Excitation creates a distributed pattern of cortical suppression due to varied recurrent input

Jonathan F O'Rawe, Zhishang Zhou, Anna J Li, Paul K LaFosse, Hannah C Goldbach, Mark H Histed

National Institute of Mental Health Intramural Program, National Institutes of Health, Bethesda MD USA · jonathan.o'rawe@nih.gov, mark.histed@nih.gov

Abstract

Sensory cortex transforms its inputs, generating neural responses that represent features of the sensory world. Dense local, recurrent connections are a major feature of cortical circuits, yet how they affect these transformations is unclear, with some studies reporting that recurrent circuits amplify certain excitatory inputs, and other studies showing instead local suppression. Here we find that input to mouse V1 excitatory neurons generates salt-and-pepper patterns of excitation and suppression. We show that neurons' responses are not strongly predicted by their optogenetic input alone, suggesting that recurrent inputs nonlinearly transform the population response. A balanced-state network model explains the observed dynamics, suppressed responses, and long tail of excited responses, but only with substantial variance in cell-to-cell connectivity. In sum, we find recurrent cortical connectivity has high variability and high average strength, which can decouple individual cells' responses from the direct input each cell receives, leading to substantial suppression generated by excitatory input.

Introduction

The cerebral cortex of mammals is specialized into areas that perform different functions[1].
Animals from rodents to primates have several different visual cortical areas, each containing
neurons with different types of selectivity[2–4]. In principle, these different representations in
different visual areas could be created purely by feedforward mechanisms, where
transformations happen via projections from one area or layer to the next, without outputs of a
neuron feeding back (directly or indirectly) to influence that neuron's activity. In fact, in a variety
of artificial neural networks, much or all computation is provided by feedforward mechanisms[5].

Yet in the brains of animals and humans, cortical recurrent connectivity is extensive. Most
excitatory connections that a cortical neuron receives originate within a few hundred microns of
their cell bodies[6–8]. Such recurrent connections can in principle have large effects on neural
computation[9], dramatically changing how cortical neurons respond to input.

How recurrent connections affect cortical computation is not fully understood, but important
aspects of the structure of cortical recurrent connectivity have been determined. Some features
of cortical network activity, such as irregular firing, are well-described by balanced-state models
which assume strong recurrent coupling between excitatory and inhibitory neurons (either
moderately strong, yielding 'loose balance', or very strong, yielding 'tight balance'[10]). Work using
inhibitory perturbations has shown that not just excitatory-inhibitory connectivity is strong, but
the average excitatory-excitatory connectivity is strong as well. More precisely, cortical recurrent
excitatory coupling is strong enough that the excitatory network is unstable and self-amplifying,
a phenomenon described by inhibition-stabilized network models (ISNs)[11–14].

While some consensus has developed on these average cortical connectivity properties (but
see[15]), the effect recurrent connections have on transforming sensory or input signals has been
less clear. For example, some recent studies have shown that certain patterns of excitatory
input can be amplified by the cortical network[16,17], consistent with some theoretical
predictions[18,19]. On the other hand, however, other studies have shown that nearby neurons can
be substantially suppressed by stimulation that excites a single or a small ensemble of
excitatory cortical neurons[20,21]. How excitatory and inhibitory neurons might interact through
recurrent connections to create such suppression has not been determined.

Here, to understand how cortical neurons' responses are shaped by the cortical recurrent
network, we stimulate excitatory cells in the visual cortex optogenetically and record responses
of local neurons with electrophysiology and two-photon imaging. First, we find that stimulation of
excitatory cells leads to a salt-and-pepper pattern of local suppression, a pattern that is
consistent with the pattern of excited and suppressed cells produced when animals see a strong
visual stimulus. To understand how this suppression effect might arise from cortical recurrent
circuitry, we examine both the patterns of firing rate changes and the dynamics of responses
and compare them to cortical models. Recent theoretical work has shown that cortical visual
responses can be "reshuffled" by additional excitatory input[22] — that is, strong average
recurrent coupling allows individual neurons' firing to change significantly in response to input
while the distribution of population activity is little-changed[23,24]. We implement this scenario in a

conductance-based simulation and find that it can explain the suppression we observe, and
additionally we find our data is consistent with substantial variability in local recurrent
connectivity, with some neurons receiving large net recurrent excitation and others smaller or
net suppressive recurrent input. Our results go beyond prior work that found strong *average*
recurrent connectivity, showing that *variance* in excitatory-excitatory connectivity must also be
substantial, and further show that this variance in recurrent connectivity can decouple neurons'
firing rate responses from the direct input they receive.

The suppression we observe during excitatory cell stimulation occurs in individual cells, but the
mean response is elevated. This increase in mean, however, seems at odds with the prior
finding that single-cell stimulation leads to inhibition on average[20]. To resolve this, we simulate
the effect of single cell stimulation and find that the difference in the two results can be
explained by the activation state of the cortical network. Increasing activity in the network with
visual stimulation results in a slight decrease in mean responses to stimulation, showing the
prior results and our current results can be described in the same model framework. Thus, we
find that a balanced-state, excitatory-inhibitory model, with strong average coupling and
variability in recurrent connectivity, explains many features of our data, including dynamics and
neural response distributions. These results show how cortical neural suppression can be
generated from excitatory input: variability in recurrent input means that firing rate responses
are decoupled from (are only weakly affected by) the level of excitatory input we provide to that
cell, as much of the input a cell receives comes from recurrent sources. Because recurrent input
varies from cell to cell, what results is many excited cells, but also a substantial fraction of
suppressed neurons.

Results

Strong visual input leads to salt-and-pepper distributed suppression in primary visual cortex

We first measured local patterns of suppression in visual cortex in response to visual stimuli. We presented small high-contrast visual stimuli to head-fixed mice while measuring activity in V1 layer 2/3 neurons via two-photon imaging (Fig. 1A). We expressed GCaMP7s in all neurons via viral injection (AAV-hSyn-GCaMP). Animals were kept awake and in an alert state[25] with occasional drops of water reward, and to ensure animals did not drift into a quiet wakefulness or quiescent state, we monitored animals during data collection to verify they continued to drink the delivered reward.

We imaged responses to two types of high-contrast visual stimuli, a fast-changing stimulus designed to minimize adaptation ("oriented noise", Fig. 1A)[26–28] and a drifting grating (data

Figure 1: V1 neurons show salt-and-pepper suppression to strong visual stimuli. (A) Experimental setup. Awake mice viewed a small (15 degree diameter) visual stimulus with rapidly changing frames of white noise, filtered to have a peak spatial frequency and orientation though pixel values changed from frame to frame (0° peak orientation, 90% contrast, frames presented at 60 Hz, total stimulus duration 3-5s, more details of stimulus in Methods). **(B)** Example 2-photon imaging data from layer 2/3 of mouse V1 in response to the oriented noise stimulus, during the transient and **(C)** steady-state periods. Time intervals used for averaging in (B-D) displayed in green and purple in (G). Note the emergence of intermixed (salt-and-pepper) elevated and suppressed responses during the steady-state period. **(D)** Deconvolved responses from (C), projected onto segmented cell masks (see Methods). **(E)** Example dF/F trace for one elevated and one suppressed cell. Shaded regions: SEM across trials. Shaded red: optogenetic stimulation duration. **(F)** Deconvolution of the traces in (D) reveals an initial transient period and then a steady-state response. **(G)** Average response for all elevated and suppressed cells in (B)**. (H)** Schematic of optogenetic and imaging experiments, to determine if suppression arises from recurrent sources. We expressed stChrimsonR selectively in excitatory neurons in V1 (AAV-FLEX-stChrimsonR; Emx1-Cre mouse line) and used two-photon imaging, extracellular physiology, and widefield imaging to measure cortical responses to excitatory cell stimulation.

shown in Fig. 1 Suppl. 1D). We found a salt-and-pepper mix of suppressed and excited cells
(Fig. 1B,C), with suppression especially visible during the later part of responses to long
stimulus presentations of the oriented noise stimuli (Fig. 1C). In other words, in response to both
types of visual stimuli, we found some cells that responded with strongly elevated steady-state
responses, and other cells that showed suppressed steady-state responses (Fig. 1C–F, Fig. 1
Suppl. 1).

Deconvolving fluorescence responses to yield a proxy measure of spike rate confirmed this salt-
and-pepper pattern with substantial numbers of suppressed neurons intermingled with excited
neurons (Fig. 1D,F-G).

Neurons showed sustained responses during the oriented noise stimuli (Fig. 1E), and
deconvolving the timecourses revealed an initial transient in excited cells (Fig. 1F, population
responses shown in Fig. 1G), followed by a steady-state period.

The viral expression strategy we used for these experiments results in both excitatory and
inhibitory neurons that express GCaMP. However, the large fraction of suppressed neurons
(Fig. 1C, 40% of responding neurons, Wald CI 29–51%; population: 44%, Wald CI 35–53%)
implies that many excitatory neurons were suppressed, and it was not just that a group of
inhibitory neurons was suppressed by stimulation. Below, we confirm with electrophysiology and
imaging that optogenetic excitatory input produces suppression in many excitatory cells.

Optogenetic excitatory drive also results in sparse and distributed suppression

To probe the influence of recurrent excitatory-inhibitory circuits on local response properties, we
next measured V1 responses while optogenetically stimulating excitatory cortical cells. Direct
stimulation allows us to exclude some feedforward mechanisms for suppression – for example,
to argue against the possibility that cortical suppression is generated principally by suppression
of thalamic inputs[29].

We injected a Cre-dependent excitatory opsin (soma-targeted ChrimsonR, or stChrimsonR) in
layer 2/3 of a mouse line expressing Cre in excitatory neurons only (*Emx1-Cre*[30]), and
expressed GCaMP7s in all neurons with a second virus (AAV-hSyn-GCaMP7s).

With optogenetic stimulation we also found a clear salt-and-pepper distribution of elevated and
suppressed responses (Fig. 2C,E,G). We saw this pattern in imaging experiments with two
different stimulation pulse lengths. First, we used short (200 ms) stimulation pulses and
measured cells' responses after stimulation onset to avoid optical artifacts (Fig. 2B,C). We also
saw the same interspersed pattern in other imaging experiments where we used long (4 sec)
stimulation pulses and imaged during stimulation (Fig. 2D,E).

Figure 2: **Salt-and-pepper distribution of elevation and suppression to optogenetic excitation shown using 2-photon imaging. (A)** Experimental setup. Animals expressing both GCaMP7s and stChrimsonR (both via AAV) were imaged using 920 nm 2-photon scanning laser, while stChrimsonR was stimulated with 595 nm light. **(B)** Example field of view, showing GCaMP fluorescence. Emx1-cre animal, injected with AAV-DIO/FLEX-stChrimsonR and non-Cre dependent AAV-GCaMP7s. **(C)** Steady-state response to 1-photon stimulation (200 ms duration) Red: elevation of firing rate relative to baseline, blue: suppression. Data is dF/F measured after stimulus offset, deconvolved, with peak response for each cell normalized to peak dF/F% for that cell. Same field as in B, stChrimsonR expression not shown. **(D)** Field of view from a second animal stimulated with long (4 sec) optogenetic pulses; stimulation during imaging flyback (gray region at edges) to avoid contaminating imaging responses and allow stimulation during imaging. Data analysis same as for C but responses collected during stimulation (stimulation delivered on horizontal imaging flyback, Methods). Red: GCaMP7s fluorescence. Gray: areas omitted from analysis due to stimulation artifact. **(E)** Deconvolved response of cells in (D) during the steady-state period, colored similarly to (C). Salt-and-pepper suppression is evident. **(F)** Increasing power leads to stronger elevation and suppression (steady-state response) in their respective populations, as expected. Shaded region: SEM across cells. **(G)** Spatial distribution of elevated and suppressed cells, N = 3 animals, elevated cells n = 246, suppressed n = 139 cells (one sample t-test, p < 0.05). **(H)** Population timecourses for all cells in (E)**.** Red shaded region: optogenetic stimulation period. Steady-state response calculated over 200 ms to 3750 ms post stimulus onset. Light lines: individual cell traces, heavy lines: population averages. Shaded region (largely obscured by thick lines): SEM across cells. **(I)** Selective expression of opsin in excitatory cells only, as expected for the double-inverted lox-site AAV vector and excitatory Cre mouse line (Emx1-Cre). The stChrimsonR opsin was fused to mRuby2, so we measured mRuby2 mRNA (red) and VGLUT1 mRNA (white), a marker of excitatory cells, via fluorescent *in situ* hybridization (RNAscope; Methods). Cell counting showed 76% of neurons are VGLUT1 positive (N = 195/257.) Arrows highlight a few example neurons. As expected, all cells that express the opsin are excitatory, but not all excitatory neurons express the opsin (59% of VGLUT1 cells are mRuby2 positive: N = 115/195). **(J)** mRuby2 mRNA (red) and VGAT mRNA (green), a marker of inhibitory neurons. 24% of neurons are VGAT positive (N = 62/257), and zero (N=0/62, Clopper-Pearson 95% CI 0-0.472), express the opsin.

