Single cell optogenetics reveals attenuation-by-suppression in visual cortical neurons

Paul K. LaFosse[1,2,3], Zhishang Zhou[1], Victoria M. Scott[1], Yanting Deng[1], Mark H. Histed[1]

1 Intramural Program, National Institute of Mental Health, National Institutes of Health, Bethesda, MD USA 20892
2 NIH-University of Maryland Graduate Partnerships Program, Bethesda, MD USA 20892
3 Neuroscience and Cognitive Science Program, University of Maryland College Park, College Park, MD USA 20742

Correspondence: mark.histed@nih.gov

Abstract

The relationship between neurons' input and spiking output is central to brain computation. Studies *in vitro* and in anesthetized animals suggest nonlinearities emerge in cells' input-output (activation) functions as network activity increases, yet how neurons transform inputs *in vivo* has been unclear. Here, we characterize cortical principal neurons' activation functions in awake mice using two-photon optogenetics and imaging. We find responses to fixed optogenetic input are nearly unchanged as neurons are excited, reflecting a linear response regime above neurons' resting point. In contrast, responses are dramatically attenuated by suppression. This attenuation is a powerful means to filter inputs arriving to suppressed cells, privileging other inputs arriving to excited neurons. These data have two major implications: first, neural activation functions *in vivo* accord with the activation functions used in recent machine learning systems, and second, neurons' IO functions can enhance sensory processing by attenuating some inputs while leaving others unchanged.

Introduction

Neurons transform their synaptic inputs into output activity. The way in which neurons, and connected networks of neurons, achieve these transformations is a central and critical component of neural computation. For single neurons, frequency-current (f-I) curves describe how aggregate synaptic input generates spiking. F-I curves (in different contexts called static nonlinearities, neural transfer functions, or what we will use in this work, input-output or IO functions (McCormick et al., 1985; Powers and Binder, 2001; Fourcaud-Trocmé et al., 2003)) thus govern how spikes are produced and passed down to postsynaptic partners. At the network level, the shape of the input-output function affects performance of the overall network. For example, IO function shape affects how well sensory networks can encode features of the world (Moskovitz et al., 2018). Machine learning (ML) work has confirmed that the shape of IO functions impacts network function, in part by exerting influence on learning speed and accuracy. For example, now-classic work has found that rectified-linear (ReLU) IO functions (or

activation functions, the name often used in the ML context) can make learning more efficient and robust than learning with sigmoidal activation functions (Glorot et al., 2011; Nair and Hinton, 2010; Krizhevsky et al., 2012).

While work in all these fields often assumes for simplicity that neural IO functions are fixed or static, in real biological networks IO function shape is critically affected by all the inputs neurons receive — that is, IO function shape is affected by activity in the local recurrent network (Chance et al., 2002; Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002; Priebe and Ferster, 2002; Agüera y Arcas and Fairhall, 2003; Rauch et al., 2003; Carandini, 2004; Ostojic and Brunel, 2011; Atallah et al., 2012; Haider et al., 2013). Thus, accurately determining neural IO functions requires measurements *in vivo* during normal network operation.

Past work has reported various shapes for neural IO functions. *In vitro* studies, where network activity levels can be low, have often found ReLU (also called threshold-linear) functions (Granit et al., 1963; McCormick et al., 1985; Jagadeesh et al., 1992; Carandini et al., 1996; Ahmed et al., 1998; Chance et al., 2002). Studies done in networks with higher levels of network activity than expected *in vitro* — in anesthetized animals, *ex vivo*, or *in vitro* with simulated noise — have found power law functions, with a supralinear regime throughout the input range (Anzai et al., 1999; Anderson et al., 2000; Priebe et al., 2004; Finn et al., 2007; Cardin et al., 2008; Petersen and Berg, 2016). On the other hand, theoretical studies of integrate-and-fire networks of active neurons have found an analytical description of the IO function in those networks, the Ricciardi input-output function, which has a supralinear regime followed by a linear, then a sublinear regime (Ricciardi and Smith, 1977; Sanzeni et al., 2020).

Those different descriptions conflict in important ways, and therefore the IO function shape seen *in vivo* during awake states of normal operation has been unclear (Yamashita et al., 2013; Tan et al., 2014; Zhao et al., 2016). This conflict arises in part because electrophysiological studies with control of neural input are difficult to perform in awake animals (Dubois, 2000; Bar-Yehuda and Korngreen, 2008; Wilson et al., 2011). Here, we use two-photon optogenetics to expand measurements of IO functions to hundreds of neurons in awake mice. We deliver the same optogenetic input to single cells at different activity levels, and thus determine neurons' input-output responses at different points along their IO function.

We use this approach to explore the IO functions of neurons in visual cortex. We stimulate selected neurons optogenetically, deliver visual stimuli that excite or suppress cells, and measure differences in response to the same optogenetic stimuli. We find neurons *in vivo* have an IO function that resembles not a ReLU or a power law function, but the Riccardi function that is created by integrate-and-fire networks. The measured IO functions, as input varies from low to high, have a supralinearity followed by a linear region, and show some saturation at the upper end of the range of normal input. That upper range is defined by normal network activation levels — neurons' firing in response to strong, high-contrast visual stimuli. A new and key observation of this work is that neurons rest near the top of the supralinear region. This means suppressed neurons produce smaller responses to the same input, suggesting cortical neurons

attenuate responses to input that arrive to suppressed cells. This attenuation-by-suppression may be a computational means by which single neurons filter or mask out certain inputs.

Results

Uncovering input-output nonlinearities via cell-specific optogenetic stimulation

We use cell-specific two-photon optogenetic stimulation (Marshel et al., 2019; Dalgleish et al., 2020; LaFosse et al., 2023) to measure cells' IO functions. By providing the same fixed input at different points along a cells' IO function, the corresponding evoked changes in output can be used to infer the IO function shape (Fig. 1a). For example, if a cell's IO function were threshold-linear (Fig. 1b, left), the response to a fixed optogenetic input would remain unchanged as the cell's activity level increases. Similarly, when that cell is suppressed below the threshold, input would elicit no response. Our stimulation bypasses dendritic nonlinearities, by delivering input to the cell at the soma (using a soma-targeted opsin, stChrimsonR), allowing us to measure the transformation of the cell's aggregate input into output spiking.

The paper proceeds as follows. We deliver visual stimulation to change cells' activity levels. We first deliver optogenetic stimuli paired with high-contrast visual stimuli that produce a wide range of responses, both suppression and activation, across populations of neurons in V1 (Fig. 2). Using these data we infer the shape of the average IO function across cells. We then show that this average shape is consistent with that seen in individual cells (Fig. 3). To do this we vary first visual stimuli to alternately suppress and activate the same cell, and then vary optogenetic stimulation intensity. We end by reconstructing this IO function, inferring the average operating point (Fig. 4; see Fig. 1c for a schematic), and discussing the major features of the IO function that we have uncovered.

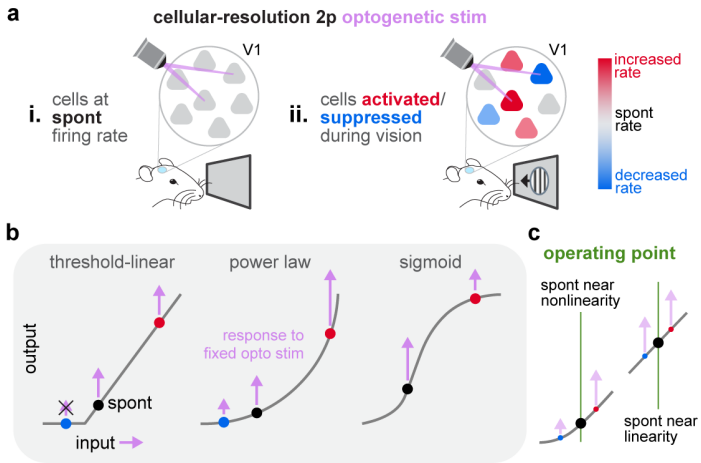

Fig. 1: Mapping neurons' input-output functions in awake mice using cellular-resolution optogenetic stimulation.
a, Schematic of experiments in V1 of awake mice during (**i**) spontaneous activity and (**ii**) vision. **(b)** Schematics of some possible input-output (IO) functions. Optogenetic responses at different points along the IO function can uncover the shape and extent of nonlinearities. (**c**), Beyond the shape of the IO function, a key unknown feature is where neurons sit on their IO curves during spontaneous activity (schematized by vertical green lines).

Cell-specific optogenetic stimulation during and without visual stimulation reveals attenuation of inputs with suppression

We expressed the soma-targeted opsin stChrimsonR and the calcium indicator jGCaMP8s via a single Cre-dependent virus (in the Emx1-Cre mouse line), and stimulated V1 excitatory pyramidal cells (Fig. 2a) at cellular resolution while imaging them. To compare responses to optogenetic stimulation at different points along cells' IO functions, we stimulated during both spontaneous conditions (i.e., no concurrent sensory stimulation), as well as during presentation of a visual stimulus (Fig. 2b). Many cells change their activity in response to drifting gratings, with many increasing their firing rates (visually-activated) and some decreasing (visually-suppressed; Fig. 2c-d; Extended Data Fig. 1).

