Behavioral/Cognitive

Mice Preferentially Use Increases in Cerebral Cortex Spiking to Detect Changes in Visual Stimuli

Jackson J. Cone, Morgan L. Bade, Nicolas Y. Masse, Elizabeth A. Page, David J. Freedman, and John H.R. Maunsell

Department of Neurobiology and Grossman Institute for Neuroscience, Quantitative Biology and Human Behavior, University of Chicago, Chicago, Illinois 60637

Whenever the retinal image changes, some neurons in visual cortex increase their rate of firing whereas others decrease their rate of firing. Linking specific sets of neuronal responses with perception and behavior is essential for understanding mechanisms of neural circuit computation. We trained mice of both sexes to perform visual detection tasks and used optogenetic perturbations to increase or decrease neuronal spiking primary visual cortex (V1). Perceptual reports were always enhanced by increments in V1 spike counts and impaired by decrements, even when increments and decrements in spiking were generated in the same neuronal populations. Moreover, detecting changes in cortical activity depended on spike count integration rather than instantaneous changes in spiking. Recurrent neural networks trained in the task similarly relied on increments in neuronal activity when activity has costs. This work clarifies neuronal decoding strategies used by cerebral cortex to translate cortical spiking into percepts that can be used to guide behavior.

Key words: **interneurons; optogenetics; perception psychophysics; visual coding; visual cortex**

Significance Statement

Visual responses in the primary visual cortex (V1) are diverse, in that neurons can be either excited or inhibited by the onset of a visual stimulus. We selectively potentiated or suppressed V1 spiking in mice while they performed contrast change detection tasks. In other experiments, excitation or inhibition was delivered to V1 independent of visual stimuli. Mice readily detected increases in V1 spiking while equivalent reductions in V1 spiking suppressed the probability of detection, even when increases and decreases in V1 spiking were generated in the same neuronal populations. Our data raise the striking possibility that only increments in spiking are used to render information to structures downstream of V1.

Introduction

Elaborating how the brain detects changes in the environment from collections of neuronal responses is essential for understanding neural circuit computations. The onset of even a small, simple visual stimulus drives activity in tens to hundreds of thousands of neurons in sensory cerebral cortex (Grinvald et al., 1994), with some neurons increasing their rate of firing and others decreasing their rate of firing. Detection strategies based on spiking increments are well supported by studies in which subjects responded to artificial activation of cerebral cortex using electrical microstimulation (Histed et al., 2013). However, increments and decrements in neuronal spiking could both provide equivalent information and be given equal weight by downstream brain regions. Alternatively, the brain might have evolved to decode cortical sensory representations using particular response profiles preferentially. For example, the energetic cost of spikes (Attwell and Laughlin, 2001; Lennie, 2003) might favor neural encoding that requires fewer spikes. That constraint might have driven the development of the visual ON and OFF pathways, where increases and decreases in luminance are converted into parallel channels with low basal rates of firing that respond selectively to changes in only one direction (Schiller, 1992). ON and OFF pathways use fewer spikes to convey the same amount of information compared with an ON-only or OFF-only system (Gjorgjieva et al., 2014, 2019). Because cortical neurons similarly have low spontaneous rates of firing, decoding of cortical representations might monitor increases in firing rates preferentially.

In addition to the question of the weight given to different signs of spike rate changes, there is also the question of the relative weight given to rapid versus sustained changes in spike rate.

Received May 8, 2020; revised July 20, 2020; accepted Aug. 26, 2020.

Author contributions: J.J.C., M.L.B., N.Y.M., D.J.F., and J.H.R.M. designed research; J.J.C., M.L.B., N.Y.M., and E.A.P. performed research; J.J.C. and N.Y.M. analyzed data; J.J.C. wrote the first draft of the paper; J.J.C., M.L.B., N.Y.M., E.A.P., D.J.F., and J.H.R.M. edited the paper.

This work was supported by an Albert O. Beckman Postdoctoral Fellowship (J.J.C.), the Vannevar Bush Faculty Fellowship (D.J.F.), and National Institutes of Health U01-NS090576 and U19-NS107464 (J.H.R.M.). We thank Dr. Vytas Bindokas and the University of Chicago Integrated Light Microscopy Core Facility for assistance with confocal imaging; Dr. Supriya Ghosh for assistance with electrophysiological recordings; and Zaina Zayyad, Dr. Supriya Ghosh, and Dr. Matthew Kaufman for critical feedback on the manuscript.

The authors declare no competing financial interests.