Deconvolving the imaging responses revealed that elevated populations showed a transient
response, then a steady-state response that lasted for the duration of the optogenetic
stimulation (Fig. 2H).

These data suggest that the network is being driven to a new steady state or fixed point by
input. While there was a slight decay in the excited population's response at high power
(perhaps due to network effects, spike rate adaptation, or opsin desensitization), at moderate
stimulation power (1 mW, Fig. 2H), deconvolved firing rates are largely constant while
stimulation is on.

We next verified that the opsin we used was expressed only in excitatory cells. We used a
mouse line that expresses Cre in excitatory neurons, and a genetic strategy (FLIP/DIO[31]) that
blocks expression in non-Cre-expressing cells. Still, to confirm expression in excitatory neurons,
we labeled excitatory, inhibitory, and stChrimsonR-expressing neurons via fluorescence in-situ
hybridization (RNAScope, ACD Inc; Fig. 2I,J). Excitatory neurons expressed the opsin (Fig. 2I),
but as expected for AAV expression[32], not all excitatory neurons were positive (59%, N =
115/195, Fig. 2I). None of the inhibitory neurons (24% of neurons in the sample, N = 62/257)
showed expression of the opsin (Fig. 2J), demonstrating, as expected from the *Cre* mouse and
floxed virus strategy, that the opsin was expressed in excitatory neurons.

Next, we examined whether this salt-and-pepper pattern exists at larger distances from the
stimulation site. Using the same opsin preparation, in awake mouse V1, we recorded neural
responses to excitatory cell stChrimsonR activation using a silicon electrode array (Fig. 3A,E).

We found both elevation and suppression across all distances (Fig. 3B–D) and depths (Fig. 3F–
H) from the stimulation site, suggesting a similar local salt-and-pepper organization of elevated
and suppressed cells that we observed directly in the 2-photon response data.

Across the population of recorded neurons, 56.6% (77 of 136) showed an elevated steady-state
response to the optogenetic stimulation, and 36.0% (49 of 136) showed a suppressed steady-
state response. Both elevated and suppressed cells, on average, showed an initial (positive)
transient followed by a (positive or negative) steady-state response (Fig. 3I,J). The
electrophysiological recordings show similar dynamics as the deconvolved imaging timecourses
(Fig. 2H), except for one feature: the recordings show an initial brief positive transient in the
suppressed cells (Fig. 3B–D, F–H, blue lines; Fig. 3J) not just the elevated cells as in the
imaging data. The imaging data likely conceal the initial transient in suppressed cells via the
slower timescale of imaging. That is, due to the imaging frame rate (30 Hz; 33 ms frames), the
transient would often be contained within one frame where it would be averaged with
suppression.

Global spatial patterns arise from trends in local salt-and-pepper suppression

The neurophysiology data showed some evidence of a larger-scale organization on top of the
local salt-and-pepper distribution of elevation and suppression. Over distances of more than a
millimeter from the stimulation site, we found that the number of elevated units gradually
decreased (Fig. 3K,L; Pearson's chi-squared test, χ^2 = 51.31, df = 3, p < 0.001) and the number
of suppressed units gradually increased (Fig. 3M,N; χ^2 = 44.83, df = 3, p < 0.001). There was
also a similar trend in neurons' firing rates (Fig. 3 Suppl. 2). Elevated single units showed less
elevated firing rate with distance from the stimulation site, and suppressed single units showed

Figure 3: **Focal stimulation of excitatory cells in V1 results in widespread steady-state elevation and suppression with a weak large-scale center-surround organization. (A)** We used arrays of four electrodes (200 µm spacing between electrodes) to measure activation across the cortex. We applied optogenetic stimulation light at an electrode on one end of the array while measuring responses from all electrodes. **(B, C, D)** Example neurons recorded at locations increasingly distant from the stimulation light (0 µm, 400 µm, 1200 µm), within a single animal, showing elevated and suppressed cells at all distances. **(E)** Multiple recording sites on each electrode allow recordings across depth. **(F, G, H)** Example neurons recorded at increasing depth (200-300 µm, 500-600 µm, and 800 µm) within a single animal, showing elevated and suppressed cells are present at different depths. **(I)** Population average timecourses of elevated cells. Blue bar: period used for calculating average firing rate during the steady-state period. Lighter colors: more intense laser stimulation. Shaded regions: SEM across cells. Same data as in Fig. 2A-H, collapsed across distance and depth. **(J)** Population time courses of suppressed cells, same display as (I). **(K)** Counts of elevated units (single and multi-units) by distance and depth, smoothed with a Gaussian kernel for display. **(L)** Distribution of elevated steady-state responses across horizontal distance, calculated by summing across depths. Shaded region: ±1 SEM given binomial counts. Note lower limit of y-axis is not zero. **(M-N)** Same as (K-L), but for units with suppressed steady-state responses.

more suppression with distance from the stimulation site, though the linear trend between
distance and population response was stronger in unit counts than in average population firing
rates (Pearson's r = -0.11, df = 29, p = 0.56, Pearson's r = 0.32, df = 46, p < 0.05; Fig. 3 Suppl.
2). Notably, however, the number of elevated neurons did not go to zero even at 1.2 mm from
the stimulation site: only the relative numbers of elevated and suppressed neurons changed.
This suggests that the salt-and-pepper organization we saw with imaging persists across the
cortex.

The trends over distance we saw with physiology, however, gave only a limited view of how
mean population responses varied with distance from the stimulation site. To characterize
population average responses more fully, we turned to widefield, mesoscale calcium imaging. In
this approach, each pixel in the imaging integrates signals from many neurons. For these
experiments, we expressed GCaMP in all excitatory cells using a genetic mouse line to
maximize consistency of GCaMP expression across cortical distance (Fig. 4A; Ai148::Cux2-
CreERT2, or Ai162::Cux2-CreERT2, see Methods). We expressed stChrimsonR also in
excitatory cells (AAV-CamKIIa-stChrimsonR) and stimulated while simultaneously imaging
responses.

We saw clear spatial patterns in widefield imaging, broadly consistent with the spatial trends we
saw in the electrophysiology data. During the initial frame of stimulation (~7 Hz imaging, 140 ms
frame period), we saw an increase in activity both at the center of the stimulation light and
extending some distance outside the center of expression (Fig. 4B,E,H). This first-frame
increase is consistent with the transient response in both elevated and suppressed cells we saw
in electrophysiology (Fig. 3I,J). Transient responses like this, before the network reaches a new
steady-state, are also consistent with neural responses being shaped by a strongly-coupled
recurrent network.

Later in the stimulation pulse, consistent with the large-scale patterns in the electrophysiological
recordings, a center-surround pattern emerged (Fig. 4C,F,I.) The area with maximum
stChrimsonR expression continued to show an elevated response, while a donut-shaped region
around it was suppressed. The activated area in the center reflected the area of expression,
measured with fluorescence imaging of the cortical surface (Fig. 4 Suppl. 2). This late
suppression also seemed consistent with the electrophysiology data, which showed
suppression not in transient responses but in the steady-state response. To examine these
timecourses (Fig. 4D, Fig 4. Suppl. 3), we deconvolved imaging responses to better reflect spike
rate changes. The deconvolution was essential to reveal the suppression because fluorescence
(dF/F) measurements, without deconvolution, integrate spiking-related calcium signals over
time. For responses with this timescale, dF/F measures without deconvolution mix initial
transients with later suppression, potentially hiding one or the other. Indeed, the dF/F signals

 showed less suppression (Fig. 4 Suppl. 1). We compared several different deconvolution

Figure 4: **Widefield imaging of excitatory neurons shows average center-surround organization during steady-state periods. (A)** Experimental setup: stChrimsonR in excitatory neurons via viral transfection (AAV-CamKIIa-stChrimsonR) with widespread expression of GCaMP via mouse line (either Ai148::Cux2-creERT2, GCaMP6f, or Ai162::Cux2-creERT2, GcaMP6s; induced with tamoxifen as adult; Methods). Right: imaging field of view for one animal. **(B-C)** Mean response during first frame (B) and during the late stimulation period (C) in an example animal. Note signal here and throughout is deconvolved (see Fig. 4 Suppl. 1 for details) to better reflect firing rate changes over time. **(D)** Response to stimulation over time, example animal (Ai148,GCaMP6f). Red bar: stimulation period, orange bar: first frame, maroon bar: late stimulation (steady-state) time period. **(E,F)** Average responses across n = 3 animals. Responses for each animal were aligned spatially to the peak during the late stimulation period, normalized by the value of the maximum response, and then averaged and smoothed (Methods.) **(G)** Response as a function of distance, averaged radially from data in E, F. Smoothing: lowess. Shaded regions: bootstrapped 95% CIs. Vertical lines: zero crossings and inflection points. Zero crossings defined by shortest distance at which 95% CI included zero, late stim local minimum determined from mean response. Black solid line represents the first zero crossing for the late stimulation period, black dotted line represents the minimum response in the late stimulation period, and dotted red line indicates the first zero crossing in the first frame of stimulation. **(H, I)** Same as (E,F) but with superimposed circles whose radii correspond to lines in (G).

methods and found suppression in all cases (Fig. 4 Suppl 1). The suppressed area of the cortex
was largest about 500 μm from the center of our laser stimulus, and extended over 1 mm from
the stimulation center (Fig. 4 G–I).

In summary, the physiology and imaging data together support the idea that suppressed and
elevated neurons are locally organized in a salt-and-pepper pattern, and that the proportion of
suppressed to elevated neurons increases with distance from the stimulation site. This change
in the proportion of suppressed cells results in a center-surround pattern that can be seen with
population-level imaging, with net suppression in excitatory cells emerging, after an initial
positive transient, about 500 μm away from the stimulation site.

Response dynamics support a balanced-state excitatory-inhibitory network that is driven to a new steady state by input

If suppression is due to local recurrent network effects, we would expect excitatory cells to be
recruited first by stimulation, and then inhibitory cells should receive inputs from excitatory cells
and respond slightly later. After this first few milliseconds, balanced-state models predict that
excitatory and inhibitory cells should later show similar response distributions[10,33]. This is in
contrast to weakly-coupled models, or a feedforward inhibition framework, where excitatory and
inhibitory populations change firing rates in opposite directions: input drives inhibitory cells to
increase their rates, inhibiting excitatory cells, which then decrease their rates.

In the data we saw support for the balanced-state recurrent model (Fig. 5A) — differences in
excitatory and inhibitory responses in the first few milliseconds, but later similar distributions of
excitatory and inhibitory rates.

We classified cells into putative excitatory and inhibitory classes by waveform (Fig. 5B). We
have previously confirmed with *in vivo* pharmacology that narrow-waveform cells are inhibitory
interneurons, likely PV-positive fast-spiking cells, while wide-waveform cells are primarily
excitatory neurons [13]. We saw that wide-waveform (excitatory) neurons have a slightly faster
onset latency than narrow-waveform inhibitory cells, by approximately 2.5 ms (Fig. 5A, inset,

Figure 5: **Response dynamics are consistent with steady-state suppression shaped by a recurrent network mechanism.** Units with wide waveforms (excitatory units) have slightly faster onset latencies than narrow (inhibitory) units, as expected when stimulating excitatory cells, but other quantities do not differ across neural populations. **(A)** Average traces for narrow- (inhibitory) and wide-waveform (excitatory) units. Narrow-waveform units are likely inhibitory, fast-spiking neurons, while wide-waveform units include mainly excitatory cells with a minority of wider-waveform inhibitory neurons. Inset: enlarged view of laser onset, highlighting latency difference between narrow and wide units (green markers: time to half peak.) **(B)** Wide/narrow sorting approach. The distribution of spike waveforms is bimodal (classification threshold: 0.445 ms.) **(C)** Onset latencies are slightly shorter for wide (excitatory) units, as expected when stimulating excitatory cells. See Results for test statistics and p-values. **(D)** No significant difference in initial transient firing rate (peak firing rate) between wide-waveform and narrow-waveform units (medians, wide: 24.5 spk/s, narrow: 31.2 spk/s.) **(E)** No significant difference in steady-state firing rate between wide and narrow units. For both putative inhibitory and excitatory cells, some units show suppressed rates and some elevated rates (medians, wide: 4.0 spk/s, narrow: 5.1 spk/s.) **(F)** No difference in onset latency for elevated and suppressed populations. Conventions as in C. **(G)** No difference in time to steady state for elevated and suppressed populations.

5C; narrow latency 7.9 ms, wide latency 5.4 ms, difference 2.5 ms, Mann-Whitney U = 1256.0, p
< 0.01; onset latencies computed via curve-fitting to rising phases, see Fig. 5 Suppl. 1A for
details).