We targeted small groups of cells for stimulation (5-12 cells) to increase yield by collecting data from several neurons in parallel, while avoiding triggering larger network interactions, which have been reported to occur at and above approximately twenty stimulation targets (Marshel et al., 2019; Dalgleish et al., 2020; O'Rawe et al., 2023). The fact that larger polysynaptic network effects do not change our results are supported by additional experiments (below, Fig. 3) where we measure the same neurons with different visual stimuli and different stimulation intensity, and find consistent results throughout.

We applied optogenetic stimulation while cells' firing rates were changed during visual stimulation, as well as during periods of spontaneous activity with no visual stimulus (visual stimulus: 3-4 seconds, opto stimulus: 300 ms, with onset 1.5-2.0 s after visual stimulation onset; Methods). Across trials, visual responses remained consistent, with minimal response adaptation (Extended Data Fig. 2). We compared the response produced by optogenetic stimulation during the visual stimulus (Opto+Vis, Fig. 2f; incremental response during visual stimulation defined as (Opto+Vis)-Vis, Fig. 2g) to the response during optogenetic stimulus alone (Opto, Fig. 2g).

If the IO function follows a power law (is concave up; Fig. 1b), the optogenetic response should increase as neurons' activity increases (schematic: Fig. 2e). If cells instead operate in a linear regime, they should produce a similar response increment to optogenetic input even as neurons' activity varies.

For many neurons whose firing rates were increased by visual stimulation, the response to stimulation was similar when optogenetic stimulation was delivered alone and with visual stimulation (example cell, Fig. 2h; population, Fig. 2n). Neurons that were moderately activated by the visual stimulus (N = 70 cells; $\Delta F/F_0 > 5\%$ and $< 60\%$) showed the same response in the two conditions ($p = 1.0$, K-S two-sample test with Bonferroni correction for multiple comparisons; Fig. 2n, top). This indicates that above neurons' firing rate at rest, neurons' average IO function is nearly linear.

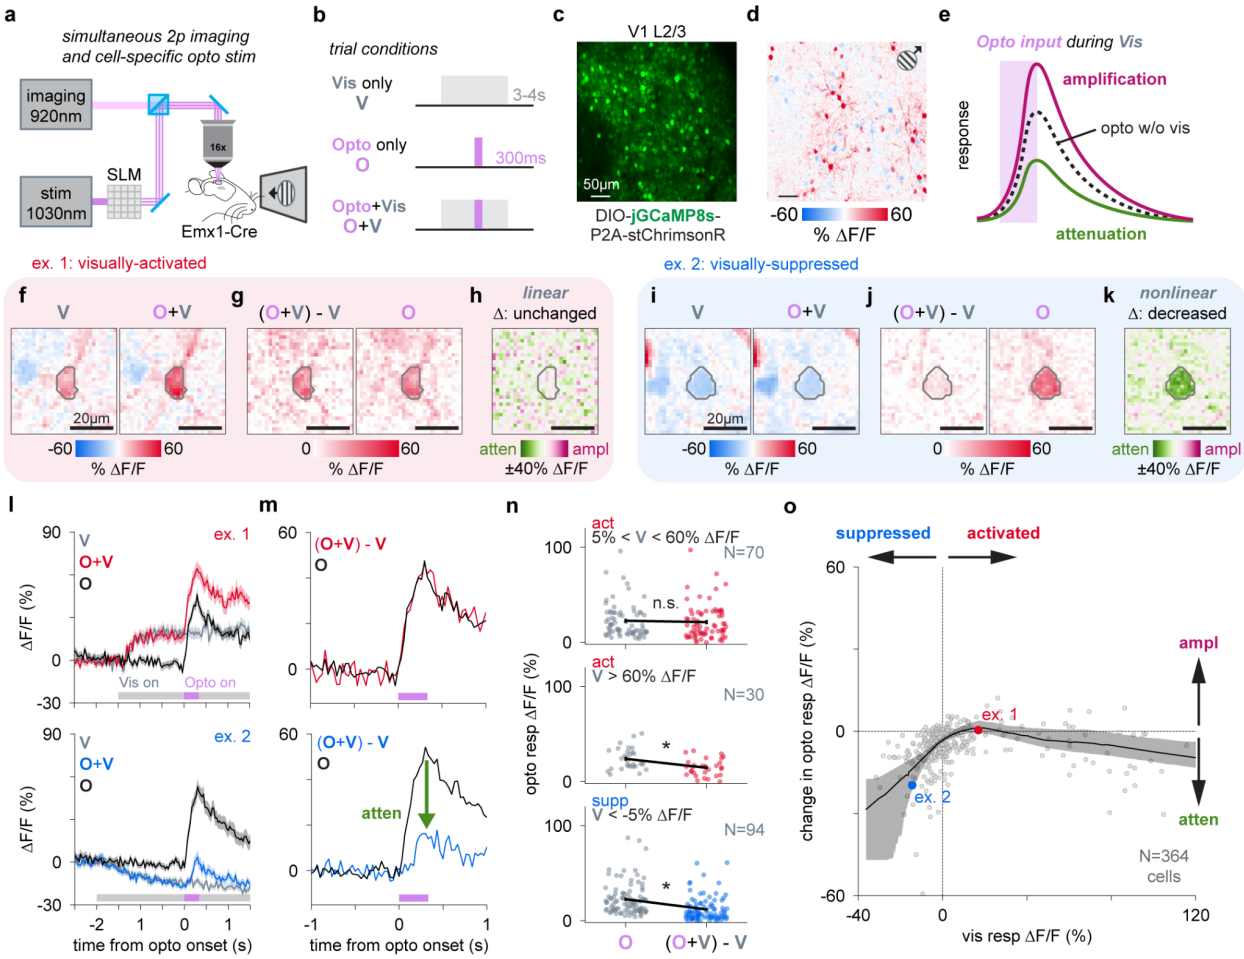

Fig. 2: V1 neurons attenuate inputs when suppressed.
a-b, Experiment schematic. Two-photon calcium imaging during visual, cell-specific opto, or paired opto+vis stimulation in awake mice. **c**, DIO-jGCaMP8s-P2A-stChrimsonR expressing in V1 layer 2/3. **d**, Visually-evoked responses to drifting grating shows suppressed (blue) and activated (red) neurons. **e**, Schematic showing increases/decreases of opto responses during vision reveal nonlinear amplification/attenuation of inputs. **f**, Example visually-activated cell. Visual response (left) and paired opto+vis response (right). **g**, Opto response during vision (left) and opto only response (right). **h**, No change in opto response during and without vision (linear regime) **i-k**, Same as **f-h** for a visually-suppressed cell with decreased opto response during vision (nonlinear regime). **l**, Trial-averaged (mean ± SEM, N = 50 reps) activity traces for the same visually-activated (top) and visually-suppressed (bottom) example cells in **f-k**. **m**, Comparison of opto responses during and without vision. Green arrow depicts attenuation in visually-suppressed cell. **n**, Visually-activated cells (vis response between 5% and 60% $\Delta F/F_0$; N = 70) show no change in opto response (top). Strongly-activated cells (vis response > 60% $\Delta F/F_0$; N = 30) show decreased opto response, suggesting saturation (middle). Visually-suppressed cells (vis response < -5% $\Delta F/F_0$; N = 94) show attenuated opto response (bottom). *p < 0.05; K-S two-sample test with Bonferroni correction for multiple comparisons. **o**, Change in optogenetic response in neurons (N = 364 total cells, 15 ± 1.3 cells responsive per pattern; N = 25 patterns across N = 7 animals) across the working range of visual responses. Black line: LOESS fit to data with 95% confidence intervals via bootstrap.

In cells exhibiting large increases in firing rate, we found some signs of saturation. Neurons more strongly activated by the visual stimulus (N = 30 cells; $\Delta F/F_0$ > 60%) showed decreased responses to optogenetic input ($p = 2*10^{-3}$, K-S two-sample test with Bonferroni correction; Fig. 2n, middle).

In contrast to the linear effects we observed for increases in firing rate, many cells suppressed by visual stimulation showed substantial attenuation of their response to optogenetic input (Fig. 2i-k). The attenuation was often dramatic, as in the example cell (Fig. 2l-m, bottom) whose response was reduced to less than 50% of its response before suppression. Attenuation was widespread across the population of suppressed neurons (N=94 suppressed cells, $\Delta F/F_0 < -5\%$, attenuation average 47.3%, $p = 6*10^{-8}$, K-S two-sample test with Bonferroni correction; Fig. 2n, bottom). These observations suggest that below neurons' resting activity levels (below spontaneous activity) the IO function shape is not linear, and instead shows a substantial nonlinearity.

The timecourses of neurons' responses to optogenetic input (Fig. 2l-m) were as expected for both suppressed and activated neurons, featuring a fast increase in fluorescence during optogenetic stimulation consistent with induced spiking, followed by a slower decay after optogenetic stimulus offset consistent with calcium decay dynamics. In agreement with these observations, deconvolving our data produced qualitatively unchanged results (Extended Data Fig. 3).

