but implies a rate-limiting step between thalamus and cortex dramatically slower than the ~10 ms CDF. This low-pass filtering was already pronounced in layer 4 (Figure 3.4b, top), V1's main thalamo-recipient layer, indicating that the filtering of the sensory response occurs very early as visually evoked activity enters the cortex.

A candidate source of low-pass filtering of the sensory response under anesthesia is frequency-dependent depression of thalamocortical (TC) afferent synapses [19-21,113,114], a progressive depression of synaptic strength as the frequency of thalamic activity increases. If present, this process would preferentially attenuate the cortical response to high-frequency thalamic activity. I measured TC depression in my preparation *in vivo* under anesthesia (Figures 3.2, 3.3d). To do this, I optogenetically

[*] Most dLGN single units showed band-pass responses to different temporal frequencies of the visual stimuli, consistent with previous work [58,136,137]. Here I combine the responses across all single units to obtain the typical response of dLGN (median response) at each temporal frequency, a summary of dLGN activity. See Figure 3.4 for the full distribution of responses in dLGN at each temporal frequency of the visual stimulus.

Figure 3.3: Relative contributions of thalamocortical synaptic depression and recurrent circuit dynamics to V1's response in anesthetized versus awake mice.

(a) Experimental configuration. Left: Presentation of visual stimulus (i.e., full-field luminance flicker) to awake mice. Right: Extracellular recordings in dLGN or V1.

(b) Example peri-stimulus time histograms (PSTH) of multi-unit (MU) activity in dLGN (top) or V1 (bottom) in response to flicker stimuli of different temporal frequencies. **Right column:** Under anesthesia. **Left column:** In awake mouse. Black is real neural response. Pink (bottom) is prediction of V1's response by convolving the CDF with dLGN response (from top) at each stimulus frequency. Note poor prediction of V1's response under anesthesia (left) but good prediction in awake mice (right).

(c) Frequency response, i.e., response amplitude as a function of stimulus frequency. Green: dLGN response (as median single-unit amplitude; error bars are 45th to 55th percentiles; for full distributions, see Figure 3.4c). Black: V1 response (as mean MU amplitude; error bars are s.e. across mice). Pink: Prediction of V1 response using CDF. (For curve calculation, normalization, and alignment, see "Frequency response", pg. 131.) **Left:** In anesthetized mice (dLGN SU n=131). **Right:** In awake mice (dLGN SU n=80). Note poor prediction of V1's frequency response (pink, compare with real V1 frequency response in black) under anesthesia (left) but good prediction in awake mice (right).

(d) Experimental configuration. **Left:** Visual stimulus train of full-field flashes presented at various temporal frequencies. **Right:** Silencing cortical excitatory recurrent circuits (as in Fig. 1b) to isolate the field excitatory post-synaptic potential of thalamic afferents (Thal. fEPSP) detected by extracellular recordings in layer 4 of V1.

(e) Top: Thalamic fEPSP (thal. fEPSP, each curve is average of 30 sweeps) recorded in layer 4 in response to repeated visual stimulation at 4.5 Hz (each stim. is 10 ms in duration) in example anesthetized (left) then awake (right) mouse (same mouse and recording site across states). Note depression of response only under anesthesia.
Center: Summary graph of thal. fEPSP amplitude versus stimulus number. Amplitudes are normalized to first response and spike rate adaptation in thalamus (see "Measuring thalamocortical synaptic depression", pg. 130). Error bars are s.e. across mice.
Bottom: Mean, steady-state thal. fEPSP amplitude (see "Measuring thalamocortical synaptic depression", pg. 130) as a function of visual stimulus frequency (n=4 mice, same 4 mice across brain states). Note strong frequency dependence of thal. fEPSP amplitude under anesthesia but no frequency dependence in awake mice.

(f) Thal. fEPSP versus V1's frequency response under anesthesia. Blue: Steady-state thal. fEPSP from (e) under anesthesia. Black: V1 frequency response from (c) under anesthesia. Curves are baseline-subtracted and normalized. Note similar attenuation with increasing stimulus frequency.

(g) Predicting V1's frequency response. Black and pink from (c) in awake mice. Gray curves show failed predictions of V1's response using fit 10 times faster (1 ms) or slower (100 ms) than correct fit to CDF (10 ms).

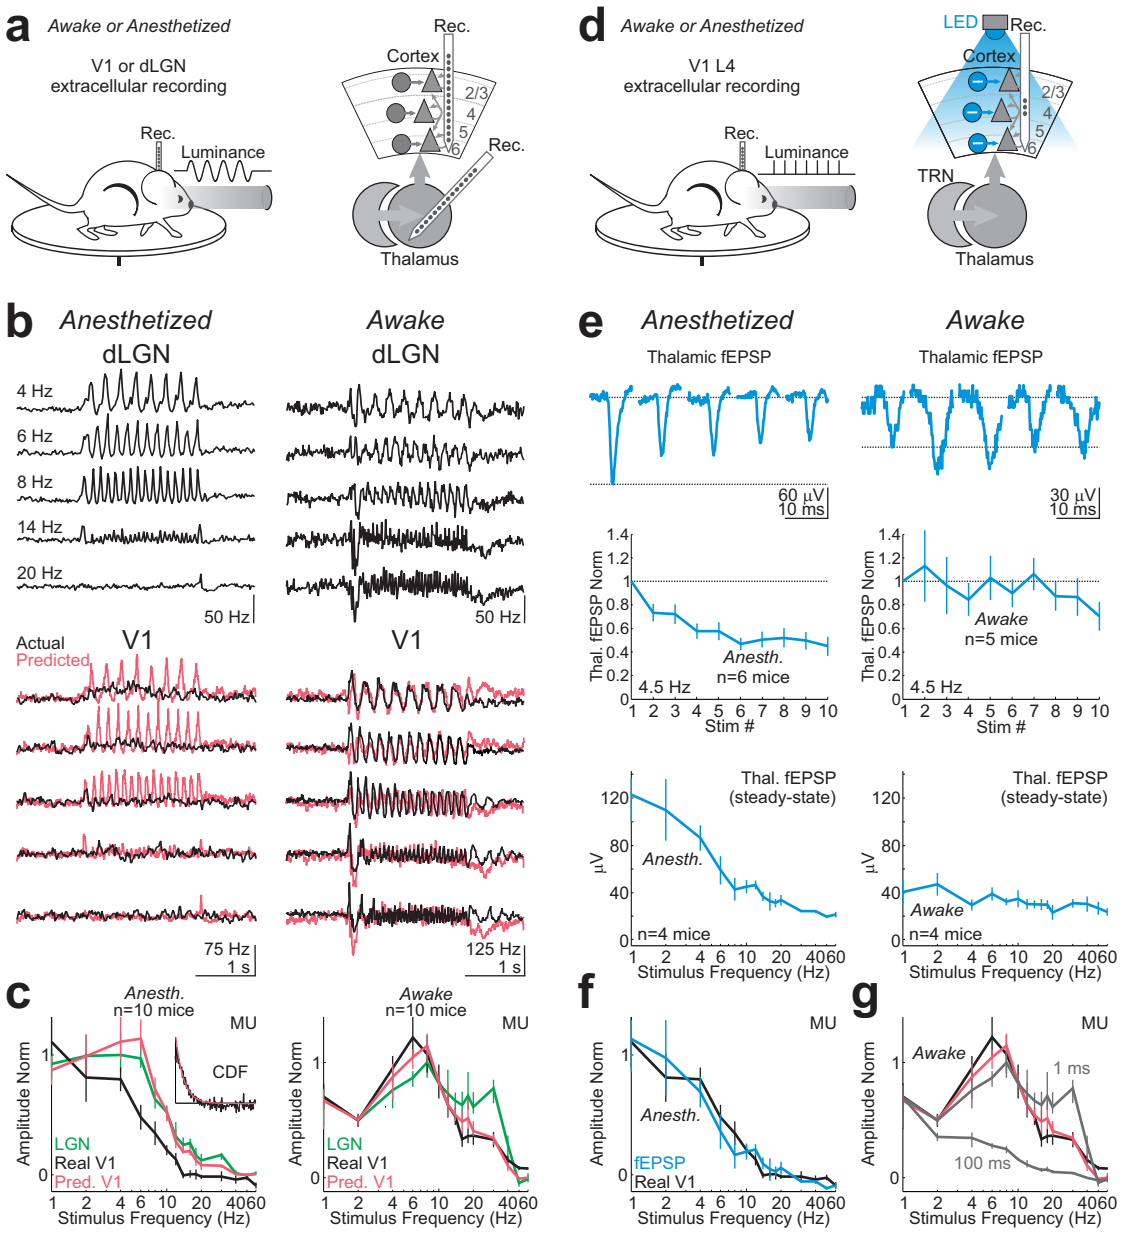
a
Awake or Anesthetized
V1 or dLGN
extracellular recording
Rec.
Luminance
Cortex
2/3
4
5
6
Thalamus
d
Awake or Anesthetized
V1 L4
extracellular recording
LED
TRN
b
Anesthetized
Awake
dLGN
4 Hz
6 Hz
8 Hz
14 Hz
20 Hz
50 Hz
Actual
Predicted
V1
75 Hz
1 s
125 Hz
e
Thalamic fEPSP
60 µV
10 ms
30 µV
Thal. fEPSP Norm
Anesth.
n=6 mice
Awake
n=5 mice
4.5 Hz
Stim #
Thal. fEPSP
(steady-state)
µV
n=4 mice
Stimulus Frequency (Hz)
c
Anesth.
n=10 mice
MU
CDF
Amplitude Norm
LGN
Real V1
Pred. V1
f
fEPSP
g
1 ms
100 ms

silenced cortical recurrent excitatory circuits, as in Figure 2.1a, while recording field excitatory post-synaptic potentials (fEPSPs) in layer 4, which under these conditions represent TC transmission (Figures 3.2, 2.1a and described previously [30]). These TC fEPSPs showed strong frequency-dependent depression, quantified as the decrease in response amplitude over a train of visual stimuli (each stimulus: flash of 10 ms duration, train of 10 stimuli) relative to the amplitude of the first response (Figure 3.3e, left; normalized by the response amplitude in thalamus, see Appendix A: Supplementary Methods, "Measuring thalamocortical synaptic depression *in vivo*", pg. 130, for details). Notably, the measured frequency dependence of the TC fEPSPs accounted quantitatively for V1's low-pass filtering of thalamic input under anesthesia (Figure 3.3f).

If this TC depression underlies the low-pass filtering of V1's response under anesthesia, conditions reducing TC depression should improve V1's ability to follow high-frequency thalamic input. I observed that frequency-dependent TC depression in layer 4 disappeared as animals woke up from anesthesia (Figure 3.3e, right; for a potential mechanism, see section 3.2.4 below), consistent with previous suggestions that TC depression depends on brain state [19-21,51,113,136]. Thus, if TC depression impairs V1's ability to follow high-frequency stimulation under anesthesia, then in awake mice, relief from TC depression should enable V1 to better follow high-frequency thalamic activity, up to the cut-off predicted by the CDF.

3.2.2 Cortical decay function (CDF) predicts the amplitude of the visual response in awake mice

Waking up the animals led to a dramatically stronger V1 response to thalamic activity at temporal frequencies above ~6 Hz (Figure 3.3b-c, right). Furthermore, consistent with my previous measurements of recurrent circuit dynamics in visual cortex, the response attenuation in V1 in awake mice precisely matched the attenuation between thalamus and cortex predicted by the CDF (Figure 3.3b-c, right). Combining (convolving) the CDF with dLGN's frequency response gave an excellent prediction of V1's frequency response measured in terms of amplitude during both high (100%) and low (20%) contrast visual flicker (Figure 3.3c, right, and Figure 3.5; for additional filtering properties of

CDF, including subtle theta-band amplification [138], also observed in cats [62], see Figure 3.6 and section 3.2.5 below). In awake mice, the cortex exhibited increasing attenuation of sensory-evoked thalamic activity above ~15 Hz, attenuating two-thirds of the amplitude of the thalamic input at 30 Hz, in accordance with the CDF prediction. Critically, combining dLGN's frequency response with a time constant ten times faster or slower than the 10 ms fit to the CDF failed to predict V1's frequency response (Figure 3.3g).

Thus, in awake mice, the dynamics of recurrent circuits predict how V1 attenuates the amplitude of a temporally modulated thalamic input. But what fraction of V1's total response in awake mice is actually explained by this simple model of input attenuation? To address this precisely, I performed dual simultaneous extracellular recordings in dLGN and V1 to measure response coherence between these structures. I found that dLGN's average multi-unit response convolved with the CDF explained 53% of V1's average multi-unit response between 1 and 30 Hz (Figure 3.7, see section 3.2.6 below).

Furthermore, both fundamental (1X temporal frequency of stimulus, 65% of total V1 response amplitude) and first harmonic (2X temporal frequency of stimulus, 35% of total) components of the visual response (Figure 3.8) were attenuated at high frequencies in V1 regular-spiking (RS) units in a manner quantitatively consistent with filtering by the CDF (Figure 3.9). Thus, these data suggest that, in awake mice and therefore in the absence of low-pass filtering by thalamic afferent synapses, dynamics of cortical recurrent circuits are the key constraint on V1's response to sensory activity from thalamus.

As further validation of how brain state affects the temporal transformation of sensory activity between thalamus and cortex, deconvolution of dLGN's spiking response to the flicker stimulus from V1's response, an analysis that measures the combined effects of both TC afferent synaptic dynamics and cortical recurrent circuit dynamics, showed slow dynamics under anesthesia but faster dynamics in the awake state (Figure 3.10).

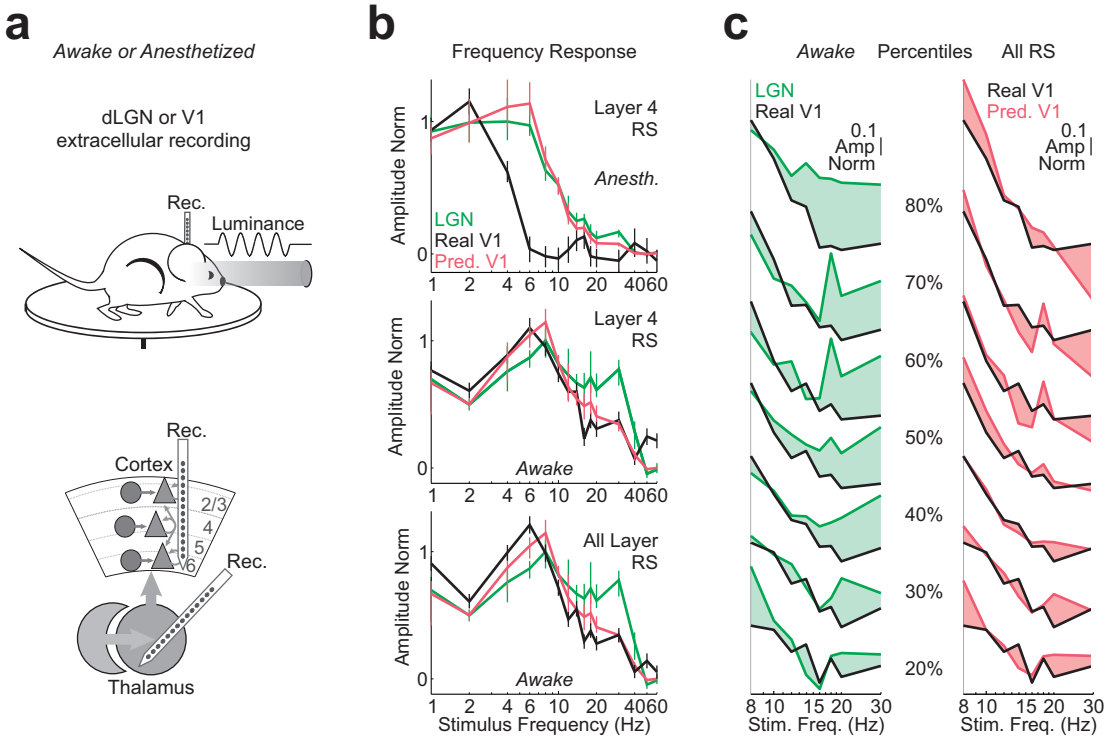

Figure 3.4: Cortical decay function (CDF) predicts low-pass filtering of responses in V1 in awake, but not anesthetized, mice in layer 4 and across cortical layers.

(a) Experimental set-up to measure frequency response by recording in either dLGN or V1.

(b) Top: Anesthetized. Frequency response under anesthesia of regular-spiking (RS) units in Layer 4 (black; n=30 units; mean and s.e.). dLGN (green) and predicted V1 (pink) frequency response from left panel of Figure 3.3c.
Middle: Awake. Frequency response (n=30 units) of RS units in Layer 4 (black; n=30 units; mean and s.e.). dLGN (green) and predicted V1 (pink) frequency response from right panel of Figure 3.3c.
Bottom: Awake. Frequency response of RS throughout all layers (black; n=270 units; mean and s.e.). dLGN (green) and predicted V1 (pink) frequency response from right panel of Figure 3.3.

(c) Attenuation predicted by CDF over full response amplitude distribution. Frequency response shown as percentiles of single-unit (SU) amplitudes (peak-normalized, see "Frequency response", pg. 131, for alignment, normalization and scaling).
Left: Percentiles of dLGN frequency response (green) versus same percentiles of V1 RS SU frequency response (black). Note that V1 amp. is below dLGN amp. at high frequencies for all percentiles of response amplitude distributions.
Right: Black as left. Pink is percentiles of prediction for V1 by convolving dLGN frequency response with CDF. Shaded areas are differences between curves. At all percentiles, prediction for V1 after filtering by CDF is much better match to real V1 response than dLGN response without filtering by CDF.

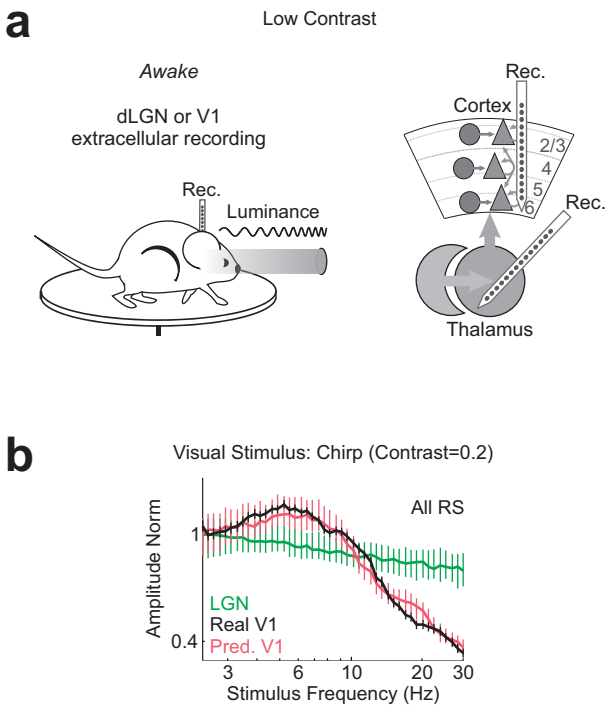

Figure 3.5: CDF predicts V1's filtering of thalamic input in response to low contrast visual flicker in awake mice.

(a) Experimental set-up to measure frequency response by recording simultaneously in dLGN and V1. Visual stimuli are low contrast frequency-modulated sweeps (chirps, logarithmic change in frequency over time).

(b) Frequency response of dLGN and V1 to low contrast (20%) frequency chirps. Green=median frequency response of dLGN units and 45th to 55th percentile error bars, as in Figure 3.3c. Black=mean, s.e. of regular-spiking (RS) cells across layers. Pink=prediction for V1, i.e., green filtered by CDF. Zero is noise (for alignment and scaling, see "Frequency response", pg. 131).

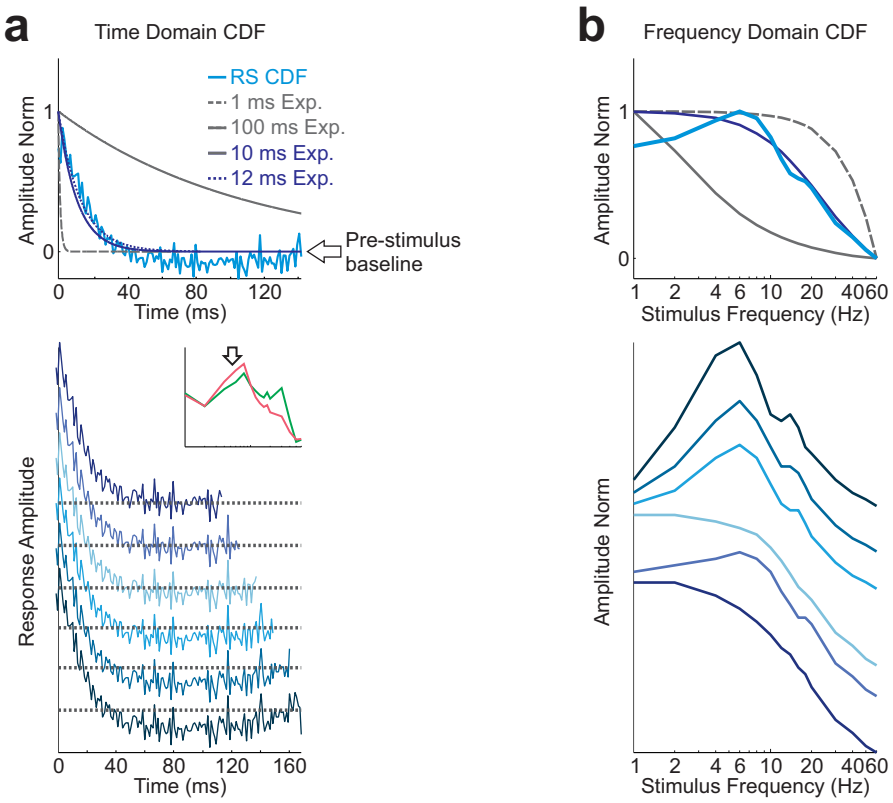

Figure 3.6: Theta-band amplification by cortical decay function (CDF).

The slow, small recovery component of the CDF (dip below baseline, see Figure 2.9c) predicts theta-band amplification in V1's frequency response.

(a) CDF in time domain. **Top:** Cyan: CDF of regular-spiking (RS) units. Curve is zeroed at pre-stimulus baseline activity (arrow). Note both fast shut-off, fit by 12 ms decaying single exponential, and second small, slow recovery component. Dark blue, dotted blue, solid gray, dotted gray: single exponential decays with time constants as indicated; 1 and 100 ms time constants shown for comparison. **Bottom:** Magnitude of slow recovery component as a function of where the baseline of CDF is set. The higher the baseline (dotted gray line), the more pronounced the slow recovery component of CDF (from top to bottom). In this study, I choose the baseline of the CDF (under anesthesia) to be the mean pre-stimulus baseline, because thalamic silencing does not affect spontaneous activity under anesthesia. However, in each experiment, the CDF relaxes to a baseline that depends on the strength of thalamic silencing; thus, the baseline of the CDF may be offset with respect to the pre-stimulus baseline. Curves are offset on ordinate for clarity. Inset: green and pink from (Figure 3.3, right). Arrow shows CDF amplification of stimulus frequencies in theta band.

(b) CDF represented in the frequency domain. **Top:** Cyan: amplitude frequency spectrum of cyan curve in (a). This cyan filter is used to predict V1's response. Colors as in (a) but spectra are in frequency domain (all scaled). **Bottom:** Frequency-domain forms of CDFs in bottom panel of (a) (colors are matched to bottom panel of (a)). Increasing the magnitude of the slow recovery component of CDF increases theta-band amplification (i.e., the 4-11 Hz component of the frequency-domain form of the CDF).

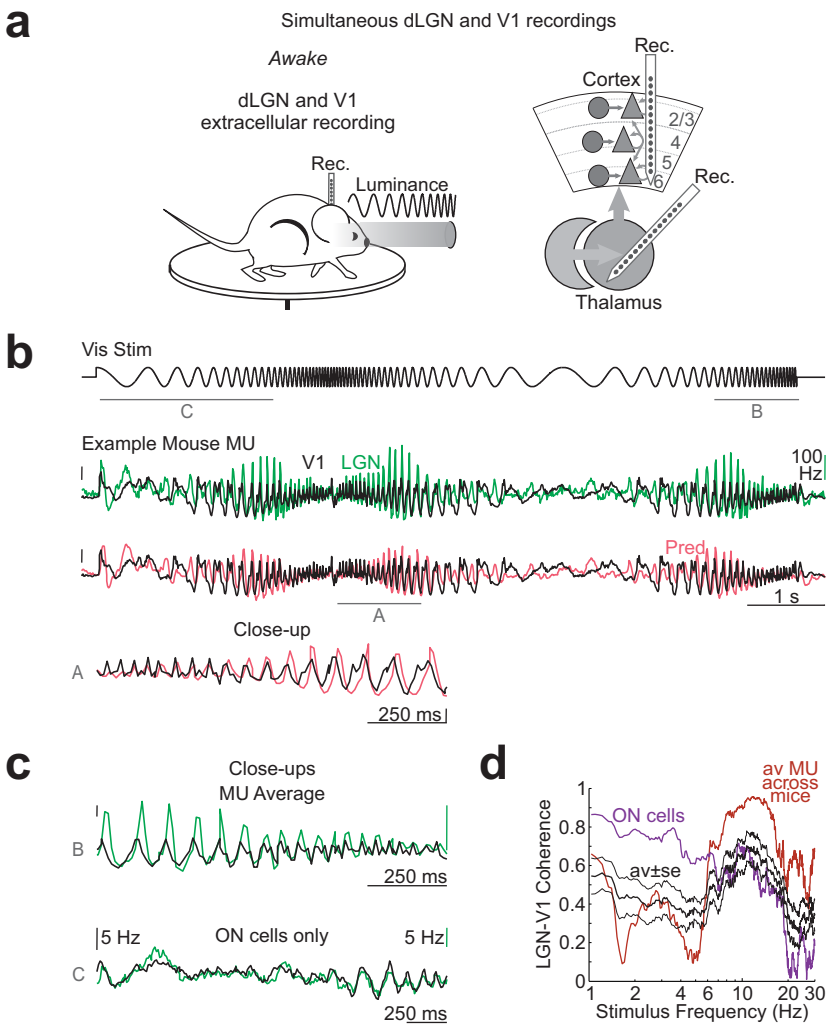

Figure 3.7: High dLGN-V1 coherence in response to visual flicker measured by simultaneous dLGN and V1 recordings.

(a) Experimental set-up to measure frequency response recording simultaneously in dLGN and V1.

(b) Unless specified, vertical scale bars in (b-c) are 100 Hz. **Top:** Vis. stim. as logarithmic frequency modulation of full-field luminance over time (chirp). **2nd Row:** Superimposed MU PSTHs from dLGN (green) and V1 (black) of example mouse in response to vis. stimulus. **3rd Row:** V1 response prediction (pink; by convolution of dLGN response with CDF) superimposed on actual V1 response (black). Note that pink predicts extent of amplitude filtering at higher frequencies. Close-up: Expanded timescale of segment A (above).

(c) Averaging MU across mice (**top**) or including only ON-responsive units (**bottom**) improves coherence between dLGN and V1 traces. Responses are to chirp segments marked B and C in (b).

(d) Coherence between dLGN and V1 recorded simultaneously. Black: mean and s.e. of coherence (between mean MU dLGN vs. V1 within each mouse) across mice. Red: Coherence between averaged-across-mice MU responses in dLGN and V1. Purple: Coherence between On-responsive dLGN and V1 units in example mouse.

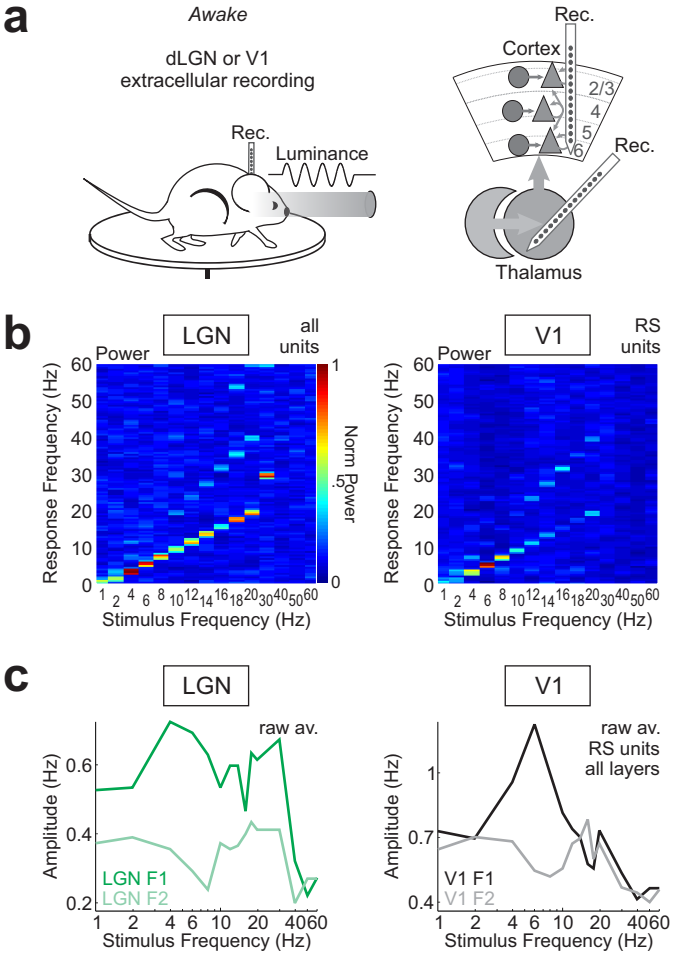

Figure 3.8: Fundamental (F1) and first harmonic (F2) components of frequency response in dLGN and V1.

(a) Experimental set-up to measure frequency response in awake mice recording in either dLGN or V1.

(b) Left: Heatmap is median power of dLGN single-unit response to flicker stimulus in awake mice as a function of stimulus frequency and response frequency. F1 response is the diagonal where response frequency matches stimulus frequency. F2 is the diagonal where response frequency is twice the stimulus frequency.
Right: Same as left but for V1 regular-spiking (RS) units.

(c) Average amplitude of F1 and F2 response components from heatmaps in (a). Average includes all units, including units at each stimulus frequency that are or are not well-driven by that stimulus frequency; thus, the low spike rate.
Left: For dLGN relay units.
Right: For V1 RS units. Consistent with data from cats [62], V1's F1 response dips between 10 and 20 Hz, while the F2 response dips around a 8 Hz.