Correspondence should be addressed to Jackson J. Cone at jackson.j.cone@gmail.com.

https://doi.org/10.1523/JNEUROSCI.1124-20.2020

Sensory responses often emphasize stimulus onsets and offsets. Spatially, the edges of stimuli are emphasized by mechanisms, such as lateral inhibition (Hartline et al., 1956). Temporally, the starts and ends of stimuli are emphasized through mechanisms, such as rapid adaptation (Müller et al., 1999; Kohn, 2007). Frequent changes in sensory input, such as those produced by normal head and eye movements, generate transient changes in spiking (Gallant et al., 1998; Martinez-Conde et al., 2000; Vinje and Gallant, 2000). Decoding of cortical representations might similarly be designed to be particularly sensitive to rapid changes in the rate of firing.

Despite the importance of the question, the relative weight assigned to rapid and slow increments and decrements in cortical spike rates remains poorly understood. Because approaches based on correlations between neuronal activity and behavioral reports cannot provide conclusive results about the weighting of different signals (Haefner et al., 2013), this issue requires direct perturbation of neuronal spiking in the cortex of behaving subjects. Optogenetic methods provide an approach for producing controlled increments or decrements in neuronal firing (Yizhar et al., 2011; Wiegert et al., 2017). Studies that have used optogenetics to produce synchronous inhibition of sensory cortex have found that it impairs perception across a range of modalities and stimuli (Glickfeld et al., 2013; Guo et al., 2014; Resulaj et al., 2018; Cone et al., 2019; Jin and Glickfeld, 2019). Because electrical excitation of neurons can produce percepts (see Histed et al., 2013), whereas optogenetic inhibition of cortical spiking suppresses perceptual reports, there may be an asymmetry in the ability of increments and decrements of cortical spiking to be used for guiding behavior. However, a direct comparison of the effects of spiking increments and decrements requires that spike rate changes of comparable magnitude be produced and subjects need to be explicitly encouraged to respond to both increments and decrements in signals.

It is critical to understand how downstream neurons readout sensory information from patterns of spiking in upstream populations. Here, we used optogenetic approaches to present increments and decrements in V1 spiking to mice trained to perform visual detection tasks. We report that perceptual detection depends on increments in V1 spiking. When spiking increments and decrements were matched for absolute magnitude in the same neuronal populations, decrements in spiking work against perceptual detection, even when they could provide a strong, behaviorally relevant signal. Furthermore, rapid increases in spiking appear to carry no special weight in producing percepts. Instead, perceptual reports appear to depend on the integrated number of spikes rather than instantaneous positive or negative changes in V1 spiking.

Materials and Methods

Mouse strains. All animal procedures were in compliance with the guidelines of the National Institutes of Health and were approved by the Institutional Animal Care and Use Committee at the University of Chicago. Mouse lines were obtained from The Jackson Laboratory. Data come from *parvalbumin-Cre* mice (PV, 11 mice, 6 female; JAX stock #017320 (Hippenmeyer et al., 2005), *somatostatin-Cre* mice (SST, 2 mice, both male; Jax stock #013044 (Taniguchi et al., 2011), and *Emx1-IRES-Cre* mice (Emx, 7 mice, 2 female, Jax stock #005628 (Gorski et al., 2002). Experimental animals were heterozygous for Cre recombinase in the cell type of interest (outbred by crossing homozygous Cre-expressing strains with WT BALB/c mice, Jax stock #000651). Mice were singly housed on a reverse light/dark cycle with *ad libitum* access to food. Mice were water scheduled throughout behavioral experiments, except for periods around surgeries. Mice used for electrophysiological recordings had *ad libitum* access to food and water.

Cranial window implant. Mice (3–5 months old) were implanted with a headpost and cranial window to give stable optical access for photostimulation during behavior (Goldey et al., 2014; Histed and Maunsell, 2014). Animals were anesthetized with ketamine (40 mg/kg, i.p.), xylazine (2 mg/kg, i.p.), and isoflurane (1.2%-2% in 100% O_2). Using aseptic technique, a headpost was secured to the skull using acrylic (C&B Metabond, Parkell) and a 3 mm craniotomy was made over the left cerebral hemisphere (3.0 mm lateral and 0.5 mm anterior to λ) to implant a glass window (0.8 mm thickness; Tower Optical).