An alternative to excitatory and inhibitory cells acting as a balanced-state network (with
excitatory and inhibitory cells responding similarly but with slightly different latency) is that
suppression was due to a subtype or subpopulation of inhibitory cells whose firing rate went up
as excitatory rates went down. If such separate subnetworks existed, we might first expect to
see differences in the dynamics of elevated and suppressed cells. But we found no significant
differences in onset time or time to steady state for elevated and suppressed cells (Fig. 5F,G,
onset time Mann-Whitney U = 1741.0, p = 0.17, time to steady state, Mann-Whitney U = 801.0,
p = 0.32.)

Moreover, if there were suppressed and elevated excitatory cells that were part of different
subnetworks, instead of acting as part of an overall coupled network, we might expect to see
differences in the dynamics of suppressed and elevated sets of wide-waveform (excitatory) cells
only. But there were no significant differences in onset time or latency for suppressed and
elevated cells when we restricted these analyses to wide-waveform cells (Fig. 5 Suppl 1C,D).

Beyond the differences in onset latency, other response dynamics were not different between
excitatory and inhibitory cells. Consistent with a recurrent network with strongly coupled
excitation and inhibition, we found that both excitatory and inhibitory cell populations increase
their *average* firing rate when excitatory cells are stimulated (wide mean Δ: 14.54 spk/s, t =
5.52, df = 93, p < 0.001, narrow mean Δ: 13.53 spk/s, t = 3.07, df = 41, p < 0.01). That is, both
excitatory and inhibitory populations contain elevated and suppressed neurons, though elevated
cells dominate both averages. Further, the initial transient and steady-state firing rate medians
were not detectably different between inhibitory and excitatory cells (transient: Mann-Whitney U
= 1816.0, p = 0.23, steady-state Mann-Whitney U = 1866.0, p = 0.31, Fig. 5D,E). Also, the time
it took for wide-waveform and narrow-waveform cells to settle into their steady states did not
differ (Fig. 5 Suppl. 1B), consistent with the idea that the steady-state dynamics emerge from
integration of both inhibitory and excitatory inputs. Overall, in a strongly-coupled network
operating as an ISN, excitatory and inhibitory mean firing rates tend to vary together[11,34],
consistent with our observation that excitatory and inhibitory populations have similar dynamics
and firing rate distributions.

A balanced-state model with strong coupling explains many features of the data

Thus far, a moderately- or strongly-coupled balanced-state network seems consistent with both
the response distributions and dynamics we observe. Such network models, with excitatory and
inhibitory neuron populations that have strong coupling across populations, explain
experimental data in many, if not all, cortical areas[10]. In balanced networks, individual neurons
each receive large excitatory and inhibitory inputs, each strong enough that by themselves they
would control the neuron — either to fire at high rates or to be completely suppressed. But the
two currents sum to a small net value, and the neurons fire based on fluctuations caused by
transient shifts in excitatory/inhibitory balance.

The large excitatory and inhibitory inputs in balanced-state cortical models could underlie the
suppression we see. Intuitively, stimulation would excite excitatory cells, the excited cells would
excite their inhibitory targets, and this would lead to some inhibitory cells excited more than
others. In turn, this would cause varying amounts of inhibition across network neurons, which in
some cases would lead to suppression. Indeed, recent theoretical work[22] has shown that this
kind of heterogeneous network response occurs in strongly-coupled cortical networks. To
understand if our experimental data could be explained by this mechanism, we examined
recurrent network models with features reflecting our data, and determined which features of the
recurrent network models were important to explain the suppression.

Figure 6: **stChrimsonR expression only weakly predicts 2-photon steady-state response. (A)** Left: Red (stChrimsonR-mRuby2) and green (GCaMP7s) fluorescence of an example cell with a suppressed steady-state response during optogenetic (1-photon) stimulation. Right: donut-shaped region of interest (ROI), inner and outer boundary calculated by shrinking or expanding the cell border (CaImAn, Methods) **(B)** Same as (A), except for an example cell that shows an elevated steady-state response. Both cells show qualitatively similar red fluorescence. **(C)** Measured red fluorescence in donut-shaped ROIs covering cell membranes, plotted against steady-state response of each cell (n = 113) in an example animal, at 2 mW stimulation power. stChrimsonR expression only weakly predicts steady-state response of each cell (least-squares linear regression, r = 0.19, df = 111, slope *p* = 0.04). Example cells (A,B) are highlighted (colored markers, letters). **(D)** Population response (n = 3 animals), same plot as in C, across two laser intensities (1 mW, 2 mW) (n = 244 neurons). LOWESS fits with bootstrapped standard error are plotted, demonstrating the same weak positive relationship, though the relationship at high fluorescence trends back toward zero, largely due to a few highly fluorescent cells, perhaps due to response saturation or overexpression of opsin. **(E)** Lowess fits from (D), zoomed in to the section of the graph indicated by the dashed box. The weak positive relationship is present.

First, we used the imaging data to determine if the suppression we observed was tightly linked
to the variability in opsin from cell to cell. In principle, it could be the case that only the cells that
are excited by the opsin would increase their firing rates, while non-opsin cells would be
suppressed. This might happen if recurrent excitatory input were small compared to optogenetic
input. That is, a tight relationship between opsin and firing rate could mean network effects were
not important.

Instead of a tight link, we found a very weak relationship between opsin expression and cells'
firing rate changes. To approximate the optogenetic drive to individual neurons, we measured
fluorescence of mRuby2 (fused to stChrimsonR) in donut-shaped regions around each cell's
membrane using two-photon imaging (Fig. 6). The amount of opsin-related fluorescence
explained little of the variance in steady-state responses (i.e. the responses were spread out
along the y axis, and did not form a tight line; Fig. 6A-C shows single animal example at 2 mW
stimulation, Pearson's r = 0.19, df = 111, slope of linear fit > 0, p < 0.05). We found that the
weak but positive relationship between opsin-related fluorescence and steady-state response is
consistent across the population (Fig. 6D), though the slope trends back towards zero at high

fluorescence, driven by a few cells that have high fluorescence but show little response,
perhaps due to overexpression of opsin in a small number of neurons (excluding cells with high
fluorescence, 1 mW: Pearson's r = 0.10, df = 193, p = 0.17, 2 mW: Pearson's r = 0.17, df = 193,
p < 0.05; Fig 6E).

Because the amount of suppression was not well-explained by stChrimsonR expression, we
next examined how network effects could lead to suppression. We first sought to confirm that
our results are consistent with a balanced-state model (as in[22]) and that we could replicate
these effects within a conductance-based network. We constructed and simulated conductance-
based spiking neural networks, varying network connectivity and opsin drive across neurons in
these models, and measured network responses to excitatory cell stimulation (Fig. 7A,B).

Each network consisted of two sparsely connected populations of conductance-based spiking
neurons, one excitatory (8000 neurons; 80%) and one inhibitory population (2000 neurons;
20%). We used a model with a minimal assumption of only one inhibitory population. The
inhibitory population may map onto fast-spiking parvalbumin-positive (PV) cells, which seem
likely to be able to play a balancing role in the cortex given their strong input to excitatory
cells[13,35], consistent with the dynamics and responses we observed in narrow-waveform, likely
PV, neurons (Fig. 5). Excitatory and inhibitory synaptic mean strengths were adjusted so the
total recurrent excitatory and inhibitory inputs were similar: for every network, we adjusted a
background input current to either excitatory or inhibitory neurons to hold the spontaneous firing
rate of the neurons at a value (4–6 spk/s) consistent with the data. We drew the opsin input
strength to each neuron from a distribution fit to the imaging data (log-normal distribution, Fig.
7C; Methods).

We first manipulated the mean connectivity strength of recurrent connections. We constructed
three different models, varying the average strength of recurrent coupling in each (schematic,
Fig. 7B). The "tightly balanced" network had the strongest recurrent coupling, and we scaled
down the synaptic weights by a factor of 2 or 4 to create more weakly-coupled "50%" and "25%
strength" networks.

We checked that the networks showed paradoxical suppression of inhibitory cells, a sign of
strong recurrent coupling within the excitatory network and the ISN regime, as observed in
visual cortex[11,13,36]. We stimulated the inhibitory cells in each network and found, as expected,
paradoxical suppression (Fig. 7 Suppl. 1).

Figure 7: **Strongly-coupled balanced-state recurrent neural network model with heterogeneous connectivity describes experimental data. (A)** Simulation design: conductance-based spiking network model with 8000 excitatory cells and 2000 inhibitory cells (2% connection probability between cells, synaptic strength variability chosen to be of same order as mean (Chapeton et al., 2012).) **(B)** We varied the overall strength of recurrent connectivity by scaling the recurrent synaptic strengths. **(C)** Optogenetic input strengths (effects of opsin on each cell) were chosen based on the fluorescence values observed in 2-photon imaging data. A distribution function (log-normal) was fit to the fluorescent data and each weight was drawn independently from this distribution. **(D)** Schematic depicting the characteristic features of the response distribution we have observed. **(E)** Electrophysiological data and simulated network response to excitatory cell optogenetic stimulation, using input weights chosen to match the data (C). The more tightly balanced networks show a long tail of cell responses similar to the long tail seen in the data. On the right we display the same data but zoomed in on the suppressed (negative) values of the distribution. The network with strongest recurrent coupling shows the most suppression. **(F)** Cumulative distribution of responses to increasingly strong 1-photon stimulation during electrophysiology experiments, **(G)** 2-photon experiments**,** and **(H)** simulations. **(I, left)** Mean time course for elevated cells in the strongest recurrent network. **(I, right)** Mean time course for suppressed cells in the strongest recurrent network. **(J)** Balance index (Ahmadian and Miller, 2021) of the 3 networks represented in (B). Values between 0–0.1 correspond to tight balance, between 0.1 and 1 loose balance, and above 1 unbalanced, where there is little recurrent contribution to neurons' responses. The strongest-coupled network (purple) has a median index corresponding to the tight balance regime. **(K)** Schematic of types of input variability. Number of recurrent connections across cells can vary (left) or the strength of those connections can vary (middle). Variability in input (here optogenetic input) can also contribute to response variability. **(L)** Simulated neural responses to optogenetic input, with (left, same as Fig. 7F) and without (center) variability in number of recurrent inputs. The relationship between optogenetic input and steady-state response strengthens with the variance in number of recurrent inputs held constant (R^2 original = 0.29, R^2 reduced conn. var. = 0.58, ΔR^2 = 0.29). Further, when recurrent strength variability is reduced, the relationship gets even tighter (L right) (R^2 original = 0.29, R^2 reduced conn + strength var. = 0.70, ΔR^2 = 0.41) Therefore, when recurrent variability is reduced, each cell's steady-state response is more strongly dependent on opsin expression. **(M)** Cumulative distributions of baseline-subtracted steady-state firing rates. Variability from either recurrent or optogenetic sources is sufficient to approximate the long tail of elevated firing rates, however having both best approximates the amount of suppression we observe. "All var:" network with parameters as in Fig. 7, including variability in input number, variability in optogenetic input weight (fit from data, Fig. 7C), and variability in synaptic strength. Pink, network with variability in number of recurrent connections removed. Orange, network with optogenetic input variability removed (all optogenetic input weights set to a constant). Yellow: both types of variability removed. (Right) Zoom of (M) to the suppressed portion of the distribution. Some suppression still exists if either source of variance is removed, but suppression is effectively extinguished when both sources of variance are removed. **(N)** Reducing the variance of the synaptic weights by a factor of 10, shows the cumulative distribution becomes strictly positive with little variance. Yellow: network with no connection number variability and no optogenetic input variability. Blue: same as yellow network but with variance in synaptic strength distributions reduced by a factor of 10.

Recurrent excitatory-inhibitory networks can be tightly or loosely balanced[10], depending on the
total amount of recurrent excitatory and inhibitory input to network neurons. To classify the
networks, we calculated the balance index, a ratio that measures how completely inhibitory
input cancels out the excitatory input for each neuron in the networks[10] (see Methods). We
found that all three networks we constructed are balanced as expected due to their irregular
spontaneous activity (balance index << 1), and the networks span a range from loose to tight
balance (Fig. 7J).

The model replicates the long tail of positive responses, suppressed responses, and
aspects of the dynamics

Two characteristic features of the response data we observed are the long-tailed positive
response and the substantial proportion of suppressed cells (Fig. 7D). To match the input
strength to the data, we scaled the input distribution until the 75th percentile of the model
response matched the electrophysiological response to the strongest optogenetic input. All
three simulated networks showed a long tail of elevated responses as in the data, with many
neurons showing increases in firing rate to stimulation, and a few showing large increases (Fig.
7H). However, the amount of suppression did depend on recurrent coupling strength. Scaling
the total excitatory and inhibitory input leads to more suppressed neurons when other network
parameters are held constant (Fig. 7E). The network that best fit the fraction of suppressed cells

we observed was the most strongly-coupled network, which was just inside the tight balance
regime (Fig. 7E,J). Excitatory and inhibitory cells' response distributions were similar in the
model, as also seen as in the data (Fig. 7 Suppl 2). Finally, to further demonstrate how well the
tightly-coupled model could reproduce the features of our response distributions, we simulated
responses to parametrically-increasing strengths of the input distribution, and compared them to
the electrophysiological and 2-photon responses to increasingly strong experimental
optogenetic stimulation. The shapes of the response distributions in the most tightly coupled
model and the data were similar (Fig. 7F–H).