To characterize neurons' IO functions across a range of activity levels, we plotted responses across many optogenetic stimulation sessions in many cells, spanning the working range of visual responses — from strongly suppressed to strongly activated (N = 11 experiments in N = 7 mice; N = 364 total cells stimulated). We plotted the average change in optogenetic response as a function of the neurons' responses to visual input just before receiving optogenetic stimulation (Fig. 2o). As cells were more suppressed, we found, on average, more attenuation (computed via LOESS regression; Methods). The effects did not depend on the size of the visual stimulus. Cells across the working range of visual responses showed the same attenuation-by-suppression, and linear responses above spontaneous activity, whether animals were presented with full-field or retinotopically-aligned Gabor (full width, half max = 15 degrees) drifting gratings (Extended Data Fig. 4).

Together, these results imply the average IO function of pyramidal cells in mouse V1 exhibits a linear regime above and a nonlinear regime below the resting activity level of neurons in the awake state. The nonlinearity below rest attenuates inputs, so that suppression leads to attenuation.

Attenuation, but not amplification, measured in individual pyramidal cells

We next asked if responses reflected by the cell-average trends were present within individual neurons. First, we identified cells whose activity was suppressed by a drifting grating of one orientation but increased by the orthogonal orientation (Fig. 3a). We then optogenetically stimulated these cells as they fired spontaneously and also while we presented either grating (Fig. 3b-c; N = 6 cells in N = 3 animals). We found results in single cells that are consistent with our previous results from populations of cells: suppression led to attenuation of responses, while optogenetic stimulation applied when cells' activity was elevated by visual input produced similar responses (i.e., a linear response; p = 0.028; student's t-test; Fig. 3d).

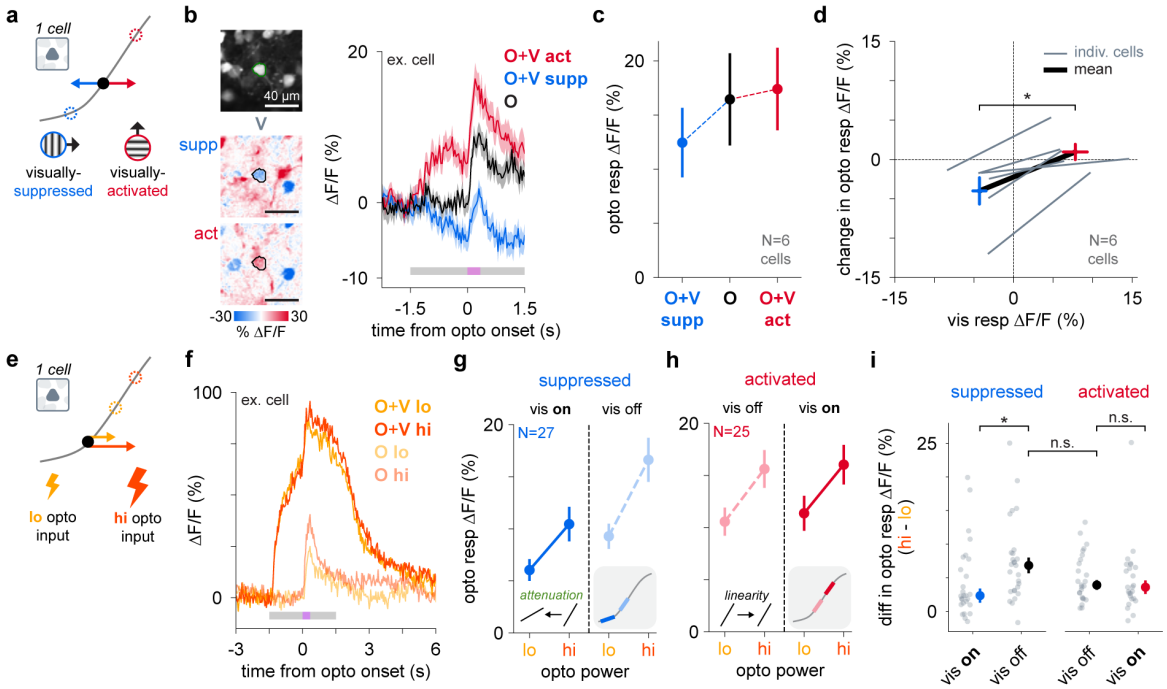

Fig. 3: Individual V1 neurons exhibit nonlinear responses when suppressed and linear responses when activated.

a, Experiment schematic showing single cells suppressed and activated by orthogonal gratings, moving them down and up along their IO function. **b**, Example neuron suppressed and activated by two different visual stimuli (left). Trial averaged traces (mean ± SEM; N = 50 trials) of opto only and opto during either visual stimulus. **c**, Cell-averaged (mean ± SEM; N = 6) opto responses during and without visual stimuli. **d**, Change in opto response plotted against visual response for all cells. Gray lines: individual cells at their suppressed and activated states. Black line: mean across cells. Crosses are SEM of change in opto response and visual response across cells. Blue: suppressed. Red: activated. **e**, Experiment schematic showing single cells stimulated with two opto powers (lo: 6.5 mW/target, hi: 8 mW/target). **f**, Example neuron stimulated with both opto powers during and without visual stimulation. **g**, Cell-averaged (mean ± SEM; N = 24) opto responses in visually-suppressed cells for vis on and vis off conditions. Inset: as cells are visually-suppressed, they move to a lower slope region of their IO function, measured between the two points of opto powers. **h**, Same as **g**, for visually-activated cells (N = 20). Inset: as cells are visually-activated, they remain in a linear regime of their IO function. **i**, Difference in opto response between low and high opto powers. Grey points: individual cells. Mean ± SEM across cells plotted for vis on (blue: suppressed, red: activated) and vis off (black) conditions. Suppressed cells show decreases in response difference between opto powers. *p < 0.05; K-S two-sample test with Bonferroni correction for multiple comparisons.

We next examined whether stimulation of a single cell with multiple optogenetic intensities (Fig. 3e-f) was consistent with the results we found in averages over cells (Fig. 2). If neurons have a nonlinear region below their resting point, then we should expect a smaller difference between responses (smaller slope) to the two stimulation intensities in suppressed cells. Consistent with this idea, we found that as cells were suppressed, the difference in responses between high and low powers decreased (Fig. 3g, N = 27 cells in N = 3 animals; Fig 3i, left, p = 0.027; K-S two-sample test with Bonferroni correction). Similarly, a linear region above cells' resting point implies there should be little or no difference in responses as cells' firing rates increase. Confirming this, we found no difference in responses at either optogenetic power (N = 25 cells in N = 3 animals; Fig. 3h) and found there was no difference (change in slope) between high and low powers (Fig. 3i, right; Extended Data Fig. 5).

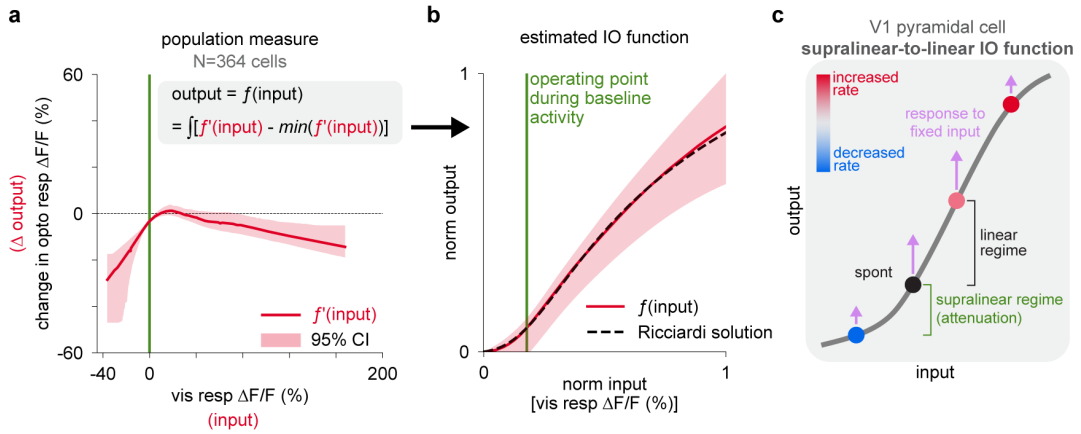

Fig. 4: V1 pyramidal cells operate at the transition point of a supralinear-to-linear IO function.
a, Change in opto response as a function of visual response. Red line: LOESS fit to cell data (N = 364 cells) with 95% confidence interval. **b**, Estimated IO function and 95% confidence intervals calculated by shifting and integrating the fit trend line in **a** (see Methods). Green line: the operating point of V1 cells, set as the point in which cells' visual response is 0% $\Delta F/F_0$ (green line in **a**). Black line: fit of the Ricciardi transfer function to the estimated IO function (see Methods). **c**, Schematic of the underlying IO function of the average V1 pyramidal cell showing a supralinear regime for low inputs and a linear regime before response saturation for high inputs. Cells operate between the supralinear and linear regimes, allowing cells suppressed below their spontaneous activity levels to attenuate inputs.

In sum, these results confirm that individual neurons show IO functions that are consistent with the population data in Fig. 2. These IO functions are nonlinear below resting activity levels, showing attenuation-by-suppression, and linear for moderate increases in firing rate — across much of the range evoked by visual input.