Intrinsic autofluorescence imaging. We located V1 by measuring changes in the intrinsic autofluorescence signal using visual stimuli and epifluorescence imaging (Andermann et al., 2011). Autofluorescence produced by blue excitation (470 ± 40 nm, Chroma) was collected using a green long-pass filter (500 nm cutoff) and a 1.0× air objective (Carl Zeiss; StereoDiscovery V8 microscope; ~0.11 NA). Fluorescence was captured with a CCD camera (AxioCam MRm, Carl Zeiss; 460 × 344 pixels; 4 × 3 mm FOV). The visual stimuli were full contrast drifting Gabors (10° SD; 30°/s; 0.1 cycles/deg) presented for 10 s followed by 6 s of mean luminance. The response to the visual stimulus was computed as the fractional change in fluorescence during the first 8 s of the stimulus presentation compared with the average of the last 4 s of the preceding blank.

Viral injections and ChR2 stimulation. Virus injections were targeted to a monocular region of V1 based on each animal's retinotopic map (+25° in azimuth; between −15° and 15° in elevation). Before virus injection, mice were anesthetized (isoflurane, 1%–1.5%), and the glass window was removed using aseptic technique. We used a volume injection system (World Precision Instruments) to inject 200–400 nl of AAV9-Flex-ChR2-tdTomato (~10^{11} viral particles; Penn Vector Core) 300 μm below the pial surface. The virus was injected at a rate of 50 nl/min through a glass capillary attached to a 10 µl syringe (Hamilton). Following the injection, a new cranial window was sealed in place. Several weeks after injection, we localized the area of ChR2 expression using tdTomato fluorescence, and attached an optical fiber (400 μm diameter; 0.48 nA; Doric Lenses) within 500 μm of the cranial window (~1.3 mm above the cortex). We delivered light though the fiber from a 455 nm LED (ThorLabs) and calibrated the total power at the entrance to the cannula. Optogenetic stimulation began no earlier than 4 weeks after injection. We prevented optogenetic stimuli from cueing the animal to respond by wrapping the fiber implant in blackout fabric (Thor Labs) that attached to the headpost using a custom mount.

Behavioral tasks. Mice were trained to respond to changes in a visual display for a water reward using a lever while head fixed (Histed et al., 2012). In the primary experiment, a static 50% contrast Gabor stimulus was continuously on the screen, presented on a uniform background with the same average luminance. Mice initiated trials by depressing a lever. Following a random delay (400-3000 ms), the contrast of the Gabor stimulus either increased or decreased (interleaved). The Gabor stimulus (SD 5–7°, 0.1 cycles/deg, odd-symmetric) changed contrast for the duration of a brief response window. The size of the contrast change varied randomly from trial to trial across a range that spanned behavioral thresholds for both change types. The mouse had to release the lever within the response window running from 100 ms to 700 or 900 ms after change onset to receive a reward. Following completion of the trial, the contrast of the Gabor stimulus returned to 50%. Stimuli for each animal were positioned at a location that corresponded to the V1 representation expressing ChR2. Early releases and misses resulted in a brief timeout before the start of the next trial. Behavioral control and data collection were done using custom software written using Objective-C, MWorks (mworks.github.io), MATLAB (MathWorks) and Python.

Optogenetic stimulation did not begin until animals worked reliably for hundreds of trials each day, and performance was stable at threshold for both increases and decreases in contrast. This typically required ~2.5 months of training. During optogenetic experiments, we activated ChR2-expressing neurons on a randomly selected half of trials for a single contrast change (~±15% for all mice). These change magnitudes were chosen for stimulation as they approximated the detection

thresholds and thus maximized our ability to resolve an impairment or facilitation of detection capability. We aligned the opsin illumination with visually evoked spiking in V1 by delaying the optogenetic stimulus by 35 ms relative to the appearance of the visual stimulus on the monitor. Opsin illumination persisted until the end of the trial to prevent mice from using the offset of optogenetic input as a task-relevant signal. The optogenetic stimulation intensity was fixed within a session and chosen for each mouse basis based on one or two preliminary testing sessions that were not included in the main analysis. Using these preliminary observations, powers were selected to avoid saturating behavioral performance (ranges for high power sessions: Emx: 0.12-0.25 mW; PV: 0.12-0.25 mW). Following data collection at high powers, we conducted additional sessions in some mice at lower optogenetic stimulus intensities to determine how changes in performance scaled with power (ranges 0.02-0.15 mW).