We next examined response dynamics to determine if the dynamics we observed were also
consistent with a balanced E-I model. We found that model responses were qualitatively similar
(Fig. 7I) to the timecourses of responses seen in the data (Figs. 2–3). Excitatory cells first
showed a brief, positive transient response before the network settled into a new steady state,
with some cells excited and some suppressed. The biphasic response in suppressed cells
(positive transient followed by suppression) is a key data point, as this biphasic response
suggests a network mechanism where input initially excites many excitatory neurons, but later
recurrent inputs lead to suppression of some excitatory cells. A second matching feature of the
dynamics is the offset dynamics in both elevated and suppressed cells: after stimulation ends,
both show a slight suppression before returning to baseline. A third matching feature is that
excitatory cells have earlier onset times, indicating that E cells were directly stimulated and I
cells were recruited just a few milliseconds later (Fig. 7 Suppl. 2), before both populations then
evolved to a new steady state.

One feature of the dynamics seen in the data but not the model is that for high stimulation
powers there is a slight decay during the tonic or steady-state period (Fig. 3I). However, this
decay effect is not seen at lower stimulation intensities, suggesting it arises from known opsin
dynamics (inactivation at high light power, e.g.[37]) or other known biophysical, non-network,
effects like spike-rate adaptation.

Neurons' responses are not tightly related to optogenetic input due to recurrent inputs

Thus, a balanced excitatory-inhibitory spiking model with one inhibitory population shows the
main features of both our two-photon and electrophysiological data. The observations reflected
in the model include the distribution of responses as a function of optogenetic input, the
dynamics, the long tail of excited neurons, and the suppression of a fraction of excitatory
neurons.

In addition to these effects, we also saw another characteristic in the data that can be
informative about the model: strikingly little correlation between opsin expression and neural
response to stimulation (Fig. 6). This implies that the amount of optogenetic input a neuron
receives is not a good predictor of a cell's response, and further that recurrent input strongly
governs neurons' responses. In fact, the tightly-coupled model showed the same pattern of
responses (Fig. 7L, left).

We therefore asked which sources of variability were important to explain why neural responses
were weakly related to optogenetic input. To do this we varied sources of input variability in the

model (Fig. 7K). First, we reduced variability in either the number or strength of recurrent inputs
and found that this created a tighter correlation between optogenetic input strength and
response (Fig. 7L, middle and right) which made the model a worse fit for the data (Fig. 6). The
original relationship was not recovered by increasing the recurrent strength of the network (Fig.
7 Supp 3), implying recurrent connection variability was required to produce this effect and
higher recurrent strength could not substitute for it. Next, we asked whether variability in opsin
input across cells was also essential to explain the suppression we observed. Removing the
variance in optogenetic input across cells (so that each network E cell with optogenetic input
received the same strength optogenetic input) significantly reduced the number of suppressed
neurons and also produced a worse fit to the data (Fig. 7M; right inset highlights suppressed
neurons).

If input variability was the primary source of variability in neural responses, then removing
variability from both kinds of input — both optogenetic input variability and recurrent input
variability — should substantially reduce the amount of suppression observed. This is what we
found. Removing or reducing both types of variability produced a set of neural responses
clustered tightly around the mean response (Fig. 7N) with no suppressed neurons. Thus, both
variability in recurrent input and in optogenetic input are required to explain the data.

In these simulations, we set several parameters as fixed based on data. Those include: the
variability of optogenetic input (estimated from Fig. 6) and the spontaneous rate of the networks
(setting the spontaneous rates to ~ 5 spk/s, Fig. 3I–J). In each network, we also adjusted the
optogenetic input strength to match the data (Fig. 7F–H, matching the 75th percentile of the data
rather than the median or mean due to the long positive tail of responses in the data). Once
those parameters were fixed, we varied only the mean strength of excitatory and inhibitory
connections (Fig. 7E,J), and the variability of recurrent input (Fig. 7K–N) to determine how each
affected network responses, finding that variability in recurrent connectivity was needed to
describe the data.

Together these results show that both optogenetic input variability and recurrent connection
variability help create the variability in the response of different neurons that receive similar
optogenetic input. Why do both contribute? Each neuron's firing rate is affected not just by the
optogenetic input that particular cell receives, but also by recurrent input received from other
neurons, which themselves receive different amounts of optogenetic and recurrent input. That
means the firing rates of the input neurons to a given cell are affected by both the variability in
their recurrent connections to the network, and the distribution of optogenetic input they receive.
The resulting recurrent input to a given neuron varies from neuron to neuron, and this is what
uncouples neurons' responses from the amount of direct optogenetic input they receive,
producing the weak relationship between optogenetic input and firing rate that we observed
(Fig. 6).

A balanced-state network model with connection variability also explains expected
responses to single cell stimulation

Past work has found that stimulating a single cell in visual cortex leads to average suppression
in the surrounding population[20] (averages across space and across orientation yield a negative

mean response, e.g. their Fig. 2D). Our results seem initially to contradict this finding, because
our data and simulations both find a net positive response across the population when we
stimulate many excitatory cells.

To determine if the effects of single-cell stimulation could also be explained by the balanced-
state simulation that describes our data, we performed simulations of single cell stimulation in
the same tightly-coupled network we examined previously, measuring the response of the non-
stimulated population (Fig 8A,B). We found that, while single cell stimulation produced a range
of individual cell responses (i.e. reshuffling, Fig. 8C), the mean response was not negative, but
instead close to 0 (Fig. 8C,D; mean firing rate 95% CI [-0.006, 0.021], t (9998) for nonzero mean
= 1.13, p = 0.26).

Different neurons in the network showed different responses based on recurrent connectivity,
but this was insufficient to explain the negative mean response that results from single-cell
stimulation. The excitatory cells that received a direct connection from the stimulated cell
(feedforward, FF, cells, n=107 neurons) had an elevated response. Those that received a
connection from the inhibitory cells which received a monosynaptic input from the 107 directly
connected E cells had a very slightly suppressed response (i.e. E-I-E connections, or feedback
(FB) cells, n=4380 neurons; Fig. 8E). The small set of strongly excited cells compensate for the
large number of weakly suppressed cells, leading to a mean response of 0 (Fig. 8F,G). Single
cell stimulation alone in our balanced-state model, therefore, could not account for the mean
suppression observed in previous studies.

We next hypothesized that this difference could be due to difference in the activation state of the
network. Chettih and Harvey (2019) stimulated during visual input, while here we delivered
optogenetic input during spontaneous activity. One explanation for this effect could be explained
by a balanced-state model: Sanzeni et al. (2022) found that increasing the firing rate of a similar
network to our model reduced the mean response of the network. Further, in models of visual
cortex with subnetwork connectivity (e.g. higher connectivity between neurons with similar
orientation tuning[38], it has also been shown that visual input can shift the network response to
be negative[39]. Therefore, to test whether additional network drive could reproduce mean
suppressive responses, we simulated single-cell stimulation on top of an input that mimics
visual drive (Fig. 8G; Methods). This indeed shifted the mean population response negative
(Fig. 8G; firing rate change: mean: -0.05; 95% CI: [-0.07, -0.04]). Further, we show that for
single-cell activation, as seen by Sanzeni et al. (2022) for optogenetic activation of many
neurons, that the relationship between baseline firing rate before input and network response to
input is systematic. The more strongly the network is activated with visual-like input, the more
negative the response to optogenetic input (Fig. 8H).

Therefore, we find that a strongly-coupled balanced-state model can explain both our results
with stimulation of many neurons (a long tail of positive responses, positive mean, and a
substantial fraction of suppressed cells) and the negative mean that results from single-cell
stimulation. In both cases, strong mean connectivity, as well as variability in recurrent
connectivity shape the responses of the network.

Figure 8: **Single cell stimulation produces elevated firing rates in a small subset of cells, but widespread weak suppression across the population. (A)** A single cell was chosen and repeatedly stimulated in the network with the strongest coupling (Fig. 7B top). **(B)** Firing rate distribution during single cell stimulation shows that single cell stimulation only strongly modulates the stimulated cell and perhaps a few others. **(C)** Cumulative density functions during single cell stimulation and without stimulation (black and green, respectively). While single cell stimulation widens the distribution of responses, the mean remains around 0. **(D)** Same as (C), but with probability density representations. **(E)** Two subsets of the E cell population were drawn: cells that receive direct connections from the stimulated cell (n = 107; brown outlined cells), and cells that receive feedback inhibition as a direct result of the stimulated cell (n = 4380; dark blue outlined cells). While cells that receive direction connections from the stimulated cells show a positively shifted distribution with a heavy positive tail, the cells that receive feedback inhibition are biased towards suppression. **(F)** Bootstrapped means of the distributions mirror the effects described in earlier panels. The whole distribution (excluding the stimulated cell) shows a mean suppressed response (black). The small subset of E cells that receive direction connections from the stimulated E cell show a significantly elevated mean response (brown). The subset of cells that receive feedback inhibition as a direct result of the stimulated cell show an even greater bias towards suppression (dark blue), while the whole population (black) is mean 0. **(G)** To simulate Chettih and Harvey (2019) conditions more closely, where single cell stimulation was carried out during visual drive, we performed our simulations with various forms of synthetic visual drive by providing the network with large scale input during single cell stimulation. Either the previously modeled 1-photon stimulus was added during single cell stimulation, or all excitatory and inhibitory cells were provided with a uniform input. Adding single cell stimulation on top of our previously modeling 1-photon stimulus (red) leads to a negative bias in comparison to the single cell stimulation condition (black). **(H)** Parametric manipulation of a uniform input to both E and I cells shows that stronger network drive leads to the effect of mean suppression during single cell stimulation.

Discussion

471 In this work, we see robust suppression in visual cortex due to direct optogenetic drive to
472 excitatory neurons. Using electrophysiology and 2-photon imaging, we found intermixed
473 elevated and suppressed neurons distributed in a salt-and-pepper pattern. This salt-and-pepper
474 distribution of responses resembles what is observed during visual input and arises without
475 input to inhibitory neurons.
476

The salt-and-pepper local pattern shows global trends over larger areas of the cortex. With
widefield calcium imaging, we saw mean elevated rates at the center of the stimulation site, and
mean suppression in a concentric circle a few hundred microns. The suppressed cortical
surround region is consistent with past data on surround suppression in vision (e.g.[40]). Our
observations show that cortical suppression can be driven by direct excitatory input to the
cortex, suggesting that surround suppression also arises from recurrent interactions. Our data
suggest recurrent connections are organized to produce a local salt-and-pepper pattern of
activation and suppression not just at the stimulation location (Fig. 2), but in the cortical
surround (Fig. 3). The salt-and-pepper distribution of responses we observe is therefore overlaid
on top of the global trends we observed in widefield imaging (Fig. 4) and electrophysiology (Fig.
3).

The response distributions and dynamics of the response to excitatory cell stimulation suggest a
network mechanism for the observed suppression. Previous work with balanced-state models
with strong coupling shows reshuffling: relatively large responses in individual neurons to
excitatory cell stimulation, while the population response distribution remains relatively invariant,
seen in both mice and monkeys[22]. We confirm this mechanism applies in a conductance-based
simulation (Fig. 7). Further, Sanzeni et al. (2022) shows that variability in optogenetic (external,
non-recurrent) input across cells is required for this reshuffling. By measuring opsin expression
(Fig. 6) we further show that variability in recurrent connectivity in the cortex is required to
explain our data. Recurrent input variability, combined with external input variability, largely
decouples external input strength and firing rate response in individual cells (Fig. 6). This is
because the input activates many different neurons, which each provide recurrent input to other
cells, and these network inputs have significant effects on neurons' responses.

Intuitively, the network mechanism that creates the salt-and-pepper excitation and suppression
is that external inputs first elevate the firing rates of excitatory cells (Fig. 5A-C), some more than
others. That activation excites inhibitory cells, also some more than others. The result is the
network settles into a new steady state (Fig. 5E–G) with a very broad distribution of excitatory
cell firing rate changes (Figs. 6,7). Each individual excitatory neuron receives a different amount
of net input (excitatory minus inhibitory) and thus many cells' rates are elevated and at the same
time a substantial fraction of neurons are suppressed by stimulation. This mechanism seems
likely to underlie variability in visual responses as well (Fig. 1), because recurrent connection
variability is present in the cortical network for all kinds of input, and thus is present both for
visual input and for the optogenetic input we provide.