Cells operate at the transition point of a supralinear-to-linear IO function at spontaneous activity levels

We next sought to estimate the IO function across the range of activity levels we measured. To do this, we used changes in optogenetic responses (Fig. 2k) to infer the slope of the underlying function. We integrated the curve fitted across all cells' change in responses (Fig. 2o) to generate the underlying IO function (Fig. 4a-b). Consistent with our individual observations, this IO function has a supralinear region followed by a linear region, before eventual saturation at high inputs (Fig. 4b).

One constraint with our approach is that we obtain the IO function by integrating, operating on changes in optogenetic response as a function of changes in ongoing neural activity. This method characterizes changes in response, but does not rigorously define a firing rate of zero (x-axis, Fig. 4b) because the smallest amount of spiking elicited by visual input must exceed zero. While the zero firing rate point is not critical to IO function shape, we also estimated neurons' firing rate in another way, based on their calcium transient rate, by deconvolving each neuron's response. This analysis revealed that there are indeed cells whose activity reaches zero, approximated by a zero rate of calcium transients, during visual stimulation (Extended Data Fig. 6). We used this value to represent zero firing rates in Fig. 4b.

A related point is that our ability to measure the shape of the IO function may be limited as activity levels approach zero. Our optogenetic input provides increments in neural activity, which

move neurons' firing rates away from zero. To ensure that the shape of the IO function we found is valid even at low activity levels, we characterized the impact of interpolated extensions to the fitted line on the resulting IO function. We found no qualitative differences (Extended Data Fig. 6) from our prior estimates. Finally, our results remained consistent whether we used estimated firing rates or measured fluorescence directly. Deconvolved responses yield the same shape of the underlying IO function as when using calcium responses (Extended Data Fig. 7).

Our results are summarized in Fig. 4c. Neurons operate during spontaneous activity near the high end of a supralinearity in their IO function. Suppression then moves neurons into this supralinear region, so that the more a neuron is suppressed, the smaller its response to a fixed input. Thus, the nonlinearity in the IO function, and the operating point of the neurons during spontaneous activity, together produce attenuation-by-suppression.

Discussion

Here, we demonstrate through direct stimulation that excitatory cells in the cortex of awake mice operate along a supralinear-to-linear IO function. Cells' f-I curves, or IO functions, have previously been suggested to exhibit smooth nonlinearities, such as power-law nonlinearities, across the working range of visual responses (Anderson et al., 2000; Carandini, 2004; Finn et al., 2007), and it was thought these nonlinearities would serve to amplify inputs (Rubin et al., 2015; Ahmadian et al., 2013). However, we find a key feature of V1 excitatory cells is not that they exhibit amplification above baseline (spontaneous) activity levels, but rather, that they show linear responses. We do find that neurons' IO functions show a smooth nonlinearity — a supralinearity — but this effect occurs below the level of spontaneous firing on average. Thus, we conclude that the prominent nonlinearity in neurons' IO functions does not allow amplification, but instead causes neurons to attenuate their responses to input as they are suppressed.

When input was strong, we observed some decreased responses. This trend may be caused by mechanisms such as sodium channel inactivation in single cells, strong inhibitory coupling in the network, and/or consequences of the refractory period (Rauch et al., 2003; Sanzeni et al., 2020). However, we cannot rule out that decreased responses were due to saturation of fluorescent signals at high activity despite the relatively linear behavior of the indicator we used (Zhang et al., 2023).

The IO function we measure is strikingly similar to the Ricciardi IO function that is derived from integrate-and-fire model networks (Ricciardi and Smith, 1977; Sanzeni et al., 2020). Based on experimental data from different systems, prior work has put forth proposals for a variety of IO function shapes, including power laws and threshold-linear functions. Our inferred IO function shape (Fig. 4b; Extended Data Fig. 8), like the Ricciardi function, has a supralinear region (which we show here is below resting or spontaneous activity), a linear region that covers most of neurons' operating range as defined by high visual stimulus contrast (100% contrast, Fig. 2c-o, Fig. 3a-d), and a weakly sublinear region for the highest activity levels.

Two major consequences for brain function arise from our results. First, the attenuation-by-suppression we observe can have direct consequences on cortical computation — in other words, the brain can filter out inputs to neurons in certain contexts, by suppressing the activity of those neurons so their incremental responses are attenuated. Second, and more speculatively, because the supralinear-to-linear function we characterized shares general features with activation functions used in modern machine learning systems, including generative AI systems like ChatGPT (Brown et al., 2020), our results may have implications for learning.

We first discuss the direct computational consequences of attenuation-by-suppression. In several visual sensory contexts neurons are suppressed; a classic example is the case of sensory surround suppression. In surround suppression (Hubel and Wiesel, 1965; Allman et al., 1985; Cavanaugh et al., 2002), sensory stimuli outside a neurons' key feature tuning (outside its receptive field) decrease that neuron's firing. Surround suppression, combined with the attenuation-by-suppression effect we observe, could lead to sharpening of responses to small visual stimuli or stimuli that change quickly in space, like edges. In this scenario, as some cells are activated by the stimulus, those activating inputs, falling on neurons' linear region, would be passed on to downstream areas. In contrast, other nearby neurons will have receptive fields that are in the surround — the small stimulus would fall outside their receptive field, causing them to be suppressed (Ozeki et al., 2009; Haider et al., 2010, 2013; Angelucci et al., 2017; Lakunina et al., 2020). These suppressed neurons would then attenuate any additional inputs that arrive, filtering out those inputs, so that surround inputs have less impact on downstream neurons. This effect could, at the population level, serve to amplify relevant responses above background noise (Ringach and Malone, 2007).

Notably, this mechanism is not just the usual conception of suppressive surrounds. Here we show not just that visual stimuli suppress some neurons' activity, but that suppressed neurons are actively filtering out the inputs they receive — that is, attenuating added inputs, in attenuation-by-suppression. We arrive at this result by going beyond what is possible with visual stimuli alone, using an *in vivo* cell-specific stimulation approach that allows us to directly measure IO functions and changes in response to fixed input as a function of activity level.

Attenuation-by-suppression may also offer an explanation for perceptual masking effects arising from optogenetic stimulation (Chen et al., 2022), as optogenetic stimulation can lead to suppression caused by recurrent excitatory connections (Dalgleish et al., 2020; O'Rawe et al., 2023). These effects, where optogenetic stimulation generates suppression and thus attenuation, may also underlie the differences in visual perception noted when visual and optogenetic inputs are combined (Lafer-Sousa et al., 2022; Azadi et al., 2023).

Another major result of this work is that our measured IO functions share important features with the activation functions used in recent AI systems. Over a decade ago, activation functions with a linear regime over much of the unit's activity range (rectified-linear units, ReLUs) were shown to yield some learning improvements compared to sigmoidal functions. Since then, it has been observed that learning performance can be additionally improved by activation functions that smooth the ReLU's sharp transition from zero activation to the linear region (Clevert et al., 2015;

Hendrycks and Gimpel, 2016). Intuitively, this smoothing avoids a zero gradient below the rectification threshold and limits the degree to which synapses drop out of learning due to a lack of available error signal. One might map this to biology by saying that neurons that “cannot fire never wire.” Thus, while learning rules *in vivo* are not fully understood, if there were a sharp IO function threshold *in vivo*, it might be that learning would be difficult. This smooth nonlinearity followed by a linear region is a shape shared by many activation functions in modern AI systems, including the softplus, Gaussian error linear (GELU), and exponential linear (ELU) functions. Our data show striking similarities to these activation functions, with a smooth transition around the nonlinearity and a largely linear regime through most of the range used for even high-contrast visual responses. Since these activation functions work well in model networks, we might speculate that biological neurons have faced selection pressure to produce IO functions with this shape, to optimize some form of learning or other computation.

In summary, our findings characterize the average IO function of neurons in the cortex during normal brain operation — in the awake state, as animals actively receive sensory input. The attenuation-by-suppression we observe suggests a new role for ongoing activity in cortical networks. Because ongoing input creates the smooth nonlinearity below spontaneous firing rates, it can perform computations by allowing neurons to selectively filter out extraneous features of sensory input.

Methods

Viral injections and cranial window implants

All procedures were approved by the NIMH Institutional Animal Care and Use Committee (IACUC) and conform to relevant regulatory standards. Emx1-Cre mice (The Jackson Laboratory) were used in all experiments to target expression of Cre to excitatory neurons. Adult mice were anesthetized (1–3% in 100% O2 at 1 L/min) and administered intraperitoneal dexamethasone (3.2 mg/kg). A custom metal head post was fixated at the base of the skull. A 3 mm diameter circular craniotomy was made over the left hemisphere of primary visual cortex (ML -3.1 mm, AP +1.5 mm relative to Lambda).