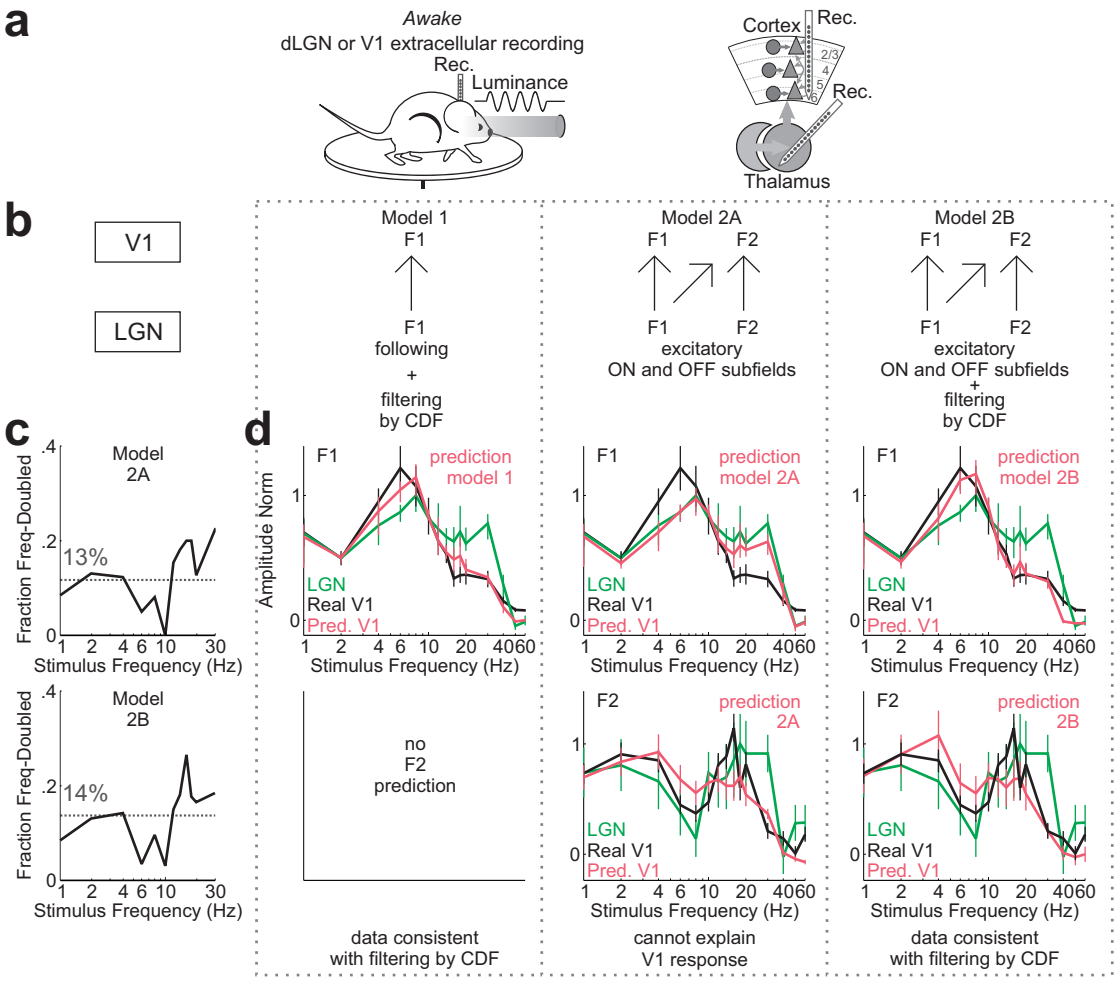

Figure 3.9: Both fundamental (F1) and first harmonic (F2) components of V1 response are consistent with filtering by Cortical Decay Function (CDF).
This figure rules out an alternate model where responses at the fundamental (F1) frequency in dLGN are converted to first harmonics (F2) in V1 without filtering by the CDF. I find that both F1 and F2 components in V1 are attenuated, consistent with filtering by the CDF. A transformation from F1 in dLGN to F2 in V1 should instead increase F2 in V1.
(a) Experimental set-up to measure frequency response in awake mice recording in either dLGN or V1.
(b) Simple models of dLGN-to-V1 connectivity affecting spatial receptive field structures. **Model 1:** F1 response in dLGN drives F1 response in V1. No consideration of F2. This is the model presented in Figure 3.3. **Model 2A:** F1 response in dLGN drives some F1 in V1 and some F2 in V1. F1-to-F2 transformation between dLGN and V1 is consistent with certain models of visual system (see "Models of F1...", pg. 132). Also, F2 response in dLGN drives some F2 in V1. In Model 2A, *no amplitude filtering by CDF*. **Model 2B:** Same as 2A, but 2B *does* include amplitude filtering by CDF. See pg. 133 for associated equations.
(c) Fraction of F1 response in dLGN that frequency-doubles in cortex to produce an F2 response in V1, derived from Model 2A (**top**) or 2B (**bottom**). See pg. 133 for associated equations and derivation. Y axis is the fraction of response that is frequency-doubled as a function of stimulus frequency, according to best fit models. Above 30 Hz enters noise (so not shown). Dotted gray is average across stimulus frequencies.
(d) Different predictions for V1 F1 and F2 response components by Models 1, 2A and 2B. **Left:** Prediction by Model 1, as in Figure 3.3. **Middle:** Best prediction for both F1 and F2 components of V1 response according to Model 2A. Prediction fails to match actual V1 F1 response. **Right:** Best prediction for both F1 and F2 components of V1 response according to Model 2B. By including filtering by the CDF in this model, this model can better account for the average amplitude of both the F1 and F2 response components in V1 RS cells.

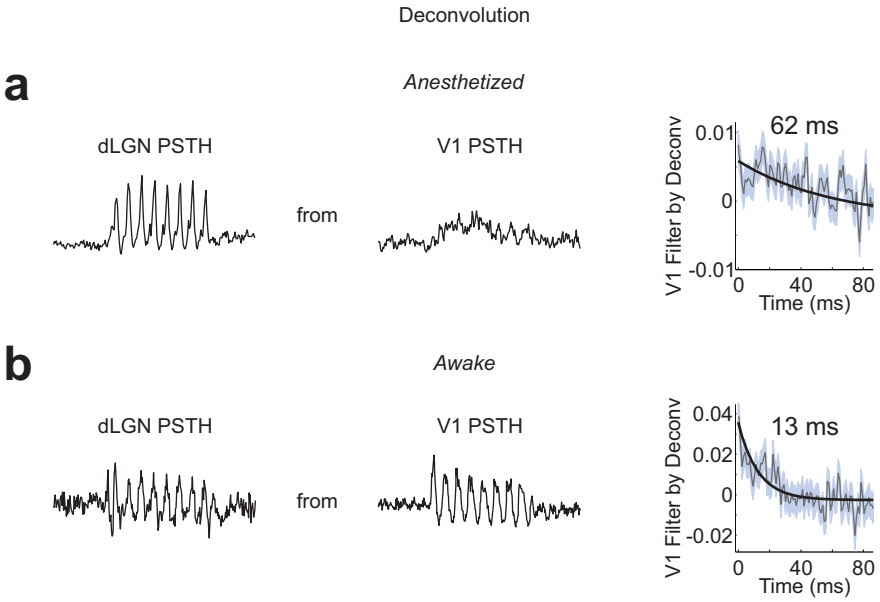

Figure 3.10: Deconvolution of dLGN response from V1 response shows that V1 response dynamics are slower under anesthesia than in awake state.

Deconvolution of dLGN PSTH response to flicker stimulus from V1 PSTH response to same stimulus (across all temporal frequencies, Methods) in anesthetized **(a)** and awake **(b)** mice.
Left: Example trial-averaged PSTH responses to 4 Hz visual flicker.
Right: Result of deconvolution. Gray line: mean deconvolution across all thalamus-V1 pairs with s.e. (gray-blue error). Black: single exponential fit and time constant value. Note slower dynamics under anesthesia as compared to awake.

(a) Anesthetized mice.

(b) Awake mice.

3.2.3 Pre-depression of thalamocortical synapses in the awake state

I measured TC depression in my preparation *in vivo* as the decrease in magnitude of the evoked response (field potential in layer 4 during cortical silencing) over a train of brief visual flashes (10 ms duration of each flash). Therefore TC depression here, as in the literature [19,20], is quantified with respect to the magnitude of the first evoked response. TC depression measured in this way disappeared as the animals woke up from anesthesia.

However, the magnitude of the first response also decreased as animals woke up from anesthesia (Figure 3.3e). The magnitude of the final, fully depressed response under anesthesia was approximately equivalent to the magnitude of the first response in the awake state. This suggests that TC synapses might already be depressed (i.e., are pre-depressed) in the awake state, even prior to the onset of visual stimulation.

3.2.4 Increase in spontaneous thalamic activity in the awake state

Consistent with this, I found that spontaneous activity in the thalamic relay cells increased dramatically, by more than a factor of five, as animals woke up from isoflurane anesthesia (Figure 3.11), simultaneous with the loss of TC depression. This increase of spontaneous activity in thalamus was pronounced in both dLGN and LP. This suggests, although it does not prove, that increased spontaneous activity in the awake state leads to a pre-depression of TC synapses prior to visual stimulation and, thus, to the loss of low-pass filtering by TC synapses in the awake state.

Furthermore, this change in spontaneous thalamic activity in the awake state was quantitatively consistent with a pre-depression of TC synapses in awake mice (Figure 3.12). Specifically, under isoflurane anesthesia, the average spontaneous firing rate of thalamic relay cells was 0.5 Hz, insufficient to depress TC synapses based on measurements *in vivo* (Figure 3.12). However, in the awake state, the average spontaneous firing rate of thalamic relay cells was 5 Hz, sufficient to fully depress TC synapses (Figure 3.12).

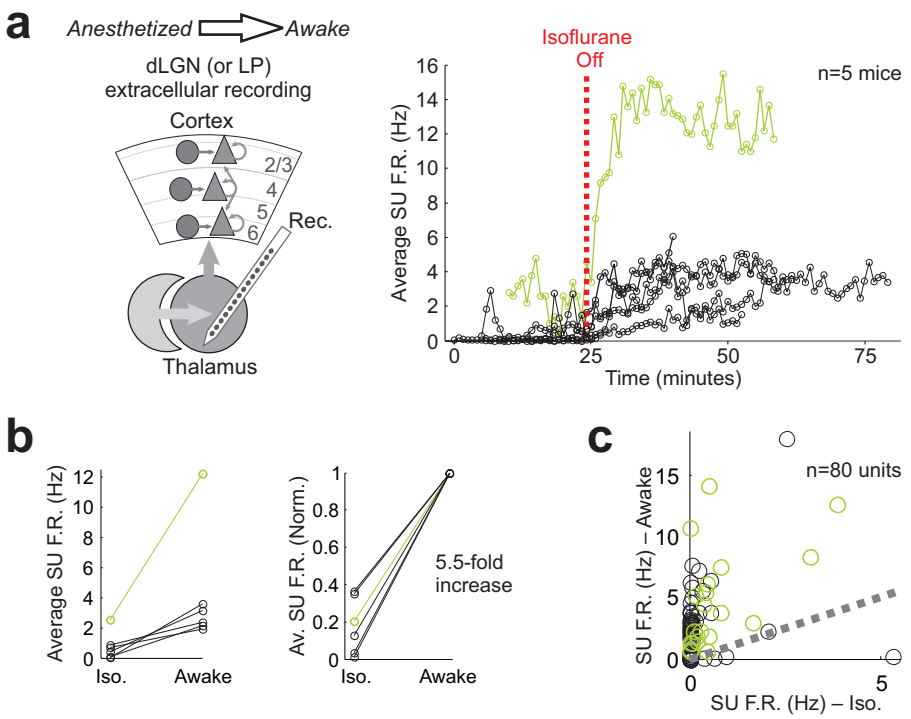

Figure 3.11: Waking up from anesthesia increases baseline spiking in the thalamus.

In all parts, light green: recording site in LP; black: 4 sites in dLGN.

(a) Left: Schematic of recordings in dLGN or LP while mice wake up from isoflurane anesthesia.
Right: Average thalamic single-unit (SU) spontaneous firing rate over time as mice wake up from isoflurane anesthesia (5 curves are 5 example recordings; anesth. average SU rate=0.8 Hz, awake average SU rate=4.2 Hz).

(b) Left: Average spontaneous SU firing rate under anesthesia and awake. **Right:** Normalized to final rate in awake.

(c) Average spontaneous SU firing rate under anesthesia (isoflurane) versus SU average rate in awake state. Dotted gray line is unity. Points above dotted line have higher firing rates in the awake state. Light green dots are units in LP. Black dots are units in dLGN.

As animals awoke from anesthesia, the increase in thalamic spontaneous activity was observed prior to the animal's first movement, suggesting that the increase in thalamic spontaneous activity is not a result of increased movement in the awake state.

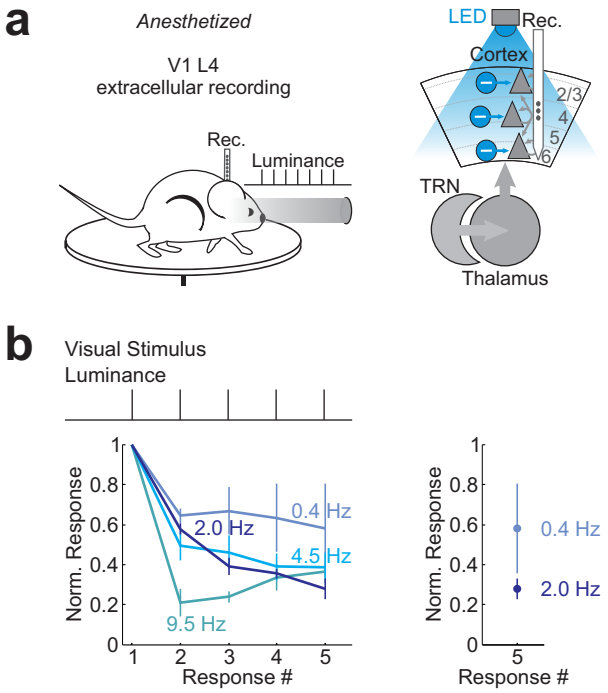

Figure 3.12: Frequency-dependence of thalamocortical synaptic depression measured *in vivo*.

(a) Schematic of experiment to measure thalamocortical synaptic depression *in vivo* by recording in layer 4 of V1 while silencing cortical excitatory circuits (LED photoactivation of cortical inhibitory interneurons). Visual stimulus is a pulse train of 10 ms-long flashes of light.

(b) Left: Mean and s.e. of amplitude of field excitatory post-synaptic potential (fEPSP) in layer 4 during cortical silencing as a function of visual stimulus pulse train frequency and response #. Note stronger depression in response to higher visual stimulus frequencies. All curves are normalized to amplitude of first response.
Right: Magnitude of steady-state (#5) fEPSP response for low versus high frequencies of activity. Note that 2 Hz (and higher, see left) input produces a strong depression of the thalamocortical synapses *in vivo*.

3.2.5 Theta-band amplification by the CDF

The Fourier transform of the CDF (Figure 3.1) shows that the CDF predicts some amplification of the low theta-frequency band (~4-8 Hz), relative to the unfiltered response at 1 Hz and the attenuation of higher frequencies (above ~15 Hz). Consistent with this prediction by the CDF, with findings in cats [62], and with proposed mechanisms for the cortical amplification of theta frequencies [138], I observed subtle theta-band amplification of the response in V1 (Figures 3.3c, right, and 3.6).

3.2.6 Limitations of the measured coherence between dLGN and V1

The 53% coherence of multi-unit activity between dLGN and V1 in response to the visual flicker stimulus (Figure 3.7) is likely an underestimate of the real coherence between connected populations in thalamus and cortex. There are several reasons why this measurement is likely an underestimate. First, I did not record from retinotopically aligned regions of dLGN and V1. Although the visual stimulus was full-field and spatially uniform, it is probable that, in most cases of simultaneous dLGN and V1 recordings, the recorded regions in dLGN and V1 were not synaptically coupled. Second, extracellular recordings sample only a small subset of neurons in a region, providing only a gross estimate of the overall activity. Consistent with these caveats, I found that averaging the multi-unit visually evoked activity across mice led to an increase in the measured dLGN-V1 coherence (in particular, at high frequencies between 6 to 30 Hz, where there are few detected spikes per cycle of the visual stimulus), probably because this across-animal average sampled more cells in each brain structure, better capturing network activity (Figure 3.7d). Moreover, comparing only On-responsive [74] cells in dLGN with only On-responsive cells in cortex further increased the observed dLGN-V1 coherence to over 70% in the temporal frequency range of 1-15 Hz (Figure 3.7d), consistent with multiple parallel pathways between dLGN and V1.

The data therefore suggest that the majority of the amplitude[*] of the network-level visually evoked response in visual cortex to these flicker stimuli can be explained by linear filtering according to

[*] The ~10 ms CDF predicts only a few ms delay in the phase or onset of the cortical response (predicted phase shift computed by Fourier transform of CDF), with respect to the timing of the thalamic response. Simultaneous recordings from synaptically connected dLGN and V1 regions would likely be required to

the CDF. However, this does not imply that the CDF predicts the response of any given single neuron in cortex.

3.2.7 Increases and decreases around baseline during visual responses in awake mice

In anesthetized mice, units in dLGN and V1 that responded to the visual stimuli (moving gratings, Chapter 2, and full-field flicker, Chapter 3) showed only stimulus-evoked *increases* in their spiking activity. However, in awake mice, unit responses to the visual stimuli included both increases above and *decreases* below the pre-stimulus spontaneous baseline (Figure 3.3b). This change in the response profile as a function of brain state appears consistent with the observed low level of thalamic spontaneous activity under isoflurane anesthesia (mean rate: 0.5 Hz), which may be expected to preclude any further decrease of activity during visual stimulation, and consistent with increased spontaneous thalamic activity in the awake state, allowing the expression of visually evoked decreases in activity in dLGN and V1.

3.2.8 Spike-timing represents response of regular-spiking cells in V1

Consistent with this observation that V1 visually evoked responses in the awake state included both increases above and decreases below the pre-stimulus spontaneous baseline, the mean firing rate of V1 cortical regular-spiking (RS) units was not altered dramatically by visual stimulation in awake mice (Figure 3.13), suggesting that increases and decreases in the neural activity cancel each other out over longer timescales. In awake mice, the pre-stimulus spontaneous firing rate of RS units was, on average, 1.8 Hz (Figure 3.13b). During visual stimulation, the average firing rate of RS units was about 1.85 Hz. Yet spiking in these RS cells time-locked to visual flicker up to 20 Hz (Figure 3.3c, right, and Figure 3.13c for example RS cell).

detect this subtle delay as a result of the CDF. Moreover, achieving millisecond-level time resolution for this measurement would require either sampling from a very large number of neurons in each brain structure simultaneously or averaging across mice (but averaging across mice destroys the comparison of simultaneous activity in dLGN and V1). Because these are technical challenges and because the time delay predicted by the CDF is expected to have only a very subtle effect (few ms) on the overall temporal structure of visually evoked activity in dLGN and cortex (over timescales of hundreds of ms), I did not pursue an analysis to detect phase or time delays predicted by the ~10 ms CDF.

This indicates that RS cells did not spike in response to every cycle of >2 Hz visual flicker. However, RS cells were equally likely to spike in response to any given cycle of, for example, 20 Hz visual flicker; thus, the trial-averaged response of each RS cell revealed power at 20 Hz, the frequency of the visual stimulus, although spike times were highly variable across trials.

The mean rate of RS cells during visual stimulation in awake mice was ~1.8 Hz across all visual stimulus temporal frequencies (Figure 3.13b). (However, note that the peak firing rate of RS units could be much higher, e.g., 10 Hz in Figure 3.13c.) This indicates that most putative pyramidal (RS) cells in V1 of awake mice express their response to the full-field flicker stimuli as an alteration of spike-timing rather than a change in firing rate over timescales of a few hundred milliseconds or more.

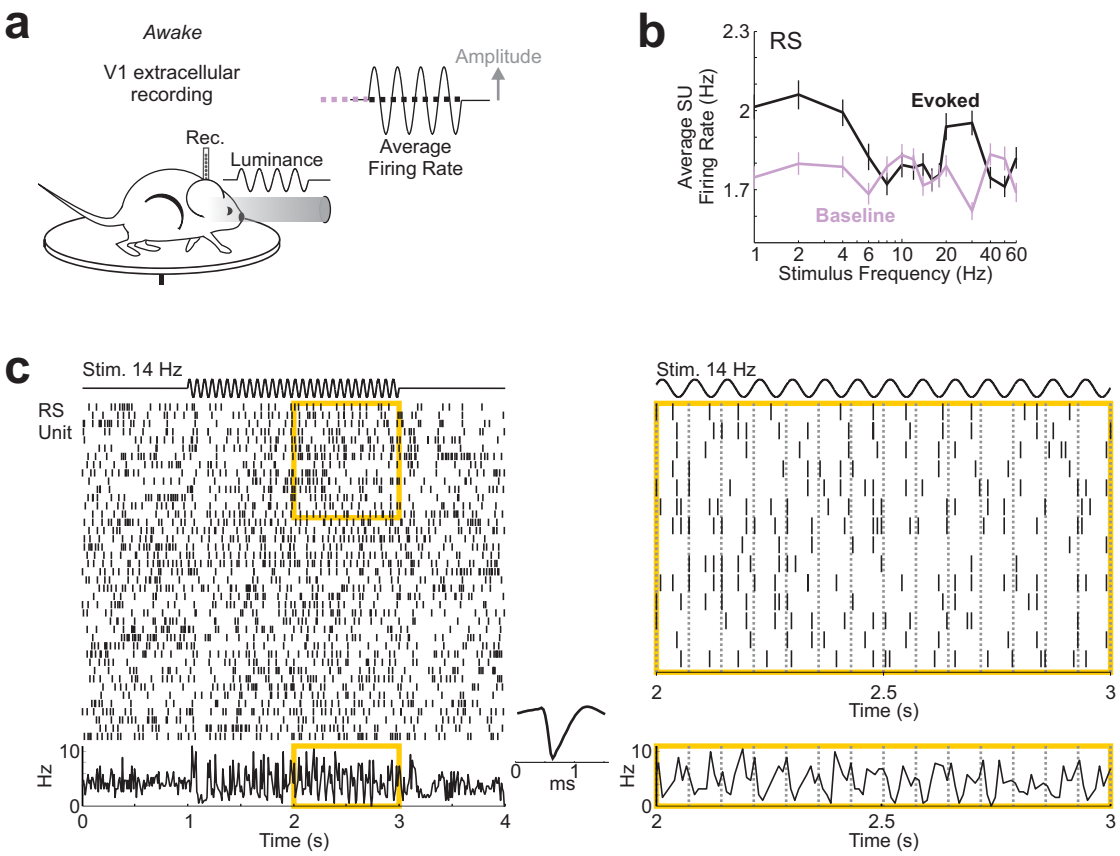

Figure 3.13: Average visually evoked firing rate of regular-spiking (RS) units in V1 during visual flicker does not deviate substantially from pre-stimulus baseline.

(a) Experimental set-up to measure frequency response in awake mice recording in V1. Inset: Computation of baseline (light purple dotted line, pre-stimulus) and visually evoked (black dotted line, mean rate during visual stimulation) firing rates. Amplitude of response indicated by gray arrow. Note that mean visually evoked firing rate may not change although amplitude of response is large.

(b) Average and s.e. of pre-stimulus baseline (light purple) and visually evoked (black) firing rates of regular-spiking (RS) single units (SU, n=80) in V1 in response to visual flicker stimulus at different temporal frequencies.

(c) Response of example RS single unit during 14 Hz visual flicker. Top: visual stimulus. Middle: raster plot. Bottom: PSTH. Inset: unit waveform. Yellow: close-ups.

3.4 Discussion

Thus, in awake mice, the CDF predicts details of how intra-cortical recurrent circuits transform the temporal structure of sensory activity from thalamus. Specifically the CDF predicts that cortical circuits will attenuate the amplitude of the response to frequencies above ~15 Hz, filtering out more than two-thirds of the amplitude of the input signal at frequencies above 30 Hz (Figure 3.1). This prediction matched the frequency response measured in visual cortex (V1) of awake mice (Figure 3.3b-c, right).

While the CDF correctly predicts that V1 will poorly follows input frequencies from the thalamus above ~30 Hz, this does not mean, necessarily, that all >30 Hz activity in the thalamus will fail to drive a response in cortex. It is possible, for instance, that frequencies in the thalamus exceeding ~30 Hz may drive an integrated response in cortex. To understand this, consider a simple example. If dLGN activity is a 50 Hz train of 15 ms-long step pulses, the CDF predicts that V1 will increase its activity in response to each step pulse. The CDF also predicts that, in the 5 ms interval between steps, activity in V1 will decrease, yet V1's activity will not fully return to the pre-stimulus baseline. Hence, V1's response to each subsequent step builds upon previous responses. In this way, V1 activity gradually builds, or integrates, in response to this dLGN input. At the end of the train of thalamic input, V1 activity still shuts off rapidly according to the CDF.

Hence, the CDF predicts a frequency range above which V1 only poorly follows fluctuations of thalamic input, but perhaps the upper cut-off of this frequency range might be more appropriately viewed as a boundary between two qualitatively different modes of thalamocortical communication (Figure 3.14). At thalamic frequencies below 30 Hz, cortex follows its input. At thalamic frequencies above 30 Hz, cortex may integrate its input. However, here I show only that, consistent with the prediction by the CDF, V1 does not follow input from the thalamus above ~30 Hz. Whether V1 also integrates >30 Hz input from thalamus in a manner consistent with the CDF will be a topic for future study. Furthermore, future work will need to address whether the V1 response is consistent with the CDF's prediction given arbitrary time-varying patterns of thalamic input.

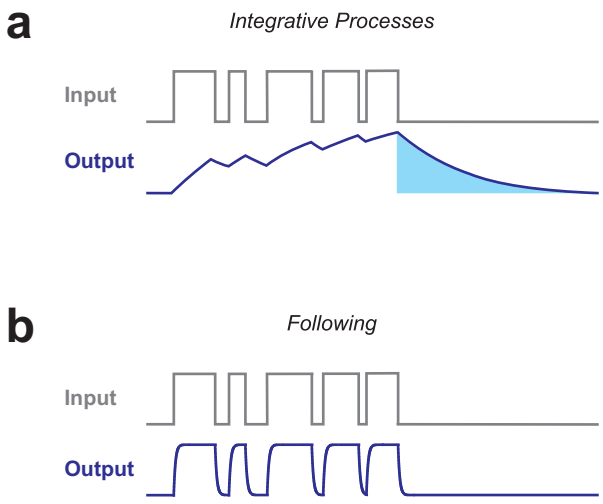

Figure 3.14: Two potential modes of thalamocortical communication.

(a) Cortical output integrates input from thalamus.

(b) Cortical output follows input from thalamus.

The fact that the CDF successfully predicts the response in primary visual cortex, given thalamic input, is not surprising if we consider the results from Chapter 2 showing that the CDF is constant across experimental conditions (Table 2.1) and is even constant in response to partial reductions of thalamic input (Figure 2.14b-c). If the CDF captures the response of cortex to fluctuations of thalamic input across all of these conditions, it should not be surprising that the CDF captures the response of cortex to a visual stimulus composed of fluctuations of thalamic input. Moreover, previous work has indicated that the response in V1 may be approximated by a linear model in various cases [2,5,31,133,134].

It has been suggested that this approximate linearity is the result of proportional amplification of thalamic input by cortical recurrent circuits and that proportional amplification is a computation that preserves information [5]. To best encode a range of levels of thalamic input, without compression or loss of a portion of that input range, every level of the thalamic input should map to a different level of cortical output. Moreover, the separation between levels of thalamic input should be maintained in the cortical output. This describes an approximately linear relationship between thalamic input and cortical output. Similarly, amplification preserves temporal information. Therefore it may not be surprising to find that early sensory areas, like visual cortex, maintain an approximately linear response over a large fraction of the normal physiologic range of thalamic input, as we observe in Figure 3.3b-c in awake mice. However, linearity in terms of average firing rate, as described here, does not imply linearity of the sensory-evoked response along other dimensions. For example, the visual system discards stimulus-specific dimensions of sensory input not relevant to the animal's behavior (e.g., the visual system seems to over-represent local contrast with respect to luminance [139]). Moreover, it is important to note that many other studies have defined regimes where cortical activity deviates from linearity, although these deviations are usually subtle [11,135,140,141]. Therefore it may be that the CDF fails to predict the response in V1 under these conditions, which produce responses in V1 that are non-linear transformations of the sensory input. However, it is also possible that the CDF is a constant, as Chapter 2 suggests, across conditions, and previously reported non-linearities in the V1 response primarily arise from dynamics of interactions of V1 with the thalamus, not from dynamics of the cortical recurrent circuits themselves.

Is the CDF measured here consistent with measurements of the frequency response in other species? In cats and mice, measurements of the frequency response that show strong low-pass filtering of thalamic input by V1 have been made under anesthesia [9-12,14,15,56-58]. As I show, anesthetics may introduce low-pass filtering as a result of thalamocortical synaptic depression. However, in primates, certain measurements of V1's frequency response have been made using sedatives rather than anesthetics (sedatives and anesthetics differ in their mechanisms and physiologic effects). Interestingly, in this case (using sedatives), the observed low-pass filtering in primate V1 is consistent with the CDF measured here in mouse visual cortex. For example, Hawken et al. [9] found that, in sedated monkeys, the visual stimulus temporal frequency eliciting the peak response in dLGN neurons was about 16 Hz and the median high cut-off temporal frequency for responses in dLGN neurons was about 45 Hz. The ~10 ms CDF measured here predicts that V1 will attenuate the amplitude of the sensory response above 30 Hz by more than two-thirds but will attenuate the sensory response at 15 Hz by only about one-fourth, with respect to transmission of the response at 1 Hz (Figure 3.1). Hence, given the measured response in dLGN in Hawken et al., the CDF predicts that V1's high cut-off frequency will be ~30 Hz and that V1's peak response will be just under 15 Hz. Consistent with this prediction by the CDF, Hawken et al. found that the median high cut-off frequency for neurons in primate V1 is ~30 Hz and the median peak response of these neurons is between 12 and 16 Hz. Therefore the CDF measured here seems to account well for the filtering properties of V1 in mice, in the absence of thalamocortical synaptic depression, and also in primates. Moreover, it seems that the differences between the temporal frequency tuning properties of mouse V1 and primate V1 may be inherited from thalamus.

The thalamocortical synaptic depression observed here under anesthesia disappeared in the awake state (Figure 3.3). Other work suggests that non-alert, or inattentive, states in awake animals also produce depression at the thalamocortical synapses [23,113], suggesting that thalamocortical synaptic depression might be an important mechanism controlling the flow of information during normal brain function (e.g., non-alert awake states, sleep, etc.). If so, thalamocortical synaptic depression may be expected to decouple thalamic and cortical activity at frequencies above ~6-10 Hz in these states.

Consistent with existing hypotheses in the literature [21-23,113], I suggest a mechanism to explain this loss of thalamocortical synaptic depression in the awake state. It may be that increased spontaneous activity in the thalamus in the awake state "pre-depresses" [22] thalamocortical synapses even prior to the onset of visual stimulation. Hence, visual stimulation to measure thalamocortical synaptic depression fails to produce any further depression. This hypothesis is also consistent with work indicating that reducing spontaneous activity in the thalamus increases the observed depression at thalamocortical synapses in the visual system [114]. Neuromodulatory changes between the anesthetized and awake states [75] may also play a role in the loss of thalamocortical synaptic depression in awake mice.

If, consistent with my data, thalamic input to cortex is indeed pre-depressed in the awake state, cortical sensitivity to this thalamic input must be higher in the awake state than it is under anesthesia. Under anesthesia, the depressed thalamic input at frequencies above ~30 Hz is insufficient to drive a cortical response, yet the same-magnitude thalamic input in the awake state is able to drive a robust response in cortex (Figure 3.3). The mechanism underlying this brain state-specific switch in cortical sensitivity to thalamic input will be a topic for future study.