Variability in recurrent connectivity in the cortex: experimental evidence

We find that variability in connection strengths between L2/3 neurons creates heterogeneous responses to input. Several observations suggest that the brain has recurrent variability at least as large, and possibly larger, than we assumed in the simulations. First, electrophysiological studies often find a long tail of synaptic strengths between pairs of neurons with a few very large connections[41–43], though the variance of individual synaptic weights may be lower[44], with the larger connection strengths due to multiple synaptic contacts between neurons (though see[45] for evidence of long-tailed synaptic bouton sizes). If there is a long tail in synaptic connection strengths, this would still support our finding of high recurrent variance, as it would increase the recurrent variability even beyond the weight distributions we used (truncated Gaussian with mean and variance equal). Second, we used a connection sparsity of 2%, and generally set the number of inputs a cell received from the recurrent network according to a binomial sum, with fixed connection probability between neurons. Connection probability in the brain may be higher than we used, as for example paired recording studies have found connection probabilities of 10% or higher[42,43]. And higher connection probability will produce greater variance in net input into different cells, as binomially-distributed sums have a larger variance as connection probability increases (in the 0–50% range). Finally, patterned or subnetwork-specific connections, which we did not include, would also only increase variance, though specific connections seem to have just a moderate effect on connection probability — shared tuning changes the connection probability from 10–20% on average to 30–50% for like-tuned neurons, in some cases[38]. Taken together, the substantial recurrent variability that explains our data is consistent with experimental measurements of recurrent connection variability.

Strong balance, loose balance: implications for models that describe cortical networks

In addition to supporting the idea that recurrent connections between neurons have substantial variability, our results also confirm that the mean V1 recurrent connectivity is strong — i.e. V1 operates an inhibitory-stabilized network, meaning that the excitatory network is unstable if inhibition could be frozen[11,13,14,22]. Within the class of balanced networks, two sorts of balance have been distinguished: "loose" and "tight" balance[10]. The best network in our results (Fig. 7) is on the border of the tight- and loose-balance regimes, with individual cells falling in either the tight or loose-balance regimes. A network near the transition from loose to tight balance is broadly consistent with past experimental data (reviewed in ref. [10]) which do not suggest a very tightly-balanced regime for the cortex (Fig. 7).

We find that a two-population excitatory-inhibitory model is sufficient to explain the data we observe. *A priori*, it could have been that a model with multiple inhibitory subtypes[15,46] would be needed to reproduce the dynamics and population statistics we saw. Recent work has argued for particular roles for cortical inhibitory subtypes: that parvalbumin-positive (PV) neurons are the primary class providing inhibition stabilization[13,35], while somatostatin-positive (SOM) cells are involved in gain control[35]. These separate roles are still consistent with our findings. PV cells are likely to be the primary inhibitory cell class in our data and model, as PV cells are the narrow-waveform cells that we identify in electrophysiology[13] (see Fig. 5). Those cells show dynamic and response firing rate changes expected for the inhibitory population in an E-I model (slightly delayed onset latency, similar distribution of firing rate change as E cells). It is also

plausible that stimulating cortical excitatory cells as we did does not cause gain to vary, so that
a separate gain role of SOM neurons was not evident in our experiments. Thus, a two-
population inhibitory model (with PV cells likely making up a large part of the I population in the
model) seems to be sufficient to explain the suppression we observe.

The mechanism we find for suppression is strikingly different than paradoxical suppression in an
ISN when inhibitory cells are stimulated[12–14]. In both cases, suppression is paradoxical: here we
excite excitatory cells and see suppression of excitatory cells, and in an ISN, exciting inhibitory
cells causes suppression in inhibitory cells. But in paradoxical inhibitory suppression, the *mean*
firing rate of the inhibitory population decreases[14]. Here with excitatory cell stimulation, the
mean firing rate change is non-paradoxical: excitatory cell average rates *increase*. It is the
substantial variability or heterogeneity of recurrent connections in combination with variability of
input that causes many cells to be suppressed as others increase their firing. However, both
types of paradoxical suppression, when excitatory or inhibitory cells are stimulated, are only
present when the network operates as an ISN – that is, both effects happen in a network with
strong average recurrent coupling[13,22]. The observed paradoxical suppression of excitatory cells,
however, requires variability around that strong average recurrent coupling.

Recurrent computations: what potential roles in visual cortical function are there for
recurrent connections?
Our data shows that recurrent connections in V1 affect the steady-state responses of neurons to
input. This is one possible role for recurrent connectivity in visual cortex: amplifying or
suppressing certain input patterns.

Artificial networks, such as deep network models, can perform very complex visual tasks without
recurrent connectivity. In advanced object recognition models, layers and areas connect to other
layers and areas, and the synaptic strengths of these feedforward connections control the
computation. This feedforward computation alone is extremely powerful — in principle, layered
feedforward networks can approximate any function, and can perform arbitrary classification
and transformation.

Why then, do recurrent connections seem to be able to have large effects on responses to input
in visual cortex? The suppression we see may be a signature of a common cortical operation,
normalization[47]. Normalization suppresses activity of other neurons within a layer when some
neurons increase their firing rates. From a statistical or informational standpoint, normalization
within a cortical layer results in improved coding — that is, units whose responses are less
correlated — in models trained to reflect natural scene statistics[48–50] (see also[51] for biological
data). And in artificial networks, normalization is used to keep the response mean or response
distribution similar when some cells are activated by input, which often improves learning[52].
Thus, normalization in biological and artificial networks tends to result in small changes in the
mean of population firing rates even as individual neurons show large changes in firing rates,
consistent with what we see (e.g. Fig. 7F-H).

If the suppression we see is a consequence of normalization that improves learning or other
functions, there is still a question of why it is provided by recurrent connections. Normalization in

deep networks for vision is often provided without using recurrent connections, by adding extra
feedforward layers (as in ResNet[53]). Our data imply biological brains, on the other hand,
collapse the normalization function into recurrent connections within a layer. Thus, brains seem
likely to perform normalization via recurrence, while retaining the ability to perform feedforward
computation via the pattern of connectivity between layers and areas. Models of visual cortex
with more biological features, like the stabilized supralinear network (SSN)[54] demonstrate this
mechanism — they perform a normalization-like function via recurrent connectivity within a
network that also transforms feedforward input.

Future: subnetworks, computation, and interactions between areas

These results could be extended in a few ways. First, here we did not consider how subnetwork
connectivity between excitatory neurons in the cortex might influence the effects. Ko and
colleagues (2011) showed approximately a 2–3 fold increase in probability of connection
between V1 excitatory neurons that had similar tuning (orientation or direction) vs those with
dissimilar tuning. Adding subnetwork connectivity would not qualitatively change our
conclusions: that suppression results from recurrent influences, and that it depends on
variability of connectivity within the network. However, future work stimulating within or across
subnetworks might change the fraction of cells suppressed, given that input patterns would drive
neuron populations with somewhat more or less connectivity with each other and the rest of the
network. Cell-specific two-photon holographic stimulation[16,55–57] seems well-placed to study how
patterned activity in one subnetwork affects activity in another subnetwork.

While local collaterals probably contribute the majority of recurrent cortical input, it is possible
that long-range, inter-areal, connections could contribute to the experimental results we
observe. We simulate a population of excitatory neurons and another population of inhibitory
neurons which are coupled to each other. One might assume that all these neurons are within a
few hundred microns of each other, as estimates of connectivity falloff show most connections
to a given neuron come from neurons within that distance[6,7]. But in principle, cells in other areas
could form part of the recurrent population. This could happen for example if projections from V1
to the thalamus recruited neurons there which connect back to the cortex. Understanding how
long-range connections contribute to the local pool of recurrent processing could be achieved
with stimulation studies with large scale silencing of other regions. However, our widefield
imaging data (Fig. 4) shows that the influence of stimulation extends only a few hundred
microns away from the stimulation site, suggesting relatively local influence. Therefore, it seems
likely that the recurrent connections in the simulations primarily reflect local connections within
V1 to nearby neurons.

Conclusion

Here we find paradoxical suppression of excitatory cells in the cortex when excitatory cells are
stimulated. A balanced-state model with strong mean recurrent connectivity, and variability in
drive across neurons, accounts for the data. These results suggest that a primary purpose of
recurrent connectivity in visual cortex is to change the steady-state firing rate of network
neurons, beyond just how inputs are transformed by feedforward connections. One purpose of
this recurrent cortical transformation may be to perform normalization, which confers

computational advantages in artificial neural networks and is found throughout many different biological networks. These results are a step forward in explaining how cortical networks fire in response to different patterns of input — a fundamental building block of neuronal computation.

Acknowledgements

We thank Kaya Matson for help with RNAscope, and Aanika Kashyap for histological analysis. We thank Nicolas Brunel and Alessandro Sanzeni for helpful comments on the manuscript. We also thank Ken Miller helpful discussion about the results. This work was funded by the National Institutes of Health (BRAIN U01 NS090576 and intramural support). This work utilized the computational resources of the NIH HPC Biowulf cluster (http://hpc.nih.gov).

Author contributions

Electrophysiology data was collected by ZZ. 2-photon data was collected by PKL, AJL, and JFO. Widefield data was collected by JFO. Histological data collection and analysis was performed by ZZ and HCG. JFO and MHH designed and implemented simulation experiments. JFO, MHH, and ZZ curated and analyzed data. JFO and MHH wrote and edited the manuscript.

References

1. Van Essen, D.C., Anderson, C.H., and Felleman, D.J. (1992). Information Processing in the Primate Visual System: An Integrated Systems Perspective. Science *255*, 419–423. 10.1126/science.1734518.

2. Glickfeld, L.L., and Olsen, S.R. (2017). Higher-Order Areas of the Mouse Visual Cortex. Annu. Rev. Vis. Sci. *3*, 251–273. 10.1146/annurev-vision-102016-061331.

3. Kravitz, D.J., Saleem, K.S., Baker, C.I., Ungerleider, L.G., and Mishkin, M. (2013). The ventral visual pathway: an expanded neural framework for the processing of object quality. Trends Cogn. Sci. *17*, 26–49. 10.1016/j.tics.2012.10.011.

4. Van Essen, D.C., and Gallant, J.L. (1994). Neural mechanisms of form and motion processing in the primate visual system. Neuron *13*, 1–10. 10.1016/0896-6273(94)90455-3.

5. Krizhevsky, A., Sutskever, I., and Hinton, G.E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In Advances in Neural Information Processing Systems (Curran Associates, Inc.).

6. Binzegger, T., Douglas, R.J., and Martin, K.A.C. (2004). A Quantitative Map of the Circuit of Cat Primary Visual Cortex. J. Neurosci. *24*, 8441–8453. 10.1523/JNEUROSCI.1400-04.2004.

7. Hellwig, B. (2000). A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. Biol. Cybern. *82*, 111–121. 10.1007/PL00007964.

8. Rossi, L.F., Harris, K.D., and Carandini, M. (2019). Excitatory and inhibitory intracortical circuits for orientation and direction selectivity. 556795. 10.1101/556795.

9. Douglas, R.J., Koch, C., Mahowald, M., Martin, K.A., and Suarez, H.H. (1995). Recurrent excitation in neocortical circuits. Science *269*, 981–985. 10.1126/science.7638624.

10. Ahmadian, Y., and Miller, K.D. (2021). What is the dynamical regime of cerebral cortex? Neuron. 10.1016/j.neuron.2021.07.031.

11. Ozeki, H., Finn, I.M., Schaffer, E.S., Miller, K.D., and Ferster, D. (2009). Inhibitory Stabilization of the Cortical Network Underlies Visual Surround Suppression. Neuron *62*, 578–592. 10.1016/j.neuron.2009.03.028.

12. Sadeh, S., and Clopath, C. (2021). Inhibitory stabilization and cortical computation. Nat. Rev. Neurosci. *22*, 21–37. 10.1038/s41583-020-00390-z.

13. Sanzeni, A., Akitake, B., Goldbach, H.C., Leedy, C.E., Brunel, N., and Histed, M.H. (2020). Inhibition stabilization is a widespread property of cortical networks. eLife *9*, e54875. 10.7554/eLife.54875.

683   14. Tsodyks, M.V., Skaggs, W.E., Sejnowski, T.J., and McNaughton, B.L. (1997). Paradoxical
684   Effects of External Modulation of Inhibitory Interneurons. J. Neurosci. *17*, 4382–4388.
685   10.1523/JNEUROSCI.17-11-04382.1997.

686   15. Mahrach, A., Chen, G., Li, N., van Vreeswijk, C., and Hansel, D. (2020). Mechanisms
687   underlying the response of mouse cortical networks to optogenetic manipulation. eLife *9*,
688   e49967. 10.7554/eLife.49967.