Follow-up experiments included retraining some PV mice ($n = 2$; both female) to detect contrast increments of a counterphasing Gabor (2 or 4 Hz). The average contrast of Gabor stimulus was held at 20%, except during contrast changes. Changes in contrast were synchronized with zero crossings of the temporal modulation to avoid generating instantaneous luminance steps that could cue the animal to respond. Optogenetic stimulation was delivered on a random subset of trials for a moderate contrast change magnitude (+30%) and the powers used with counterphase-modulated stimuli were identical to those used in the main experiments. As above, optogenetic stimulation was delivered from stimulus onset until the end of the trial. Other task variations included shortening the duration of visual and optogenetic stimulation or ramping and stepping optogenetic stimuli up or down during contrast changes.

Histology. Mice were perfused with 10% pH-neutral buffered formalin (MilliporeSigma), after which the brain was removed and submerged in fixative for 24 h. The brain was subsequently rinsed with PBS and placed in a 30% sucrose PBS solution until it sank. Brains were sectioned at 40 μm on a freezing microtome, mounted, and coverslipped with DAPI Fluoromount-G (Southern Biotechnology). tdTomato expression and DAPI labeling were visualized with 561 and 405 nm excitation light, respectively, using an SP5 Confocal Microscope (Leica Microsystems).

Electrophysiological recordings. We recorded extracellularly from V1 in awake, head-fixed mice ($n = 4$ Emx, 1 female; $n = 4$ PV, 3 female) using multisite silicon probes (Neuronexus; 32-site model 4x8-100-200-177). Electrophysiological recordings were done using passively viewing mice, outside of behavioral sessions. Some mice were first used for behavioral experiments (2 PV, both female), while the rest were untrained but injected with opsins before recording. Electrodes were electroplated with a gold solution mixed with carbon nanotubes (Keefer et al., 2008; Ferguson et al., 2009) to impedances between 200 and 500 kΩ.

At the start of recording sessions, mice were anesthetized with isoflurane (1.2%-2% in 100% O_2), placed in a sled, and head-fixed. While anesthetized, the eyes were kept moist with 0.9% saline. We visualized ChR2-expressing areas of monocular visual cortex by imaging tdTomato fluorescence with a fluorescence microscope and camera (Carl Zeiss). The cranial window was then removed and the electrodes lowered through a slit in the dura. We then positioned an optic fiber above the cortex at a distance comparable to that used during behavioral experiments (1.0-1.5 mm). The craniotomy was then covered with 3% agarose dissolved in aCSF (MilliporeSigma and Tocris Bioscience, respectively). Following the recovery period of 1 h, anesthetic was removed and we waited at least an additional hour for recovery from anesthesia before recording.

The electrode was advanced to locate responsive units, and was allowed to settle for 30 min before collecting data. Delivery of visual and optogenetic stimuli and data acquisition was computer controlled. Concurrent visual and optogenetic stimuli matched those used during behavioral experiments, except that visual and optogenetic changes were presented for 500 ms rather than 700-900 ms, and the visual stimuli filled the video display. As in behavioral experiments, the optogenetic stimulus was delayed by 35 ms relative to the onset of the visual stimulus to account for the neuronal response latencies in V1. We recorded at least 25 repetitions of each stimulus condition in a given stimulus set. For optogenetic stimuli presented in isolation, optogenetic input ramped up for 250 ms at the beginning of each stimulus to a moderate baseline power (0.5 mW), where it remained for 750 ms, and stepped up or down in intensity (randomly interleaved) for 250 ms before returning to the baseline for the remainder of the trial. Electrode signals were amplified, bandpass-filtered (750 Hz to 7.5 kHz), sampled around threshold crossings (Blackrock), and spikes were sorted offline (OfflineSorter, Plexon).

Neural network models. We trained recurrent neural network (RNN) models on a task similar to the one that the mice perform to test whether *in silico* networks adopt the same strategies as *in vivo* networks. RNNs were trained and simulated using the Python machine learning framework TensorFlow (Abadi et al., 2016), and the network architecture was based on our previous study (Masse et al., 2019). Briefly, all networks consisted of orientation and contrast-selective input neurons (whose firing rates are represented as $\boldsymbol{u}(t)$) that projected onto 100 recurrently connected neurons (whose firing rates are represented as $\boldsymbol{h}(t)$, which in turn projected onto the output layer. Recurrently connected neurons never sent projections onto themselves.