Chapter 3, in part, has been submitted for publication of the material as it may appear in Reinhold, K., Lien, A.D. and Scanziani, M., 2015, Distinct Recurrent Versus Afferent Dynamics in Cortical Visual Processing, Nature Neuroscience. The dissertation author was the primary investigator and author of this paper.

Chapter 4:

Role of Inhibition in Dynamics of Cortical Recurrent Circuits

4.1 Introduction

How do cortical recurrent circuits maintain dynamics fast enough to lock cortical sensory-evoked activity to the timing of thalamic input, over a wide range of temporal frequencies (approximately 1 to 30 Hz)?

Models of cortex have proposed that inhibition acts to counteract strong recurrent excitation to maintain these dynamics of the cortical response to sensory input and to prevent prolonged reverberation of sensory activity in the excitatory recurrent circuits of cortex [5,7,8,55]. According to these models, the rapid cortical response across a range of levels of thalamic input depends on an effective balance between excitation and inhibition* [5,63,131,142]. As recurrent excitation in cortex increases with increasing feed-forward drive from the thalamus, local inhibition in cortex also increases to prevent the increasing recruitment of self-sustaining positive feedback loops maintaining sensory activity within cortical recurrent circuits†. This regulation of positive feedback within cortical recurrent circuits enforces the speed of the cortical response and its sensitivity to high-frequency input.

The assumption underlying these models, that inhibition is necessary to maintain the fast sensory-evoked response in cortex, needs to be tested empirically. The prediction of these models is that

* An effective balance between excitation and inhibition in recurrent circuits is a foundation of cortical models exhibiting fast responses to changing thalamic input [5,7,55]. These models require that the effect of synaptic inhibition onto the post-synaptic neuron is sufficient to counteract the effect of recurrent excitation. This does not, however, imply that the actual magnitude of the inhibitory post-synaptic current must match the magnitude of the excitatory post-synaptic current. The excitation-inhibition balance is defined functionally as a balance between the effect of recurrent excitation and the effect of local inhibition within the network. Henceforth I will use this definition for "excitation-inhibition balance" in cortical circuits.

† Here I refer to a class of cortical models wherein cortical response dynamics are independent of the strength of thalamic input. Not all models of cortex exhibit this behavior [72,111,143]. However, my empirical results in Chapter 2 show that the CDF is indeed independent of the strength of thalamic input. Thus, sensory-evoked dynamics in visual cortex are independent of the strength of thalamic input. Hence, here I consider only the class of cortical models that matches this empirical data and reproduces robust sensory-evoked dynamics independent of the absolute level of feed-forward input.

disrupting the normal balance of excitation and inhibition in cortex will alter the rate of the cortical response to fluctuations in thalamic input.

Disrupting the normal balance of excitation and inhibition in cortex does not dramatically affect the orientation tuning of cortical neurons [144] or the centers of their spatial receptive fields (SRFs) [29]. Therefore it is believed that these cortical response properties are largely inherited from the synaptic wiring of specific feed-forward connections [29]. Hence not all properties of the cortical response are set by the balance of excitation and inhibition in cortex. Whether dynamics of the cortical response are sensitive to the balance of excitation and inhibition within recurrent circuits remains to be seen. If inhibition within recurrent circuits of visual cortex does enforce the rapid and invariant cortical response to thalamic input, characterized by the CDF, suppressing cortical inhibition should prolong the CDF.

4.2 Methods: Test whether inhibition enforces cortical dynamics

To test whether cortical inhibition controls the CDF, I designed a multi-focal optogenetic approach to simultaneously silence the thalamus and modulate inhibition within visual cortex (Figure 4.1a). These two independent optogenetic manipulations were performed in the same mouse and on interleaved trials, enabling a direct comparison of the CDF during conditions of normal cortical inhibition with the CDF in conditions of reduced activity in cortical inhibitory interneurons.

Stereotactic viral injections and targeted photo-illumination were used to confine each optogenetic manipulation to the appropriate brain structure. In the thalamus, injections of a Cre-dependent ChR2 into the GABAergic TRN in Gad2-Cre mice enabled photo-activation of the inhibitory TRN, as described previously (Chapter 2), to silence thalamus. In the same brain, a separate injection of Cre-dependent ArchT, a suppressive opsin, into the visual cortex of Gad2-Cre mice enabled the photo-activation of ArchT in these cortical inhibitory cells and thus the photo-suppression of their spiking activity. Expression of ArchT in cortical inhibitory interneurons has previously been shown to lower the levels of inhibition within cortical circuits [127,144]. Together, these two optogenetic manipulations in the same brain allow me to ask whether the CDF (at the offset of thalamic input to cortex) depends on the level of cortical inhibition.

4.3 Results

4.3.1 Decay function of cortical recurrent circuits depends on inhibition

What mechanism regulates the dynamics of recurrent circuits to prevent prolonged reverberation of cortical activity after thalamic input ends? Certain models [5,7,8,55] propose that cortical inhibition balances strong recurrent excitation to damp sensory activity. To test the role of cortical inhibitory interneurons in setting the fast CDF, I expressed the light-activated proton pump, ArchT [145], in V1's inhibitory interneurons (Cre-dependent ArchT, Gad2-Cre mice [119]) to suppress their activity by application of amber light (595 nm; Figure 4.1a-b). During the suppression of cortical GABAergic interneurons, I photo-activated the ChR2-expressing TRN in the same mice to silence the thalamus (473 nm light from fiber above TRN). I performed these experiments in anesthetized mice, because disinhibition of cortex in awake mice led to run-away activity (as in [144]).

Suppressing cortical interneurons on interleaved trials increased the gain of V1's visually evoked response* by 47% (Figure 4.1c), consistent with previous reports [127,144], and strikingly slowed the CDF by 263% (single exponential fit changes from 8±2 ms to 21±7 ms, n=5 mice, paired p=0.01; Figure 4.1d-f and Figures 4.2-4.4). This lengthening of the CDF originated in cortex†, because suppressing cortical interneurons did not affect the rate of dLGN silencing (paired p=0.78, Figure 4.1g; or cortical sensitivity to thalamic input differentially at low levels of input, see section 4.3.2 below). Furthermore, this lengthening of the CDF was not due to increased gain of the V1 response, per se, because under conditions of normal cortical inhibition the CDF is not affected by the gain of the V1 response (i.e., across stimulus contrasts or behavioral states, Figures 2.14 and 2.19, Table 2.1). Hence,

* V1 disinhibition also led to a sustained increase in spontaneous activity in visual cortex. The level of spontaneous activity in visual cortex (trial-averaged) in the absence of visual stimulation did not vary after the first 500 ms of disinhibition (Figures 4.1 and 4.4). Thus, fluctuations of spontaneous activity time-locked to the onset of the visual stimulus during cortical disinhibition cannot account for the prolonged CDF.

† Also, lengthening of the CDF during cortical disinhibition was correlated with the widespread expression of ArchT in visual cortex observed post-mortem (Figure 4.3 and Table 4.1).

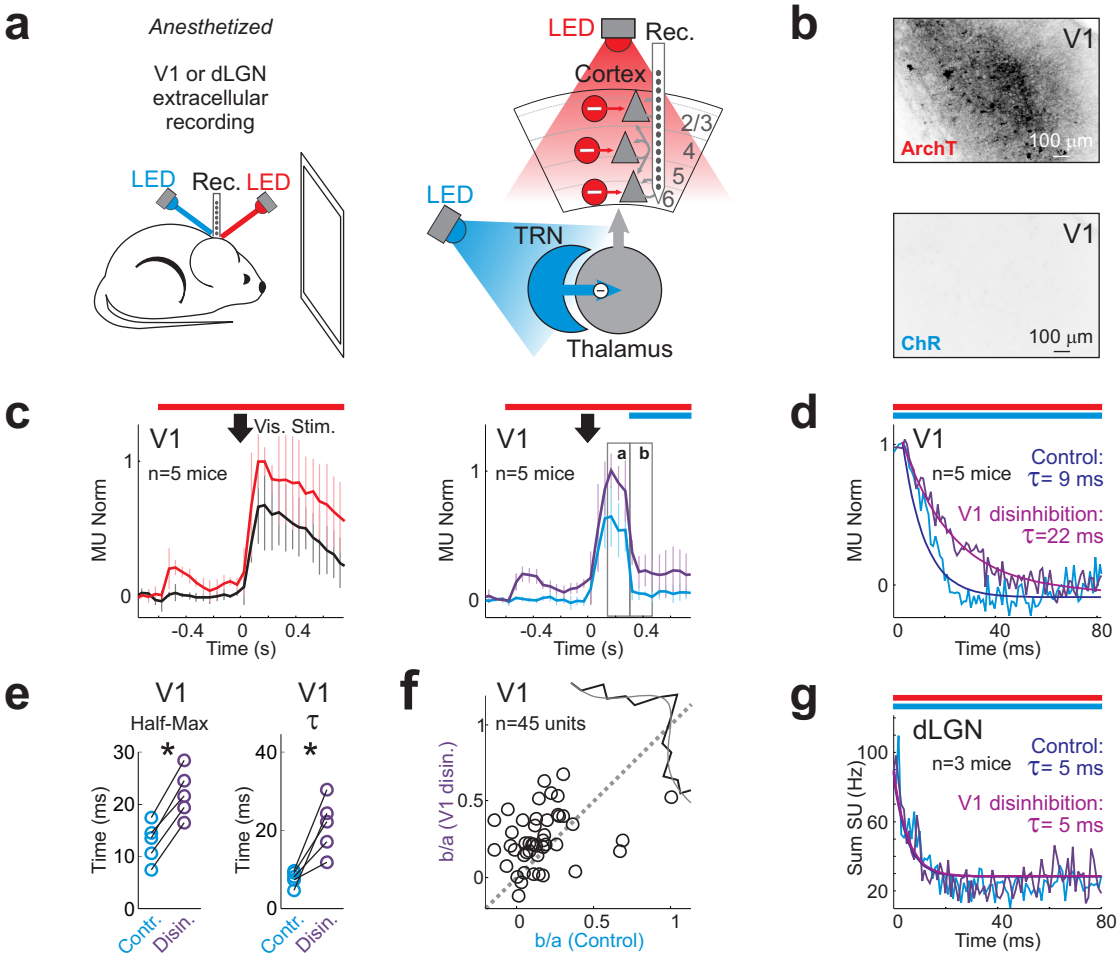

Figure 4.1: Cortical Decay Function is modulated by inhibition.
(a) Experimental configuration. **Left:** Two simultaneous optogenetic manipulations within same brain of anesthetized mouse. **Right:** Thalamic silencing with blue LED. Cortical disinhibition (reduction of inhibition) with red LED by optogenetically suppressing cortical inhibitory interneurons expressing ArchT. Extracellular recordings in V1.
(b) Expression of ArchT fused to GFP in V1. No expression of ChR2 fused to TdTomato in V1.
(c) Left: V1 multi-unit (MU) response to visual stimulus (arrow) lasting 3 s with (red) or without (black) cortical disinhibition. Red bar: amber LED illumination of V1. **Right:** During the same experiment (as left), on interleaved trials, the thalamus was silenced while cortex was (purple) or was not (blue) disinhibited. MU normalized to 150 ms window surrounding peak response in red or purple. Error bars are s.d. across mice. Blue bar: blue LED illumination of TRN. Windows a and b refer to (f).
(d) Expanded timescale at blue LED onset comparing CDF during normal (blue) or reduced (purple, disinh.) cortical inhibition. Dark blue, medium purple: single exponential fits to CDFs.
(e) Quantification of prolonged CDF during cortical disinhibition. Time to half-max (left) or time constant of single exponential fit (right) across mice (two-sided paired t-tests; time to half-max: p=0.0002, tau: p=0.01).
(f) Prolonged shut-off of single-unit (SU) population during cortical disinhibition as ratio of mean SU firing rate in window b over window a (windows refer to right panel of (c)). (See Figure 4.2 for raw SU examples.) Broken line is unity; above unity line indicates prolonged shut-off. Note significant slowing of SU shut-offs (two-sided paired t-test: p=0.0068).
(g) Time course of summed SU activity in dLGN at blue LED onset, comparing normal (blue) versus reduced (purple, disinh.) cortical inhibition. Reducing cortical inhibition does not affect shut-off time course in dLGN. The 3 mice here are a subset of the 5 mice in (d).

cortical inhibition contributes to enforcing fast recurrent circuit dynamics, which are essential to lock the duration of recurrent excitation to the timing of thalamic input.

4.3.2 V1 disinhibition does not differentially increase cortical sensitivity to low versus high levels of thalamic input

This data suggests that suppressing the activity of inhibitory interneurons in cortex leads to prolonged dynamics (observed as the prolonged CDF) within cortical recurrent circuits. To show that the site of the prolonged CDF is cortical recurrent circuits, I needed to rule out a change in cortical sensitivity to thalamic input as an alternate explanation for the prolonged CDF. Suppressing the activity of cortical inhibitory interneurons changes cortical sensitivity to thalamic input (Figure 4.1). However, this change at most time points manifests as a constant gain scaling of the V1 activity. A constant gain scaling of cortical activity will not alter the dynamics of the cortical response. However, if the effect of cortical disinhibition on V1's response is not gain scaling by a constant factor but is, instead, gain scaling that changes with the level of thalamic input, then differentially increased cortical sensitivity to low versus high levels of thalamic input could explain why the CDF appears slower during cortical disinhibition, as the level of thalamic input is quite low during photo-activation of the TRN to suppress thalamus (Figure 4.5).

To test whether cortical disinhibition differentially increases the steady-state gain of V1's response to low versus high levels of thalamic input, I performed a control experiment that measured the steady-state gain of V1's response to different levels of thalamic input. Therefore this control dissociates a change in V1's steady-state gain (V1's sensitivity to a constant level of thalamic input) from a change in V1 dynamics (what happens when thalamic input changes in time). I found that cortical disinhibition increased V1's steady-state gain equally for low and high levels of thalamic input, within the tested range (Figure 4.6). Thus, the source of prolonged cortical dynamics (the prolonged CDF) is not a change in cortical sensitivity to thalamic input and is likely mechanisms of cortex independent of thalamic input.

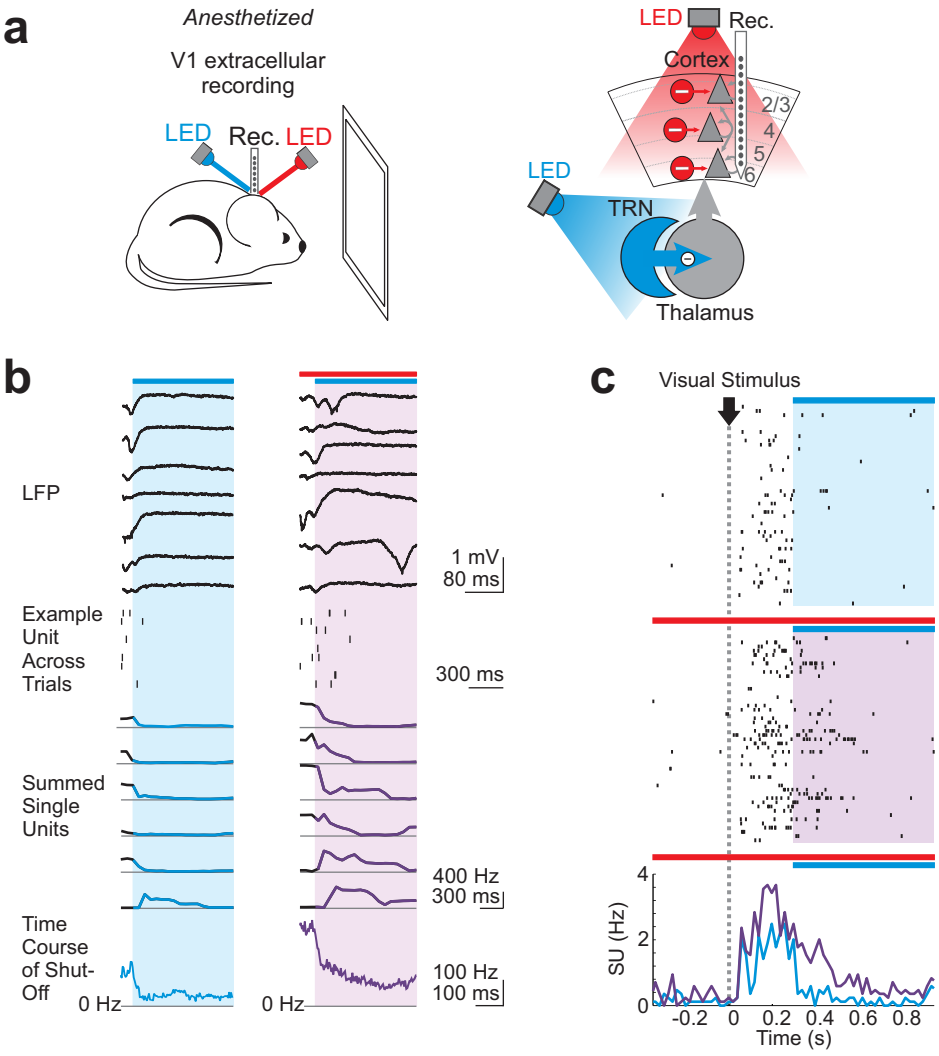

Figure 4.2: Example data upon silencing thalamus during V1 disinhibition.

In all parts, blue bar and blue shading indicate thalamic silencing during normal cortical inhibition; blue bar plus red bar and purple shading indicate thalamic silencing during cortical disinhibition.

(a) Schematic of dual optogenetic manipulation to simultaneously silence thalamus (blue) and disinhibit cortex (red) while recording from V1.

(b) Raw single trials from an example recording at moment of silencing thalamus showing (top to bottom) local field potential (LFP), example single-unit raster plot (each line is a spike), and summed single-unit activity across cortical layers on single trials. Bottom: trial-averaged summed single-unit activity (i.e., time course of shut-off) in this mouse. Horizontal gray lines indicate 0 Hz.

(c) Activity of example regular-spiking (RS) single unit (SU) at moment of silencing thalamus with (purple) or without (blue) V1 disinhibition. Visual stimulus is moving grating (3 s duration).

Figure 4.3: Slowing of V1 cortical decay function (CDF) correlates with widespread expression of ArchT in inhibitory interneurons of visual cortex.

(a) All mice tested, subdivided according to ArchT expression profile in V1. **Top Row:** Expression of ArchT-GFP in visual cortex. Overlay of one coronal section from each mouse in Group. Grp. 1: widespread ArchT (n=5 mice; these are 5 mice in Figure 4.1), Grp. 2: deep layers (n=8 mice), Grp. 3: superficial layers (n=2 mice), Grp. 4: radially restricted (n=2 mice).
Row 2: Increase in gain of V1 multi-unit (MU) activity by ArchT photo-stimulation. Black: control, Red: ArchT photo-stimulation.
Row 3: V1 MU activity with thalamic silencing. Blue: control, Purple: ArchT photo-stimulation. Striped horizontal bar is moving grating. Blue horizontal bar (blue LED) is photo-activation of TRN (thalamic silencing). Red horizontal bar (amber LED) is ArchT photo-stimulation. Error bars are s.d. across PSTH.
Row 4: Time-course of V1 MU (baseline-subtracted) after silencing thalamus. Blue: control, Purple: ArchT photo-stimulation.
Bottom: Quantification of shut-off time-course (stats. are two-sided paired t-tests).

(b) Coverage (fraction of pixels expressing ArchT within visual cortex) versus change in time to half-max of V1 shut-off as a result of ArchT photo-stimulation. Grp. 1 includes the 5 mice with greatest ArchT coverage of vis. cortex. R is correlation coefficient; p is p-value of correlation. Note strong correlation between expression of ArchT and slowing of CDF.

(c) Same as (b), but metric is time constant of single exponential fit to V1 shut-off. R is correlation coefficient; p is p-value of correlation. Note strong correlation.

(d) Quantification across all mice tested. Significant slowing of V1 shut-off with ArchT photo-stimulation across all mice ($p=0.00018$) and even across mice when excluding the 5 with the greatest ArchT coverage (i.e., excluding 5 with the strongest effect on CDF, $p=0.02$). Two-sided paired t-tests comparing interleaved conditions.

In Group 4, the lack of a significant lengthening of the CDF during disinhibition is likely due to the very focal expression of ArchT. Interneurons beyond the region of ArchT expression may be able to compensate for locally reduced inhibition. Indeed, interneurons in cortex can have extensive axonal arbors spanning horizontal distances of many hundreds of microns [157].

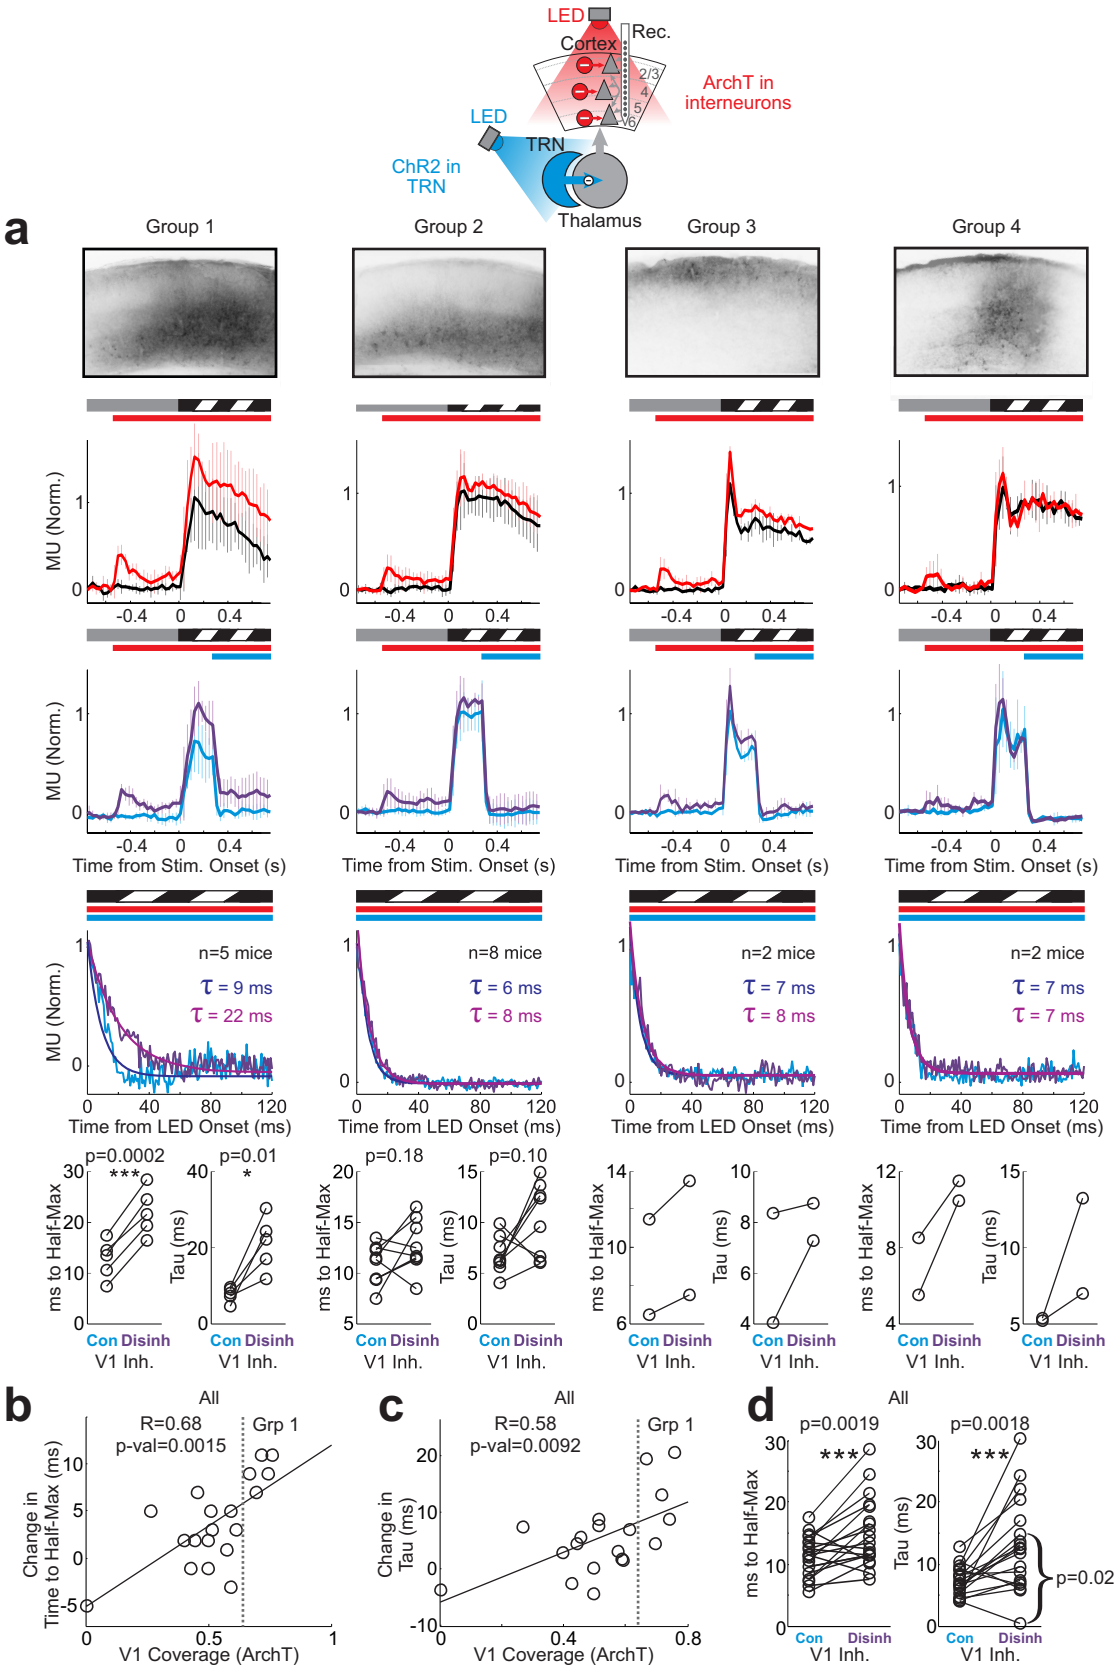
LED
Rec.
Cortex
2/3
4
5
6
ArchT in interneurons
LED
TRN
ChR2 in TRN
Thalamus
a
Group 1
Group 2
Group 3
Group 4
MU (Norm.)
Time from Stim. Onset (s)
n=5 mice
τ = 9 ms
τ = 22 ms
n=8 mice
τ = 6 ms
τ = 8 ms
n=2 mice
τ = 7 ms
τ = 8 ms
n=2 mice
τ = 7 ms
τ = 7 ms
Time from LED Onset (ms)
p=0.0002
***
p=0.01
*
p=0.18
p=0.10
ms to Half-Max
Tau (ms)
Con Disinh
V1 Inh.
b
All
R=0.68
p-val=0.0015
Grp 1
Change in Time to Half-Max (ms)
V1 Coverage (ArchT)
c
All
R=0.58
p-val=0.0092
Grp 1
Change in Tau (ms)
V1 Coverage (ArchT)
d
All
p=0.0019
***
p=0.0018
***
p=0.02
ms to Half-Max
Tau (ms)
Con
Disinh
V1 Inh.

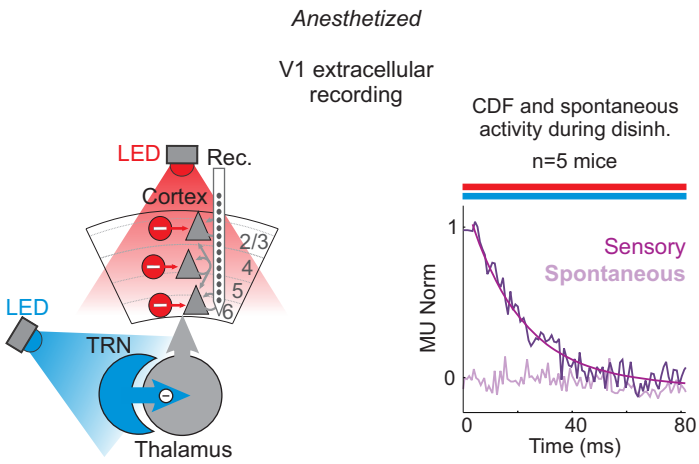

Figure 4.4: No further effect of silencing thalamus during cortical disinhibition on spontaneous activity in V1.

Left: Schematic of dual optogenetic manipulation to simultaneously silence thalamus (blue) and disinhibit cortex (red) while recording from V1.

Right: No further effect of silencing thalamus during cortical disinhibition on spontaneous activity in V1. Dark purple curves: CDF during cortical disinhibition from Figure 4.1d and single exponential fit. Lightest purple: Same time window as dark purple, i.e., activity in V1 at moment of silencing thalamus and during cortical disinhibition, but here without visual stimulus (spontaneous activity).

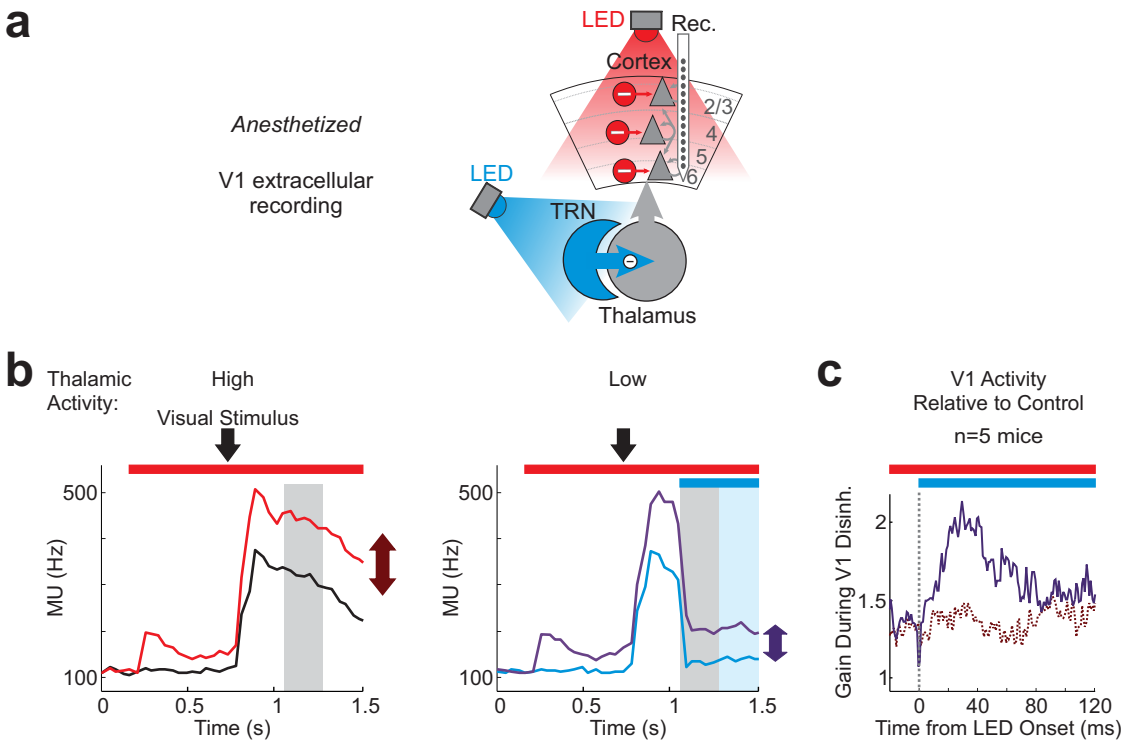

Figure 4.5: Scaling of V1 activity by a constant factor fails to explain time course of V1 activity at moment of silencing thalamus during cortical disinhibition.