689   16. Marshel, J.H., Kim, Y.S., Machado, T.A., Quirin, S., Benson, B., Kadmon, J., Raja, C.,
690   Chibukhchyan, A., Ramakrishnan, C., Inoue, M., et al. (2019). Cortical layer–specific critical
691   dynamics triggering perception. Science *365*, eaaw5202. 10.1126/science.aaw5202.

692   17. Peron, S., Pancholi, R., Voelcker, B., Wittenbach, J.D., Ólafsdóttir, H.F., Freeman, J., and
693   Svoboda, K. (2020). Recurrent interactions in local cortical circuits. Nature, 1–4.
694   10.1038/s41586-020-2062-x.

695   18. Goldman, M.S. (2009). Memory without Feedback in a Neural Network. Neuron *61*, 621–
696   634. 10.1016/j.neuron.2008.12.012.

697   19. Murphy, B.K., and Miller, K.D. (2009). Balanced Amplification: A New Mechanism of
698   Selective Amplification of Neural Activity Patterns. Neuron *61*, 635–648.
699   10.1016/j.neuron.2009.02.005.

700   20. Chettih, S.N., and Harvey, C.D. (2019). Single-neuron perturbations reveal feature-specific
701   competition in V1. Nature *567*, 334–340. 10.1038/s41586-019-0997-6.

702   21. Oldenburg, I.A., Hendricks, W.D., Handy, G., Shamardani, K., Bounds, H.A., Doiron, B., and
703   Adesnik, H. (2022). The logic of recurrent circuits in the primary visual cortex.
704   2022.09.20.508739. 10.1101/2022.09.20.508739.

705   22. Sanzeni, A., Palmigiano, A., Nguyen, T.H., Luo, J., Nassi, J.J., Reynolds, J.H., Histed, M.H.,
706   Miller, K.D., and Brunel, N. (2022). Mechanisms underlying reshuffling of visual responses
707   by optogenetic stimulation in mice and monkeys. 2022.07.13.499597.
708   10.1101/2022.07.13.499597.

709   23. Histed, M.H. (2018). Feedforward Inhibition Allows Input Summation to Vary in Recurrent
710   Cortical Networks. eNeuro *5*, ENEURO.0356-17.2018. 10.1523/ENEURO.0356-17.2018.

711   24. Nassi, J.J., Avery, M.C., Cetin, A.H., Roe, A.W., and Reynolds, J.H. (2015). Optogenetic
712   Activation of Normalization in Alert Macaque Visual Cortex. Neuron *86*, 1504–1517.
713   10.1016/j.neuron.2015.05.040.

714   25. McGinley, M.J., Vinck, M., Reimer, J., Batista-Brito, R., Zagha, E., Cadwell, C.R., Tolias,
715   A.S., Cardin, J.A., and McCormick, D.A. (2015). Waking State: Rapid Variations Modulate
716   Neural and Behavioral Responses. Neuron *87*, 1143–1161. 10.1016/j.neuron.2015.09.012.

717   26. Beaudot, W.H.A., and Mullen, K.T. (2006). Orientation discrimination in human vision:
718   Psychophysics and modeling. Vision Res. *46*, 26–46. 10.1016/j.visres.2005.10.016.

27. Bondy, A.G., Haefner, R.M., and Cumming, B.G. (2018). Feedback determines the structure
of correlated variability in primary visual cortex. Nat. Neurosci. *21*, 598–606.
10.1038/s41593-018-0089-1.

28. Rolfs, M., and Carrasco, M. (2012). Rapid Simultaneous Enhancement of Visual Sensitivity
and Perceived Contrast during Saccade Preparation. J. Neurosci. *32*, 13744–13752a.
10.1523/JNEUROSCI.2676-12.2012.

29. Nakajima, M., Schmitt, L.I., and Halassa, M.M. (2019). Prefrontal Cortex Regulates Sensory
Filtering through a Basal Ganglia-to-Thalamus Pathway. Neuron *103*, 445-458.e10.
10.1016/j.neuron.2019.05.026.

30. Gorski, J.A., Talley, T., Qiu, M., Puelles, L., Rubenstein, J.L.R., and Jones, K.R. (2002).
Cortical Excitatory Neurons and Glia, But Not GABAergic Neurons, Are Produced in the
Emx1-Expressing Lineage. J. Neurosci. *22*, 6309–6314. 10.1523/JNEUROSCI.22-15-
06309.2002.

31. Sohal, V.S., Zhang, F., Yizhar, O., and Deisseroth, K. (2009). Parvalbumin neurons and
gamma rhythms enhance cortical circuit performance. Nature *459*, 698–702.
10.1038/nature07991.

32. Watakabe, A., Ohtsuka, M., Kinoshita, M., Takaji, M., Isa, K., Mizukami, H., Ozawa, K., Isa,
T., and Yamamori, T. (2015). Comparative analyses of adeno-associated viral vector
serotypes 1, 2, 5, 8 and 9 in marmoset, mouse and macaque cerebral cortex. Neurosci.
Res. *93*, 144–157. 10.1016/j.neures.2014.09.002.

33. Ahmadian, Y., Rubin, D.B., and Miller, K.D. (2013). Analysis of the Stabilized Supralinear
Network. Neural Comput. *25*, 1994–2037. 10.1162/NECO_a_00472.

34. Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., and Harris, K.D.
(2010). The Asynchronous State in Cortical Circuits. Science *327*, 587–590.
10.1126/science.1179850.

35. Bos, H., Oswald, A.-M., and Doiron, B. (2020). Untangling stability and gain modulation in
cortical circuits with multiple interneuron classes. 2020.06.15.148114.
10.1101/2020.06.15.148114.

36. Li, N., Chen, S., Guo, Z.V., Chen, H., Huo, Y., Inagaki, H.K., Chen, G., Davis, C., Hansel,
D., Guo, C., et al. (2019). Spatiotemporal constraints on optogenetic inactivation in cortical
circuits. eLife *8*, e48622. 10.7554/eLife.48622.

37. Schneider, F., Gradmann, D., and Hegemann, P. (2013). Ion Selectivity and Competition in
Channelrhodopsins. Biophys. J. *105*, 91–100. 10.1016/j.bpj.2013.05.042.

38. Ko, H., Hofer, S.B., Pichler, B., Buchanan, K.A., Sjöström, P.J., and Mrsic-Flogel, T.D.
(2011). Functional specificity of local synaptic connections in neocortical networks. Nature
473, 87–91. 10.1038/nature09880.

39. Podlaski, W., Russell, L.E., Roth, A., Bicknell, B., Häusser, M., and Machens, C. (2022).
The dynamical regime of mouse visual cortex shifts from cooperation to competition with

increasing visual input | Cosyne 2022. https://www.world-wide.org/cosyne-22/dynamical-
regime-mouse-visual-cortex-5af75de1/.

40. Angelucci, A., and Bressloff, P.C. (2006). Contribution of feedforward, lateral and feedback
connections to the classical receptive field center and extra-classical receptive field
surround of primate V1 neurons. In Progress in Brain Research Visual Perception., S.
Martinez-Conde, S. L. Macknik, L. M. Martinez, J.-M. Alonso, and P. U. Tse, eds. (Elsevier),
pp. 93–120. 10.1016/S0079-6123(06)54005-1.

41. Arellano, J., Benavides-Piccione, R., DeFelipe, J., and Yuste, R. (2007). Ultrastructure of
dendritic spines: correlation between synaptic and spine morphologies. Front. Neurosci. *1*,
10. 10.3389/neuro.01.1.1.010.2007.

42. Holmgren, C., Harkany, T., Svennenfors, B., and Zilberter, Y. (2003). Pyramidal cell
communication within local networks in layer 2/3 of rat neocortex. J. Physiol. *551*, 139–153.
10.1111/j.1469-7793.2003.00139.x.

43. Song, S., Sjöström, P.J., Reigl, M., Nelson, S., and Chklovskii, D.B. (2005). Highly
Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits. PLOS Biol. *3*, e68.
10.1371/journal.pbio.0030068.

44. Chapeton, J., Fares, T., LaSota, D., and Stepanyants, A. (2012). Efficient associative
memory storage in cortical circuits of inhibitory and excitatory neurons. Proc. Natl. Acad.
Sci. *109*, E3614–E3622. 10.1073/pnas.1211467109.

45. Loewenstein, Y., Kuras, A., and Rumpel, S. (2011). Multiplicative Dynamics Underlie the
Emergence of the Log-Normal Distribution of Spine Sizes in the Neocortex In Vivo. J.
Neurosci. *31*, 9481–9488. 10.1523/JNEUROSCI.6130-10.2011.

46. Litwin-Kumar, A., Rosenbaum, R., and Doiron, B. (2016). Inhibitory stabilization and visual
coding in cortical circuits with multiple interneuron subtypes. J. Neurophysiol. *115*, 1399–
1409. 10.1152/jn.00732.2015.

47. Carandini, M., and Heeger, D.J. (2012). Normalization as a canonical neural computation.
Nat. Rev. Neurosci. *13*, 51–62. 10.1038/nrn3136.

48. Schwartz, O., and Simoncelli, E.P. (2001). Natural signal statistics and sensory gain control.
Nat. Neurosci. *4*, 819–825. 10.1038/90526.

49. Valerio, R., and Navarro, R. (2003). Optimal coding through divisive normalization models of
V1 neurons. Network *14*, 579–593. 10.1088/0954-898X/14/3/310.

50. Coen-Cagli, R., and Schwartz, O. (2013). The impact on midlevel vision of statistically
optimal divisive normalization in V1. J. Vis. *13*, 13. 10.1167/13.8.13.

51. Snyder, A.C., Morais, M.J., Kohn, A., and Smith, M.A. (2014). Correlations in V1 Are
Reduced by Stimulation Outside the Receptive Field. J. Neurosci. *34*, 11222–11227.
10.1523/JNEUROSCI.0762-14.2014.

52. Ioffe, S., and Szegedy, C. (2015). Batch Normalization: Accelerating Deep Network Training
by Reducing Internal Covariate Shift. In Proceedings of the 32nd International Conference
on Machine Learning (PMLR), pp. 448–456.

53. He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep Residual Learning for Image
Recognition. ArXiv151203385 Cs.

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

55. Dalgleish, H.W., Russell, L.E., Packer, A.M., Roth, A., Gauld, O.M., Greenstreet, F.,
Thompson, E.J., and Häusser, M. (2020). How many neurons are sufficient for perception of
cortical activity? eLife *9*, e58889. 10.7554/eLife.58889.

56. Packer, A.M., Peterka, D.S., Hirtz, J.J., Prakash, R., Deisseroth, K., and Yuste, R. (2012).
Two-photon optogenetics of dendritic spines and neural circuits. Nat. Methods *9*, 1202–
1205. 10.1038/nmeth.2249.

57. Sadeh, S., and Clopath, C. (2020). Theory of neuronal perturbome in cortical networks.
Proc. Natl. Acad. Sci. *117*, 26966–26976. 10.1073/pnas.2004568117.

58. Daigle, T.L., Madisen, L., Hage, T.A., Valley, M.T., Knoblich, U., Larsen, R.S., Takeno,
M.M., Huang, L., Gu, H., Larsen, R., et al. (2018). A Suite of Transgenic Driver and Reporter
Mouse Lines with Enhanced Brain-Cell-Type Targeting and Functionality. Cell *174*, 465-
480.e22. 10.1016/j.cell.2018.06.035.

59. Franco, S.J., Gil-Sanz, C., Martinez-Garay, I., Espinosa, A., Harkins-Perry, S.R., Ramos, C.,
and Müller, U. (2012). Fate-Restricted Neural Progenitors in the Mammalian Cerebral
Cortex. Science *337*, 746–749. 10.1126/science.1223616.

60. Goldbach, H.C., Akitake, B., Leedy, C.E., and Histed, M.H. (2021). Performance in even a
simple perceptual task depends on mouse secondary visual areas. eLife *10*, e62156.
10.7554/eLife.62156.

61. Kelly, R.C., Smith, M.A., Samonds, J.M., Kohn, A., Bonds, A.B., Movshon, J.A., and Lee,
T.S. (2007). Comparison of Recordings from Microelectrode Arrays and Single Electrodes in
the Visual Cortex. J. Neurosci. *27*, 261–264. 10.1523/JNEUROSCI.4906-06.2007.

62. Nordhausen, C.T., Maynard, E.M., and Normann, R.A. (1996). Single unit recording
capabilities of a 100 microelectrode array. Brain Res. *726*, 129–140. 10.1016/0006-
8993(96)00321-6.

63. Wissig, S.C., and Kohn, A. (2012). The influence of surround suppression on adaptation
effects in primary visual cortex. J. Neurophysiol. *107*, 3370–3384. 10.1152/jn.00739.2011.

64. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D.,
Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. (2020). SciPy 1.0: fundamental
algorithms for scientific computing in Python. Nat. Methods *17*, 261–272. 10.1038/s41592-
019-0686-2.

65. Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B.L., Koay, S.A., Taxidis, J.,
Najafi, F., Gauthier, J.L., Zhou, P., et al. (2019). CaImAn an open source tool for scalable
calcium imaging data analysis. eLife *8*, e38173. 10.7554/eLife.38173.

66. Friedrich, J., Zhou, P., and Paninski, L. (2017). Fast online deconvolution of calcium
imaging data. PLOS Comput. Biol. *13*, e1005423. 10.1371/journal.pcbi.1005423.

67. Stern, M., Shea-Brown, E., and Witten, D. (2020). Inferring the Spiking Rate of a Population
of Neurons from Wide-Field Calcium Imaging 10.1101/2020.02.01.930040.

68. Stimberg, M., Brette, R., and Goodman, D.F. (2019). Brian 2, an intuitive and efficient neural
simulator. eLife *8*, e47314. 10.7554/eLife.47314.

69. Vogels, T.P., and Abbott, L.F. (2005). Signal Propagation and Logic Gating in Networks of
Integrate-and-Fire Neurons. J. Neurosci. *25*, 10786–10795. 10.1523/JNEUROSCI.3508-
05.2005.

70. Song, S., Sjöström, P.J., Reigl, M., Nelson, S., and Chklovskii, D.B. (2005). Highly
Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits. PLOS Biol. *3*, e68.
10.1371/journal.pbio.0030068.

Methods

Reagent Table

Reagent type (species) or resource	Designation	Source or reference	Identifiers	Additional information
Chemical compound, drug	Tamoxifen	Sigma-Aldrich	T5648-5G	
Genetic reagent (M. musculus)	Ai148	The Jackson Laboratory	RRID:IMSR_JAX:030328	
Genetic reagent (M. musculus)	Cux2-CreERT2	MMRRC	RRID:MMRRC_032779-MU	
Recombinant DNA reagent	AAV9-hSyn-jGCaMP7s-WPRE	Addgene	RRID:Addgene_104487	
Recombinant DNA reagent	AAV9-Syn-DIO-stChrimsonR-mRuby	Addgene	RRID:Addgene_105448	
Software	Mworks	The Mworks Project		mworks.github.io
Other	C and B Metabond	Parkell	S380	
Other	Kwik-sil	World Precision Instruments	KWIK-SIL	

Animals

All procedures were approved by the NIMH Institutional Animal Care and Use Committee (IACUC). Emx1-cre animals[30] (N = 10; https://www.jax.org/strain/005628) were used for 2-photon and electrophysiology experiments. For widefield imaging experiments, Ai162 (N = 2; https://www.jax.org/strain/031562) and Ai148 (N = 1; https://www.jax.org/strain/030328) animals[58] were crossed with the Cux2-CreERT2 line[59] (https://www.mmrrc.org/catalog/sds.php?mmrrc_id=32779), and GCaMP6f or GCaMP6s was induced via tamoxifen injection during adulthood (P22 or later, tamoxifen 2 mg intraperitoneally daily for 3 days). All animals were singly housed on a reversed light cycle.

Implants and injections

Details of the headplate and window procedures are described in previous studies[13,60]. Optical glass windows (3 mm diameter) were placed over V1 (center: -3 mm ML, +1.5 mm AP, relative to lambda) for 2p and widefield imaging. Windows were also used before electrophysiology for imaging to localize V1.

For Emx1-cre animals, 300 nL of AAV9-syn-jGCaMP7s-WPRE (RRID:Addgene_104487) and/or AAV9-Syn-DIO-stChrimsonR-mRuby (RRID:Addgene_105448) were injected 250 μm below the dura (200 nL/min) prior to cementing the cranial window. For Ai148 and 162 animals, AAV9-CamKIIa-stChrimsonR-mRuby2 was generated by cloning the CaMKIIa promoter (RRID:Addgene_120219) into a pAAV backbone containing stChrimsonR-mRuby2 (RRID:Addgene_105447) and packaged into an AAV (Vigene, Inc.). This was injected at the same depth as the hSyn-DIO-stChrimsonR virus, but with 100 nL volume at 100 nL/min.

Electrophysiology

Electrophysiological methods are described in detail in previous studies[13], and are summarized here. Animals were headfixed during recording. Before the first session of electrophysiology, the animal's cranial window was removed and the craniotomy was flushed with saline. Between recording sessions, the craniotomy was covered using Kwik-Sil polymer (WPI, Inc.). A fiber optic cannula (400 µm diameter, 0.39 NA, Thorlabs) was placed to center light output at the center of stChrimsonR expression. For light intensity calculations, spot area was defined as the area inside the 50% contour of light spot intensity on the cortex, measured with a camera by imaging the spot on the brain surface. 1–2% agarose (Type IIIA, Sigma) was placed over the dura at the start of each session, and an array of four electrodes (4 probes, 32 sites in total, part #A4x8-5mm-100-400-177-A32, NeuroNexus, Inc.) were lowered into the cortex using a micromanipulator (Sutter MPC-200). One probe was placed at the center of the light spot. Probes were advanced 600–1000 µm below the point in which the first probe touched the dura. Probes were not moved for 1 hour prior to recording, as we found this improved recording stability. Recording data was sampled at 30 kHz (Cerebus, Blackrock Microsystems.)

Optogenetic stimulation was performed with randomly interleaved stimulation light pulses with several intensities over the range 0.2 mW/mm^2 to 15 mW/mm^2. Stimulation pulses were 600 ms long and delivered with a 4 s period.

For spike recordings, waveforms (bandpass filtered, 750 Hz – 7.5 kHz) were digitized and saved by storing a short data section around points where amplitude exceeded 3 times the RMS noise on that channel. Single units were identified (OfflineSorter, Plexon, Inc) based on clusters in the waveform PCA that were separate from noise and other clusters, had unimodal spike width distributions, and inter-spike intervals consistent with cortical neuron absolute and relative refractory periods. A single-unit score was assigned to each unit manually based on these factors [13,23]. To compare these populations quantitatively, we calculated SNR for both single and multiunits[61,62]. Median SNR for single units was larger than median SNR for multiunits (SU: 3.32, MU: 2.26; Fig. 3 Suppl. 1A), consistent with prior reports[13,61,63].

Visual stimulation

Visual stimuli were presented using MWorks (https://mworks.github.io/). Grating stimuli (sinusoidal contrast variation, 0.1 cyc/deg, orientation = 0 deg) were masked with a circular raised-cosine envelope (15 deg FWHM.) Visual stimuli were displayed on an LCD display, with center positioned 0-10 degrees of visual angle temporal to the central meridian. The oriented noise stimuli were generated by filtering white noise pixel arrays (each pixel drawn independently from a uniform distribution) with a spatial band-pass filter (peak orientation = 0 deg, orientation bandwidth = 10 deg, peak spatial frequency = 0.05, frequency bandwidth = 0.05). Frames were generated at 60 Hz and the noise pattern was independent from frame to frame[26–28]. Visual stimuli were presented for 3 or 5 seconds, depending on the experiment.

2-photon imaging

During 2-photon experiments, animals were awake, water-scheduled, and given periodic water rewards (20% probability per trial, reward once every 30 s on average). If animals stopped

licking in response to the rewards, data collection was ended. We imaged GCaMP7s responses
(920 nm excitation) with either a galvo-galvo (5 Hz) or resonant scanning (30 Hz) two-photon
microscope. stChrimsonR-mRuby2 expression was imaged at 1000 nm. The microscopes used
for imaging were built using MIMMS components (https://www.janelia.org/open-science/mimms-
21-2020) and other custom components, built in-house or provided by Sutter Inc. A second light
path, combined with the 2p stimulation light path before the tube lens using a dichroic, was used
to stimulate stChrimsonR using 530nm light (CoolLED, pE-4000). For 200 ms long optogenetic
pulses, we measured responses in the first frame after stimulation. For 4 s long optogenetic light
pulses, we imaged while stimulation was ongoing. To do this, we avoided stimulation artifacts by
stimulating only during horizontal flyback (approximate pulse duration 19 µs, off time 44 µs, duty
cycle 30%, line rate 8kHz). Optogenetic pulse period was 6 s.

Widefield imaging
For widefield imaging experiments, we used Ai162;Cux2-creERT2 or Ai148;Cux2-creERT2
animals, expressing GCaMP6f or 6s in L2/3 excitatory cells. Animals were headfixed and awake
during widefield imaging experiments. Prior to imaging, a fiber optic cannula was aimed at the
center of the focal stChrimsonR expression. Images were collected using a Zeiss microscope
(Discovery V12) with a 1.0x objective using excitation light with wavelength centered at 475 nm
(Xylis X-Cite XT720L). A Zyla 4.2 sCMOS camera (Oxford Instruments) collected images (100
ms exposure time, approximately 140 ms frame period) with 4-pixel binning. Laser powers were
randomly interleaved, with 50 repetitions per laser power. Laser pulses were 600 ms long, and
presented with 6 s period.

Analysis of electrophysiology data
For spike rate plots, spike counts were binned (1 ms bins), and smoothed via LOWESS[64]. To
classify units as having elevated or suppressed responses, we measured spike rate over 145–
400 ms after stim onset, relative to baseline (-1020 ms–0 ms relative to stim onset) for 6
mW/mm^2 stimulation intensity. To classify cells as wide- or narrow-waveform, we used a spike
width threshold of 0.445 ms based on the bimodal distribution of waveform widths (Fig. 5B).
This threshold is consistent with pharmacological segregation of inhibitory and excitatory cells[13].
940
941   For analysis of onset times, we fit a sigmoid (logistic function) to each cell's response from 100
942   ms before to 100ms after laser pulse onset:

$$f(x) = \frac{L}{1 + e^{-k(x-x_0)}} + b$$

944
L: upper asymptote, b: lower asymptote, k: slope, x_0: onset latency. X_0 was constrained to the
range [onset+0.5 ms, onset + 30 ms]. We defined onset latency as x_0, the time to half-max. To
estimate the time to steady-state, the same function was fit to data from 500 ms before and after
the laser onset, with the spike rates within a 50 ms window around the initial transient blanked
by setting to the baseline firing rate. Each cell's time to steady-state was computed as the
difference between the steady-state onset and the initial onset (difference between the x_0
parameters of the two fits).

Analysis of 2-photon data

For short optogenetic stimulation (200 ms pulses) during two-photon imaging, we avoided
stimulation light influencing imaging responses by measuring responses in the frame after the
stimulus offset. For long pulses (4 s), we stimulated during imaging frames by restricting
stimulation to imaging line flyback and intensities we give are the average intensity, corrected
for the 30% stimulation duty cycle. Because we found that the LED device we used for
stimulation (pE-4000, CoolLED Ltd; specified bandwidth 100 kHz) had some variability in
onset/offset for each line, we removed pixels (~40% of frame) at left and right edges of field of
view to ensure no stimulation light could affect images. Image frames were motion corrected
using NoRMCorre through the CaImAn processing suite[65]. Deconvolution was done with
OASIS[66] via CaImAn. To ease interpretability of the deconvolution signals, each neuron's
deconvolved signal was normalized to have the same maximum value as the dF/F of the
corresponding fluorescence trace. To segregate populations into elevated and suppressed cells,
we performed a one sample t-test ($\alpha = 0.05$; testing if cell response mean was larger or smaller
than 0) on the deconvolved dF/F during the stimulation period (long pulses) or the frame just
after the stimulation period (short pulses). For the short pulses, we used the frame just after
stimulation to estimate responses for each neuron per trial. For the long pulses, we averaged
data within the period 750 ms after stimulus onset to the stimulus offset in order to capture the
steady-state response.

For visual response data, data were preprocessed in the same manner as the short optogenetic
stimulation experiments. In order to characterize steady-state responses, the averaging period
started 750 ms after stimulus onset and ended at stimulus offset.

Analysis of widefield imaging data

Widefield fluorescence images were motion corrected for rigid translation, and any linear trend
across the full imaging session was estimated via regression and subtracted. Deconvolution
was done via Widefield Deconvolution[67], which differs from single-neuron deconvolution
algorithms like OASIS by dropping the sparsity assumption that is useful for deconvolving spike
trains of single neurons. This algorithm produces better results for aggregated signals, such as
that from a single pixel during widefield imaging[67]. We rescaled the deconvolved signals to the
maximum dF/F of the imaging data, as with the two-photon deconvolution. Comparison of
Widefield Deconvolution, OASIS, and first-differencing is given in Fig. 4 Suppl. 1. To average
images across animals, images were aligned on the basis of their maximum response during
the first frame of the stimulation period. For quantification of spatial falloff (Fig 4G–I), we
averaged pixels in concentric circles around the peak response and fit a LOWESS line to these
radial distributions. Crossing points are the minimum distance at which the 95% confidence
interval contains zero.