A bicistronic virus expressing both opsin and GCaMP indicator, AAV9-hSyn-DIO-jGCaMP8s-P2A-stChrimsonR (LaFosse et al., 2023) (Addgene, 174007) was diluted in phosphate-buffered saline (final titer: $2.6x10^{12}$ GC/mL, N = 3 mice; $2.9x10^{12}$ GC/mL, N = 1 mouse; $3.4x10^{12}$ GC/mL, N = 2 mice; $4.7x10^{12}$ GC/mL, N = 2 mice) and injected 200 µm below the surface of the brain (multiple injections, 300 nL/injection, 0.1 µL/min) to target expression to layer 2/3 neurons. In N = 2 mice, separate viruses of AAV9-hSyn-jGCaMP8s-WPRE ($5.6x10^{12}$ GC/mL or $1.0x10^{13}$ GC/mL) and AAV9-hSyn-DIO-stChrimsonR-mRuby2 ($2.7x10^{12}$ GC/mL) were instead injected.

A 3 mm optical window (Tower Optical) was implanted over the craniotomy. Both the optical window and metal head post were fixed to the skull using C&B metabond (Parkell) cement dyed black. Animals were individually housed after surgery. Mice were imaged three or more weeks post-injection, and water-scheduled (1.0 mL/day) to promote alertness during experiments that used intermittent water delivery.

Two-photon calcium imaging

Two-photon calcium imaging was performed with a custom-built microscope (MIMMS, Modular In vivo Multiphoton Microscopy System, components, Sutter Instruments; controlled by ScanImage, MBF Biosciences; bidirectional 8 kHz resonant scanning, 512 lines, ~30 Hz frame rate). Calcium responses were imaged with 920 nm laser pulses (Chameleon Discovery NX laser, Coherent, Inc., pulse rate 80 MHz, pulse energy 0.19–0.25 nJ/pulse; imaging power 15–20 mW measured at the front aperture of the objective; 16x objective, 0.8 NA, Nikon Inc; immersion with clear ultrasound gel, ~1 mL; 100-200 µm below pia, L2/3 of V1, field of view 400 x 400 µm).

Cell-specific two-photon stimulation

Cell-specific photostimulation was performed using either 1030 nm (Satsuma laser, Amplitude Laser; Fig. 2, Fig. 3f-i) or 1035 nm light (Monaco laser, Coherent, Inc.; Fig. 3b-d). The laser wavefront was shaped by a spatial light modulator (1920 x 1152 pixels; Meadowlark) to target patterns of cells for stimulation (10 µm diameter disks; 300 ms, 8 mW/target, 6.5 mW/target for low power in Fig. 3e-i, 500 kHz pulse rate; TTL logic gate, Pulse Research Lab, used to invert imaging line clock to gate laser during pixel acquisition and stimulate during reversal of bidirectional scanning: on time 19 µs, off time 44 µs, 30% duty cycle; stimulation power

accounting for blanking: 2.4 mW/target, or 2 mW/target for low power; see LaFosse et al., 2023 for more details).

Visual stimulation in awake mice

Visual responses were measured in awake, head-fixed mice viewing sinusoidal drifting grating stimuli (full-field or Gabor patches filtered with a 15° full-width half-max Gaussian mask; spatial frequency of 0.1 cycles/degree; 100% contrast; 3-4 seconds; stimuli presented on an LCD monitor with neutral gray background). For experiments using orthogonal visual stimuli to switch between suppressed or elevated responses in the same neurons (Fig. 3a-d), to choose the optimal orientations we measured responses to drifting gratings at multiple orientations (8 orientations 45° apart, 2 seconds on, 6 seconds off). To ensure mice were awake and alert during imaging experiments, small volumes of water were delivered on 20% of trials (3μL, ~300μL total).

To align Gabor patches to the receptive fields of neurons in the imaging FOV, we performed retinotopy via hemodynamic intrinsic imaging using 530 nm light (delivered via fiber-coupled LED, Thorlabs; 2 Hz imaging) (Goldbach et al., 2021). Retinotopic maps were calculated by measuring hemodynamic-related changes in light absorption across the cortex to square wave drifting gratings (0.1 cycles/degree, 10° diameter, 6 locations, 5 seconds; 5 second baseline window and 2.5 second response window starting 3 seconds after stimulus onset used to measure responses).

Analysis of two-photon imaging data

Images acquired via two-photon imaging were first rescaled from 512 x 512 pixels to 256 x 256 to ease data processing. Background correction was done by subtracting the minimum pixel value of the average intensity image and setting all remaining negative pixels (due to noise) to zero. We used the CaImAn toolbox (Giovannucci et al., 2019) for motion correction and Suite2p (Pachitariu et al., 2017) for cell segmentation to allow manual selection of cell masks.

We measured neurons' activity to both visual and optogenetic stimuli ($\Delta F/F_0$; F: average intensity across all pixels within a cell's segmented mask; F_0: average fluorescence across the 45 imaging timepoints, 1.5 s, prior to stimulus presentation). "Opto" trials or "Vis" trials are where the corresponding stimulus was delivered by itself. "Opto+Vis" trials, (O+V), are when visual stimulus and optogenetic stimulus are delivered on the same trial. The "Opto during vis" response, (O+V) - V, (Fig. 2m) was found by subtracting the average Vis trace from the average Opto+Vis trace. Deconvolved activity was estimated using the OASIS method (Friedrich et al., 2017). To compare responses across experiments and animals, deconvolved estimates were normalized by setting the mean activity equal to the variance under Poisson-like assumptions.

For visual display of responses across the FOV (Fig. 2d,f-k; Fig. 3b), we measured trial-averaged activity at every pixel ($\Delta F/F_0$; F: average frame across a time window, Fig. 2d, entire visual stimulus period, Fig. 2f,i and Fig. 3b, entire visual stimulus period except the first 1 second, so as to exclude the initial transient response at stimulus onset, Fig. 2g,j, 1 second

period following opto stimulus onset; F_0: average frame across the preceding 1.5 seconds to the onset of the respective stimulus).

Analysis of stimulus-evoked responses

Cell responses (points in Fig. 2n-o; Fig. 3c-d, g-i) were computed by averaging timecourses in 0.33 second windows (10 frames), just before (baseline) or after (response) opto stim onset. “Vis resp” (x-axis, Fig. 2o, 3d, 4a): response on Opto+Vis trials, in window just before opto stim onset (when only the visual stimulus was present). “Opto” (Fig. 2n): response on Opto only trials, in response period after opto stim onset, minus baseline period just before opto stim onset. “Opto during vis”, “(O+V) - V” (Fig. 2n): response period on Opto+Vis trials minus response period on Vis only trials, with baseline period difference on the same trials subtracted. Change in opto response between visual and spontaneous conditions (y-axis, Fig. 2o, 3d, 4a): difference between “Opto during vis” and the “Opto” responses.

Stimulated cells during our experiments were defined as those whose centroid fell within a 15 micron radius of any stimulation target’s center (5-12 targets per pattern, average 9 ± 0.5) and whose “Opto” response value exceeded a 5% $\Delta F/F_0$ threshold (N = 15 ± 1.3 stimulated cells per pattern). Stimulation of small groups of cells typically induces smaller network effects, such as weak, broad suppression (Jouhanneau et al., 2018; Chettih and Harvey, 2019; Oldenburg et al., 2022; LaFosse et al., 2023), though these small influences are unlikely to significantly impact the much larger evoked responses we measure in targeted cells.

IO function estimation

To estimate the underlying IO function of V1 pyramidal cells, output = f(input), we first performed locally estimated scatterplot smoothing, or LOESS (Cleveland, 1979). We fit LOESS curves (rotated axis to maximize variance on x-axis, using the loess Python package (Cappellari et al., 2013): pypi.org/project/loess/) to the data (Fig. 2o, Fig. 4a; confidence intervals via bootstrap; N = 1000 repetitions, N = 364 samples).

The resulting fit represents a change in output versus the input level of a cell, or the derivative of the IO function, shifted vertically by a constant:

$$f(\text{input}) = \int [f'(\text{input}) + y_{shift}].$$

The value of y_{shift} is dependent on two factors (Extended Data Fig. 6a): any difference between maximum visual suppression and zero firing rate ($y_{supp,0}$), and the size of the optogenetic response (y_{opto}): $y_{shift} = y_{supp,0} + y_{opto}$. Deconvolved spiking estimates of calcium activity (Extended Data Fig. 6b-d) suggest the most-suppressed cells did indeed reach zero response, so we estimate the shift as the minimum of the fit line: $y_{supp,0} = -min(f'(\text{input}))$. Assuming suppression to zero completely attenuates small optogenetic inputs (as expected for an activation function that goes to zero), we set y_{opto} to zero. While we cannot directly measure this value because any optogenetic activation moves responses away from zero, we confirm that our conclusions are not changed by various extrapolations from our observed data points (Extended Data Fig. 6e-g).

To fit the Ricciardi transfer function to our estimated IO function (Fig. 4b), we used least squares regression. Initial values for the parameters taken from Sanzeni et al. (2020): τ_i = 20 ms, τ_{rp}= 2 ms, V_r = 10 mV, θ = 20 mV, σ = 10 mV. The resulting fitted parameters and their standard errors are: τ_i = 16.8 ± 0.5 ms, τ_{rp}= 3.1 ± 0.02 ms, V_r = 10.9 ± 0.02 mV, θ = 20.0 ± 0.001 mV, σ = 10.0 ± 0.01 mV.

Author contributions

PKL performed all experiments and data analysis. ZZ, VS performed viral injections and surgical implants. YD, MH developed two-photon microscope and two-photon holographic stimulation equipment. PKL, MH wrote the manuscript.