(a) Schematic of dual optogenetic manipulation to simultaneously silence thalamus (blue, blue LED) and disinhibit cortex (red, amber LED) while recording from V1.

(b) Cortical disinhibition produces a gain scaling of the V1 response. n=5 mice from Figure 4.1. The amber LED turns on before the visual stimulus begins in this experiment, but the blue LED turns on only after the build-up of a visually evoked response in cortical recurrent circuits. Colored lines as in Figure 4.1.
Left: Effect of amber LED (red bar) on gain of visually evoked response (average, non-normalized multi-unit activity, MU) in V1 (vis. stim. is moving grating, 3 s duration). Note that cortical disinhibition both increases the gain of V1's visually evoked response and produces a higher spontaneous baseline in V1. Dark red arrow shows gain scaling in V1 within gray time window as a result of cortical disinhibition.
Right: Effect of amber LED (red bar) and thalamic silencing by blue LED (blue bar) on gain of visually evoked response in V1. Note that silencing thalamus during cortical disinhibition still leads eventually to a disappearance of visually evoked activity (i.e., return to disinhibited spontaneous baseline). Dark purple arrow shows gain scaling in V1 (during gray) as a result of cortical disinhibition during thalamic silencing.

(c) Throughout most of the visually evoked response in V1, the effect of cortical disinhibition is a constant gain scaling of visually evoked activity (see (b)). However, at the moment of silencing thalamus, activity in V1 during cortical disinhibition is no longer explained by a constant gain scaling of the time course of V1 activity under conditions of normal cortical inhibition. This is because the CDF is prolonged during cortical disinhibition. Thus, a model of constant gain scaling of activity in V1 by cortical disinhibition fails to account for the transiently increased gain of the response in V1 at the moment of silencing thalamus, as the prolonged CDF during cortical disinhibition decays more slowly than the CDF during normal cortical inhibition. This transient increase in V1 gain at the moment of silencing thalamus is shown in this plot. Here the Y-axis (gain) is the ratio of visually evoked activity in V1 during cortical disinhibition to visually evoked activity in V1 during normal cortical inhibition. Specifically, V1 gain as a function of time indicated by the dark purple line is the purple curve in (b, right) (minus its own spontaneous baseline 500 ms before vis. stim. onset) divided by the blue curve in (b, right) (minus its own spontaneous baseline). The dark red dotted line is the red curve in (b, left) (minus its own spontaneous baseline) divided by the black curve in (b, left) (minus its own spontaneous baseline).

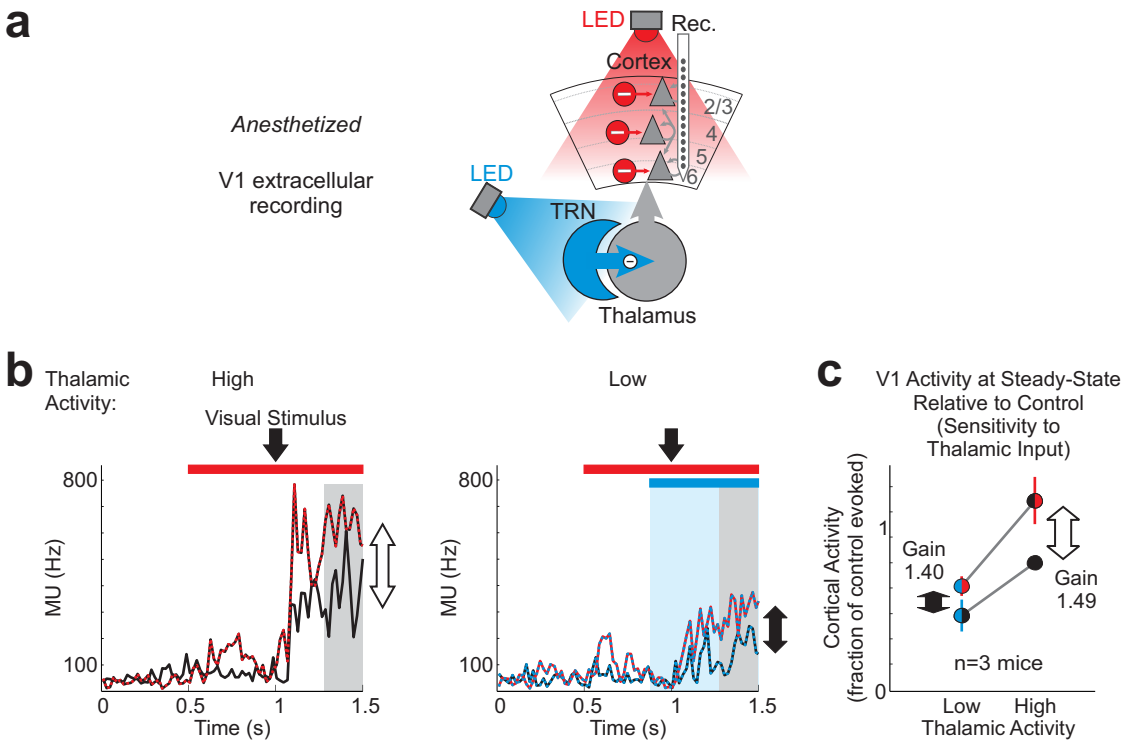

Figure 4.6: V1 disinhibition does not differentially increase cortical gain at steady-state in response to low versus high levels of thalamic input.

(a) Schematic of dual optogenetic manipulation to simultaneously silence thalamus (blue, blue LED) and disinhibit cortex (red, amber LED) while recording from V1.

(b) Cortical disinhibition produces a gain scaling at the onset of the V1 response. n=3 mice, a subset of 5 in Figure 4.1. Multi-unit (MU) PSTH responses recorded in V1 during different interleaved LED conditions. Both the blue LED and the amber LED turn on before the visual stimulus begins in this experiment, in order to measure the gain of the V1 response prior to the build-up of a sensory-evoked response in cortical recurrrent circuits. Black: Control (no V1 disinhibition). Red-Black: Amber LED to suppress inhibitory interneurons in cortex (V1 disinh.). Blue-Black: Low thalamic activity is achieved by low-intensity blue LED illumination of TRN over first 500 ms of visually evoked response. Red-Blue: Both amber and blue LEDs (to partially suppress thalamic activity during V1 disinhibition).
Left: No suppression of visually evoked response in thalamus; thus, response to high-contrast visual stimulus (moving grating onset) is high thalamic activity. Two-headed white arrow indicates increase in visually evoked activity at steady-state (gray shaded area) upon V1 disinhibition (red bar indicates amber LED illumination of V1).
Right: Lowering thalamic activity through partial suppression of visually evoked response in thalamus by blue LED (blue bar indicates blue LED illumination of TRN). Two-headed black arrow indicates increase in visually evoked activity at steady-state (gray shaded area) upon V1 disinhibition.

(c) V1 disinhibition does not produce a greater gain increase at low, relative to high, levels of thalamic activity at steady-state. Amplitude of cortical activity during low or high thalamic activity. Circle colors same as colors in plots to left. Gain, computed during steady-state (gray shaded area) is (black arrow) Red-Black over Black or (white arrow) Red-Blue over Blue-Black. Circles show mean response amplitude at steady-state (gray shaded areas), normalized to the response in control (black circle), across 3 mice.

Table 4.1: Expression of Flexed-ChR2 and Flexed-ArchT in brain areas of Gad2-Cre mice tested for effect of cortical inhibition on CDF.

For each Gad2-Cre transgenic mouse tested for slowing of the CDF in V1 as a result of cortical disinhibition, this table documents non-targeted locations of expression of ChR2-TdTomato and ArchT-GFP in post-mortem histology after physiology. y=expression observed in brain structure. n=no expression observed in brain structure. In all 19 mice, targeted ChR2 expression was observed in the TRN, and targeted ArchT expression was observed in V1 (not included in table). The five mice in Figure 4.1 (Group 1 in Figure 4.3) are represented by the five gray columns. In cases of observed leak of ChR2 expression to V1, expression levels of TdTomato fused to ChR2 in V1 were extremely low. Inset picture at upper left shows the mouse with the greatest leak of ChR2 expression to V1 -- note extremely low level of expression of ChR2 in V1. If anything, expression of ChR2 in Gad2+ interneurons in V1 should counteract the disinhibiting effect of ArchT expression in these inhibitory interneurons; thus, I do not make an attempt to exclude from my experiments the subset of mice with extremely low levels of a leak of ChR2 to V1 if these mice show cortical disinhibition measured physiologically.

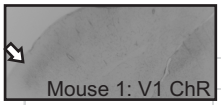

Brain Area		Mouse #																		
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
V1 posterior	ChR2	**y**	n	n	n	y	n	y	n	n	y	n	n	n	n	y	n	n	**y**	n
V1 middle	ChR2	**y**	**y**	n	n	y	y	n	y	y	y	n	n	n	n	y	y	y	n	n
V1 anterior	ChR2	**y**	n	n	n	y	n	n	n	n	n	n	n	n	n	y	y	y	n	n
Superior colliculus	ArchT	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n
Dorsal LGN	ArchT	n	n	n	n	y	n	n	n	n	n	n	**y**	n	**y**	y	n	y	**y**	y
Ventral LGN	ArchT	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n	n
Lateral posterior nuc.	ArchT	n	n	n	n	y	n	n	n	n	n	n	**y**	n	**y**	n	n	y	**y**	y
S1 cortex	ArchT	**y**	n	n	n	n	y	y	n	y	n	y	n	y	**y**	y	y	y	**y**	y
Auditory cortex (A1)	ArchT	n	n	n	n	n	y	n	n	y	n	y	n	n	**y**	n	n	y	**y**	n
Hippocampus	ArchT	**y**	**y**	y	y	y	y	y	y	y	y	y	**y**	y	**y**	y	y	y	**y**	y
TRN	ArchT	n	n	n	n	n	n	n	n	n	n	n	**y**	n	**y**	n	n	y	**y**	y
V2 Lateral	ArchT	**y**	**y**	n	y	y	y	y	y	y	y	y	**y**	n	**y**	y	n	y	**y**	y
Retrosplenial	ArchT	**y**	**y**	n	y	y	y	y	y	y	y	y	**y**	y	**y**	y	y	y	**y**	n
V2 Medial	ArchT	**y**	**y**	n	y	y	y	y	y	n	y	y	**y**	y	**y**	y	y	y	**y**	y

4.3.3 Changing the CDF of cortical recurrent circuits predictably alters V1's visually evoked response

The CDF predicts that V1 can readily follow thalamic frequencies up to ~15 Hz but attenuates thalamic frequencies above ~15 Hz (Figure 3.1). If V1's ability to rapidly follow high-frequency thalamic input is constrained by the CDF, prolonging the CDF should reduce V1's response to high-frequency fluctuations of thalamic activity. To test a causal role for the CDF in rate-limiting V1's response to high-frequency input, I prolonged the CDF by suppressing cortical interneurons in visual cortex of anesthetized mice (as described above, Figure 4.1a). (I performed this experiment in anesthetized mice, because suppressing cortical interneurons in the visual cortex of awake mice led to run-away, seizure-like activity.)

However, because in anesthetized mice thalamocortical depression strongly attenuates the cortical response to repeated high-frequency visual stimulation, there is little steady-state response in V1 to visual stimulus frequencies above 6 Hz. Yet prolonging the CDF is expected to affect V1's response measurably only at these higher visual stimulus frequencies. Therefore I could not test the effect of prolonging the CDF on V1's steady-state response under anesthesia. However, thalamocortical depression has not yet developed at the onset of the visual response in V1 under anesthesia. Therefore I could study the high-frequency increase in visual activity at the onset of the visual response under anesthesia to test whether prolonging the CDF led to a predictable slowing (or filtering) of V1's response to this high-frequency onset of sensory activity.

Prolonging the CDF lengthened V1's onset response to an extent predicted by the measured change in the CDF (Figure 4.7a-c). Control experiments to change the gain of V1's response without prolonging the CDF did not lengthen V1's onset response (Figure 4.7d). Thus, changing the CDF changes cortical dynamics. Taken together, my results show that the CDF of visual cortex is fast enough to temporally lock the cortical response to the timing of sensory-evoked thalamic activity.

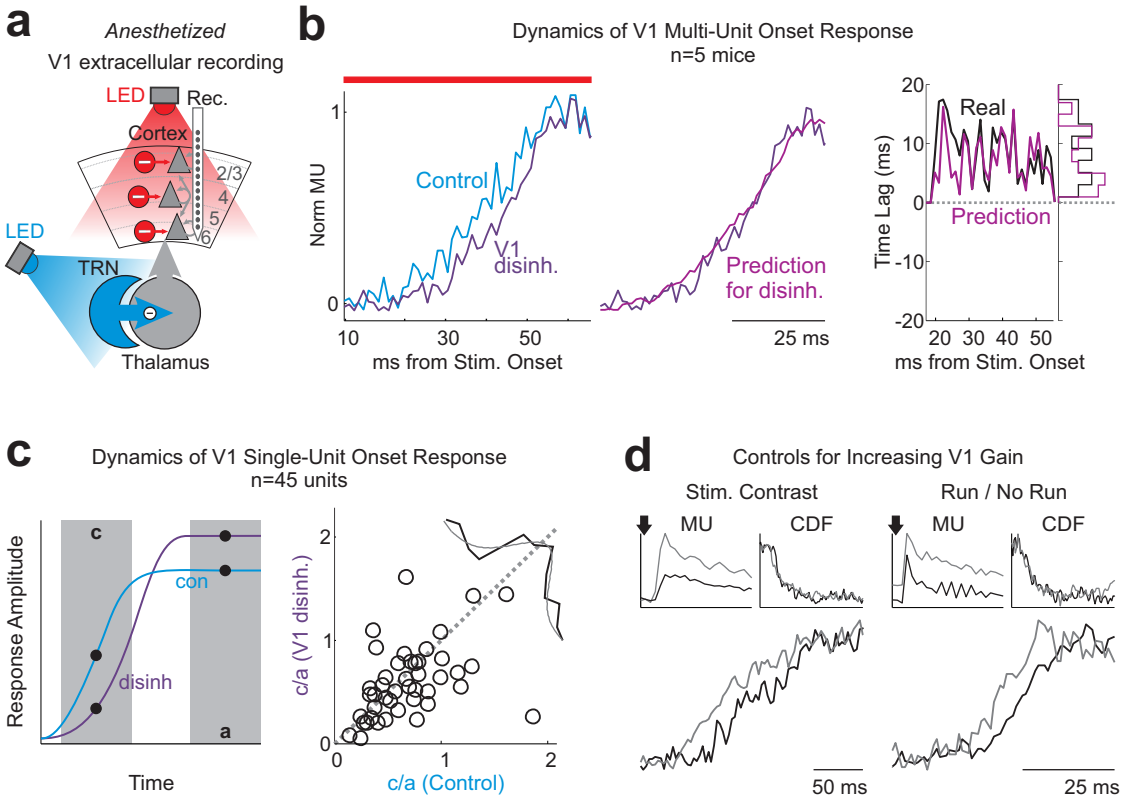

Figure 4.7: Slowing the CDF slows V1's onset response.

(a) Schematic of dual optogenetic manipulation to simultaneously silence thalamus (blue) and disinhibit cortex (red) while recording from V1.

(b) Left: Effect of control or slowed CDF on time course of response in V1 (MU) to visual stimulus onset. Blue is onset response in conditions of control CDF. Purple is onset response in conditions of slowed CDF (V1 disinh.). Same mice as in Figure 4.1. Note lagged onset response following V1 disinhibition.
Middle: Purple curve is from left. Lighter purple: Prediction of lagged V1 response to stimulus onset given the measured change in the CDF during V1 disinhibition ("Predicting slowed cortical onset ...", pg. 136).
Right: Quantification of average lag (n=5 mice) over the initial 50 ms of the onset response (black) and of predicted lag (light purple). Positive values indicate delayed onset response. Note good match between experimental (black) and predicted lag (light purple). Left axis y-axis shows histogram of experimental (black) and predicted (light purple) lags.

(c) Slowing of single-unit (SU) onsets upon slowing the CDF, computed in a manner analogous to Figure 4.1f scatter plot. Scatter plot at right shows the ratio of SU firing rate during initial 50 ms of the onset response, called time window c, divided by firing rate during the subsequent 50 ms, i.e., 50-100 ms after response onset, called time window a. At left, the cartoon of onset response shows how higher values of this ratio (c over a) for control (blue) indicate faster rise in control conditions vs. V1 disinh. (purple). p-value (two-sided paired t-test) for SU population slowing is p=0.04.

(d) Increasing the gain of V1 response by varying vis. stim. contrast (**Left**, gray=high contrast, black=low contrast) or behavioral state (**Right**, gray=running, black=non-running) leads to faster, not slower, rise of onset response. Insets at top: MU is multi-unit PSTHs, and CDF is CDF for each condition.

4.4 Discussion

Consistent with models of cortex [5,7,8,55], suppressing the activity of cortical inhibitory interneurons prolonged the cortical decay function (CDF), a measure of cortical dynamics. Thus, cortical dynamics depend on inhibition. Not all properties of cortical activity depend on cortical inhibition (e.g., orientation tuning [144] and spatial receptive field centers [29] in visual cortex are largely inhibition-independent).

Although cortical disinhibition increased the time constant fit to the CDF by about 250%, a large effect, the time constant of the CDF during cortical disinhibition was still just 22 ms, rather fast with respect to some previous estimates of intra-cortical recurrent dynamics (hundreds of ms or more [6,16,17,44,46,59,60,72,109-112]). Perhaps I did not observe an even longer-lasting CDF during cortical disinhibition, because I could not fully suppress the activity of cortical inhibitory interneurons (strong suppression led to run-away, seizure-like activity in cortex, which I avoided). Or, perhaps, additional cortical mechanisms (e.g., membrane time constants or the number of feed-forward synaptic steps in the network) also contribute to enforcing robust dynamics in visual cortex.

Moreover, here I modulate all GABAergic interneurons in visual cortex. Future work will need to determine which type or types of interneurons are primarily responsible for enforcing the CDF.

Understanding the mechanistic details of how inhibition impacts the dynamics intrinsic to cortical recurrent circuits will therefore require additional future work. However, it seems clear that the results of this chapter empirically implicate inhibition as one critical component of the mechanism deciding cortical dynamics in response to thalamic input.

Chapter 4, in part, has been submitted for publication of the material as it may appear in Reinhold, K., Lien, A.D. and Scanziani, M., 2015, Distinct Recurrent Versus Afferent Dynamics in Cortical Visual Processing, Nature Neuroscience. The dissertation author was the primary investigator and author of this paper.

Chapter 5:

Final Conclusions

Cortical dynamics are fundamental to sensation and cognition. These dynamics are sculpted by intra-cortical recurrent circuits, which dominate the anatomy of both associative and sensory cortical areas [24,146]. In sensory areas, these recurrent circuits provide the majority of sensory-evoked excitation to cortical neurons [29-32]. How these recurrent circuits impact the duration and dynamics of the response in sensory areas is still not known. Here I measure directly, for the first time, the dynamics intrinsic to recurrent circuits in visual cortex in isolation from dynamics at thalamic afferents. I demonstrate that the dynamics of these cortical recurrent circuits are on the order of the integration time window of a single cortical neuron (~10 ms [53]). Furthermore, I establish that these dynamics are a fundamental constraint on the temporal response of cortex to visual stimulation and precisely predict the amplitude attenuation of the cortical response in awake mice (Figure 5.1) over a wide range of stimulus frequencies.

By optogenetically dissecting intra-cortical and thalamic contributions to the sensory response, I discover that the prolonged timescale of continuing sensory-evoked activity in visual cortex after removal of the sensory stimulus is a product, not of cortical recurrent circuits, but of thalamocortical communication. Therefore thalamocortical interactions are necessary for long-lasting responses in visual cortex. In higher cortical areas, even longer-lasting sensory responses have been observed [126,147] and are often attributed to intrinsic dynamics of cortical recurrent circuits [110,111,148,149]. My work thus raises the possibility that long-lasting responses in higher cortical areas arise from thalamocortical interactions rather than cortical recurrent circuits [150]. However, consistent with the data of Chapter 4 suggesting that inhibition modulates intrinsic dynamics of cortical recurrent circuits, it is also plausible that different relative properties of inhibition and excitation lead to different intrinsic dynamics of the recurrent circuits in these higher cortical areas. My approach should allow researchers to measure these intrinsic dynamics across cortical areas.

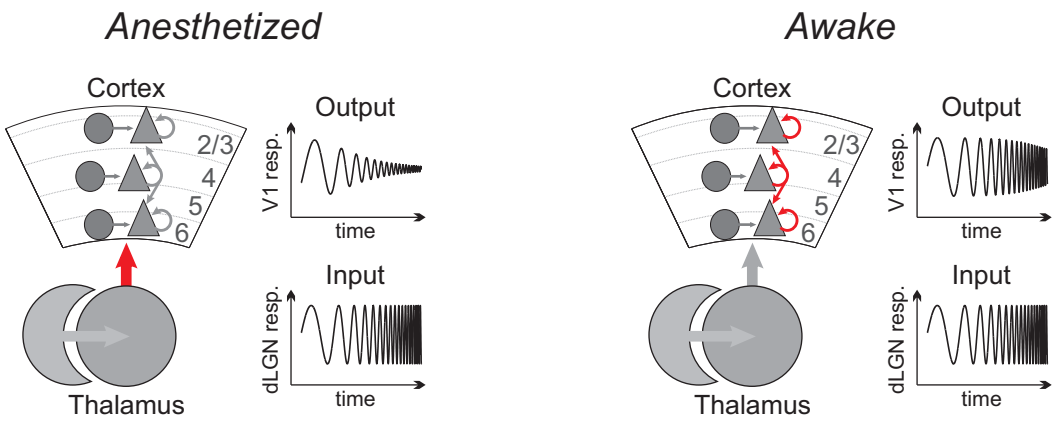

Figure 5.1: Rate-limiting afferent or recurrent dynamic control of responses in visual cortex.

Sites of rate-limiting processes are colored in red. Under anesthesia (left), depression at thalamic afferent synapses is rate-limiting and strongly low-pass filters input from dLGN. In awake mice (right), both low- and high-frequency input from dLGN pass equally well through thalamic afferent synapses. In the awake state, the dynamics of intra-cortical recurrent circuits become rate-limiting.

My results reveal that thalamocortical communication is necessary to sustain sensory-evoked activity in visual cortex over timescales longer than a few tens of milliseconds. The duration of the response in thalamus may be set by input from peripheral sensory structures [69] or by corticothalamic feedback [150].

In collaboration with Dr. Tony Lien, I find that cortical recurrent circuits provide the majority of synaptic excitation to cortical neurons only after an initial transient excitation from thalamus, which dominates the first 43 ms of the response (Figure 2.1), at least under anesthesia. If this also holds true in the awake state, this initial transient excitation from thalamus might be particularly relevant to natural viewing conditions producing rapid transitions of visual input [151,152].

Notably, the decay of sensory-evoked activity in intra-cortical recurrent circuits after silencing thalamus (CDF) was invariant to changes in stimulus strength (e.g., stimulus contrast, size), consistent with the CDF's ability to predict the amplitude of the cortical response across stimulus contrasts and temporal frequencies. Existing models of cortex achieve dynamics invariant to stimulus strength by assuming that local inhibition counteracts recurrent excitation to preserve the speed of the response [5,7,8]. Here I empirically test this assumption and find evidence that an appropriate ratio of excitation and inhibition contributes to the invariance of the CDF. However, future experiments that modulate specific types of interneurons in cortex will be required to understand exactly how inhibition controls the CDF.

My approach also provides a mechanism for how anesthesia disrupts the timing of the cortical response, despite the CDF's invariance to anesthesia. I find that marked depression of thalamocortical synapses under anesthesia disrupts the transmission of high-frequency sensory activity to cortex. In the awake state, however, thalamocortical synapses do not depress over a train of repeated visual stimuli and, thus, transmit both low and high frequencies of sensory activity equally well. My data suggest a possible explanation for this difference between the anesthetized and awake states. I observed increased spontaneous activity in thalamus in the awake state (Figure 3.11), which likely "pre-depresses" the thalamocortical synapses prior to visual stimulation, so no further depression is observed during visual stimulation [22]. If this brain state-specific switch in thalamocortical communication generalizes to other cortical areas, it might underlie the distinction between the awake and non-conscious states. Consistent

with the effects of anesthesia, non-alert awake states increase thalamocortical synaptic depression relative to alert awake states [51,113], suggesting that reduced thalamocortical transmission of high temporal frequencies is not just specific to anesthesia. It likely represents a normal mode of brain function correlated with reduced alertness or inattention [136].

Moreover, the increase in spontaneous thalamic activity as animals wake up from anesthesia likely also explains my finding that the majority of spontaneous cortical activity in the awake state depends on thalamic input, either from dLGN or higher-order thalamic nuclei, like LP, which were also suppressed by photo-activation of the TRN.

Thus, pre-depression of thalamocortical synapses in the awake state along with an appropriate balance of excitation and inhibition in recurrent circuits seems to enable the sensitivity of visual cortex to stimulus frequencies up to 30 Hz, consistent with the perceptual cut-off in mice [153] (and humans [154]). Many neurologic disorders with strong perceptual alterations, such as schizophrenia and autism, have been considered to be, in part, disruptions of the balance between excitation and inhibition in cortex [155]. The resulting change in the intrinsic dynamics of cortical recurrent circuits may account, in part, for the observed perceptual alterations. Understanding the mechanistic source of such alterations in brain dynamics should allow us to better target interventions. More generally, a mechanistic understanding of cortical dynamics will be vital to understanding various processes underlying cognition.

Chapter 5, in part, has been submitted for publication of the material as it may appear in Reinhold, K., Lien, A.D. and Scanziani, M., 2015, Distinct Recurrent Versus Afferent Dynamics in Cortical Visual Processing, Nature Neuroscience. The dissertation author was the primary investigator and author of this paper.

Appendix A:
Supplementary Methods

A.1 Specific methods for Chapter 2

A.1.1 Surgeries and animal preparation

Stereotactic viral injections of TRN

I stereotactically injected AAV2/1.CAGGS.flex.ChR2.tdTomato.SV40 into the thalamic reticular nucleus (TRN) of adult Gad2-Cre x C57Bl6 transgenic mice. Mice were anesthetized with 2% isoflurane and placed into a Kopf stereotax. Core body temperature was maintained at just above 35˚ C for the duration of the surgery with an FHC rectal probe/heating pad. Lubricating ointment (Artificial Tears) was applied to the eyes, the head was shaved, and the skin was sterilized with alcohol and povidone-iodine before exposing and stereotactically flattening the skull. I made a small craniotomy (approx. 50 µm in diameter above the somatosensory cortex, [1540 µm posterior, 2235 µm lateral] of bregma), inserted a thin pipette containing virus to 400 µm beneath the final target injection site (final target site: [1540 µm posterior, 2235 µm lateral, 3158 µm ventral] of bregma), returned the pipette tip to the final target injection site, and pressure injected 200 nL of the virus (titer: 6.86e12 genome copies/ml) into the TRN, at an injection speed of 30 nL/min. I removed the pipette only 10-15 min after completion of the injection to reduce the spread of virus along the pipette track. I administered a single dose of 10% buprenorphrine as a post-operative analgesic, sutured the skin, applied povidone-iodine, and waited 2 weeks before recording.

Terminal experiments under anesthesia

During my surgical preparation of a mouse for electrophysiological recording, I maintained an anesthetic depth such that the mouse did not respond to a toe pinch using either 2% gas isoflurane plus 2 mg/kg chlorprothixene (intraperitoneal injection), 2% gas isoflurane alone, or 0.5-1% isoflurane plus 1.5 g/kg urethane diluted to 0.1 g/ml in phosphate-buffered saline (PBS; i.p. injection). I covered the animal's eyes in transparent seed oil. After shaving the head and sterilizing the skin, I removed the skin and fascia, scored the skull with a bone scraper, and secured a metal head-frame to the skull using dental cement (Ortho-Jet) mixed with black paint, which prevents light during optogenetic stimulation from penetrating the dental cement. I allowed the cement to set for at least 30 min before drilling the skull.

Craniotomies/Thinned Skull In Anesthetized Preparation:

- Photo-Activation of TRN – I made a large circular craniotomy (diameter approx. 1.25 mm, [1540 µm posterior, 2235 µm lateral] of bregma) above the TRN. I then used a cautery to make a hole in the brain tissue above the TRN and inserted an optical fiber (1 mm diameter) to an approximate depth of 1.25 mm, through the somatosensory cortex and part of the hippocampus at coordinates [1.54 mm posterior, 2.235 mm lateral, 1.5 mm ventral] of bregma. I verified that this acute insertion of the large fiber optic did not affect my results by 1. confirming all results in a subset of anesthetized mice in which I used a smaller 250 micron-diameter fiber optic implanted more than a week before recording (see pg. 113, "Animal Surgery – Awake Recordings") and 2. verifying that the CDF and extent of shut-off of V1 activity did not change over the course of long anesthetized recordings (>1 hr).
- V1 Extracellular Recordings – I made a small craniotomy (diameter approx. 50 µm) above V1.
- Thalamic Extracellular Recordings – I made a craniotomy above dLGN/LP (diameter: 200-300 µm; approx. 1.8 mm posterior and 2 mm lateral of bregma).

I then covered the brain in artificial cerebrospinal fluid (ACSF: 142 mM NaCl, 5 mM KCl, 10 mM D-glucose, 10 mM Hepes Na-salt, 3.1 mM $CaCl_2$, 1.3 mM MgCl, pH 7.4) and inserted the recording electrode. During the recording, I reduced the level of isoflurane to 0.9-1% (or less when combined with urethane).

Awake recordings

I fitted animals with a head-frame for awake recordings more than a week before the day of the recording. I then performed the following relevant surgical procedures to prepare the animal for head fixation, optogenetic stimulation, and recording:

- Head-frame Surgery – I anesthetized animals with 2% isoflurane. After I shaved and sterilized the head, I cut off the skin and scored the skull using a bone scraper and surgical blade, which removes the soft upper layer of bone. I glued the edge of the skin to the skull using VetBond. I then inserted two bone screws bilaterally at 0.75 mm anterior and 2 mm lateral of bregma. I used black dental cement to attach a metal head-frame to the skull.
- Cannulation – For optogenetic silencing of the thalamus, during the head-frame surgery, I also inserted a metal cannula (a guide for the optical fiber) with an outer diameter of 460 µm through the skull and brain to a depth of 2375 µm beneath bregma, just above the TRN at coordinates [1.54 mm posterior, 2.235 mm lateral] of bregma. The bottom of this cannula was angled (45 deg.) and beveled to provide a clean, sharp edge for penetrating the brain tissue. Also, the bottom of the cannula was sealed with a clear window of Kwik-Sil that permanently separated the inside of the cannula from the brain. (At the time of the recording, I inserted an optical fiber into this cannula and optically stimulated through this clear window, without re-damaging the brain.) I cemented this windowed cannula to the skull and capped it with Kwik-Cast.
- Recording Well – During the head-frame surgery, I used the black dental cement to build up a recording well surrounding visual cortex.