The activity of the recurrent neurons was modeled to follow the dynamical system (Song et al., 2016) as follows:

$$\tau \frac{d\boldsymbol{h}}{dt} = -\boldsymbol{h} + f\left(W^{rec}\boldsymbol{h} + W^{in}\boldsymbol{u} + \boldsymbol{b}^{rec} + \sqrt{2\tau}\sigma_{rec}\zeta\right)$$

where τ is the neuron's time constant (set to 50 ms), $f(\cdot)$ is the activation function, W^{rec} and W^{in} are the synaptic weights between recurrent neurons, and between input and recurrent neurons, respectively, $\boldsymbol{b}^{rec}$ is a bias term, ζ is independent Gaussian white noise with zero mean and unit variance applied to all recurrent neurons, and σ_{rec} is the strength of the noise (set to 0.05). To ensure that neuron's firing rates were non-negative and nonsaturating, we chose the rectified linear (ReLu) function as our activation function: $f(x) = \max(0, x)$.

To simulate the network, we used a first-order Euler approximation with time step Δt as follows:

$$\boldsymbol{h}_t = (1-\alpha)\boldsymbol{h}_{t-1} + \alpha f\left(W^{rec}\boldsymbol{h}_{t-1} + W^{in}\boldsymbol{u}_t + \boldsymbol{b}^{rec} + \sqrt{\frac{2}{\alpha}}\sigma_{rec}N(0,1)\right)$$

where $\alpha = \frac{\Delta t}{\tau}$ and $N(0,1)$ indicates the standard normal distribution.

The decision to release the lever (when there was a contrast change) or to hold the lever (when there was no contrast change) was mediated by a competition between two output units. The 80 excitatory neurons linearly projected onto the output unit associated with releasing the lever as follows:

$$z_t^{release} = W^{out}\boldsymbol{h}_t + b^{out}$$

where W^{out} are the synaptic weights between the excitatory neurons and the output unit, and b^{out} is a bias term. The activity of the output unit associated with holding the lever was simply the negative of the activity of the unit associated with releasing the lever: $z_t^{hold} = -z_t^{release}$.

We then calculated the network policy, $\boldsymbol{\pi}_t$, which was the probability of holding or releasing the lever, by taking the softmax of these two values as follows:

$$\boldsymbol{\pi}_t = \text{softmax}\left(\left[z_t^{hold}, z_t^{release}\right]\right)$$

To maintain separate populations of 80 excitatory and 20 inhibitory neurons, we decomposed the recurrent weight matrix, W^{rec} as the product between a matrix for which all entries are non-negative, $W^{rec,+}$ whose values were trained, and a fixed diagonal matrix, D, composed of 1 s and –1 s, corresponding to excitatory and inhibitory neurons, respectively (Song et al., 2016), as follows:

$$W^{rec} = W^{rec,+}D$$

$$D = \begin{bmatrix} 1 & & \\ & \ddots & \\ & & -1 \end{bmatrix}$$

Initial connection weights between excitatory neurons were randomly sampled from a gamma distribution with shape parameter of 0.1 and scale parameter of 1.0, and then multiplied by 0.25. Initial connections weights projecting to or from inhibitory neurons were sampled from a γ distribution with shape parameter of 0.1 and scale parameter of 1.0 and then multiplied by 0.5. Initial bias values were set to 0.

Networks consisted of 24 orientation- and contrast-selective input neurons. The tuning of the input neurons followed a von Mises' distribution, such that the activity of the input neuron i was as follows:

$$u^i = g^i(contrast)Aexp(\kappa\cos\left(\theta - \theta^i_{pref}\right)) + \sqrt{\frac{2}{\alpha}}\sigma_{in}N(0,1)$$

where θ is the orientation of the stimulus (always fixed at 0°), θ^i_{pref} is the preferred direction of input neuron i, κ was set to 2, and A was set to $\frac{1}{\exp(\kappa)}$. The strength of the input activity noise, σ_{in}, was set to 0.05. The function $g^i(contrast)$ determined how neuron i responded to different contrasts. Half (12) of the input neurons were contrast increasing, defined as $g^i(contrast) = contrast$, while the other half were contrast decreasing, defined as $g^i(contrast) = 1/contrast$.

Contrast change detection task for network model. The networks were trained to indicate whether the stimulus contrast changed by responding within a fixed interval. Trials lasted 3000 ms, divided into 10 ms steps. A Gabor patch with an orientation of 0° and a baseline contrast level was presented from the start of the trial, and at a random time, the contrast either doubled or was halved for 100 ms, before returning to baseline. The time of the contrast change was randomly sampled from an exponential distribution with a time constant of 1300 ms, plus 400 ms. If the contrast change did not occur before the end of the trial, the network was rewarded for maintaining hold of the lever throughout the trial. The network received a reward of 1 if it chose to release the lever during the 100 ms duration contrast change, a reward of –0.1 if it chose to either release the lever before the contrast change (i.e., false alarm), or not release during the contrast change (i.e., miss).