Acknowledgements

We thank Georg Jaindl and others from Vidrio/MBF Biosciences for expert technical assistance. This work was supported by the NIH BRAIN Initiative, grant U19NS107464, and the NIMH Intramural Research Program, ZIAMH0020967.

References

Agüera y Arcas B, Fairhall AL (2003) What causes a neuron to spike? Neural Comput 15:1789–1807.

Ahmed B, Anderson JC, Douglas RJ, Martin KA, Whitteridge D (1998) Estimates of the net excitatory currents evoked by visual stimulation of identified neurons in cat visual cortex. Cereb Cortex 8:462–476.

Allman J, Miezin F, McGuinness E (1985) Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. Annu Rev Neurosci 8:407–430.

Anderson JS, Lampl I, Gillespie DC, Ferster D (2000) The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. Science 290:1968–1972.

Angelucci A, Bijanzadeh M, Nurminen L, Federer F, Merlin S, Bressloff PC (2017) Circuits and Mechanisms for Surround Modulation in Visual Cortex. Annu Rev Neurosci 40:425–451.

Anzai A, Ohzawa I, Freeman RD (1999) Neural mechanisms for processing binocular information I. Simple cells. J Neurophysiol 82:891–908.

Atallah BV, Bruns W, Carandini M, Scanziani M (2012) Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. Neuron 73:159–170.

Azadi R, Bohn S, Lopez E, Lafer-Sousa R, Wang K, Eldridge MAG, Afraz A (2023) Image-dependence of the detectability of optogenetic stimulation in macaque inferotemporal cortex. Curr Biol 33:581–588.

Bar-Yehuda D, Korngreen A (2008) Space-clamp problems when voltage clamping neurons expressing voltage-gated conductances. J Neurophysiol 99:1127–1136.

Brown T, Mann B, Ryder N (2020) Language models are few-shot learners. Adv Neural Inf Process Syst.

Cappellari M, McDermid RM, Alatalo K, Blitz L, Bois M, Bournaud F, Bureau M, Crocker AF, Davies RL, Davis TA, Others (2013) The ATLAS3D project--XX. Mass--size and mass--σ distributions of early-type galaxies: bulge fraction drives kinematics, mass-to-light ratio, molecular gas fraction and stellar initial mass function. Mon Not R Astron Soc 432:1862–1893.

Carandini M (2004) Amplification of trial-to-trial response variability by neurons in visual cortex. PLoS Biol 2:E264.

Carandini M, Mechler F, Leonard CS, Movshon JA (1996) Spike train encoding by regular-spiking cells of the visual cortex. J Neurophysiol 76:3425–3441.

Cardin JA, Palmer LA, Contreras D (2008) Cellular mechanisms underlying stimulus-dependent gain modulation in primary visual cortex neurons in vivo. Neuron 59:150–160.

Cavanaugh JR, Bair W, Movshon JA (2002) Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. J Neurophysiol 88:2530–2546.

Chance FS, Abbott LF, Reyes AD (2002) Gain modulation from background synaptic input. Neuron 35:773–782.

Chen SC-Y, Benvenuti G, Chen Y, Kumar S, Ramakrishnan C, Deisseroth K, Geisler WS, Seidemann E (2022) Similar neural and perceptual masking effects of low-power optogenetic stimulation in primate V1. Elife 11.

Chettih SN, Harvey CD (2019) Single-neuron perturbations reveal feature-specific competition in V1. Nature 567:334–340.

Cleveland WS (1979) Robust Locally Weighted Regression and Smoothing Scatterplots. J Am Stat Assoc 74:829–836.

Clevert D-A, Unterthiner T, Hochreiter S (2015) Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs). arXiv [csLG] Available at: http://arxiv.org/abs/1511.07289.

Dalgleish HW, Russell LE, Packer AM, Roth A, Gauld OM, Greenstreet F, Thompson EJ, Häusser M (2020) How many neurons are sufficient for perception of cortical activity? Elife 9.

Dubois JM (2000) What is the true resting potential of small cells? Gen Physiol Biophys 19:3–7.

Finn IM, Priebe NJ, Ferster D (2007) The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. Neuron 54:137–152.

Fourcaud-Trocmé N, Hansel D, van Vreeswijk C, Brunel N (2003) How spike generation mechanisms determine the neuronal response to fluctuating inputs. J Neurosci 23:11628–11640.

Friedrich J, Zhou P, Paninski L (2017) Fast online deconvolution of calcium imaging data. PLoS Comput Biol 13:e1005423.

Giovannucci A, Friedrich J, Gunn P, Kalfon J, Brown BL, Koay SA, Taxidis J, Najafi F, Gauthier JL, Zhou P, Khakh BS, Tank DW, Chklovskii DB, Pnevmatikakis EA (2019) CaImAn an open source tool for scalable calcium imaging data analysis. Elife 8.

Glorot X, Bordes A, Bengio Y (2011) Deep Sparse Rectifier Neural Networks. In: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (Gordon G, Dunson D, Dudík M, eds), pp 315–323 Proceedings of Machine Learning Research. Fort Lauderdale, FL, USA: PMLR.

Goldbach HC, Akitake B, Leedy CE, Histed MH (2021) Performance in even a simple perceptual task depends on mouse secondary visual areas. Elife 10.

Granit R, Kernell D, Shortess GK (1963) Quantitative aspects of repetitive firing of mammalian motoneurones, caused by injected currents. J Physiol 168:911–931.

Haider B, Häusser M, Carandini M (2013) Inhibition dominates sensory responses in the awake cortex. Nature 493:97–100.

Haider B, Krause MR, Duque A, Yu Y, Touryan J, Mazer JA, McCormick DA (2010) Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. Neuron 65:107–121.

Hansel D, van Vreeswijk C (2002) How noise contributes to contrast invariance of orientation tuning in cat visual cortex. J Neurosci 22:5118–5128.

Hendrycks D, Gimpel K (2016) Gaussian Error Linear Units (GELUs). arXiv [csLG].

Hubel DH, Wiesel TN (1965) Receptive fields and functional architecture in two nonstriate visual areas (18 and 19) of the cat. J Neurophysiol 28:229–289.

Jagadeesh B, Gray CM, Ferster D (1992) Visually evoked oscillations of membrane potential in cells of cat visual cortex. Science 257:552–554.

Jouhanneau J-S, Kremkow J, Poulet JFA (2018) Single synaptic inputs drive high-precision action potentials in parvalbumin expressing GABA-ergic cortical neurons in vivo. Nat Commun 9:1540.

Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. Adv Neural Inf Process Syst 25.

Lafer-Sousa R, Wang K, Afraz A (2022) Detectability of optogenetic stimulation of inferior temporal cortex depends significantly on visibility of visual input. J Vis 22:3768–3768.

LaFosse PK, Zhou Z, Friedman NG, Deng Y, Li AJ, Akitake B, Histed MH (2023) Bicistronic Expression of a High-Performance Calcium Indicator and Opsin for All-Optical Stimulation and Imaging at Cellular Resolution. eNeuro 10.

Lakunina AA, Nardoci MB, Ahmadian Y, Jaramillo S (2020) Somatostatin-Expressing Interneurons in the Auditory Cortex Mediate Sustained Suppression by Spectral Surround. J Neurosci 40:3564–3575.

Marshel JH, Kim YS, Machado TA, Quirin S, Benson B, Kadmon J, Raja C, Chibukhchyan A, Ramakrishnan C, Inoue M, Shane JC, McKnight DJ, Yoshizawa S, Kato HE, Ganguli S, Deisseroth K (2019) Cortical layer–specific critical dynamics triggering perception. Science 365:eaaw5202.

McCormick DA, Connors BW, Lighthall JW, Prince DA (1985) Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. J Neurophysiol 54:782–806.

Miller KD, Troyer TW (2002) Neural noise can explain expansive, power-law nonlinearities in neural response functions. J Neurophysiol 87:653–659.

Moskovitz TH, Roy NA, Pillow JW (2018) A comparison of deep learning and linear-nonlinear cascade approaches to neural encoding. bioRxiv:463422.

Nair V, Hinton GE (2010) Rectified linear units improve restricted Boltzmann machines.

Oldenburg IA, Hendricks WD, Handy G, Shamardani K, Bounds HA, Doiron B, Adesnik H (2022) The logic of recurrent circuits in the primary visual cortex. bioRxiv:2022.09.20.508739.

O'Rawe JF, Zhou Z, Li AJ, LaFosse PK, Goldbach HC, Histed MH (2023) Excitation creates a distributed pattern of cortical suppression due to varied recurrent input. Neuron *in press*.

Ostojic S, Brunel N (2011) From spiking neuron models to linear-nonlinear models. PLoS Comput Biol 7:e1001056.

Ozeki H, Finn IM, Schaffer ES, Miller KD, Ferster D (2009) Inhibitory stabilization of the cortical network underlies visual surround suppression. Neuron 62:578–592.

Pachitariu M, Stringer C, Harris KD (2017) Robustness of spike deconvolution for calcium imaging of neural spiking. bioRxiv:156786.