Finally, I applied a thin layer of clear cement to the remaining exposed skull to prevent infection. I sterilized the edges of the skin with povidone-iodine again before allowing the animals to wake up and administered a single dose of 10% buprenorphrine as a post-operative analgesic. I checked on the mice daily after the head-frame surgery.

On the day of recording, I anesthetized the mice using 1.8-2% isoflurane, drilled off the clear dental cement covering the recording site, and made a small craniotomy (diameter approx. 50 µm) over either V1 or visual thalamus. I covered the animal's eyes in Artificial Tears lubricant. I fixed the head-frame to a post, positioned the body of the mouse on the circular treadmill, and inserted the recording electrode into the brain before allowing the animal to wake up from anesthesia. Immediately after waking up, the mouse groomed to remove the lubricant from its eyes. When I recorded in the thalamus both before and after the animals woke up from anesthesia, I applied only a thin layer of the relatively transparent eye lubricant to be able to provide visual stimulation through this protective layer.

Habituating awake mice to the recording set-up

I habituated mice to the head-fixed electrophysiological set-up prior to recording in the awake state. Four or more days after the head-frame surgery, I began a habituation protocol that exposed a mouse to the recording set-up for 30-60 min each day for at least three days. During each habituation session, the mouse was head-fixed (by fixing the head-frame to a metal post) and allowed to run freely during visual and "fake" optogenetic stimulation (light-emitting diode, LED, pulses near the head but pointed away from the brain). I recorded using the same set-up.

A.1.2 Electrophysiology

Extracellular recordings from cortex and thalamus were performed using a NeuroNexus silicon probe (A series) with 16 linear recording sites. In cortex, the 50 µm spacing between these sites provided an array that spanned 800 µm in total, enabling simultaneous recording from all cortical layers. In three of my thalamic recording experiments, I instead used an array with 25 µm spacing to improve unit isolation. The recording electrode was connected to an AM Systems pre-amplifying head-stage (20X) through a Plexon adaptor, and then the voltage signals were further amplified 500X and filtered between 0.1 Hz and 10 kHz by the AM Systems model 3600 extracellular amplifier. Data were digitized with a National Instruments Data Acquisition card and acquired with custom Matlab data acquisition software written by Dr. Shawn Olsen.

V1 recordings

I targeted the center of the monocular zone when recording from V1.

Recordings in thalamus

I targeted the dLGN by advancing the electrode straight down at coordinates [1.8 mm posterior, 2 mm lateral] of bregma. When post-mortem reconstruction of this thalamic recording track revealed penetration of both dorsal and ventral LGN (in about 10% of the animals), I included only the top-most visually responsive channels on the recording array to exclude from the analysis visually responsive units in ventral LGN. I targeted LP at a site 300-400 µm medial of this dLGN site.

V2 recordings

I stereotactically targeted V2 at >3.25 mm lateral of bregma. I only included V2L data sets showing an obvious visually evoked response and for which I was able to confirm by post-mortem histology a recording site more than 50 µm lateral of a clear V1-V2L border (see pg. 116, "Verification of V2 Recording Sites").

A.1.3 Post-mortem histology

Preparation of histologic samples

Animals were not perfused. I fixed the brains in 4% paraformaldehyde (PFA) in phosphate-buffered saline (PBS). After tissue fixation, I rinsed each brain in PBS overnight, placed the brain into 30% refrigerated sucrose solution for 2 days, and then cold-sectioned the tissue using a freezing microtome. I mounted the 50 µm-thick sections on slides with mounting medium including DAPI. For V2 track reconstructions, I stained the slices with blue NeuroTrace (fluorescent Nissl) before mounting. To stain with NeuroTrace, I rinsed each slice in PBS plus 0.1% Triton X-100 for 20 min, washed twice for 5 min in PBS, incubated the slices in 1:200 of NeuroTrace:PBS for 30 min at room temperature, and washed the slices overnight at 4˚ C. When using the NeuroTrace stain, I did not include DAPI in the mounting medium. Viral expression (ChR2-TdTomato or ArchT-GFP) and DiI tracks were imaged on a fluorescent microscope.

Excluding animals with ChR2 expression in V1

All mice with stereotactic injections of AAV2/1.CAGGS.flex.ChR2.tdTomato.SV40 into the TRN were tested post-mortem for off-target expression of ChR2 in visual cortex. I excluded all the data from animals showing ChR2 expression of >25% "fractional coverage" (fraction of pixels presenting any detectable ChR2-TdTomato reporter fluorescence) in any slice of visual cortex. Furthermore, I verified that <25% "fractional coverage" of V1 by ChR2 had no effect on V1 activity during TRN photo-activation (Figure 2.12). Such off-target expression of ChR2 was observed in approximately 5-10% of the animals. In addition, when I pointed an LED-coupled fiber directly at V1 rather than at the TRN, I found that direct V1 illumination did not affect V1 activity. Finally, I ruled out contamination of my results by off-target expression of ChR2 in V1 interneurons by noting that spontaneous activity in V1 under anesthesia was unaffected by thalamic silencing, whereas direct photo-activation of ChR2-expressing V1 interneurons produced a strong suppression of both visually evoked and spontaneous activity in V1 (Figure A.1).

Excluding ChR2 expression in dLGN interneurons

In about 20% of the mice, I observed some ChR2-TdTomato reporter fluorescence in local GABAergic interneurons of dLGN. Excluding these mice had no effects on the results (i.e., no change in the strength of V1 shut-off or its time course, the CDF). Furthermore, ChR2 expression in local dLGN interneurons, if anything, should lead to stronger and faster silencing of the dLGN, not slower and weaker silencing of dLGN. Therefore I included all mice in the final figures.

Labeling of electrophysiology recording track

I verified recording sites in the thalamus and V2 by post-mortem inspection of the electrophysiology recording tracks. To label a track, at the time of the experiment but after conclusion of the electrophysiological recordings, I retracted the recording electrode along the axis of entry, applied a drop of 1 mM DiI to the probe, quickly wicked away the artificial cerebrospinal fluid (ACSF) on the surface of the brain, and then reinserted the recording electrode to the site of data collection. I

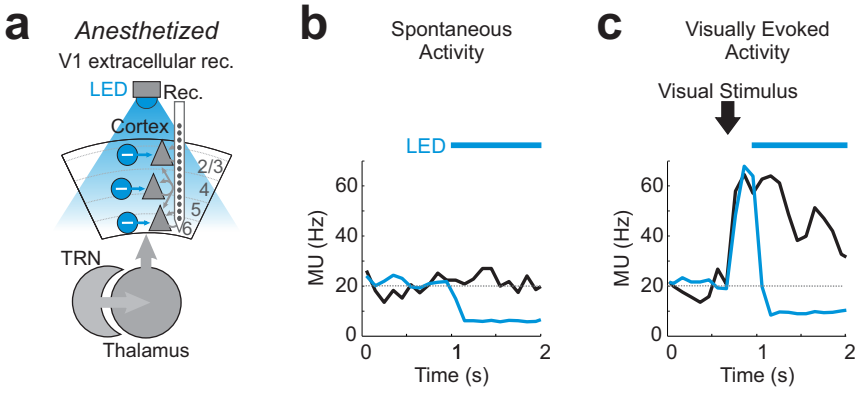

Figure A.1: Photo-activation of Gad2-positive cortical interneurons suppresses both spontaneous and visually evoked activity in V1.

In all parts, blue bar indicates LED illumination of ChR2-expressing cortical inhibitory interneurons.

(a) Schematic of photo-activating cortical inhibitory interneurons expressing ChR2 to suppress activity in cortical pyramidal neurons while recording in V1.

(b) Effect of photo-activation of cortical inhibitory interneurons on trial-averaged spontaneous multi-unit (MU) activity in V1 in example mouse. Without (black) or with (blue) photo-activation of interneurons for duration of blue bar.

(c) Effect of photo-activation of cortical inhibitory interneurons on trial-averaged visually evoked MU activity in V1 in same mouse as (b). Without (black) or with (blue) photo-activation of interneurons for duration of blue bar. Visual stimulus was moving grating (3 s duration). Trials including visual stimulation (c) were interleaved with trials lacking visual stimulation (spontaneous activity in (b)). Note that photo-activation of ChR2 in V1 in anesthetized mice suppresses both visually evoked and spontaneous activity, unlike photo-activation of ChR2 in the TRN in anesthetized mice, which only suppresses visually evoked activity (see Figure 2.3).

allowed the DiI-coated electrode to remain in the brain for at least 10 min before removing the electrode, sacrificing the animal, and fixing the brain.

Verification of V2 recording sites

After slicing the brain (50 µm sections), I identified the section containing the fluorescent DiI recording track in lateral secondary visual cortex (V2L). I identified the boundary between V1 and V2L (Figure A.2), first, by imaging the section in bright-field prior to mounting the tissue, which reveals the V1-V2L border as a difference in darkness between the middle layers of V1 and V2L (likely due to differences between V1 and V2L in the density of myelinated thalamic afferents terminating in these layers), then, by Nissl staining and mounting to reveal the dense cytoarchitectonic band of layer 4 in V1, and, finally, by registering the images from both of these methods to each other and to the Paxinos mouse atlas [156]. I excluded data obtained from animals in which the boundary between V1 and V2L appeared inconsistent or ambiguous.

A.1.4 Visual stimulation

I used an LCD computer monitor for visual stimulation (gamma-corrected, mean luminance 50 cd/m^2, refresh rate 75 Hz, 16x24 inches) at a distance of 25 cm from the eye contralateral to the V1 recording site. I began the recording with the computer monitor at about a 60 degree angle from the anterior-posterior axis of the head. I then re-positioned the monitor as necessary to place the apparent spatial receptive field of active units on the center of screen. The Psychophysics Toolbox in Matlab supported my presentation of various stimuli. Unless otherwise noted, stimulus contrast was 1, and the mean luminance of each full-contrast stimulus matched that of the interleaved blank gray screen.

Visual stimuli during V1 extracellular recordings

Visual stimuli consisted of static gratings (duration: 3 s) of 8 interleaved and randomized orientations. Static gratings (spatial frequency: 0.04 cyc/deg) appeared following a gray screen (6 s inter-trial interval of gray screen) of mean luminance matching the mean luminance of the static grating. I also used, as visual stimuli, as specified in the text, full-field luminance changes (transition from black to white screen, duration of white screen: 3 s, duration of inter-trial black screen: 5 s), reversing checkerboards (one-time contrast reversal of checkerboard pattern, each square side length: 3 degrees, inter-reversal interval: 5.5 s), moving gratings (transition from a gray screen to a pattern of black and white oriented bars, 8 interleaved and randomized orientations, moving across screen at a constant speed, spatial frequency: 0.04 cyc/deg, temporal frequency: 2 to 3 cyc/s) and full-field flashes of white or blue light lasting between 10 and 100 ms.

A.1.5 Optogenetic manipulations

Photo-activation of TRN

Anesthetized animals:

The optical fiber coupled to a blue LED (1 mm diameter; max power output 80 mW) was inserted and targeted as described above (pg. 112, “Surgeries and animal preparation”). I was able to use <10 mW step functions to drive maximal silencing of the dLGN in most mice, but in some animals I increased the power output of the LED to achieve maximal thalamic silencing. LED intensities below 10 mW typically provided partial thalamic silencing. In anesthetized animals, in the presence of continuous TRN illumination, dLGN silencing lasted for more than one second.

The ability of LED illumination to drive sustained spiking of the TRN units wore off over the course of about 1-2 hours of recording under anesthesia, when I frequently photo-activated the TRN (LED on for 1 s every 7-13 s; Figure A.3). (I recorded TRN population activity as LED-driven axonal signals in the dLGN – see pg. 121 below, “Sorting single units”.) Less frequent TRN photo-activation eliminated this run-down effect (LED on for <0.5 s every 10 s). This run-down was independent of anesthetic depth (Figure A.4) and time from penetration of the electrode. Thus, it was a function of illumination frequency (Figure A.5) and not of brain state or tissue recovery. For strong, reliable and sustained TRN photo-activation leading to strong thalamic silencing, I increased the inter-trial interval

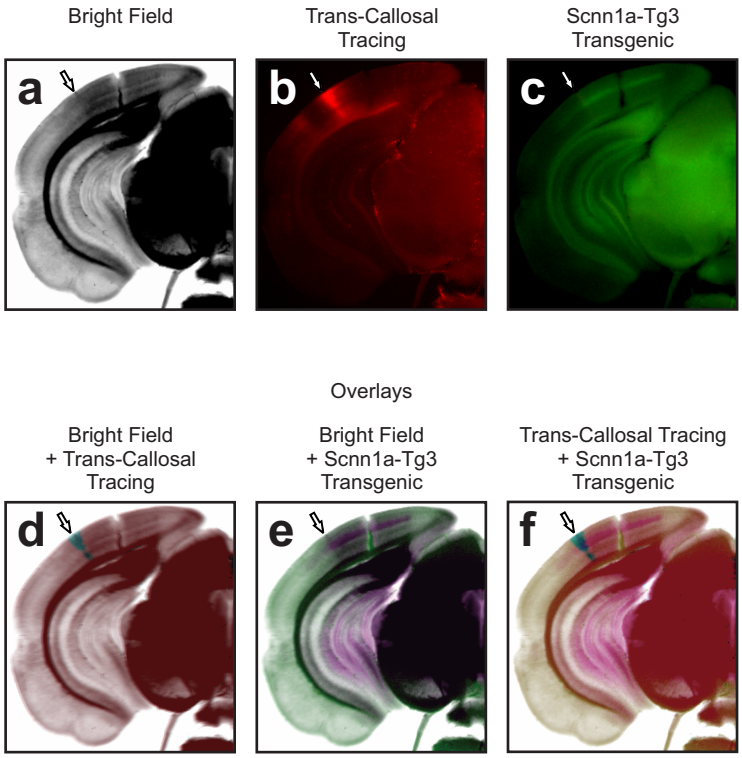

Figure A.2: Identification of V1-V2L border by trans-callosal tracing and bright field image registration in Scnn1a-Tg3 transgenic mice.

All images are of the same post-mortem coronal section through visual cortex (-3.8 mm posterior of bregma) from one mouse.

(a) Bright field of coronal section through V1. Arrow shows lateral border of V1 with V2 lateral (V2L).

(b) Fluorescence image of fluororuby (red, anterograde tracer) in trans-callosal axons from visual cortex of opposite hemisphere (25 injections of 15 nL to sites in 5x5 grid in visual cortex of opposite hemisphere). Arrow as in (a).

(c) Expression of GFP under control of Scnn1a-Tg3 (cross of Scnn1a-Tg3 Cre transgenic mouse line with GFP reporter mouse line). Arrow as in (a).

(d) Overlay of (a) (here red) and (b) (here blue).

(e) Overlay of (a) (here green) and (c) (here purple).

(f) Overlay of (b) (here blue) and (c) (here red).

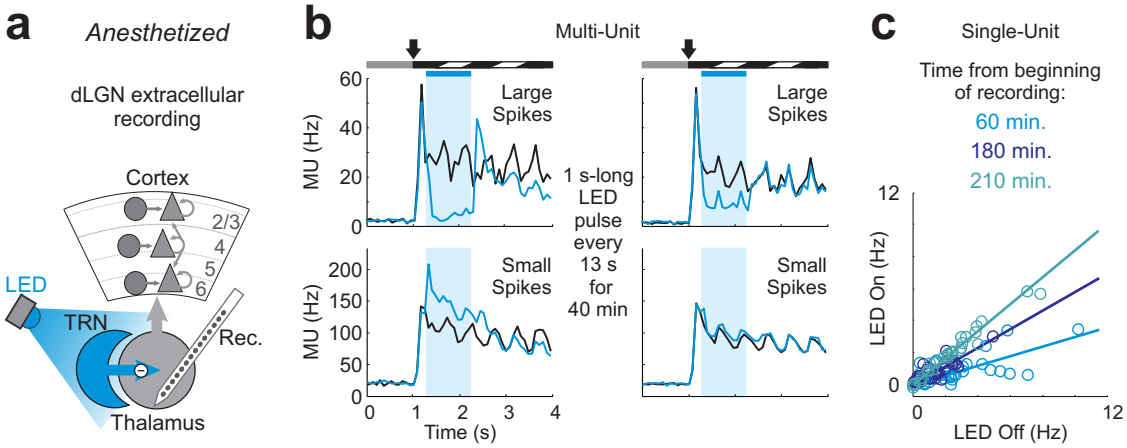

Figure A.3: Optogenetic silencing of thalamus wears off with repeated photo-stimulation of TRN.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in dLGN.

(b) Multi-unit (MU) activity in dLGN separated according to amplitude of spike waveform and time into experiment. **Left:** Activity in dLGN early in experiment. **Right:** Activity in dLGN 40 min later and after repeated photo-stimulation of TRN (1 s-long LED pulse every 13 s). Power of LED is 30 mW. Small-amplitude MU hash (**Bottom**) represents putative TRN axons recorded in dLGN (see Figure A.8), driven by LED. Large-amplitude MU activity (**Top**) represents putative thalamic relay cells, suppressed by LED. Photo-stimulation of TRN at this frequency leads to the loss of an ability to drive spikes in putative TRN axons in dLGN.

(c) Effect of TRN photo-activation on activity of single-unit population in dLGN over course of experiment. Each point is a unit. Well-isolated single units, here, are putative thalamic relay cells (Figure A.8). TRN photo-activation continues for 1 s every 13 s over course of experiment. As experiment proceeds (longer time from start of recording), suppression of putative thalamic relay cells wears off. Lines are fits to population activity measured at different times. X axis is visually evoked firing rate of unit (not baseline-subtracted) in control conditions. Y axis is visually evoked firing rate of same unit during TRN photo-activation. Flatter slope of line indicates stronger suppression.

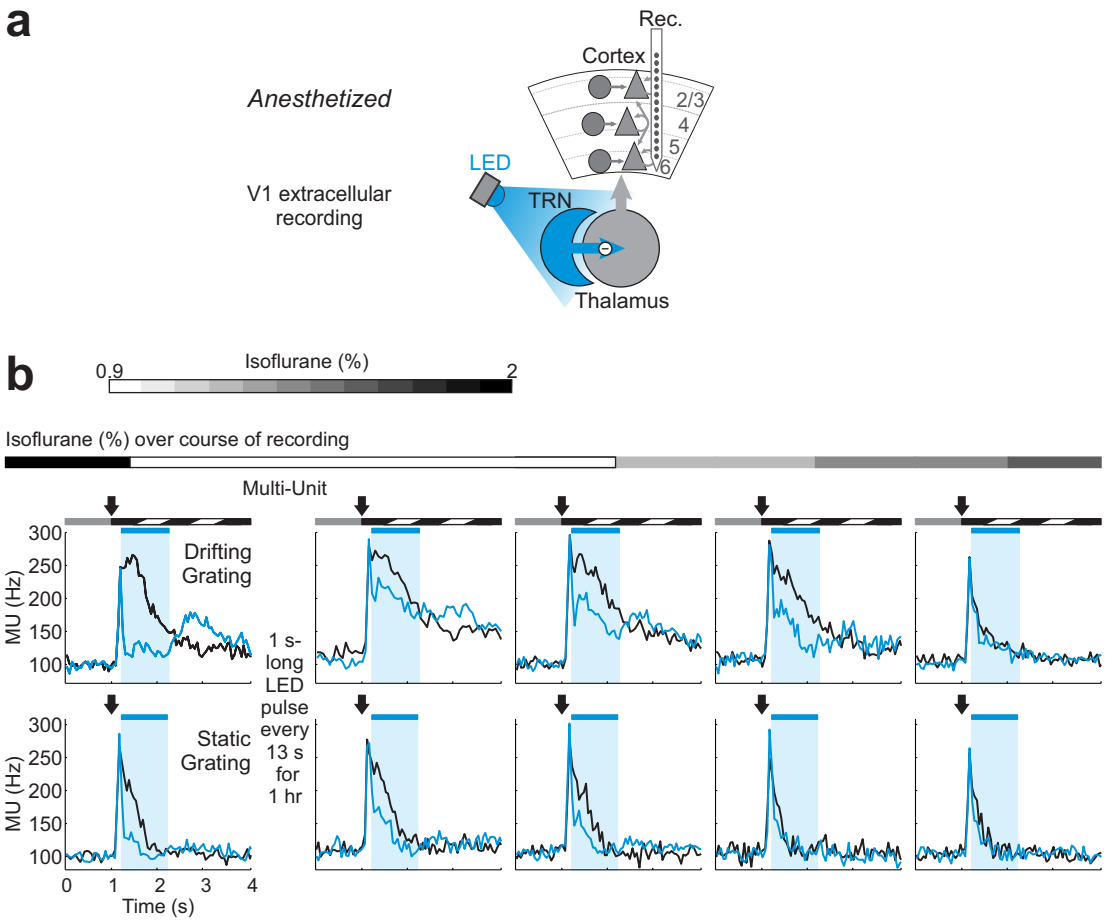

Figure A.4: Depth of isoflurane anesthesia does not correlate with magnitude of effect of thalamic silencing on V1 activity.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) Effect of depth of isoflurane anesthesia on silencing of visually evoked activity produced by photo-activation of TRN. Traces show V1 response to onset of moving grating (3 s duration, **top**) or static grating (3 s duration, **bottom**).
Top: Gray scale gives depth of isoflurane anesthesia at various time points in experiment.
Bottom: Multi-unit peri-stimulus time histograms (PSTHs) of V1 activity at different time points in the experiment as the depth of isoflurane anesthesia is altered. Left to right: During surgery to prepare the animal for recording, isoflurane is maintained at 2%. At the beginning of the recording, isoflurane is reduced to 0.9% (leftmost graphs). The strength of suppression of visually evoked activity in V1 produced by photo-activation of the TRN is strong at this early time point. Is this strong suppression an effect of lingering deep anesthesia? To test this, % isoflurane was maintained at 0.9% for 1 hr, while TRN photo-activation continued regularly. Suppression wore off (2nd column from left). Then, % isoflurane was increased to test whether deeper levels of anesthesia recovered the strength of suppression of the V1 response (3rd column from left to rightmost column). Strong suppression was not recovered. Therefore anesthetic depth does not explain strength of silencing of visually evoked response.

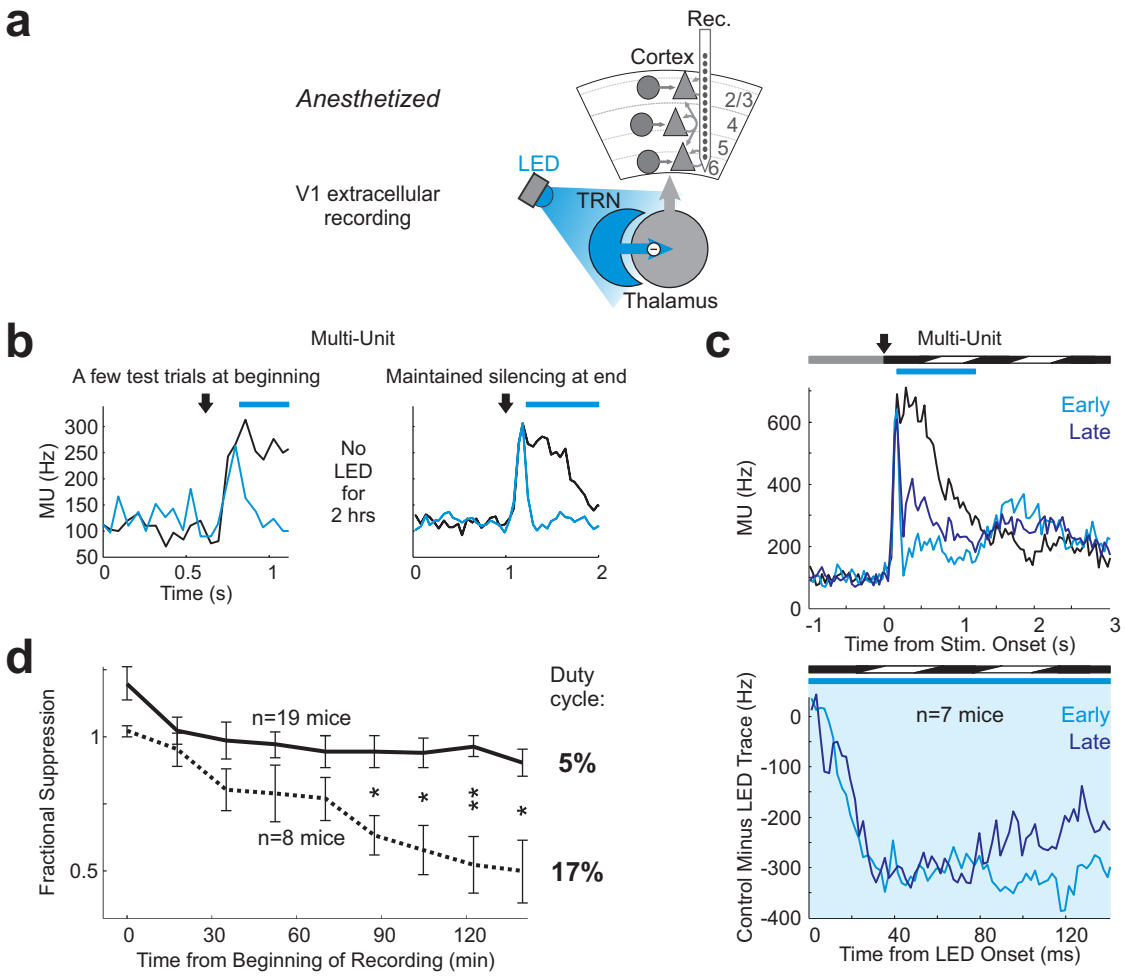

Figure A.5: Frequency and duration of TRN photo-stimulation correlates with whether or not thalamic silencing wears off over course of experiment.

(a) Schematic of thalamic silencing by photo-activation of TRN and recording in V1.

(b) Thalamic silencing by LED illumination of TRN does not wear off in the absence of repeated TRN photo-stimulation. Visually evoked responses to moving grating (duration 3 s) as multi-unit (MU) PSTHs, either in control (black) or during thalamic silencing (blue; throughout figure, blue bar/shading indicates photo-illumination of TRN). **Left:** A small number of test trials to measure effect of TRN illumination at the beginning of recording in example mouse. **Right:** Measurement of effect of TRN illumination after 2 hrs in same mouse without repeated photo-stimulation of TRN. Over this time, in the absence of photo-stimulation of TRN, no run-down of effect of TRN illumination.

(c) Top: In a different example mouse, repeated photo-illumination of TRN leads to run-down of optogenetic effect measured as shut-off of MU visually evoked activity in V1. Black trace is control (no optogenetic manipulation). Light blue is with TRN photo-illumination early in experiment. Dark blue is with TRN photo-illumination after repeated TRN photo-illumination (note some reduction in effect size). **Bottom:** Close-up of (top) at moment of silencing thalamus. Note that immediate shut-off time course in V1 (CDF) not changed by reduction in the strength of shut-off of visually evoked activity, i.e., sustained shut-off is lost, but immediate fast shut-off is maintained.

(d) Mean and std. dev. of fractional suppression of visually evoked response in V1 across mice as a function time and frequency (duty cycle) of TRN photo-illumination. More frequent TRN photo-stimulation (17% duty cycle) leads to run-down of effect, but 5% duty cycle does not. Thus, all subsequent experiments use a duty cycle of 5% or less for TRN photo-stimulation, which does not lead to run-down. * $p<0.05$, ** $p<0.01$.

to eliminate this run-down effect and chose to include only data from the first hour of each anesthetized recording session.

Awake animals:

I inserted an optical fiber (diameter 200 μm) coupled to a blue laser (power output 10 mW for photo-activation) into a cannula chronically implanted above the TRN, as described above ("Awake recordings", pg. 113). In contrast to what was observed in anesthetized animals (see section above), continuous illumination of the TRN in awake animals provided sustained silencing of the dLGN for only about 250 ms (Figure A.6). In fact, dLGN silencing in awake animals was invariably followed by a rebound of activity (recorded in both dLGN and V1) approximately 250 ms following laser onset, although the laser remained on (Figure A.7). This rebound was the first cycle of a thalamocortical oscillation (4-8 Hz) lasting for a second or more. Thus, in awake mice, I only considered the thalamus to be silenced during the first 125 ms immediately following TRN photo-activation. Given the approx. 10 ms CDF, this 125 ms time window is more than long enough for a complete characterization of the fast time-course of V1 shut-off following TRN photo-stimulation in awake mice.

A.1.6 Data analysis

Calculating single-unit response durations (Figure 2.1b)

To measure the response duration for each single unit individually, I plotted the trial-averaged PSTH of the unit's response and then measured the time it took for this PSTH to re-enter the noise (i.e., unit's firing rate after the stimulus returns to within one standard deviation of mean of unit's activity in the absence of a visual stimulus). I report this time as the response duration.

Sorting single units

I used UltraMegaSort from D. N. Hill, S. B. Mehta, and D. Kleinfeld [157] to cluster spike waveforms into putative units and then manual sorting to verify the quality of the isolated units. My units were well-isolated clusters with large spikes (Gaussian fit to the distribution of spike amplitudes suggested >85% of the spikes in every cluster were greater than 4X the standard deviation of the high-frequency noise), contained fewer refractory period violations than 1% of the total spikes, and did not appear similar in waveform shape or amplitude to any neighboring clusters, thus excluding units that could be confused.

Sorting thalamic relay units:

I separated the signals obtained from my dLGN recordings into spikes originating from putative relay cells and spikes originating from putative TRN axons (Figure A.8). Spikes from putative relay cells had larger, broader waveforms and exhibited significantly higher F1 modulation. Furthermore, these spikes were suppressed by TRN photo-activation, and they showed a clear post-inhibitory rebound. Activity from the putative TRN axons, in contrast, represented a very thin, low-amplitude spike population (hash). The spike rate of this hash was increased by TRN photo-stimulation with very low latency. Thus, I separated the spikes of the putative relay cells from the LED-driven hash by isolating single units and then excluding any units with an average spike waveform width-at-half-max of less than 0.22 ms.

Separating regular-spiking (RS) and fast-spiking (FS) units in the cortex:

I separated the regular-spiking (RS, putative excitatory) and fast-spiking (FS, putative inhibitory) units in the cortex by evaluating the average waveform of each unit by-eye. The mean width at half-max of the spike waveform for RS units was 0.44 ms (std. dev. 0.16 ms). The mean width at half-max of the spike waveform for FS units was 0.25 ms (std. dev. 0.067 ms).