Spiking network model

We simulated a conductance-based neural network model with 10000 neurons (8000 excitatory,
2000 inhibitory) to understand the recurrent features that contribute to the response properties
we observe during excitatory cell stimulation. Simulations were performed using Brian2[68].

Membrane and synaptic dynamics evolve according to the following equations:

$$(1)\ C\frac{dV}{dt} = g_L(E_L - V) + g_E(E_E - V) + g_I(E_I - V) + I_{background} + c \cdot I_{opto}(t)$$

$$(2)\ \frac{dg_E}{dt} = -g_E/\tau_E$$

$$(3)\ \frac{dg_I}{dt} = -g_I/\tau_I$$

Each synapse was stepped by its corresponding connection weight for every presynaptic spike.

Connections between neurons were made with 2% probability, independently for each potential
connection. This sparsity is lower than that observed *in vivo* (2% vs ~10%)[42,43], though this
value has been used before in studies of spiking neural networks[69]. Probability of connection
can vary dramatically across distance[42] (though see[70]), and unconditional estimates vary across

Table 1: Spiking neural network model parameters

Parameter	Value		Parameter	Value
τ_E	5 ms		Mean W_{II}	4.0 nS
τ_I	10 ms		Variance W_{II}	0.4 nS
E_L	-60 mV		Mean W_{IE}	5.0 nS
E_I	-80 mV		Variance W_{IE}	0.4 nS
E_E	0 mV		E cell $I_{background}$ Full Network	250 pA
g_L	10.0 nS		I cell $I_{background}$ Full Network	150 pA
Mean W_{EE}	0.4 nS		E cell $I_{background}$ 50% Network	222 pA
Variance W_{EE}	0.4 nS		I cell $I_{background}$ 50% Network	178 pA
Mean W_{EI}	0.8 nS		E cell $I_{background}$ 25% Network	207 pA
Variance W_{EI}	0.4 nS		I cell $I_{background}$ 25% Network	193 pA

different studies[38,42]. We chose a low probability to balance what has been observed in cortex
with producing networks that show important physiological characteristics, such as
asynchronous firing and low signal correlations.

Synaptic weights were drawn from truncated (rectified) Gaussian distributions. Allowing synaptic
weights to vary allows us to further assess how variability in recurrent input may impact the
variance in response to stimulation of excitatory cells. Mean connectivity parameters were
selected to produce irregular and asynchronous firing. Background currents were chosen for
inhibitory and excitatory cell populations to fix baseline firing-rates for each constructed network
to be between 4–6 Hz. Network parameters are shown in Table 1.

1-photon stimulation simulations

Optogenetic stimulation was provided as an additional constant current for the length of the
stimulation period, with onset and offset ramped linearly over 3 ms. The strength of the
optogenetic stimulation (c in Eq. 1) could vary across cells, and was initially chosen from a
lognormal distribution fit, scaled between 0 and 1 (Fig. 7C) to our in-vivo fluorescence estimates
of stChrimsonR expression, and later held constant (Fig. 7M,N). For each network, this
stimulation distribution was scaled by a constant in order to reproduce the same response rate
at the 75th percentile across excitatory cells.

Steady-state response was measured for each cell as their firing rate during the 1 s baseline
period subtracted from the firing rate during the last 500 ms of the stimulation period.

To reduce connection strength variability (Fig. 7N), we reduced the variability of the truncated
Gaussians that define connection strength by a factor of 100 (setting both the synaptic strength
variability and connection number variability to zero produced a network that was less stable).

Single cell stimulation simulations

Single cell stimulation simulation experiments were performed in the network we constructed
with the strongest recurrent coupling (described above). A single cell was chosen and
stimulated at the maximum value of the stimulation distribution previously used to push the
network to a matched 75th percentile point. This was done for 10,000 trials. A set of control
experiments were performed with the same parameters but with no stimulation. We then
compared response distributions across the stimulated distribution and the control distribution.

To further examine the effects of single cell stimulation, we segregated the E cell population into
two subsets: cells that receive a projection stimulated cell, that is: cells that have a non-zero
connection weight from the stimulated cell, and cells that receive projections from the inhibitory
cells that the stimulated E cell projects to.

To simulate single cell stimulation under similar conditions to Chettih and Harvey (2019) with
visual input, we provided single cell stimulation during either the log-normal 1-photon
stimulation, as previously described, or during uniform input of both excitatory and inhibitory
cells.

Balance index

We computed the balance index as described by Ahmadian and Miller (2021). For each neuron,
we computed this index as the net current (excitatory + inhibitory) divided by the excitatory
current. The index becomes smaller as balance becomes tighter, with component currents
becoming larger, and the index becomes larger as inhibitory input from the network shrinks.

1043

Supplemental Figures

1045   Fig. 1 Supplement 1

1046

Fig. 1 Suppl. 1: **Both grating patches and oriented noise stimuli produce steady-state**
elevation and suppression in layer 2/3 of V1. (A) dF/F response at each imaging pixel to
oriented noise stimuli (small stimulus, FWHM = 15 deg, stimulus approximately aligned to cells'
receptive fields measured outside this experiment, same animal as in Fig. 1), corresponding to
the deconvolved cell responses shown in Fig. 1B. Here and in Fig. 1B, responses are measured
beginning 750 ms after stimulus onset to focus on steady-state response (Methods.) Evidence
of suppression is seen here but is more evident when data is deconvolved (compare this panel
to Fig. 1B), as expected for sustained suppression preceded by a transient, as the initial
transient seen in Figs. 2, 3, 7. **(B)** dF/F response to a drifting grating (Gabor patch, spatial
frequency 0.1 cpd, FWHM 15 deg), from the same animal, showing cells that are elevated and
suppressed in response to drifting gratings. Overall pattern of responses to noise stimulus and
grating is similar. **(C)** Deconvolved population response to oriented noise stimulus (replicated
from Fig. 1F for comparison.) Stimulus on during time indicated by light red shaded box. **(D)**
Deconvolved cell responses to Gabor patch, same data as in (B). Gabor patches drive both
steady- state elevation and suppression, though show signs of stronger off responses and
potentially a larger onset transient. **(E)** Population deconvolved response to oriented noise
stimulus in two additional animals, consistent with effects from example animal.

Fig. 3 Supplement 1

1065

1066 Fig. 3 Suppl. 1: **Sorting and quality of electrophysiology data (A)** Single units demonstrate
1067 higher SNR (N=136, median = 3.32) than multi-units (N=184, median = 2.26). **(B)** Mean spike-
1068 waveform of putative excitatory units (wide) in red (N = 94), mean spike-waveform of putative
1069 inhibitory units (narrow) in blue (N = 42.) Bimodal histogram of spike widths is shown in Fig. 5B.

1070

Fig. 3 Supplement 2

1072
Fig. 3 Suppl. 2: **Steady-state firing rates of neurons in layer 2/3 follow a weak spatial**
gradient with similar trends as the spatial distribution observed in cell counts. (A)
Elevated cell steady-state rates, with the highest and lowest powers for comparison. Rate is the
difference in firing rate during stimulation relative to baseline. **(B)** Suppressed cell steady-state
rates, with the highest and lowest powers for comparison, measured in relation to decreases
from baseline.

1079

1081

Fig. 4 Suppl. 1: **Center-surround organization is present regardless of deconvolution method. (A)** Mean whole-frame dF/F GCaMP response in an example animal. **(B)** We tested 3 different methods of deconvolution, OASIS [66], first-differences (i.e. subtracting one frame from the previous), and Widefield Deconvolution [67]. Widefield Deconvolution is expected to be the best method, as it is designed for data like this and does not incorporate the sparse-event constraints of OASIS, which is designed for single neurons. We found similar time-series results for each of the methods. The first-differences method (i.e. deconvolution with an kernel that decays immediately) seems to overestimate decreases in firing rate, as might be expected. All panels use the same dataset. **(E, F, G)** Spatial distribution of responses during the early laser period. All deconvolution methods produce a qualitatively similar excitatory response during this early period. **(H, I, J)** Spatial distributions of responses during late laser period demonstrates slight differences in size of surround, but overall a qualitatively similar center-surround organization with all methods. **(K, L, M)** Spatial distributions of response during the late laser period, but with dashed contours depicting the manually-drawn regions of interest (ROIs) that we used to produce the time-series data in **(B, C, D)**, with red dashed contours representing the center ROIs, and blue dashed contours representing the surround ROIs.

1098   Fig. 4 Supplement 2

1099

1100

1101   Fig. 4 Suppl. 2: **Stimulation response correlates to the pattern of stChrimsonR expression.**
1102   **(A)** Example animal's response during the late stim period. **(B)** Example animal's stChrimsonR
1103   expression pattern (gray: fluorescence) with overlaid contours of fluorescence intensity. **(C)**
1104   Example animal's response during the late stim period overlaid with their stChrimsonR contours.

1105

1106

1107

1108   Fig. 4 Suppl. 3: **Population time course of normalized center and surround deconvolved**
1109   **GCaMP response to laser stimulation (N = 3).** Error represents the standard error to the
1110   mean across the 3 animals.

1111

 Fig. 5 Supplement 1

1113

Fig. 5 Suppl. 1: **Differences in dynamics are restricted to those seen between the onset of wide- and narrow-waveform cells.** The excitatory and inhibitory (wide- and narrow-) onset latency difference is shown in Fig. 7C. Other quantities shown here do not differ: wide- vs narrow (excitatory vs inhibitory) time to steady state (B), and onset time and time to steady state (C,D) for elevated and suppressed groups of wide-waveform excitatory cells. **(A)** Example single neuron firing rate with fits. To obtain the onsets for individual cells, each cell's mean timecourse was smoothed with width dependent on the detectability of the transient signal (SNR; Methods), then a logistic function was fit to data from time range [-100ms, 100ms]. The onset time (latency) was defined as the time to half-max of the logistic function. **(B)** No difference in median time to steady state was found across narrow-spiking and wide-spiking cells. **(C)** No difference in median onset time for elevated and suppressed groups of wide-waveform (excitatory) cells. **(D)** Same as B, but difference in median time to steady-state.

1127

Fig. 7 Suppl. 1: **Networks at all tested recurrent strengths operate within the ISN regime.**
(A) To examine paradoxical suppression, we record the steady-state responses of the I cells in
response to different levels of stimulation. We performed this experiment on all networks
presented in Figure 8. **(B)** Steady-state responses of the I cell population to 5 levels of
stimulation. Error bars are standard error to the mean. Graph has been zoomed into the region
which clearly shows paradoxical suppression in all 3 networks. This paradoxical suppression is
predicted for both loosely and tightly balanced networks. Our simulations used three recurrent
strength values, one in the tight-balance regime and two in the loose-balance regime, and we
confirmed that all three showed paradoxical effects of suppression when I cells are stimulated
(C) Same as (B) but input normalized by the input value calculated in Fig. 8 to drive each
network to the same firing rate (input level that achieves same value of the 75th percentile of
evoked rates; see Fig. 8.)

1140

1141   Fig. 7 Supplement 2

1142

1143

Fig. 7 Suppl. 2: **Inhibitory neurons in balanced state model show similar responses to**
excitatory neurons but are recruited after initial stimulation. (A) Mean timecourses for
elevated and suppressed inhibitory cells (left and right, respectively) show the same
characteristic transient response followed by steady-state responses. **(B)** Cross-correlation
analysis of E- and I-cell response. Network has no synaptic delays built into the model. E-cells
respond to direct stimulation, and then I-cells are recruited after. **(C)** Population distribution of
steady-state responses is similar across E- and I-cells, though excitatory cells show a slightly
longer-tailed positive response (true in the data as well; Fig. 5E), as seen through the
distribution of responses or their corresponding CDFs **(D)**.

1153

1154

1155

1156

 Fig. 7 Supplement 3

1158

1159

Fig. 7 Suppl. 3: **Increasing strength of recurrent connections does not substitute for recurrent connection variability. (A)** Cumulative distribution of responses to optogenetic stimulation in model with 1x recurrent strength, matched to the 75th percentile of the response measured using electrophysiology. Negligible recurrent variability in this simulation (same number of recurrent connections to each neuron, variability in recurrent strength ~1% of mean, see Methods), and so spread in responses as a function of input is due to optogenetic input variability. Distribution of input is inferred from data in Fig. 6 (lognormal fit; Methods.) **(B)** Same as (A) but in model with 2.5x recurrent strength. **(C)** Relationship between input and steady-state response in the model with 1x recurrent strength. Marginal distribution of response show on the right. **(D)** Same as (C), but in model with 2.5x recurrent strength. Note that both stimulations produce similar variability between input strength and firing rates. This variability is seen as spread in the red cloud of points around an imagined curve that could be fit through the points. **(E)** Estimated mean absolute error of the relationship between the input and output as measured by a LOWESS fit across all recurrent strength manipulations.