Petersen PC, Berg RW (2016) Lognormal firing rate distribution reveals prominent fluctuation–driven regime in spinal motor networks. Elife 5:e18805.

Powers RK, Binder MD (2001) Input-output functions of mammalian motoneurons. Rev Physiol Biochem Pharmacol 143:137–263.

Priebe NJ, Ferster D (2002) A new mechanism for neuronal gain control (or how the gain in brains has mainly been explained). Neuron 35:602–604.

Priebe NJ, Mechler F, Carandini M, Ferster D (2004) The contribution of spike threshold to the dichotomy of cortical simple and complex cells. Nat Neurosci 7:1113–1122.

Rauch A, La Camera G, Lüscher H-R, Senn W, Fusi S (2003) Neocortical Pyramidal Cells Respond as Integrate-and-Fire Neurons to In Vivo–Like Input Currents. J Neurophysiol 90:1598–1612.

Ricciardi LM, Smith CE (1977) Diffusion Processes and Related Topics in Biology. Springer Berlin Heidelberg.

Ringach DL, Malone BJ (2007) The operating point of the cortex: neurons as large deviation detectors. J Neurosci 27:7673–7683.

Sanzeni A, Histed MH, Brunel N (2020) Response nonlinearities in networks of spiking neurons. PLoS Comput Biol 16:e1008165.

Tan AYY, Chen Y, Scholl B, Seidemann E, Priebe NJ (2014) Sensory stimulation shifts visual cortex from synchronous to asynchronous states. Nature 509:226–229.

Wilson JR, Clark RB, Banderali U, Giles WR (2011) Measurement of the membrane potential in small cells using patch clamp methods. Channels 5:530–537.

Yamashita T, Pala A, Pedrido L, Kremer Y, Welker E, Petersen CCH (2013) Membrane potential dynamics of neocortical projection neurons driving target-specific signals. Neuron 80:1477–1490.

Zhang Y et al. (2023) Fast and sensitive GCaMP calcium indicators for imaging neural populations. Nature 615:884–891.

Zhao W-J, Kremkow J, Poulet JFA (2016) Translaminar Cortical Membrane Potential Synchrony in Behaving Mice. Cell Rep 15:2387–2399.

Extended Data Figures

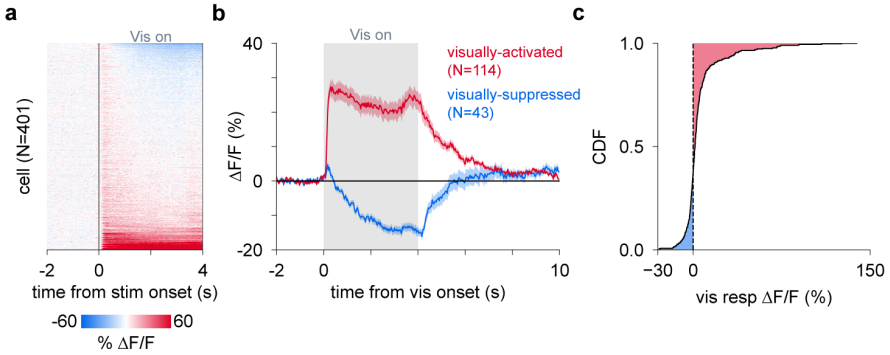

Extended Data Fig. 1
Drifting gratings evoke a pattern of visually-activated and visually-suppressed cells.
a, Example visual responses in one animal (see Fig. 2d) to a drifting grating stimulus. **b**, Cell-averaged (mean ± SEM) activity traces to the visual stimulus for visually-activated (red) and visually-suppressed cells (blue). Cells were determined as activated or suppressed by averaging the activity across the visual stimulation period (excluding the first ~330 ms to account for initial transient response). Activated: > 5% $\Delta F/F_0$, suppressed: < -5% $\Delta F/F_0$. **c**, Cumulative distribution function of visual responses across all cells.

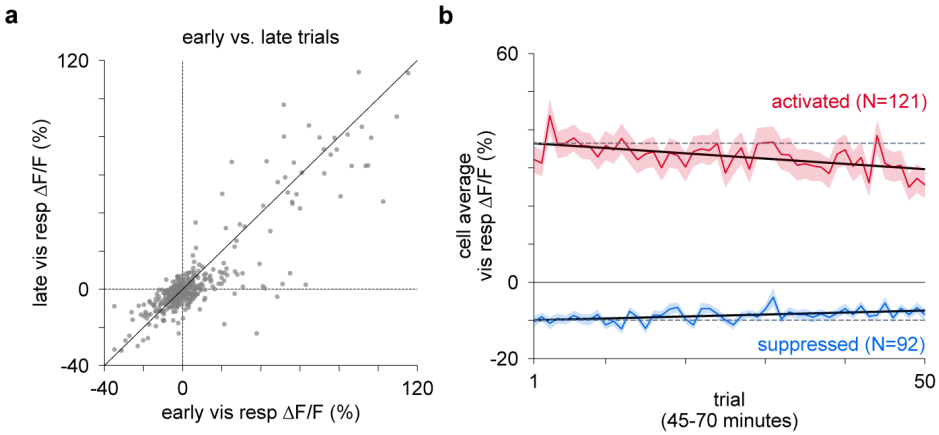

Extended Data Fig. 2

Minimal visual response adaptation across trials.

a, Visual responses (N = 364 cells; Fig. 2n) averaged across early (first 15 trials) and late (last 15 trials) periods. **b**, Cell-averaged (mean ± SEM) visual responses across trials showing minimal response adaptation for both visually-activated (red) and visually-suppressed (blue) cells. Black lines: linear fits. Grey dashed horizontal lines: value of the fitted response on the first trial.

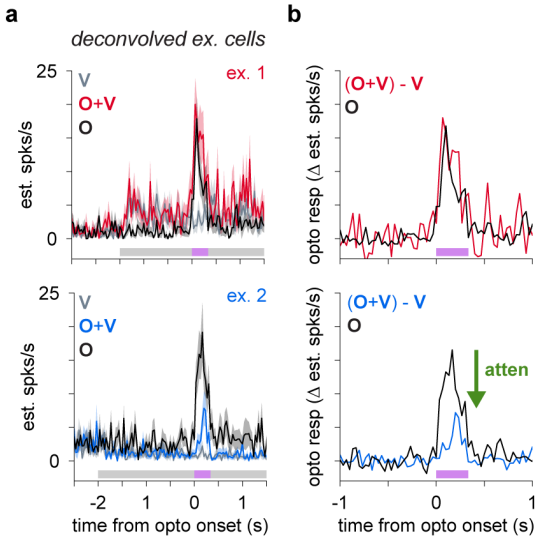

Extended Data Fig. 3

Deconvolved responses show attenuation during suppression.

a, Example deconvolved responses (mean ± SEM) in a visually-activated (top) and visually-suppressed (bottom) neuron (from Fig. 2f-m) across Visual only, Opto only, and Opto+Vis trials. **b**, Comparison of deconvolved opto responses during and without vision showing attenuation in a suppressed neuron.

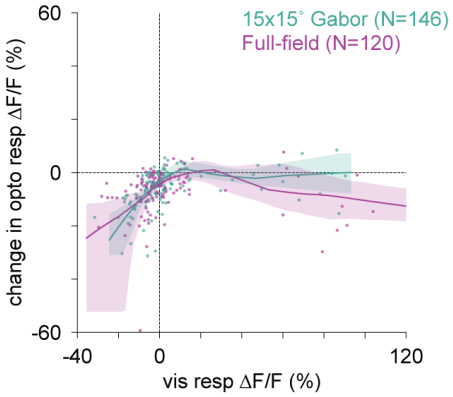

Extended Data Fig. 4

Changes in opto response are not impacted by size of visual stimulus.

Change in opto response as a function of visual response across cells measured with both full-field drifting gratings (purple) and 15 x 15° Gabor drifting gratings (teal). Lines: LOESS fits with 95% confidence intervals.

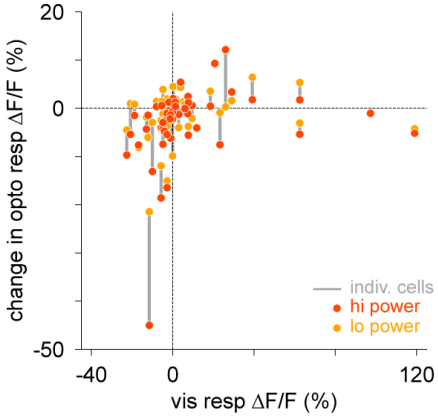

Extended Data Fig. 5

Changes in opto responses during vision are similar for different opto intensities.

Change in opto response as a function of average visual response measured using two opto powers. Gray lines: individual cells.