V1 active states under anesthesia

I identified V1 active states according to the local field potential (LFP) ratio, a method that quantifies the ratio of the amplitude in the LFP between 5 and 30 Hz over the amplitude between 30 and 100 Hz (Figure A.9). Increases in this LFP ratio, as calculated from single-trial spectrograms (Gabor-Morlet wavelets, as in [68], code from http://dxjones.com/matlab/timefreq, or Chronux [158]) are indicative of Up states (cortical active states under anesthesia). I found a threshold for the LFP ratio able to separate Up (high activity) and DOWN (low activity) states by plotting the bimodal distribution of

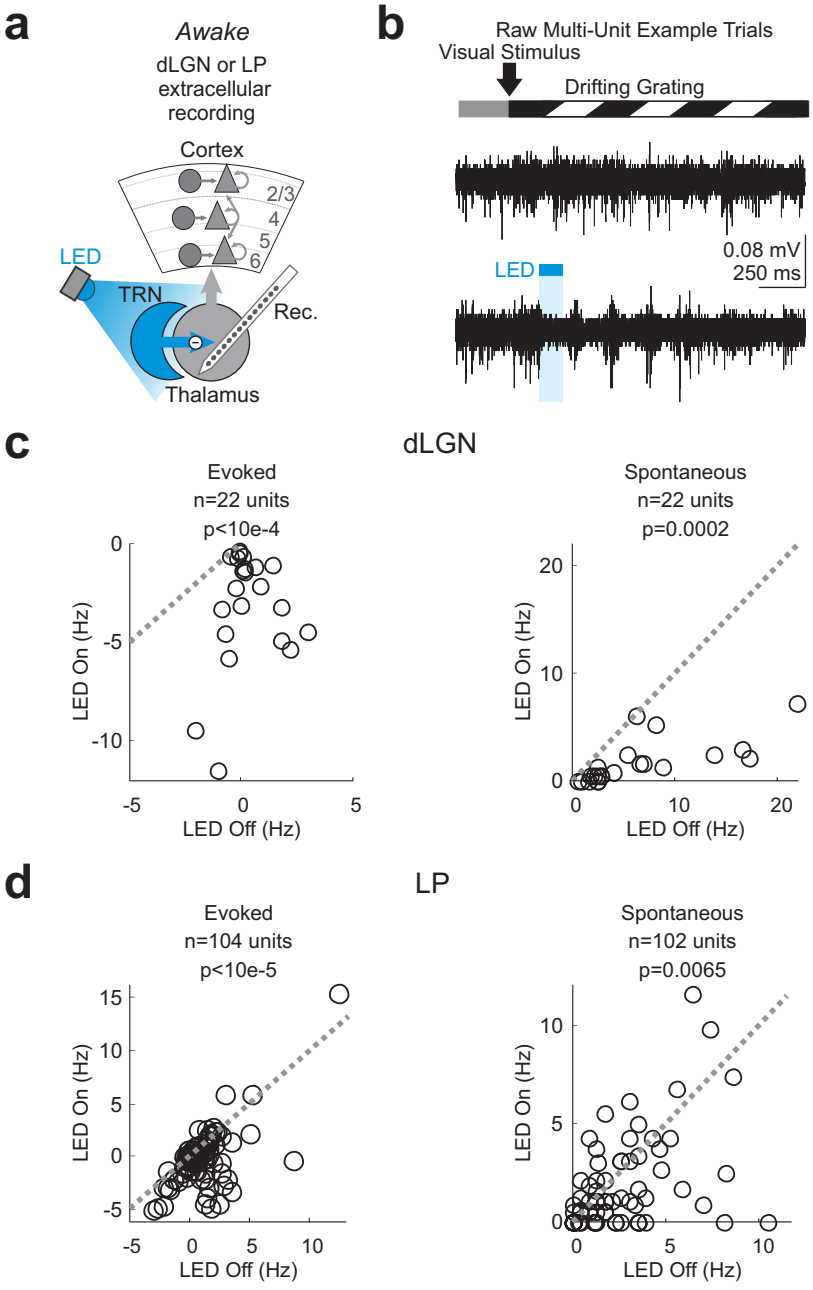

Figure A.6: TRN photo-activation suppresses thalamic activity in awake mice.
(a) Schematic of thalamic silencing by photo-activation of TRN and recording in dLGN or LP of awake mice.
(b) Example single trials of multi-unit recordings in dLGN. Blue bar indicates LED illumination of TRN. Vis. stim. in (b-c) is moving grating.
(c) Effects of TRN photo-activation on visually evoked (**left**) or spontaneous (**right**) activity in populations of well-isolated single units (SU) in dLGN (**top**) or LP (**bottom**). Each point gives SU's visually evoked firing rate without (X axis) or with (Y axis) TRN photo-activation. p=paired t-test of values on X and Y axes. Dotted lines are unity. Points below dotted lines are suppressed.

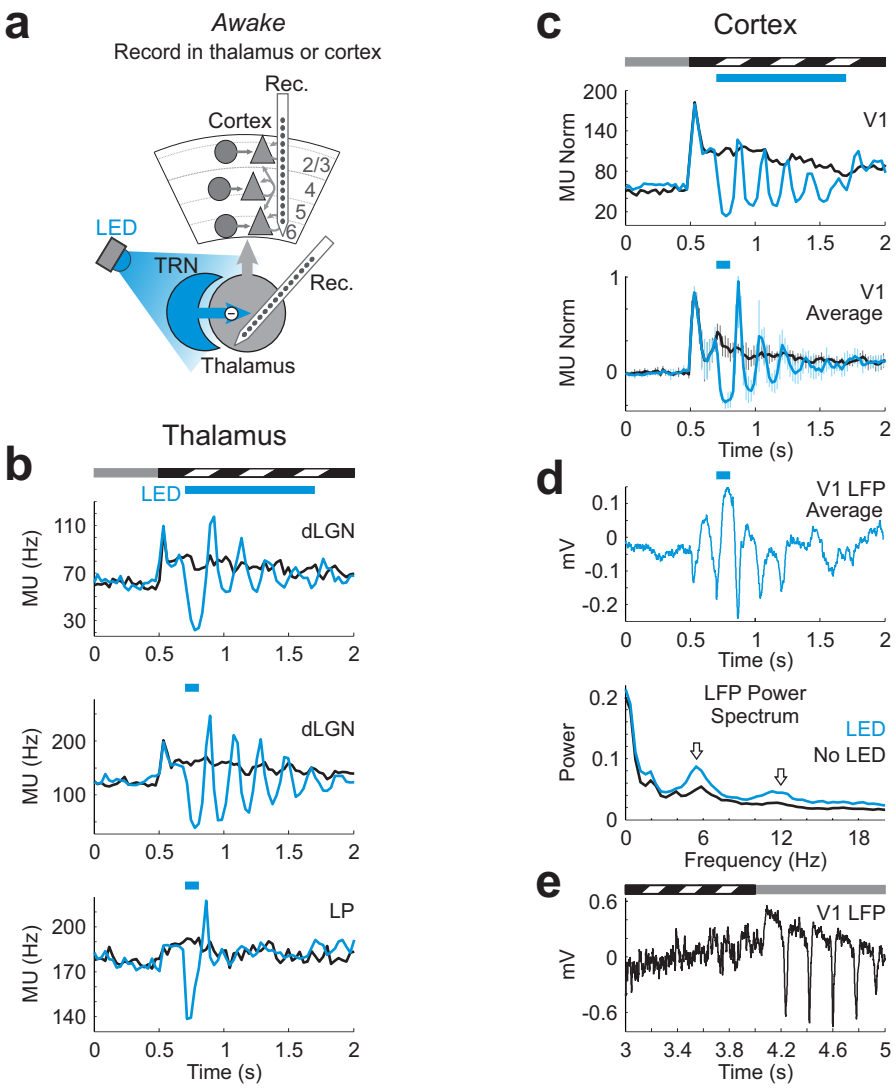

Figure A.7: TRN photo-activation in awake mice triggers ~6 Hz thalamo-cortical oscillation after initial suppression of relay thalamus.
For all parts, blue bar indicates duration of TRN photo-illumination. Angled black and white bars indicate moving grating visual stimulation (duration 3 s).
(a) Schematic of thalamic silencing by photo-activation of TRN and recording in thalamus or cortex of awake mice.
(b) ~6 Hz oscillation of multi-unit (MU) activity in dLGN (**top rows**) or LP (**bottom**) triggered by TRN photo-activation after initial suppression of thalamic activity in awake mice. Each row is PSTH from an example mouse.
(c) ~6 Hz oscillation of MU activity in V1 triggered by TRN photo-activation. **Top row:** example mouse. **Bottom:** average and std. dev. of normalized PSTH across mice (n=8, norm. to visually evoked response over first 100 ms).
(d) Top: ~6 Hz oscillation in V1 local field potential (LFP) after onset of visual stimulus (vis. stim. timing same as in (c)). **Bottom:** power spectrum of LFP after TRN photo-activation (blue) vs. spectrum of same time window in trial without TRN photo-activation (black). Arrows indicate that TRN photo-activation triggers strongest increases in LFP power at 6 and 12 Hz.
(e) ~6 Hz oscillation occurring spontaneously in the awake cortex (no TRN photo-activation) at the offset of moving grating stimulus. Trace is single trial of LFP.

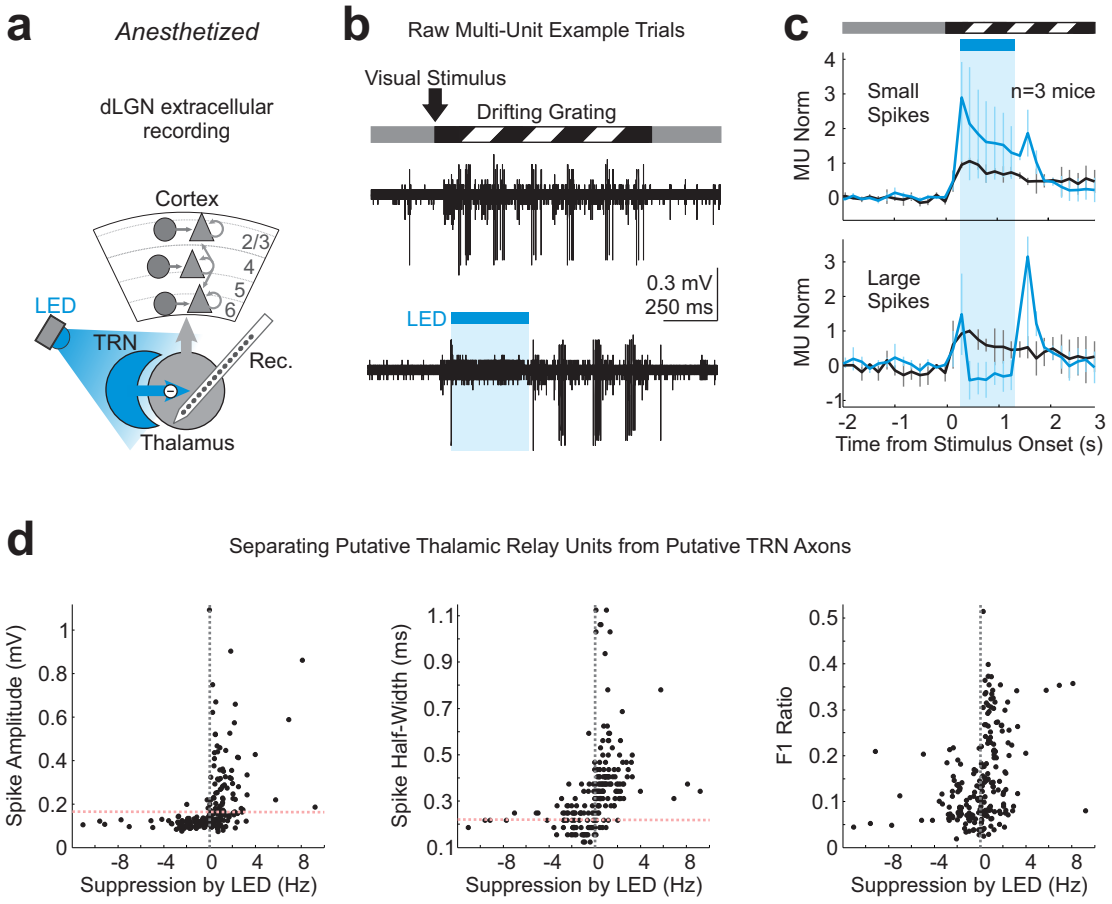

Figure A.8: Spike amplitude and width separate units recorded in dLGN suppressed by TRN photo-activation from driven hash (putative TRN axons) recorded in dLGN.
(a) Schematic of thalamic silencing by photo-activation of TRN and recording in dLGN.
(b) Example raw multi-unit data from recording site in dLGN in response to the appearance of a visual stimulus (moving grating, 3 s duration). **Top** trace: single trial of multi-unit (MU) activity in response to visual stimulus. **Bottom** trace: single trial of MU activity in response to visual stimulus with LED illumination of TRN for duration of blue bar.
(c) Mean and std. dev. of peri-stimulus time histogram (PSTH) of normalized (to height of visually evoked response across mice, see "PSTH normalization", pg. 126) MU activity in dLGN separated into small-amplitude (**top**) and large-amplitude (**bottom**) spike waveforms (see (d) for separation threshold). Blue bar is LED illumination of TRN. Angled black and white bars show moving grating visual stimulus.
(d) Two populations of spikes recorded in dLGN with qualitatively different responses to TRN photo-activation: a population of small-amplitude spikes driven by TRN photo-activation (consistent with the idea that these spikes are from TRN axons in dLGN, see time period of blue shading in (b) for example hash) and a population of large-amplitude spikes suppressed by TRN photo-activation (consistent with relay cells as source of large spikes). Scatter plots show metrics used to separate the populations driven or suppressed by TRN photo-activation. Each point is a cluster of spikes (clustered by k-means), including clusters that are not well-isolated. **Left to right:** spike amplitude, spike half-width-at-half-max, and F1 ratio (Methods) of clusters versus increase or decrease in that cluster's spiking activity during LED illumination of TRN. To right of vertical gray dotted line, clusters are suppressed by LED; to left, clusters are driven. Pink horizontal dotted line shows the cut-off threshold used to isolate putative dLGN relay cells suppressed by TRN photo-activation. Only well-isolated units with amplitudes and spike half-widths greater than the pink cut-off thresholds were used as measures of dLGN SU putative relay cells.

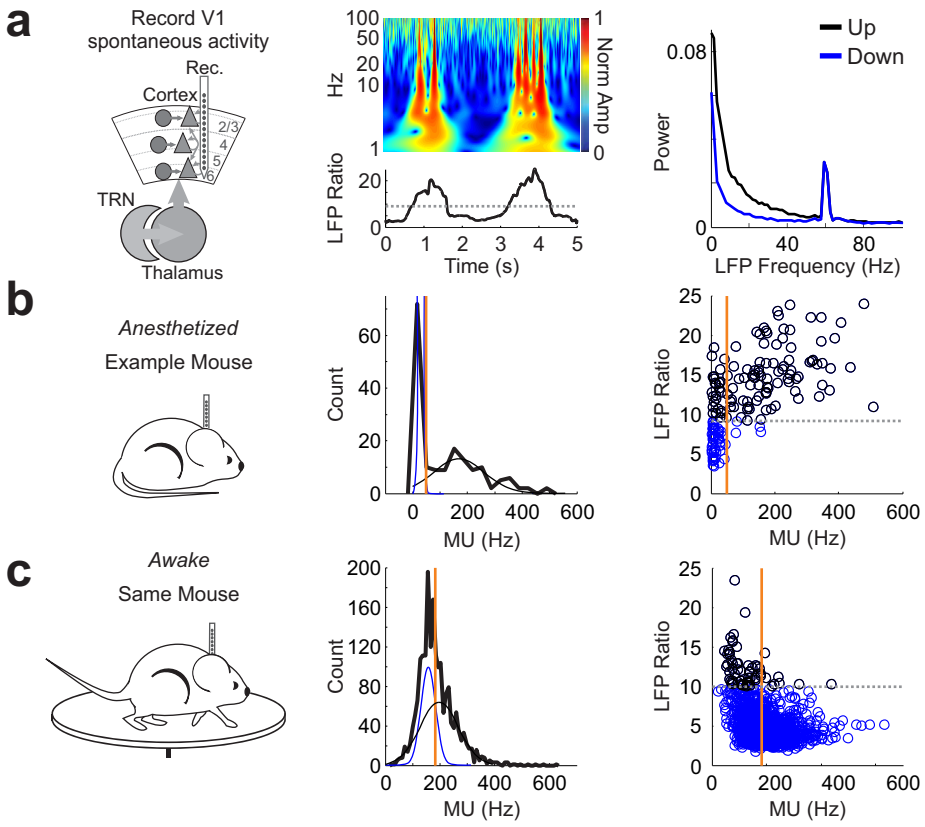

Figure A.9: Detection of Up states in V1 under isoflurane anesthesia, but not in the awake state, based on frequency content of LFP.
(a) Left: Schematic of recording in V1 in the absence of visual stimulation. **Middle:** Example of spectrogram of local field potential (LFP) activity in V1 on example single trial (top, whitened) and plot of corresponding LFP ratio (bottom, from non-whitened spectrogram) as a function of time. LFP amplitude ratio calculated as amplitude in the frequency band 5 to 30 Hz divided by amplitude in the frequency band 30 to 100 Hz [68]. Dotted line is threshold in this mouse used to detect Up states (see rest of figure for method of finding this threshold). In Up states, LFP ratio exceeds threshold. **Right:** Power spectrum of Up (black) versus Down (blue) states separated by LFP ratio.
(b) Right: Schematic of V1 recording under anesthesia. **Middle:** Histogram of multi-unit activity (MU) in 250 ms time bins during spontaneous activity in example anesth. mouse. Note bimodal distribution of V1 MU. Histogram fit by a sum of two gaussians, i.e., black and blue curves corresponding to Up and Down states, respectively. Orange line is point of intersection of these two gaussian curves. **Right:** Plot of MU versus LFP ratio in same example mouse. Note positive correlation. Orange line is threshold from (middle). Gray dotted line is the LFP ratio threshold that best distinguishes the population to the left of the orange line from the population to the right of the orange line. Gray dotted line is taken as the LFP ratio threshold for the subsequent detection of all Up states in this mouse. Appropriate LFP ratio threshold was determined in this way for each recording separately.
(c) As in (b) but showing the same analysis in an example awake mouse. Note that the distribution of spontaneous firing rates is not bimodal in the awake state. Moreover, no correlation between the LFP ratio and MU activity during spontaneous activity in the awake state. These observations are consistent with a lack of Up states in awake mice. Thus no Up states detected in awake mice.

cortical multi-unit (MU) firing rates across 500 ms time bins, using Receiver Operating Characteristic (ROC) analysis to identify the MU threshold best separating the two modes in this histogram (Up and DOWN states), and then converting this MU threshold to an LFP ratio threshold using the least-squares linear fit to the relationship between MU firing rate and the LFP ratio. Thus, consistent with the identification of Up states in previous work, I selected time periods of cortical activity characterized by a high LFP ratio. I applied this LFP ratio threshold to all the cortical LFP data to pick out the Up states beginning within 300 ms of silencing of the thalamus, aligned the Up states to their onsets and to the LED onset, and calculated the average firing rate over the course of the Up state (Figure 2.16).

PSTH normalization

In Figures 2.3, 2.5, 2.18, 2.6, 4.1 and A.8, error bars of PSTHs are standard deviation across mice.

dLGN PSTH (Figures 2.3, 2.5 and A.8):

- Figure 2.3, anesthesia: To show the effects of TRN photo-activation on thalamic relay cell activity (Figures 2.3c), I isolated putative relay cells as described above (pg. 121, "Sorting thalamic relay units") and plotted, for each mouse, a peri-stimulus time histogram (PSTH) of summed single-unit activity, analogous to multi-unit activity but excluding high-frequency noise and putative axonal hash (see above). I then normalized the PSTH for each mouse by the mean summed single-unit activity over the full duration of the moving grating (3 seconds) during control conditions (black). Summed single-unit activity in the thalamus was not baseline-subtracted. In all figures with a normalized PSTH across mice comparing control conditions with thalamic silencing, error bars are standard deviation of the normalized PSTH across mice.
- Figure 2.5, anesthesia: In Figure 2.5a, I show dLGN activity as the PSTH of baseline-subtracted multi-unit (MU) activity normalized by the mean activity in the 250 ms window surrounding the peak of the evoked response. In Figures 2.3 and 2.5, visual responses are to moving gratings (3 s), but silencing of the evoked response in dLGN was not different between responses to static or moving visual stimuli. At 300 ms into the response, as in Figure 2.3c, the percentage of the evoked response suppressed in dLGN for static versus moving stimuli was: (static, n=2) 93±57% (mean±s.e.), (moving, n=12) 95±28%.
- Figure A.8, anesthesia: In Figure A.8, I investigate the types of activity I record in dLGN using extracellular probes. This figure shows that well-isolated single units have a profile consistent with relay cells, but a low-amplitude hash has a profile consistent with TRN axons in dLGN. In Figure A.8c, I separate these two populations and plot the baseline-subtracted MU activity across mice as the average PSTH normalized to the evoked response under control conditions over the 250 ms surrounding the peak of the evoked response.

V1 PSTH (Figures 2.3, 2.5, 2.18 and 4.1):

- Figure 2.3, anesthesia: For visually evoked activity in V1 (Figure 2.3f), I show the baseline-subtracted multi-unit PSTH in response to brief visual stimuli, normalized by the mean multi-unit firing rate over the first 500 ms of the control evoked response. This PSTH includes responses to brief flashes, reversing checkerboards, full-field luminance transitions, and static gratings. The results were the same for all these brief stimuli and also for moving gratings. I show the cortical decay function (CDF) during brief stimuli in Figure 2.3f. Again, there was no difference in the CDF between brief visual stimuli and moving gratings (Table 2.1). The PSTH showing spontaneous activity in V1 during Up states (Figure 2.3e) is multi-unit activity normalized by the mean firing rate in the control condition over the 500 ms window surrounding the peak of the Up state, similar to the PSTH showing visually evoked activity. In all cases, error bars are standard deviation of the normalized PSTH across mice.
- Figure 2.5, anesthesia: In Figure 2.5b, I show V1 activity as the PSTH of baseline-subtracted MU activity normalized by the mean activity in the 250 ms window surrounding the peak of the evoked response.
- Figure 2.18, awake: I also show the visually evoked PSTH in awake mice (Figure 2.18). In Figure 2.18, I show the PSTH for visual stimuli (i.e., moving gratings) baseline-subtracted and normalized by the rate over the first 150 ms of the control evoked response. I also show the

CDF curves separately for brief visual stimuli (i.e., brief flashes, reversing checkerboards, full-field luminance transitions, and static gratings) and moving gratings. As under anesthesia, there was no change in the CDF across any of the visual stimuli tested (Table 2.1). In Figure 2.18c, I show spontaneous activity in V1 as the mean and standard deviation of the PSTH across mice (not baseline-subtracted). Here the PSTH in each mouse is normalized by the mean spontaneous firing rate under control conditions (no thalamic silencing).

- Figure 4.1, anesthesia: In Figure 4.1c, PSTH in each mouse is normalized to evoked rate in 150 ms window surrounding peak of response.

V2 PSTH (Figure 2.18):

- Figure 2.18, awake: For visually evoked activity in V2 (Figure 2.18d), I show the baseline-subtracted MU PSTH normalized in each mouse by the evoked rate under control conditions over the first 150 ms of the evoked response.

LP PSTH (Figure 2.6):

- Figure 2.6, anesthesia: To show the effects of TRN photo-activation on thalamic relay cell activity in LP (Figures 2.6b), I isolated putative relay cells and plotted the PSTH of summed single-unit activity normalized in each mouse by the mean summed single-unit activity over the first 200 ms of the evoked response during control conditions (black). Error bars were standard deviation of the normalized PSTH across mice. Summed single-unit activity was not baseline-subtracted.

Accounting for shut-off delay in dLGN

To calculate the exact time course of sensory activity decay intrinsic to the cortical recurrent circuits (in linear systems terms, the cortical impulse response function), I needed to remove (deconvolve) the time course of shut-off in the thalamus from the time course of shut-off in the cortex. However, deconvolution of real neural signals with noise introduces error. Therefore I decided to study the shape of the decay curve in V1 only after the thalamus was nearly off. The single exponential time constant fit to the shut-off of activity in dLGN was 3.7 ms. Thus, the dLGN activity has already decreased substantially (>60%) by about 3 ms after the LED onset. Furthermore, 3 ms is the average delay before activity begins to obviously decay in V1. As I show in Figure 2.9, beyond 3 ms, a single exponential fit well approximates the result of the deconvolution. Furthermore, the V1 decay time course approximates an exponential function, and an exponential function has the same shape at all scales, meaning theoretically I can measure any window of the exponential decay and get the same fit. Finally, considering this time window beyond 3 ms after LED onset is more than sufficient to allow me to study response frequencies between 1 and 60 Hz. For these reasons, in all cases, I fit the single exponential decay to cortical shut-off starting 3 ms after LED onset. This accounts for the time it takes the dLGN to shut off.

Single exponential fits to cortical shut-off

I applied the following protocol to fit the CDF. 1) I took the baseline-subtracted peri-stimulus time histogram (PSTH) of V1 shut-off between 3 and 50 ms after the onset of the LED (or longer, see time window in each figure). The 3 ms delay from LED onset was excluded from the shut-off curve to account for the time it takes the dLGN to shut off (see preceding section). 2) Avoiding any assumptions about the expected fractional suppression of the evoked response, I used Matlab to fit a single exponential plus a constant offset to this PSTH. This method optimized the fit to all points in the trace, without forcing the initial or final values. 3) If this method failed to return a fit, I forced the constant offset to equal the mean value of the last 5 ms of the shut-off PSTH. 4) In the case of noisy data, this second method could also fail to return a fit, in which case I forced both the initial value of the exponential fit to equal the mean of the PSTH over the 20 ms preceding the LED onset and the final value of the exponential fit to zero. 5) If all of these methods failed to fit the shut-off curve, I excluded the data. 6) All fits were verified by eye.

A.1.7 Identifying cortical layers

To correct for potential variation in the position of the 16 linear electrode channels relative to cortical layers across experiments, I used the current source density (CSD) to identify the site of the L4 current sink at stimulus onset and used this site as a reference. I verified the stability of this CSD-defined L4 sink relative to the cortical anatomy using conditional expression of virally injected ChR2 in the Scnn1a-Tg3-Cre mouse line (Figure A.10), which expresses Cre recombinase mainly in L4 excitatory neurons [127]. I first measured the CSD-defined L4 sink at stimulus onset and then photo-stimulated the ChR2-expressing L4 population on interleaved trials to determine the relationship between these two reference points (functionally-defined L4 sink and anatomically-defined Cre+ neuronal population). I found that the depth relationship was consistent across awake and anesthetized mice, thus validating the utility of the CSD-defined L4 sink as an alignment point. Finally, I assigned putative cortical layers (importantly, these are only rough estimates of the cortical layer boundaries) to the aligned average by stretching a template of putative cortical layer widths from slice experiments (quantification by Dr. Dante Bortone) across the depth of cortex, with L4 at the L4 CSD sink.

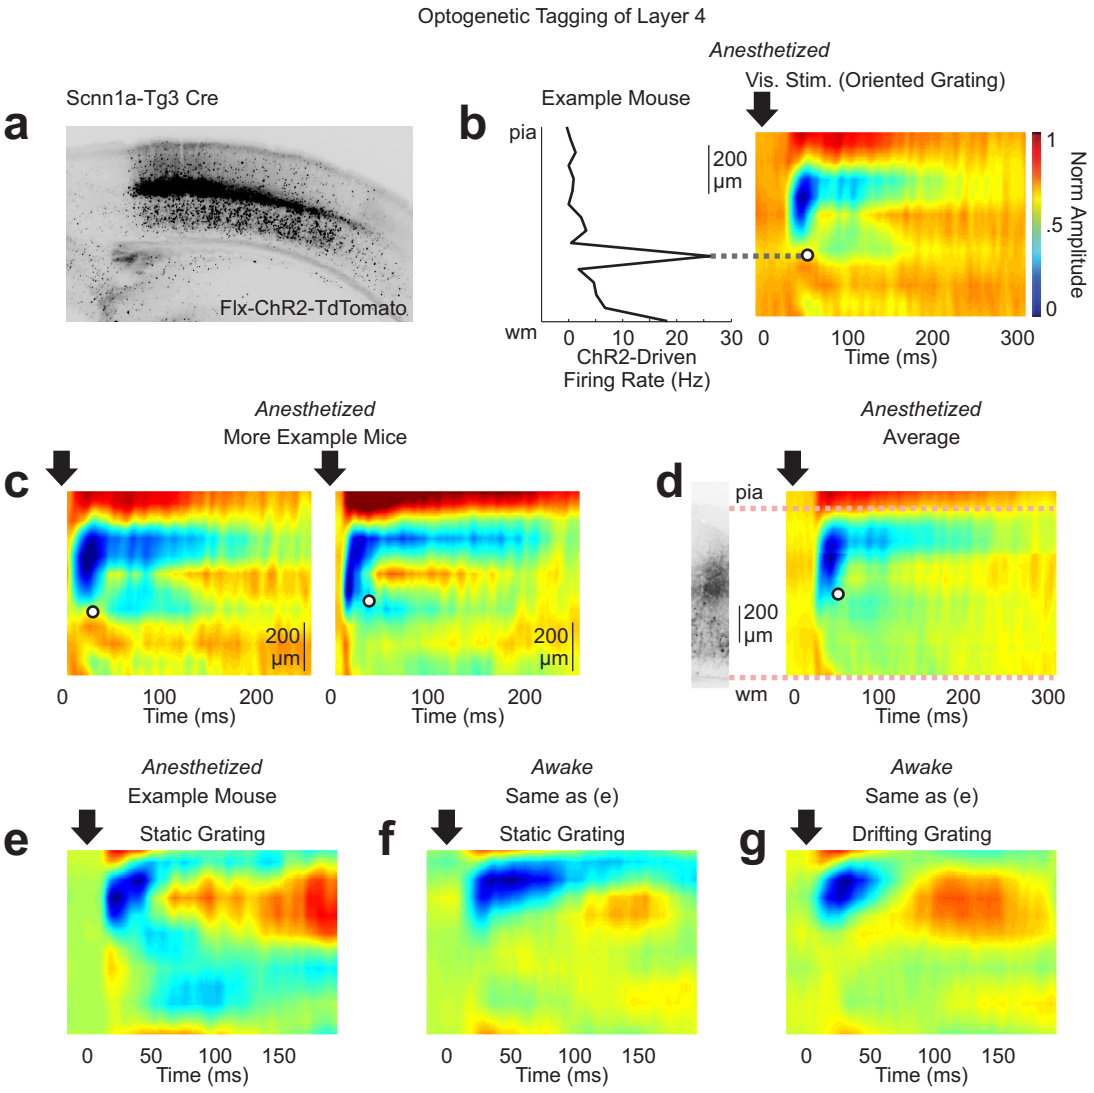

Figure A.10: Calibrating the current source density (CSD) with respect to anatomically defined layer 4 of visual cortex.
(a) Expression of Flx-ChR2-TdTomato in example Scnn1a-Tg3 Cre mouse (in this line, Cre specific to layer 4 in V1). Note strong labeling in V1 layer 4 cells. Neuropil (not cell bodies) label in other layers.
(b) Left: Change in multi-unit firing rate across cortical depths as a result of blue LED illumination of V1 of example Scnn1a-Tg3 mouse injected with Flx-ChR2-TdTomato. **Right:** Current source density (CSD) in same mouse in response to appearance of a moving grating (arrow). Current sinks are blue; current sources are red (CSD amplitude normalized for all heatmaps; same scale applies to all). Depths aligned with plot of ChR2-driven firing rates at left. Note that the depth of the peak ChR2-driven firing rate (dotted gray line) corresponds to the deeper edge of the earliest CSD current sink (open circle).
(c) As in (b) but showing 2 other example mice. Open circles are depths of maximal ChR2-driven firing rate in each mouse. Note relationship of this depth to CSD.
(d) Right: Average CSD across 3 mice in (a-b). CSDs aligned to depth of max ChR2-driven firing rate (open circles). **Left:** histology of Flx-ChR2 in L4 aligned to av. CSD at pia and white matter (wm) .
(e) CSD in example anesthetized mouse in response to appearance of static grating (SG).
(f) As in (e) but mouse is awake. Same mouse as in (e).
(g) As in (e) but mouse is awake and CSD is in response to appearance of moving grating. Same mouse as in (e). Note constant location of early current sink (blue) in CSDs in (e),(f) and (g).

A.2 Specific methods for Chapter 3

Unless described below, all experiments and analyses were performed as in Chapter 2.

A.2.1 Surgeries and animal preparation

Viral injections into V1 of Gad2-Cre mice

AAV2/1.CAGGS.flex.ChR2.tdTomato.SV40 was injected into V1 of adult Gad2-Cre x C57Bl6 transgenic mice. The virus injection surgery was the same as in Chapter 2. For AAV2/1.CAGGS.flex.ChR2.tdTomato.SV40 injections, I made a single craniotomy over the center of V1 (<50 µm in diameter) and injected 200 nL (titer: 6.86e12 genome copies/ml) of the virus at a depth of 400-600 µm. I recorded 2 weeks after injection.

A.2.2 Electrophysiology

Measuring the frequency response

Extracellular recordings from dLGN and V1 were performed as in Chapter 2.

Measuring thalamocortical synaptic depression *in vivo*

I isolated the thalamic component of the field excitatory postsynaptic potential ($fEPSP_{thal}$) recorded in Layer 4 as described in detail previously [30]. Briefly, I optogenetically silenced visual cortex by photo-activating GABAergic cortical interneurons conditionally expressing ChR2 (in Gad2-Cre mice with viral injections, see preceding section, or in vGat-ChR2 reporter mice) with a blue LED positioned above visual cortex while recording the $fEPSP_{thal}$ with a 16-channel linear extracellular array (NeuroNexus, A series). Visual stimuli consisted of brief pulse trains of blue light flashed at the contralateral eye ranging in frequency between 1 and 60 Hz (see section below "Visual stimulation: measuring the depression of thalamocortical synapses *in vivo*"). To separate the depression occurring at the thalamocortical synapse from any potential frequency-dependent changes occurring up-stream of the dLGN-V1 synapse (e.g., adaptation of the spike rate response in the dLGN), I also measured the spike rate response in the dLGN to the same visual pulse trains and normalized the $fEPSP_{thal}$ in V1 at each response number by the amplitude of that response in the dLGN multi-unit response. I also plotted the raw steady-state amplitude of the $fEPSP_{thal}$ response to the visual stimulus pulse train (Figure 3.3e-f), comparing anesthetized and awake recordings.

A.2.3 Post-mortem histology

Post-mortem histologic procedures were performed as in Chapter 2.

A.2.4 Visual stimulation

Visual stimulation: measuring the frequency response

I used the light of a blue light-emitting diode (LED) collimated through an objective lens (5x; 0.15 NA). The back aperture of this objective was pointed at the eye of the mouse to obtain a spatially unstructured, illuminated circular field at the eye (diameter approx. 3 cm at eye). The LED light source allowed me to produce precisely time-varying visual stimuli (frequencies between 1 and 60 Hz, or logarithmically modulated frequency sweeps, called chirps, see below), uniformly illuminating much of the visual field of the eye contralateral to the recording site. I modulated the LED with a computer-generated voltage signal. The power output of the LED varied sinusoidally between 0 and 25 mW (3.5 mW/cm^2) as the input signal varied in time.

I used two types of temporally modulated, spatially uniform full-field visual stimuli to drive activity in dLGN and V1:

1. Sine function flicker – Full-field modulations of luminance, varying sinusoidally in time over two seconds at a constant frequency, interleaved with time periods of constant illumination at mean luminance.
2. Logarithmic chirps – Full-field sinusoidal modulations of luminance whose frequency varied over two seconds, called frequency chirps. The frequency of the sine function describing luminance increased or decreased logarithmically over time. The logarithmic

chirp contained all frequencies between 1 and 30 Hz, each appearing instantaneously. To analyze the response to this chirp, I used Chronux to measure the amplitude of the neural response at each frequency, then took the amplitude of this response at the frequency of the stimulus as the fundamental response.

Visual stimulation: measuring the depression of thalamocortical synapses *in vivo*

I used the same LED visual stimulation set-up described above (preceding section) to present pulse trains of brief (10 ms duration, 3.5 mW/cm^2) flashes to the contralateral eye. The frequency of the pulse train varied between 1 and 60 Hz.

A.2.5 Optogenetic manipulations

Photo-activation of cortical interneurons to silence V1 recurrent excitatory circuits

A blue LED (455 nm) coupled to a 1 mm fiber was positioned several millimeters above the pial surface. Total power from the fiber was 25 mW. In LED blocks, the LED was turned on 645 ms prior to the onset of each block of visual stimulus and lasted the duration of the stimulus block.

A.2.6 Data analysis

Frequency response

Frequency response heat-map:

To compare the responses of dLGN and V1 to high-contrast, sinusoidally modulated, full-field luminance stimuli (called "flicker" stimuli) presented at various temporal frequencies (15 stimulus frequencies between 1 and 60 Hz), I used the frequency response. I computed the frequency response as follows. I used Chronux [158] to compute the amplitude spectrum (or power spectrum) of the trial-averaged steady-state PSTH (between 0.4 and 2 s after stimulus onset) in response to each stimulus frequency.

With this method, I picked out the frequency component of the response to a given stimulus as long as that component was phase-locked to the onset of the stimulus (i.e., apparent in the trial-averaged PSTH). I created heat-map matrices of the power at each frequency to summarize the neural responses in dLGN and V1 as a function of stimulus frequency (Figure 3.8b). When combining single units, each single unit's heat-map was first normalized to its peak along the diagonal, to equally weight all units in the combined heat-map, thus avoiding a larger contribution from cells with higher firing rates. I also exclude units that were not visually responsive. For dLGN, I show the median rather than the mean heat-map and frequency response to reduce biases from very high firing rate outlier units observed in dLGN (about 3% of total units). For V1, I show the mean (in Figure 3.3c,f,g) and the median (in Figure 3.4). For V1 units, the mean and median are very similar. However, I also show the full distributions for dLGN and V1 amplitudes as percentiles in Figure 3.4 to show that the full distributions are consistent with filtering by the CDF.

Frequency response 2D plot:

To obtain a simple 2D plot that captured how well the dLGN and V1 followed the stimulus input frequencies (Figures 3.3c,f,g, 3.4, 3.5 and 3.9), I plotted the amplitude of the neural response at each frequency of the stimulus. To obtain a meaningful zero value for this 2D plot, I subtracted the average amplitude of the non-specific response at all response frequencies above 70 Hz. This "non-specific response" was equivalent to the amplitude at frequencies other than the frequency of the input, when I did not observe harmonics. In the power heat-maps, for example, the non-specific power corresponds to the dark blue background. This non-specific response subtraction eliminates noise contributions to the final frequency response.

Frequency response prediction:

- Prediction in the time domain – I predicted the V1 PSTH (Figure 3.3b, pink traces) by convolving the dLGN PSTH traces (Figure 3.3b, top) with the regular-spiking (RS) empirical CDF of V1 measured during moving gratings (Figure 3.3c inset, Figure 2.15, Figure 3.1a). The CDF does not change with stimulus type for all stimulus types tested, so I chose this very clean measurement of the regular-spiking units' CDF to attempt to predict the regular-spiking units' frequency response in V1. For the example data in Figure 3.3b, to maintain the relative

amplitudes of the predicted responses across temporal frequencies of the visual stimulus, I scaled all the pink traces by the same amount to best match the actual V1 PSTH responses (black traces), by-eye. I then aligned each pink trace to each black trace by-eye.

- Prediction in the frequency domain – I predicted V1's frequency response by multiplying the dLGN's frequency response (Figure 3.3c, green) with the amplitude spectrum of the CDF (Figure 3.1b, pink). For the summary 2D single-line graphs (Figure 3.3c), I plotted the cortical prediction (pink) relative to the thalamic frequency response (green) such that an ideal frequency response curve representing "perfect following" of the thalamic input (i.e., filtered by an infinitesimally fast CDF) would lie on top of the plotted thalamic response. Because the proportionality factor relating firing rates in the dLGN to firing rates in V1 is not known, the absolute amplitude of the actual V1 frequency response (black) relative to the predicted V1 frequency response (pink) is arbitrary. Thus, in Figure 3.3 and elsewhere, where I compare the prediction to the actual frequency response in V1, I aligned the prediction for V1 (pink) to V1's actual frequency response (black) by scaling the heights of the curves to minimize the difference between the curves by eye, while also aligning these curves at the non-specific, baseline-subtracted zero (see above for explanation of zero in the frequency response). In Figure 3.3g, I use the same method just described to predict the frequency response curve that would result from replacing the real CDF with a 1 ms or 100 ms time constant. I use this same method to align the dLGN data, V1 data and prediction by the CDF in Figures 3.4, 3.5 and 3.9.

Models of fundamental (F1) and first harmonic (F2) responses in V1 (Figures 3.8 and 3.9)

In order to test whether a simple transformation of spatial receptive field (SRF) structures between dLGN and V1 could account for the decrease in V1 response amplitudes to high temporal frequencies with respect to the response in dLGN, I considered an alternate simple model: responses at the fundamental (F1) in dLGN can become responses at the first harmonic (F2) in V1. I found that this model alone, in the absence of low-pass filtering by the cortical decay function (CDF), was not able to account for the observed F1 and F2 responses in V1. Figures 3.8 and 3.9 show this test.

First, I measured the fundamental (response frequency matches stimulus frequency, called F1 modulation) and first harmonic (response frequency is twice the stimulus frequency, called F2 modulation) amplitudes in dLGN and V1 (Figure 3.8). I then attempted to fit three models, inspired by previous work on spatial receptive fields in the visual system [2] (Figure 3.9), to this data.

1. Model 1: Model 1 allowed F1 activity in dLGN to drive F1 activity in V1. F2 activity was not considered. Model 1 is the simple model presented in Figure 3.3. For discussion of Model 1 and filtering of the F1 response by the cortical decay function (CDF), see above.
2. Model 2A: Model 2A allowed F1 activity in dLGN to drive either F1 or F2 activity in V1. F2 activity in dLGN could also drive F2 activity in V1. In this model, *there is no filtering by the cortical decay function (CDF)*.
3. Model 2B: Model 2B allows the same F1-to-F2 transformations as Model 2A but also includes filtering by the CDF.

Equations for models of F1 and F2 Responses in dLGN and V1:

Let $d_{F1}(f)$ be the amplitude of dLGN's F1 response, as a function of the visual stimulus temporal frequency, f. Let $d_{F2}(f)$ be the amplitude of dLGN's F2 response, as a function of stimulus frequency. Similarly, let $V_{F1}(f)$ and $V_{F2}(f)$ be the amplitudes of V1's F1 and F2 responses, respectively, as functions of stimulus frequency. Let c_0 be a constant scaling factor between dLGN spiking and V1 spiking that is independent of stimulus frequency. Let $F_{NTC}(f_{resp})$ be the amplitude scaling produced by filtering thalamic activity at frequency f_{resp} with the CDF. Finally, let $D(f)$ be the fraction of the F1 amplitude in the dLGN at stimulus frequency f that is transformed to an F2 response in V1, without a loss of amplitude during this transformation. This is the simplest model of an F1-to-F2 transformation between dLGN and V1. Let $D(f)$ vary as a function of stimulus frequency. Then, we can write Models 2A and 2B as two systems of equations:

Model 2A – no filtering by CDF

(1) $V_{F1}(f) = c_0(1 - D(f))d_{F1}(f)$

(2) $V_{F2}(f) = c_0(d_{F2}(f) + D(f)d_{F1}(f))$

Model 2B – includes filtering by CDF

(1) $V_{F1}(f) = c_0F_{NTC}(f)(1 - D(f))d_{F1}(f)$

(2) $V_{F2}(f) = c_0F_{NTC}(2f)(d_{F2}(f) + D(f)d_{F1}(f))$

To test Models 2A and 2B, I solved each system of equations above for the fraction of the F1 response in dLGN that is frequency-doubled to produce F2 in V1 (a parameter I call $D(f)$, see above). I took the average estimate of $D(f)$ from each equation, forcing $D(f) \geq 0$ (Figure 3.9c). I then used this estimate of $D(f)$ to predict the F1 and F2 amplitudes for V1, and I compared these predictions to the real F1 and F2 responses in V1 (Figure 3.9). The important point is: if a simple solution exists to Model 2A or 2B, I should be able to find a value of $D(f)$ that accurately predicts both V1's F1 response and V1's F2 response. I was only able to find such a value for $D(f)$ in the context of Model 2B, which includes filtering by the CDF, indicating: first, Model 2A does not explain V1's response, and, second, V1's F1 and F2 responses are better explained by filtering with the CDF than by a simple, loss-less (no loss of amplitude) F1-to-F2 transformation between dLGN and V1.

Deconvolution of dLGN PSTH from V1 PSTH

In Figure 3.10, I approximate the temporal transformation of the sensory response between dLGN and V1 directly by computing the deconvolution of the dLGN response from the V1 response. I used Matlab to compute this deconvolution for each pair of responses in dLGN and V1 to the same visual stimulus (same temporal frequency of the flicker stimulus), across all possible pairs of recordings across all mice. I then computed the average and standard error of this deconvolution across all the pairwise dLGN-V1 comparisons. As Figure 3.10 shows, the result of the deconvolution is noisy, because the deconvolution of real neural signals is not ideal. Therefore this deconvolution should be viewed as a rough approximation. To summarize the difference in this result between anesthesia and awake, I fit a single exponential decay function to each deconvolution (black fits in Figure 3.10).

A.3 Specific methods for Chapter 4

Unless described below, all experiments and analyses were performed as in Chapters 2 and 3.

A.3.1 Surgeries and animal preparation

Viral injections into V1 of Gad2-Cre mice

AAV2/9.CAG.flex.Arch.GFP was injected into V1 of adult Gad2-Cre x C57Bl6 transgenic mice. The virus injection surgery was the same as in Chapters 2 and 3. However, for the AAV2/9.CAG.flex.Arch.GFP injections, I made 3 small craniotomies (<50 µm each in diameter) in a triangular pattern tiling the extent of mouse V1 and performed 3 pressure injections of 125 nL each (titer: 6x10e12 molecules/mL), at a depth of between 400 and 600 µm, at a speed of 30 nL/min, and waiting 10-15 min between injection sites. I recorded 4-6 weeks after injection.

Thinning the skull to enable photo-activation of ArchT in V1

Before performing the craniotomy to provide access to V1 for electrophysiology, I thinned the skull in the approximately 1.5 x 1.5 mm square overlying visual cortex. The thinned skull was then covered in ACSF. Together, the thinned skull and ACSF were transparent, allowing optical access of amber light to V1.

A.3.2 Electrophysiology

Extracellular electrophysiology to record from dLGN or V1 was performed as in Chapters 2 and 3.

A.3.3 Post-mortem histology

Post-mortem histologic procedures were performed as in Chapters 2 and 3.

A.3.4 Visual stimulation

Visual stimuli were presented as in Chapters 2 and 3. In Chapter 4, unless otherwise specified, the visual stimulus was the appearance of moving drifting gratings (contrast=1) following a mean-luminance gray screen. The moving grating drifted across the screen at a constant speed for 3 s before the gray screen returned (4.5 s inter-trial interval, ITI, between presentations of moving grating).

A.3.5 Optogenetic manipulations

Photo-activation of ArchT in V1 inhibitory interneurons

I used an amber LED (1 mm diameter, approx. 10-30 mW output, see following explanation) positioned above visual cortex to photo-activate ArchT expressed in cortical inhibitory interneurons. I titrated the intensity of the amber LED to maximally disinhibit the visually evoked response without triggering ictal events. This led to the clear suppression of a fraction of fast-spiking (FS) cells (putative interneurons [159]) and disinhibition of most regular-spiking (RS) cells (predominantly putative pyramidal neurons) (Figure A.11). I performed this experiment in anesthetized mice, because even slight disinhibition of the visual cortex in awake animals triggered ictal-like activity. To determine the CDF, I presented a visual stimulus (full-field moving grating, orientation varied at random) and silenced relay thalamus by activating the TRN with a blue LED (as in Chapter 2). I used 4 interleaved LED conditions: 1. No LED, 2. Blue LED, 3. Amber LED, and 4. Amber + Blue LED. I presented these LED conditions in the order "1, 2, 3, 4, 3, 2, 1, 4" to verify that any observed change in the CDF was not affected by the LED condition on the previous trial. The effect of the amber LED on the CDF was consistent over the course of the V1 recording. I observed the greatest slowing of the CDF in mice showing widespread expression of ArchT across visual cortex. I quantified the spread of ArchT expression as the "fractional coverage" of V1, that is, the fraction of V1 pixels presenting any detectable ArchT-GFP reporter fluorescence in post-mortem sections. I selected the 5 mice with the greatest ArchT fractional coverage of V1 for Figure 4.1. Importantly, when I included all the experiments in which I illuminated ArchT in V1, irrespective of fractional coverage (n=19 mice), suppression of inhibitory interneurons still produced a significant slowing of the cortical CDF (Figure

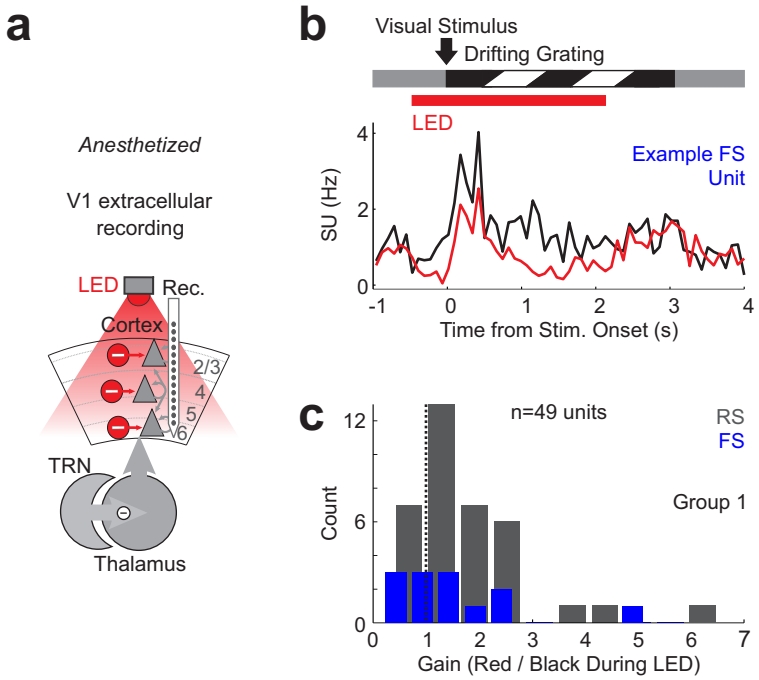

Figure A.11: Photo-activation of ArchT in V1 inhibitory interneurons suppresses a fraction of fast-spiking (FS) cells and disinhibits regular-spiking (RS) cells in V1.

(a) Schematic of cortical disinhibition by suppressing inhibitory interneurons expressing ArchT in V1. Red bar indicates amber LED illumination of V1.

(b) Response of example fast-spiking (FS) single unit. FS cells are putative inhibitory interneurons. Amber LED illumination of V1 (red bar) suppresses spiking of this unit.

(c) Histogram across single units showing the effects of ArchT photo-activation, by amber LED illumination of V1, on FS (blue) and regular-spiking (RS, gray) units in V1. Gain is the change in a unit's firing rate (not baseline-subtracted) during visual stimulation (moving grating, 3 s duration) as a result of ArchT photo-activation (i.e., firing rate during ArchT photo-activation divided by firing rate in control). Gain=1 indicates no change. Gain<1 indicates suppression. Gain>1 indicates disinhibition. Note that only a subset of FS units are suppressed, but most RS units are disinhibited. Most RS units are putative excitatory pyramidal cells.

4.3). I recorded in V1 for an hour before moving the electrode to the dLGN of the same mouse to verify no change in shut-off in the dLGN.

A.3.6 Data analysis

Predicting slowed cortical onset response when CDF is slowed

In Figure 4.7, I predict the change in the cortical response to the onset of a visual stimulus (moving grating) when the CDF is slowed by suppressing cortical inhibitory interneurons. To do this, I first consider the onset of the sensory-evoked response in V1 (blue line in Figure 4.7b marked "Control") to be the combination (convolution) of the time course of the response in the thalamus with the CDF under control conditions. Then I can ask: how does changing the CDF without changing the time course of response in the thalamus change the time course of response in the cortex? To analyze this, I first deconvolve the CDF under control conditions from the onset response in V1 under control conditions to arrive at the approximate onset response in the dLGN. Then I recombine (convolve) this onset response in the dLGN with the slower CDF under conditions of reduced cortical inhibition to predict V1's onset response under conditions of reduced cortical inhibition. This convolution gives the prediction (lighter purple line in Figure 4.7b, middle, marked "Prediction for disinhibition"). This prediction matches the actual time course (darker purple line in Figure 4.7b marked "V1 disinh.") of the V1 response to the onset of the visual stimulus under conditions of reduced cortical inhibition. I then plot how the onset response is slowed by V1 disinhibition as the time delay between the blue and dark purple onset curves in ms. If these curves are the same and not shifted in time, the time delay should be centered around 0. However, we see that the time delay is centered at a few ms lag, consistent with the prediction by the CDF.

Appendix B:
List of Visual Stimuli in Figures

- Figure 2.1a (center) – 1.7 s-long static oriented gratings and luminance steps
- Figure 2.1a (right) – 1.7 s-long static oriented gratings and luminance steps, and 250-ms long static oriented gratings
- Figure 2.1b (center) – 3 s-long static oriented gratings
- Figure 2.1b (right, top) – 3 s-long static oriented gratings
- Figure 2.1b (right, bottom, scatter plot) – cyan is 3 s-long static oriented grating, red is 3 s-long luminance step (black to white), dark blue is 10 to 100 ms-long flash of light, purple is reversal of checkerboard pattern every 3 s
- Figure 2.3c – 3 s-long moving oriented gratings
- Figure 2.3f-g – 3 s-long static oriented gratings and luminance steps
- Figure 2.3h – combine all stimulus types
- Figure 2.3i – 3 s-long moving oriented gratings vs. spontaneous activity
- Figure 2.4 – 3 s-long moving oriented gratings
- Figure 2.5 – onset of 3 s-long moving oriented gratings
- Figure 2.6 – 3 s-long moving oriented gratings
- Figure 2.8 (top) – 3 s-long static oriented gratings
- Figure 2.8 (2^{nd} row) – 3 s-long moving oriented gratings
- Figure 2.8 (3^{rd} row) – 3 s-long static oriented gratings and 3 s-long luminance steps
- Figure 2.8 (bottom, left) – 3 s-long luminance step
- Figure 2.8 (bottom, right) – 10 ms-long luminance step
- Figure 2.9 – 3 s-long moving oriented gratings
- Figure 2.10 – 3 s-long moving oriented gratings
- Figure 2.11 – 3 s-long moving oriented gratings; static stimuli including 3 s-long static oriented grating, 3 s-long luminance step and reversal of checkerboard pattern; or no visual stimulus
- Figure 2.13 – 3 s-long moving oriented gratings
- Figure 2.14 – 3 s-long moving oriented gratings
- Figure 2.15b – 3 s-long moving oriented gratings
- Figure 2.15c – 3 s-long moving oriented gratings, 3 s-long static oriented gratings, and 3 s-long luminance steps
- Figure 2.16 – 3 s-long moving oriented gratings; static stimuli including 3 s-long static oriented grating, 3 s-long luminance step and reversal of checkerboard pattern; or no visual stimulus
- Figure 2.17 – 3 s-long moving oriented gratings or no stimulus
- Figure 2.18b (top) – 3 s-long moving oriented gratings
- Figure 2.18b (bottom) – 3 s-long moving oriented gratings or 3 s-long static gratings
- Figure 2.18d – 3 s-long moving oriented gratings
- Figure 2.19 – 3 s-long moving oriented gratings
- Figure 2.20 – 3 s-long moving oriented gratings
- Figure 2.21 (left, right) – 3 s-long moving oriented gratings
- Figure 2.21 (center) – no visual stimulus
- Figure 2.22 – reversal of checkerboard pattern
- Figure 2.23 – 3 s-long moving oriented gratings or no stimulus
- Figure 2.24 – 3 s-long moving oriented gratings; static stimuli including 3 s-long static oriented grating, 3 s-long luminance step and reversal of checkerboard pattern; or no visual stimulus

- Figure 2.25 – 3 s-long moving oriented gratings; static stimuli including 3 s-long static oriented grating, 3 s-long luminance step and reversal of checkerboard pattern; or no visual stimulus
- Figure 2.26 – 3 s-long moving oriented gratings
- Figure 2.27 – 3 s-long moving oriented gratings
- Figure 3.2 – 10 ms-long luminance step
- Figure 3.3b-c – full-field flicker lasting for 2 s
- Figure 3.3e – repetitive train of full-field flashes (10 ms duration for each flash)
- Figure 3.3f-g – full-field flicker lasting for 2 s
- Figure 3.4 – high-contrast (contrast=1) full-field luminance flicker lasting 2 s
- Figure 3.5 – low-contrast (contrast=0.2) full-field luminance logarithmic frequency chirp, see Figure 3.7b "Vis. Stim."
- Figure 3.7 – high-contrast (contrast=1) full-field luminance logarithmic frequency chirp, see Figure 3.7b "Vis. Stim."
- Figure 3.8 – full-field high-contrast (contrast=1) luminance flicker lasting 2 s
- Figure 3.9 – full-field high-contrast (contrast=1) luminance flicker lasting 2 s
- Figure 3.10 – high-contrast (contrast=1) full-field luminance flicker lasting 2 s
- Figure 3.12 – repetitive train of full-field flashes (10 ms duration for each flash)
- Figure 3.13 – high-contrast (contrast=1) full-field luminance flicker lasting 2 s
- Figure 4.1 – 3 s-long moving oriented gratings
- Figure 4.2 – 3 s-long moving oriented gratings
- Figure 4.3 – 3 s-long moving oriented gratings
- Figure 4.4 – 3 s-long moving oriented gratings or no visual stimulus
- Figure 4.5 – 3 s-long moving oriented gratings
- Figure 4.6 – 3 s-long moving oriented gratings
- Figure 4.7 – 3 s-long moving oriented gratings (low=0.2 or high=1 contrast in (d), high contrast in rest of figure)
- Figure A.1 – 3 s-long moving oriented gratings or no visual stimulus
- Figure A.3 – 3 s-long moving oriented gratings
- Figure A.4 – 3 s-long moving oriented gratings
- Figure A.5 – 3 s-long moving oriented gratings
- Figure A.6 – 3 s-long moving oriented gratings or no visual stimulus
- Figure A.7 – 3 s-long moving oriented gratings
- Figure A.8 – 3 s-long moving oriented gratings
- Figure A.10 – 3 s-long static or moving oriented gratings
- Figure A.11 – 3 s-long moving oriented gratings

Appendix C:
Materials

Optogenetic stimulation

- Blue LED: LEDP-B_PF960-0.50-1m-FC_R2 from Doric Lenses, 470 nm
- Amber LED: LEDP_HB01-A_PF1000-0.50(-) from Doric Lenses, 595 nm
- Laser: Omicron, 80 mW max. output, 470 nm

Viruses

- AAV2/1.CAGGS.flex.ChR2.tdTomato.SV40 [124] (Addgene 18917) from the University of Pennsylvania viral vector core
- AAV2/9.CAG.flex.Arch.GFP [145] from the University of North Carolina viral vector core

Transgenic mouse lines

- Gad2-Cre (Jackson Labs Stock number: 010802) x C57Bl6 [119]
- vGat-ChR2 (Jackson Labs Stock number: 014548) [160]
- Scnn1a-Tg3-Cre (Jackson Labs Stock number: 009613) [161]
- PV-Cre (Jackson Labs Stock number: 008069) [162]

Appendix D:
Animal Use Statement

All procedures were conducted in accordance with the National Institutes of Health guidelines and with the approval of the Committee on Animal Care at UCSD (protocol S02160M). Animals were housed on a reverse light cycle in cages of 4 mice or less.

Appendix E:
Author Contributions

Kimberly Reinhold and Massimo Scanziani designed the study. Kimberly Reinhold conducted all experiments and analyses throughout the thesis, except the whole-cell recordings shown in Figure 2.1 of Chapter 2. Anthony Lien performed and analyzed the whole-cell recordings shown in Figure 2.1.

Bibliography

1 Ko, H., Hofer, S. B., Pichler, B., Buchanan, K. A., Sjostrom, P. J. & Mrsic-Flogel, T. D. Functional specificity of local synaptic connections in neocortical networks. *Nature* **473**, 87-91, doi:10.1038/nature09880 (2011).

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

3 Haider, B., Krause, M. R., Duque, A., Yu, Y., Touryan, J., Mazer, J. A. & McCormick, D. A. Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron* **65**, 107-121, doi:10.1016/j.neuron.2009.12.005 (2010).

4 Cossell, L., Iacaruso, M. F., Muir, D. R., Houlton, R., Sader, E. N., Ko, H., Hofer, S. B. & Mrsic-Flogel, T. D. Functional organization of excitatory synaptic strength in primary visual cortex. *Nature*, doi:10.1038/nature14182 (2015).

5 Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A. & Suarez, H. H. Recurrent excitation in neocortical circuits. *Science* **269**, 981-985 (1995).

6 Druckmann, S. & Chklovskii, D. B. Neuronal circuits underlying persistent representations despite time varying activity. *Current biology : CB* **22**, 2095-2103, doi:10.1016/j.cub.2012.08.058 (2012).

7 van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724-1726 (1996).

8 Ozeki, H., Finn, I. M., Schaffer, E. S., Miller, K. D. & Ferster, D. Inhibitory stabilization of the cortical network underlies visual surround suppression. *Neuron* **62**, 578-592, doi:10.1016/j.neuron.2009.03.028 (2009).

9 Hawken, M. J., Shapley, R. M. & Grosof, D. H. Temporal-frequency selectivity in monkey visual cortex. *Visual neuroscience* **13**, 477-492 (1996).

10 Movshon, J. A., Thompson, I. D. & Tolhurst, D. J. Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. *The Journal of physiology* **283**, 101-120 (1978).

11 DeAngelis, G. C., Ohzawa, I. & Freeman, R. D. Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *Journal of neurophysiology* **69**, 1091-1117 (1993).