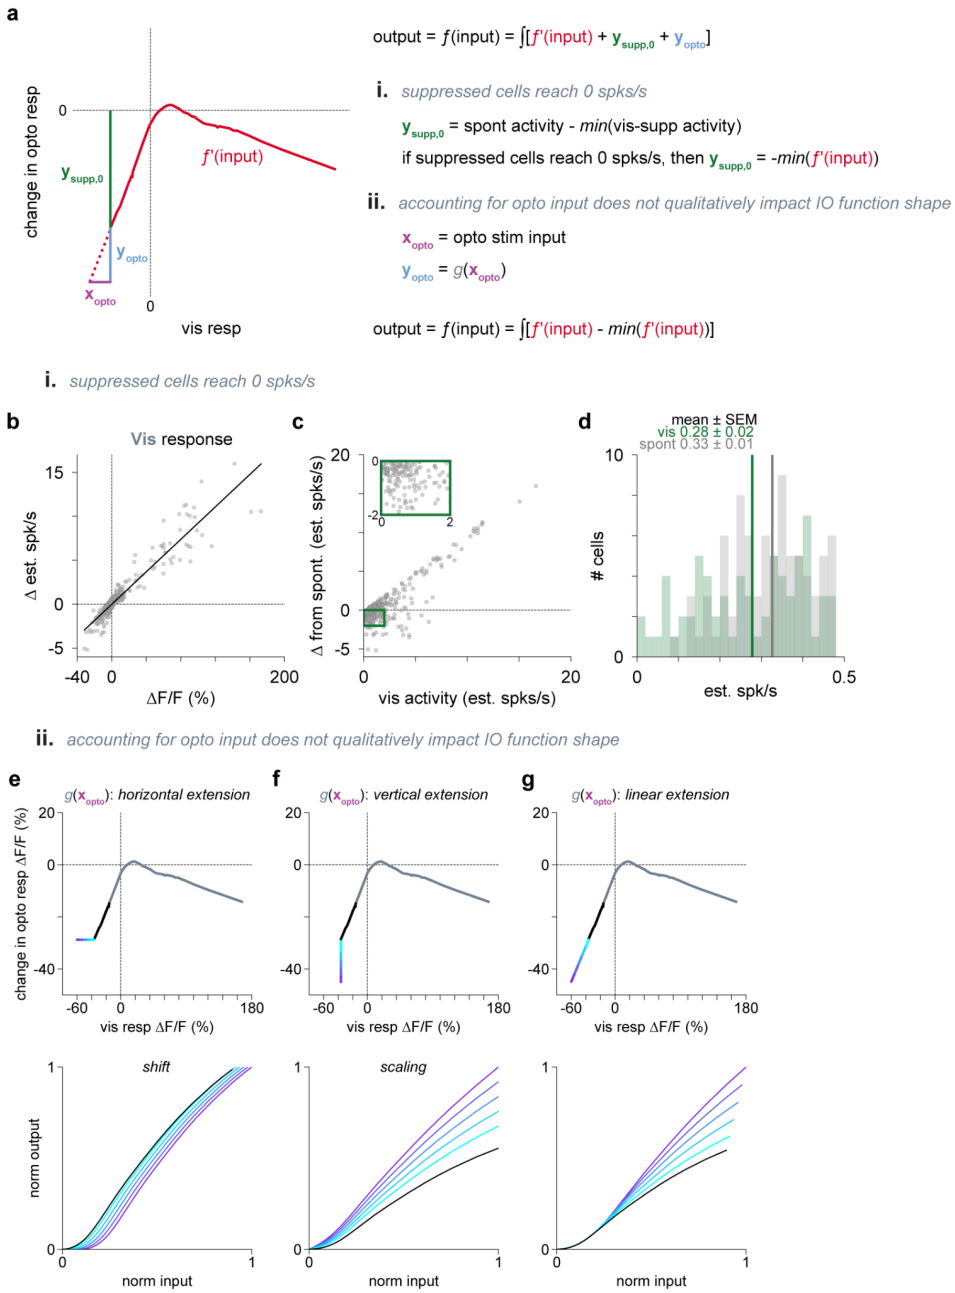

Extended Data Fig. 6

Estimating the IO function of V1 pyramidal cells.

a, Schematic showing the relationship between change in opto response (change in output) as a function of visual response (input). This relationship is the derivative of the underlying IO function, but shifted down since measures of change in opto response are taken relative to spontaneous activity levels (see Methods for details). The degree of shift is dependent on (**i**) the distance visually-suppressed cells are from zero activity during vision and (**ii**) the distance opto input pushes the response away from zero activity. **b**, Deconvolved spiking estimates of visual responses match corresponding $\Delta F/F_0$ values. **c-d**, Deconvolved spiking estimates of visually-evoked activity. **c**, Visually-suppressed cells reach 0 spks/s. Inset shows zoomed plot near zero visual activity. **d**, Histogram showing deconvolved cell activity reaches zero during vision but not during spontaneous activity, implying cells are suppressed to zero activity. **e-g**, We are unable to account for the impact of the opto input (x_{opto}) into cells as the underlying shape of this influence is some unknown function dependent on this input: $y_{opto} = g(x_{opto})$. At its extremes, this function (g) is either vertical (assuming zero opto input) or is horizontal (assuming the most-suppressed cells are well below spiking threshold), with following the trajectory of the LOESS fit line as an intermediate option. We tested

the impact of various extensions on our data to assess the qualitative impact on the estimated IO function. **e**, horizontal extensions (out to -60% $\Delta F/F_0$) shift the IO function. **f**, Vertical extensions (out to -45% $\Delta F/F_0$) scale the IO function. **g**, Linear extensions of the LOESS fit (estimated using least squares regression on the region where $x < -15\%$ $\Delta F/F_0$) do not impact nonlinearities at low input but scale saturating nonlinearities.

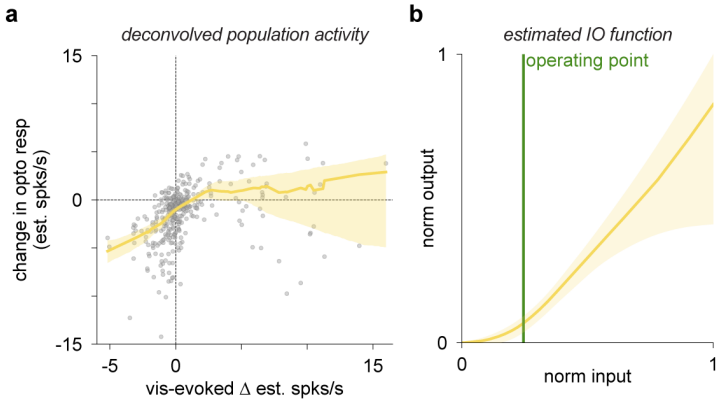

Extended Data Fig. 7

Deconvolved responses exhibit a supralinear-to-linear IO function.

a, Population (N = 364 cells) measure of change in deconvolved estimated spiking responses to opto stimulation across the working range of visual responses. Yellow line: LOESS fit with 95% confidence interval, axis not rotated prior to fit. **b**, Estimated IO function using deconvolved data. Dark green line: operating point determined as zero visual response from **a**.

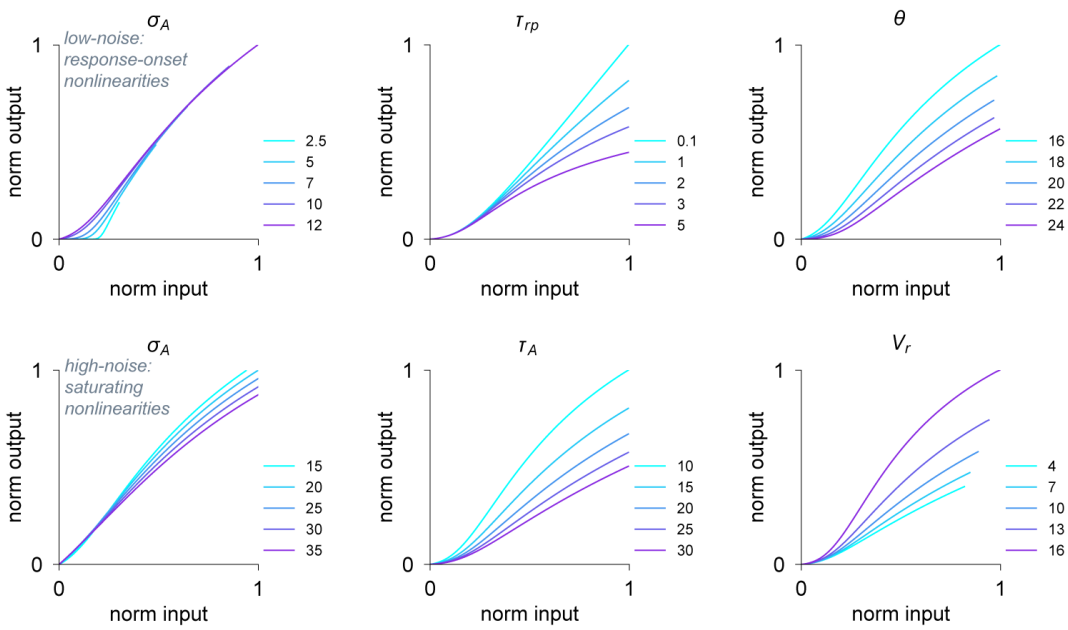

Extended Data Fig. 8

Ricciardi transfer function shape is qualitatively preserved across parameter space.

Ricciardi transfer function solution for ranges of different parameters. Changing parameter values does not qualitatively change the shape of the transfer function solution. Parameters: σ_A = 10, τ_{rp} = 2, τ_A = 20, θ = 20, V_r = 10, where not stated.