12 Holub, R. A. & Morton-Gibson, M. Response of Visual Cortical Neurons of the cat to moving sinusoidal gratings: response-contrast functions and spatiotemporal interactions. *Journal of neurophysiology* **46**, 1244-1259 (1981).

13 Saul, A. B. & Humphrey, A. L. Temporal-frequency tuning of direction selectivity in cat visual cortex. *Visual neuroscience* **8**, 365-372 (1992).

14 O'Keefe, L. P., Levitt, J. B., Kiper, D. C., Shapley, R. M. & Movshon, J. A. Functional organization of owl monkey lateral geniculate nucleus and visual cortex. *Journal of neurophysiology* **80**, 594-609 (1998).

15 Foster, K. H., Gaska, J. P., Nagler, M. & Pollen, D. A. Spatial and temporal frequency selectivity of neurones in visual cortical areas V1 and V2 of the macaque monkey. *The Journal of physiology* **365**, 331-363 (1985).

16 MacLean, J. N., Watson, B. O., Aaron, G. B. & Yuste, R. Internal dynamics determine the cortical response to thalamic stimulation. *Neuron* **48**, 811-823, doi:10.1016/j.neuron.2005.09.035 (2005).

17 Miller, J. E., Ayzenshtat, I., Carrillo-Reid, L. & Yuste, R. Visual stimuli recruit intrinsically generated cortical ensembles. *Proceedings of the National Academy of Sciences of the United States of America* **111**, E4053-4061, doi:10.1073/pnas.1406077111 (2014).

18 Xing, D., Yeh, C. I., Burns, S. & Shapley, R. M. Laminar analysis of visually evoked activity in the primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America* **109**, 13871-13876, doi:10.1073/pnas.1201478109 (2012).

19 Gil, Z., Connors, B. W. & Amitai, Y. Differential regulation of neocortical synapses by neuromodulators and activity. *Neuron* **19**, 679-686 (1997).

20 Kloc, M. & Maffei, A. Target-Specific Properties of Thalamocortical Synapses onto Layer 4 of Mouse Primary Visual Cortex. *The Journal of Neuroscience* **34**, 15455-15465, doi:10.1523/jneurosci.2595-14.2014 (2014).

21 Castro-Alamancos, M. A. & Oldford, E. Cortical sensory suppression during arousal is due to the activity-dependent depression of thalamocortical synapses. *The Journal of physiology* **541**, 319-331 (2002).

22 Borst, J. G. The low synaptic release probability in vivo. *Trends in neurosciences* **33**, 259-266, doi:10.1016/j.tins.2010.03.003 (2010).

23 Swadlow, H. A., Gusev, A. G. & Bezdudnaya, T. Activation of a cortical column by a thalamocortical impulse. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **22**, 7766-7773 (2002).

24 da Costa, N. M. & Martin, K. A. The proportion of synapses formed by the axons of the lateral geniculate nucleus in layer 4 of area 17 of the cat. *The Journal of comparative neurology* **516**, 264-276, doi:10.1002/cne.22133 (2009).

25 McGuire, B. A., Hornung, J. P., Gilbert, C. D. & Wiesel, T. N. Patterns of synaptic input to layer 4 of cat striate cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **4**, 3021-3033 (1984).

26 Peters, A. & Payne, B. R. Numerical relationships between geniculocortical afferents and pyramidal cell modules in cat primary visual cortex. *Cerebral cortex* **3**, 69-78 (1993).

27 Feldmeyer, D. Excitatory neuronal connectivity in the barrel cortex. *Front Neuroanat* **6**, 24, doi:10.3389/fnana.2012.00024 (2012).

28 Gilbert, C. D. & Wiesel, T. N. Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature* **280**, 120-125 (1979).

29 Ferster, D., Chung, S. & Wheat, H. Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature* **380**, 249-252, doi:10.1038/380249a0 (1996).

30 Lien, A. D. & Scanziani, M. Tuned thalamic excitation is amplified by visual cortical circuits. *Nature neuroscience* **16**, 1315-1323, doi:10.1038/nn.3488 (2013).

31 Li, Y. T., Ibrahim, L. A., Liu, B. H., Zhang, L. I. & Tao, H. W. Linear transformation of thalamocortical input by intracortical excitation. *Nature neuroscience* **16**, 1324-1330, doi:10.1038/nn.3494 (2013).

32 Li, L. Y., Li, Y. T., Zhou, M., Tao, H. W. & Zhang, L. I. Intracortical multiplication of thalamocortical signals in mouse auditory cortex. *Nature neuroscience* **16**, 1179-1181, doi:10.1038/nn.3493 (2013).

33 Gilbert, C. D. & Wiesel, T. N. Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **9**, 2432-2442 (1989).

34 Yabuta, N. H. & Callaway, E. M. Cytochrome-oxidase blobs and intrinsic horizontal connections of layer 2/3 pyramidal neurons in primate V1. *Visual neuroscience* **15**, 1007-1027 (1998).

35 Malach, R., Amir, Y., Harel, M. & Grinvald, A. Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex. *Proceedings of the National Academy of Sciences of the United States of America* **90**, 10469-10473 (1993).

36 Douglas, R. J. & Martin, K. A. Neuronal circuits of the neocortex. *Annual review of neuroscience* **27**, 419-451, doi:10.1146/annurev.neuro.27.070203.144152 (2004).

37 Callaway, E. M. Feedforward, feedback and inhibitory connections in primate visual cortex. *Neural Netw* **17**, 625-632, doi:10.1016/j.neunet.2004.04.004 (2004).

38 Yabuta, N. H. & Callaway, E. M. Functional streams and local connections of layer 4C neurons in primary visual cortex of the macaque monkey. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **18**, 9489-9499 (1998).

39 Hubel, D. H. & Wiesel, T. N. Sequence regularity and geometry of orientation columns in the monkey striate cortex. *The Journal of comparative neurology* **158**, 267-293, doi:10.1002/cne.901580304 (1974).

40 Ringach, D. L., Hawken, M. J. & Shapley, R. Dynamics of orientation tuning in macaque primary visual cortex. *Nature* **387**, 281-284, doi:10.1038/387281a0 (1997).

41 Douglas, R. J. & Martin, K. A. Recurrent neuronal circuits in the neocortex. *Current biology : CB* **17**, R496-500, doi:10.1016/j.cub.2007.04.024 (2007).

42 Callaway, E. M. Local circuits in primary visual cortex of the macaque monkey. *Annual review of neuroscience* **21**, 47-74, doi:10.1146/annurev.neuro.21.1.47 (1998).

43 Kara, P., Reinagel, P. & Reid, R. C. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* **27**, 635-646 (2000).

44 Cossart, R., Aronov, D. & Yuste, R. Attractor dynamics of network UP states in the neocortex. *Nature* **423**, 283-288, doi:10.1038/nature01614 (2003).

45 Chubykin, A. A., Roach, E. B., Bear, M. F. & Shuler, M. G. A cholinergic mechanism for reward timing within primary visual cortex. *Neuron* **77**, 723-735, doi:10.1016/j.neuron.2012.12.039 (2013).

46 Nauhaus, I., Busse, L., Carandini, M. & Ringach, D. L. Stimulus contrast modulates functional connectivity in visual cortex. *Nature neuroscience* **12**, 70-76, doi:10.1038/nn.2232 (2009).

47 Krukowski, A. E. & Miller, K. D. Thalamocortical NMDA conductances and intracortical inhibition can explain cortical temporal tuning. *Nature neuroscience* **4**, 424-430, doi:10.1038/86084 (2001).

48 Crowder, N. A., van Kleef, J., Dreher, B. & Ibbotson, M. R. Complex cells increase their phase sensitivity at low contrasts and following adaptation. *Journal of neurophysiology* **98**, 1155-1166, doi:10.1152/jn.00433.2007 (2007).

49 Niell, C. M. & Stryker, M. P. Highly selective receptive fields in mouse visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **28**, 7520-7536, doi:10.1523/JNEUROSCI.0623-08.2008 (2008).

50 Andermann, M. L., Kerlin, A. M., Roumis, D. K., Glickfeld, L. L. & Reid, R. C. Functional specialization of mouse higher visual cortical areas. *Neuron* **72**, 1025-1039, doi:10.1016/j.neuron.2011.11.013 (2011).

51 Zhuang, J., Bereshpolova, Y., Stoelzel, C. R., Huff, J. M., Hei, X., Alonso, J. M. & Swadlow, H. A. Brain state effects on layer 4 of the awake visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **34**, 3888-3900, doi:10.1523/JNEUROSCI.4969-13.2014 (2014).

52 Reinagel, P. & Reid, R. C. Precise firing events are conserved across neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **22**, 6837-6841, doi:20026633 (2002).

53 Gabernet, L., Jadhav, S. P., Feldman, D. E., Carandini, M. & Scanziani, M. Somatosensory integration controlled by dynamic thalamocortical feed-forward inhibition. *Neuron* **48**, 315-327, doi:10.1016/j.neuron.2005.09.022 (2005).

54 Bair, W. & Koch, C. Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput* **8**, 1185-1202 (1996).

55 Pinto, D. J., Hartings, J. A., Brumberg, J. C. & Simons, D. J. Cortical damping: analysis of thalamocortical response transformations in rodent barrel cortex. *Cerebral cortex* **13**, 33-44 (2003).

56 Saul, A. B. & Humphrey, A. L. Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *Journal of neurophysiology* **64**, 206-224 (1990).

57 Niell, C. M. & Stryker, M. P. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**, 472-479, doi:10.1016/j.neuron.2010.01.033 (2010).

58 Grubb, M. S. & Thompson, I. D. Quantitative characterization of visual response properties in the mouse dorsal lateral geniculate nucleus. *Journal of neurophysiology* **90**, 3594-3607, doi:10.1152/jn.00699.2003 (2003).

59 Shu, Y., Hasenstaub, A. & McCormick, D. A. Turning on and off recurrent balanced cortical activity. *Nature* **423**, 288-293, doi:10.1038/nature01616 (2003).

60 McCormick, D. A., Shu, Y., Hasenstaub, A., Sanchez-Vives, M., Badoual, M. & Bal, T. Persistent cortical activity: mechanisms of generation and effects on neuronal excitability. *Cerebral cortex* **13**, 1219-1231 (2003).

61 Arabzadeh, E., Zorzin, E. & Diamond, M. E. Neuronal encoding of texture in the whisker sensory pathway. *PLoS biology* **3**, e17, doi:10.1371/journal.pbio.0030017 (2005).

62 Rager, G. & Singer, W. The response of cat visual cortex to flicker stimuli of variable frequency. *The European journal of neuroscience* **10**, 1856-1877 (1998).

63 Wehr, M. & Zador, A. M. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* **426**, 442-446, doi:10.1038/nature02116 (2003).

64 Hasenstaub, A., Shu, Y., Haider, B., Kraushaar, U., Duque, A. & McCormick, D. A. Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron* **47**, 423-435, doi:10.1016/j.neuron.2005.06.016 (2005).

65 Buracas, G. T., Zador, A. M., DeWeese, M. R. & Albright, T. D. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* **20**, 959-969 (1998).

66 Haider, B. & McCormick, D. A. Rapid neocortical dynamics: cellular and network mechanisms. *Neuron* **62**, 171-189, doi:10.1016/j.neuron.2009.04.008 (2009).

67 Li, R. & Wang, Y. Neural mechanism for sensing fast motion in dim light. *Scientific reports* **3**, 3159, doi:10.1038/srep03159 (2013).

68 Goard, M. & Dan, Y. Basal forebrain activation enhances cortical coding of natural scenes. *Nature neuroscience* **12**, 1444-1449, doi:10.1038/nn.2402 (2009).

69 Bignall, K. E. & Rutledge, L. T. Origin of a Photically Evoked Afterdischarge in Cat Visual Cortex. *Journal of neurophysiology* **27**, 1048-1062 (1964).

70 Benucci, A., Ringach, D. L. & Carandini, M. Coding of stimulus sequences by population responses in visual cortex. *Nature neuroscience* **12**, 1317-1324, doi:10.1038/nn.2398 (2009).

71 Honkanen, R., Rouhinen, S., Wang, S. H., Palva, J. M. & Palva, S. Gamma Oscillations Underlie the Maintenance of Feature-Specific Information and the Contents of Visual Working Memory. *Cerebral cortex*, doi:10.1093/cercor/bhu263 (2014).

72 Wang, X. J. Synaptic reverberation underlying mnemonic persistent activity. *Trends in neurosciences* **24**, 455-463 (2001).

73 Nikonov, S. S., Kholodenko, R., Lem, J. & Pugh, E. N., Jr. Physiological features of the S- and M-cone photoreceptors of wild-type mice from single-cell recordings. *J Gen Physiol* **127**, 359-374, doi:10.1085/jgp.200609490 (2006).

74 Cleland, B. G., Dubin, M. W. & Levick, W. R. Sustained and transient neurones in the cat's retina and lateral geniculate nucleus. *The Journal of physiology* **217**, 473-496 (1971).

75 Steriade, M., McCormick, D. A. & Sejnowski, T. J. Thalamocortical oscillations in the sleeping and aroused brain. *Science* **262**, 679-685 (1993).

76 Steriade, M., Timofeev, I. & Grenier, F. Natural waking and sleep states: a view from inside neocortical neurons. *Journal of neurophysiology* **85**, 1969-1985 (2001).

77 Olsen, S. R., Bortone, D. S., Adesnik, H. & Scanziani, M. Gain control by layer six in cortical circuits of vision. *Nature* **483**, 47-52, doi:10.1038/nature10835 (2012).

78 Briggs, F. & Usrey, W. M. Emerging views of corticothalamic function. *Curr Opin Neurobiol* **18**, 403-407, doi:10.1016/j.conb.2008.09.002 (2008).

79 Cudeiro, J. & Sillito, A. M. Looking back: corticothalamic feedback and early visual processing. *Trends in neurosciences* **29**, 298-306, doi:10.1016/j.tins.2006.05.002 (2006).

80 Crandall, S. R., Cruikshank, S. J. & Connors, B. W. A corticothalamic switch: controlling the thalamus with dynamic synapses. *Neuron* **86**, 768-782, doi:10.1016/j.neuron.2015.03.040 (2015).

81 von Krosigk, M., Monckton, J. E., Reiner, P. B. & McCormick, D. A. Dynamic properties of corticothalamic excitatory postsynaptic potentials and thalamic reticular inhibitory postsynaptic potentials in thalamocortical neurons of the guinea-pig dorsal lateral geniculate nucleus. *Neuroscience* **91**, 7-20 (1999).

82 Takahashi, T. The organization of the lateral thalamus of the hooded rat. *The Journal of comparative neurology* **231**, 281-309, doi:10.1002/cne.902310302 (1985).

83 Caviness, V. S., Jr. & Frost, D. O. Tangential organization of thalamic projections to the neocortex in the mouse. *The Journal of comparative neurology* **194**, 335-367, doi:10.1002/cne.901940205 (1980).

84 Frost, D. O. & Caviness, V. S., Jr. Radial organization of thalamic projections to the neocortex in the mouse. *The Journal of comparative neurology* **194**, 369-393, doi:10.1002/cne.901940206 (1980).

85 Bourassa, J. & Deschenes, M. Corticothalamic projections from the primary visual cortex in rats: a single fiber study using biocytin as an anterograde tracer. *Neuroscience* **66**, 253-263 (1995).

86 Sherman, S. M. The thalamus is more than just a relay. *Curr Opin Neurobiol* **17**, 417-422, doi:10.1016/j.conb.2007.07.003 (2007).

87 Johnson, R. R. & Burkhalter, A. Evidence for excitatory amino acid neurotransmitters in the geniculo-cortical pathway and local projections within rat primary visual cortex. *Exp Brain Res* **89**, 20-30 (1992).

88 Sanchez-Vives, M. V. & McCormick, D. A. Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nature neuroscience* **3**, 1027-1034, doi:10.1038/79848 (2000).

89 Steriade, M., Nunez, A. & Amzica, F. A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **13**, 3252-3265 (1993).

90 Lampl, I., Reichova, I. & Ferster, D. Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron* **22**, 361-374 (1999).

91 Constantinople, C. M. & Bruno, R. M. Effects and mechanisms of wakefulness on local cortical networks. *Neuron* **69**, 1061-1068, doi:10.1016/j.neuron.2011.02.040 (2011).

92 Timofeev, I., Grenier, F., Bazhenov, M., Sejnowski, T. J. & Steriade, M. Origin of slow cortical oscillations in deafferented cortical slabs. *Cerebral cortex* **10**, 1185-1199 (2000).

93 Compte, A., Sanchez-Vives, M. V., McCormick, D. A. & Wang, X. J. Cellular and network mechanisms of slow oscillatory activity (<1 Hz) and wave propagations in a cortical network model. *Journal of neurophysiology* **89**, 2707-2725, doi:10.1152/jn.00845.2002 (2003).

94 Poulet, J. F. & Petersen, C. C. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* **454**, 881-885, doi:10.1038/nature07150 (2008).

95 Harris, K. D. & Thiele, A. Cortical state and attention. *Nat Rev Neurosci* **12**, 509-523, doi:10.1038/nrn3084 (2011).

96 Rigas, P. & Castro-Alamancos, M. A. Thalamocortical Up states: differential effects of intrinsic and extrinsic cortical inputs on persistent activity. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **27**, 4261-4272, doi:10.1523/JNEUROSCI.0003-07.2007 (2007).

97 Cruikshank, S. J., Urabe, H., Nurmikko, A. V. & Connors, B. W. Pathway-specific feedforward circuits between thalamus and neocortex revealed by selective optical stimulation of axons. *Neuron* **65**, 230-245, doi:10.1016/j.neuron.2009.12.025 (2010).

98 Cruikshank, S. J., Lewis, T. J. & Connors, B. W. Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex. *Nature neuroscience* **10**, 462-468, doi:10.1038/nn1861 (2007).

99 Beltramo, R., D'Urso, G., Dal Maschio, M., Farisello, P., Bovetti, S., Clovis, Y., Lassi, G., Tucci, V., De Pietri Tonelli, D. & Fellin, T. Layer-specific excitatory circuits differentially control recurrent network dynamics in the neocortex. *Nature neuroscience* **16**, 227-234, doi:10.1038/nn.3306 (2013).

100 Chauvette, S., Volgushev, M. & Timofeev, I. Origin of active states in local neocortical networks during slow sleep oscillation. *Cerebral cortex* **20**, 2660-2674, doi:10.1093/cercor/bhq009 (2010).

101 Amzica, F. & Steriade, M. Cellular substrates and laminar profile of sleep K-complex. *Neuroscience* **82**, 671-686 (1998).

102 Amzica, F. & Steriade, M. Short- and long-range neuronal synchronization of the slow (< 1 Hz) cortical oscillation. *Journal of neurophysiology* **73**, 20-38 (1995).

103 Mountcastle, V. B. Modality and topographic properties of single neurons of cat's somatic sensory cortex. *Journal of neurophysiology* **20**, 408-434 (1957).

104 Blasdel, G. G. & Salama, G. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature* **321**, 579-585, doi:10.1038/321579a0 (1986).

105 Wiesel, T. N., Hubel, D. H. & Lam, D. M. Autoradiographic demonstration of ocular-dominance columns in the monkey striate cortex by means of transneuronal transport. *Brain Res* **79**, 273-279 (1974).

106 Jones, E. G. Microcolumns in the cerebral cortex. *Proceedings of the National Academy of Sciences of the United States of America* **97**, 5019-5021 (2000).

107 Yuste, R., Peinado, A. & Katz, L. C. Neuronal domains in developing neocortex. *Science* **257**, 665-669 (1992).

108 Adesnik, H. & Scanziani, M. Lateral competition for cortical space by layer-specific horizontal circuits. *Nature* **464**, 1155-1160, doi:10.1038/nature08935 (2010).

109 Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D. & Yuste, R. Synfire chains and cortical songs: temporal modules of cortical activity. *Science* **304**, 559-564, doi:10.1126/science.1093173 (2004).

110 Camperi, M. & Wang, X. J. A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *Journal of computational neuroscience* **5**, 383-405 (1998).

111 Miller, P., Brody, C. D., Romo, R. & Wang, X. J. A recurrent network model of somatosensory parametric working memory in the prefrontal cortex. *Cerebral cortex* **13**, 1208-1218 (2003).

112 Tsodyks, M., Kenet, T., Grinvald, A. & Arieli, A. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* **286**, 1943-1946 (1999).

113 Swadlow, H. A., Bezdudnaya, T. & Gusev, A. G. Spike timing and synaptic dynamics at the awake thalamocortical synapse. *Prog Brain Res* **149**, 91-105, doi:10.1016/S0079-6123(05)49008-1 (2005).

114 Boudreau, C. E. & Ferster, D. Short-term depression in thalamocortical synapses of cat primary visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **25**, 7179-7190, doi:10.1523/JNEUROSCI.1445-05.2005 (2005).

115 Wu, L. G. & Borst, J. G. The reduced release probability of releasable vesicles during recovery from short-term synaptic depression. *Neuron* **23**, 821-832 (1999).

116 Tsodyks, M. V. & Markram, H. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences of the United States of America* **94**, 719-723 (1997).

117 Chance, F. S., Nelson, S. B. & Abbott, L. F. Synaptic depression and the temporal response characteristics of V1 cells. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **18**, 4785-4799 (1998).

118 Steriade, M. The GABAergic reticular nucleus: a preferential target of corticothalamic projections. *Proceedings of the National Academy of Sciences of the United States of America* **98**, 3625-3627, doi:10.1073/pnas.071051998 (2001).

119 Taniguchi, H., He, M., Wu, P., Kim, S., Paik, R., Sugino, K., Kvitsiani, D., Fu, Y., Lu, J., Lin, Y., Miyoshi, G., Shima, Y., Fishell, G., Nelson, S. B. & Huang, Z. J. A resource of Cre driver

lines for genetic targeting of GABAergic neurons in cerebral cortex. *Neuron* **71**, 995-1013, doi:10.1016/j.neuron.2011.07.026 (2011).

120 Madisen, L., Mao, T., Koch, H., Zhuo, J. M., Berenyi, A., Fujisawa, S., Hsu, Y. W., Garcia, A. J., 3rd, Gu, X., Zanella, S., Kidney, J., Gu, H., Mao, Y., Hooks, B. M., Boyden, E. S., Buzsaki, G., Ramirez, J. M., Jones, A. R., Svoboda, K., Han, X., Turner, E. E. & Zeng, H. A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nature neuroscience* **15**, 793-802, doi:10.1038/nn.3078 (2012).

121 Tanahira, C., Higo, S., Watanabe, K., Tomioka, R., Ebihara, S., Kaneko, T. & Tamamaki, N. Parvalbumin neurons in the forebrain as revealed by parvalbumin-Cre transgenic mice. *Neurosci Res* **63**, 213-223, doi:10.1016/j.neures.2008.12.007 (2009).

122 Contreras, D. & Llinas, R. Voltage-sensitive dye imaging of neocortical spatiotemporal dynamics to afferent activation frequency. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **21**, 9403-9413 (2001).

123 Poulet, J. F., Fernandez, L. M., Crochet, S. & Petersen, C. C. Thalamic control of cortical states. *Nature neuroscience* **15**, 370-372, doi:10.1038/nn.3035 (2012).

124 Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G. & Deisseroth, K. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature neuroscience* **8**, 1263-1268, doi:10.1038/nn1525 (2005).

125 Halassa, M. M., Siegle, J. H., Ritt, J. T., Ting, J. T., Feng, G. & Moore, C. I. Selective optical drive of thalamic reticular nucleus generates thalamic bursts and cortical spindles. *Nature neuroscience* **14**, 1118-1120, doi:10.1038/nn.2880 (2011).

126 Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D. & Wang, X. J. A hierarchy of intrinsic timescales across primate cortex. *Nature neuroscience* **17**, 1661-1663, doi:10.1038/nn.3862 (2014).

127 Adesnik, H., Bruns, W., Taniguchi, H., Huang, Z. J. & Scanziani, M. A neural circuit for spatial summation in visual cortex. *Nature* **490**, 226-231, doi:10.1038/nature11526 (2012).

128 Wang, Q., Gao, E. & Burkhalter, A. In vivo transcranial imaging of connections in mouse visual cortex. *Journal of neuroscience methods* **159**, 268-276, doi:10.1016/j.jneumeth.2006.07.024 (2007).

129 Buzsaki, G. & Draguhn, A. Neuronal oscillations in cortical networks. *Science* **304**, 1926-1929, doi:10.1126/science.1099745 (2004).

130 Cardin, J. A., Carlen, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L. H. & Moore, C. I. Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature* **459**, 663-667, doi:10.1038/nature08002 (2009).

131 Fregnac, Y., Monier, C., Chavane, F., Baudot, P. & Graham, L. Shunting inhibition, a silent step in visual cortical computation. *J Physiol Paris* **97**, 441-451, doi:10.1016/j.jphysparis.2004.02.004 (2003).

132 Browning, P. G., Chakraborty, S. & Mitchell, A. S. Evidence for Mediodorsal Thalamus and Prefrontal Cortex Interactions during Cognition in Macaques. *Cerebral cortex*, doi:10.1093/cercor/bhv093 (2015).

133 Movshon, J. A., Thompson, I. D. & Tolhurst, D. J. Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *The Journal of physiology* **283**, 53-77 (1978).

134 Carandini, M. & Ferster, D. A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science* **276**, 949-952 (1997).

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

136 Bezdudnaya, T., Cano, M., Bereshpolova, Y., Stoelzel, C. R., Alonso, J. M. & Swadlow, H. A. Thalamic burst mode and inattention in the awake LGNd. *Neuron* **49**, 421-432, doi:10.1016/j.neuron.2006.01.010 (2006).

137 Piscopo, D. M., El-Danaf, R. N., Huberman, A. D. & Niell, C. M. Diverse visual features encoded in mouse lateral geniculate nucleus. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **33**, 4642-4656, doi:10.1523/JNEUROSCI.5187-12.2013 (2013).

138 Stark, E., Eichler, R., Roux, L., Fujisawa, S., Rotstein, H. G. & Buzsaki, G. Inhibition-induced theta resonance in cortical circuits. *Neuron* **80**, 1263-1276, doi:10.1016/j.neuron.2013.09.033 (2013).

139 Troy, J. B. & Enroth-Cugell, C. X and Y ganglion cells inform the cat's brain about contrast in the retinal image. *Exp Brain Res* **93**, 383-390 (1993).

140 Maffei, L. & Fiorentini, A. The visual cortex as a spatial frequency analyser. *Vision Res* **13**, 1255-1267 (1973).

141 Morrone, M. C., Burr, D. C. & Maffei, L. Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. *Proc R Soc Lond B Biol Sci* **216**, 335-354 (1982).

142 Higley, M. J. & Contreras, D. Balanced excitation and inhibition determine spike timing during frequency adaptation. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **26**, 448-457, doi:10.1523/JNEUROSCI.3506-05.2006 (2006).

143 Brody, C. D., Romo, R. & Kepecs, A. Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Curr Opin Neurobiol* **13**, 204-211 (2003).

144 Atallah, B. V., Bruns, W., Carandini, M. & Scanziani, M. Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. *Neuron* **73**, 159-170, doi:10.1016/j.neuron.2011.12.013 (2012).

145 Han, X., Chow, B. Y., Zhou, H., Klapoetke, N. C., Chuong, A., Rajimehr, R., Yang, A., Baratta, M. V., Winkle, J., Desimone, R. & Boyden, E. S. A high-light sensitivity optical neural silencer: development and application to optogenetic control of non-human primate cortex. *Frontiers in systems neuroscience* **5**, 18, doi:10.3389/fnsys.2011.00018 (2011).

146 Fuster, J. M. Memory networks in the prefrontal cortex. *Prog Brain Res* **122**, 309-316 (2000).

147 Fuster, J. M. & Alexander, G. E. Neuron activity related to short-term memory. *Science* **173**, 652-654 (1971).

148 Romo, R., Brody, C. D., Hernandez, A. & Lemus, L. Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* **399**, 470-473, doi:10.1038/20939 (1999).

149 Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annual review of neuroscience* **30**, 535-574, doi:10.1146/annurev.neuro.29.051605.113038 (2007).

150 Sherman, S. M. Thalamocortical interactions. *Curr Opin Neurobiol* **22**, 575-579, doi:10.1016/j.conb.2012.03.005 (2012).

151 David, S. V., Vinje, W. E. & Gallant, J. L. Natural stimulus statistics alter the receptive field structure of v1 neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **24**, 6991-7006, doi:10.1523/JNEUROSCI.1422-04.2004 (2004).

152 Muller, J. R., Metha, A. B., Krauskopf, J. & Lennie, P. Information conveyed by onset transients in responses of striate cortical neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **21**, 6978-6990 (2001).

153 Nathan, J., Reh, R., Ankoudinova, I., Ankoudinova, G., Chang, B., Heckenlively, J. & Hurley, J. B. Scotopic and photopic visual thresholds and spatial and temporal discrimination evaluated by behavior of mice in a water maze. *Photochem Photobiol* **82**, 1489-1494, doi:10.1562/2006-02-27-RA-818 (2006).

154 Levinson, J. Z. Flicker fusion phenomena. *Science* **160**, 21-28 (1968).

155 Murray, J. D., Anticevic, A., Gancsos, M., Ichinose, M., Corlett, P. R., Krystal, J. H. & Wang, X. J. Linking microcircuit dysfunction to cognitive impairment: effects of disinhibition associated with schizophrenia in a cortical working memory model. *Cerebral cortex* **24**, 859-872, doi:10.1093/cercor/bhs370 (2014).

156 Paxinos, G. & Franklin, K. B. J. *The mouse brain in stereotaxic coordinates*. Compact 2nd edn, (Elsevier Academic Press, 2004).

157 Fee, M. S., Mitra, P. P. & Kleinfeld, D. Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *Journal of neuroscience methods* **69**, 175-188, doi:10.1016/S0165-0270(96)00050-7 (1996).

158 Mitra, P. & Bokil, H. *Observed brain dynamics*. (Oxford University Press, 2008).

159 Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G. & Wu, C. Interneurons of the neocortical inhibitory system. *Nat Rev Neurosci* **5**, 793-807, doi:10.1038/nrn1519 (2004).

160 Zhao, S., Ting, J. T., Atallah, H. E., Qiu, L., Tan, J., Gloss, B., Augustine, G. J., Deisseroth, K., Luo, M., Graybiel, A. M. & Feng, G. Cell type-specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. *Nature methods* **8**, 745-752 (2011).

161 Madisen, L., Zwingman, T. A., Sunkin, S. M., Oh, S. W., Zariwala, H. A., Gu, H., Ng, L. L., Palmiter, R. D., Hawrylycz, M. J., Jones, A. R., Lein, E. S. & Zeng, H. A robust and high-throughput Cre reporting and characterization system for the whole mouse brain. *Nature neuroscience* **13**, 133-140, doi:10.1038/nn.2467 (2010).

162 Hippenmeyer, S., Vrieseling, E., Sigrist, M., Portmann, T., Laengle, C., Ladle, D. R. & Arber, S. A developmental switch in the response of DRG neurons to ETS transcription factor signaling. *PLoS biology* **3**, e159, doi:10.1371/journal.pbio.0030159 (2005).