Network training. The RNNs using the actor-critic reinforcement learning method (Barto et al., 1983), in which the networks were trained to maximize the discounted cumulative future reward, are as follows:

$$R_\tau = \sum_{t=\tau}^{T} \gamma^{t-\tau} r_t$$

where $\gamma\epsilon(0,1]$ is the discount factor and r_t is the reward given at time t. The network was trained to estimate this discounted future reward as a linear projection from the recurrent units as follows:

$$V_t = W^{value}\boldsymbol{h}_t + b^{value}$$

by minimizing the loss function as follows:

$$\mathcal{L}_V = \frac{1}{2T}\sum_{t=1}^{T}\left[V_t - r_t - \gamma V_{t+1}\right]^2$$

Concurrently, the network adjusts the network policy, $\boldsymbol{\pi}_t$ (described above) to select the actions that would lead to the greatest cumulative reward by minimizing the loss function as follows:

$$\mathcal{L}_P = -\frac{1}{T}\sum_{t=1}^{T}[r_t + \gamma V_{t+1} - V_t]\log\boldsymbol{\pi}_t$$

The network was also encouraged to explore different strategies by maximizing the entropy of the policy output as follows:

$$\mathcal{L}_H = -\frac{1}{T}\sum_{t=1}^{T}\boldsymbol{\pi}_t\log\boldsymbol{\pi}_t$$

Finally, the network was encouraged to solve the task using low levels of neural activity by minimizing the L2 norm of the recurrent neuron firing rates as follows:

$$\mathcal{L}_{sc} = \frac{1}{N\times T}\sum_{t=1}^{T}\sum_{i=1}^{N}h^2_{i,t}$$

where $h_{i,t}$ is the neural activity of the ith recurrent neuron at time t.

Together, the overall loss function is the weighted sum of all four terms as follows:

$$\mathcal{L} = \mathcal{L}_P + \beta\mathcal{L}_V - \alpha\mathcal{L}_H + \gamma\mathcal{L}_{sc}$$

where α and β were set to 0.01. To understand different network solutions across various metabolic constraints, γ was randomly sampled for each network from a logarithmically uniform distribution between $10^{-6.5}$ and $10^{-0.5}$.

We trained all network parameters using the Adam version of stochastic gradient descent, with first and second moment decay rates set to their default values (0.9 and 0.999, respectively). All networks were trained for 50,000 batches, with a batch size of 1024 trials and a learning rate of 0.001.

Experimental design and statistical analysis. Behavioral and electrophysiological data were collected from mice of both sexes. We did not observe any systematic differences between genders, so the data were combined. Details on mouse numbers and genders are included separately for each subsection in the Results. Generally speaking, we were interested in how neuronal spiking or behavioral detection differed when visual stimuli were paired with optogenetic stimulation compared with trials without optogenetic perturbations. In other cases, we compared the probability of behavioral responses following increases or decreases in optogenetic excitation relative to trials in which optogenetic input remained constant. Details regarding statistical tests used in each experiment are listed in the Results corresponding to each figure. Behavioral and electrophysiological data were analyzed in MATLAB, while RNN data were analyzed with Python. $p < 0.05$ was considered the threshold for statistical significance.

For electrophysiological experiments, visually responsive units were taken as those with a 10% change in the average firing rate during the 50–250 ms after stimulus onset (stimulus period) relative to the average firing rate during the baseline epoch (baseline period −250 to −50 ms before stimulus onset) for the largest stimulus intensities. Optogenetically responsive units were defined as any unit with a significant difference ($p < 0.05$; Wilcoxon signed-rank test) in firing rate during the stimulus period for the same visual stimulus with and without optogenetic stimulation. For optogenetic stimuli presented in isolation (see Fig. 5), optogenetically responsive units were classified based on significant changes in firing for optogenetic increments or decrements (from 0 to 250 ms following optogenetic steps) relative to prechange firing rates (from −250 to 0 ms; $p < 0.05$; Wilcoxon signed-rank test). We recorded both single- and multi-units but did not differentiate between them because our primary interest was how optogenetic manipulations affect visually evoked responses across the V1 population.

For behavioral data, the proportion correct for each stimulus condition was determined using trials in which the subject either responded correctly (hit) or failed to respond (miss). Trials in which the animal

released the lever before stimulus onset (false alarm) were not considered in performance analyses. Sessions in which the false alarm rate or the miss rate was >50% were excluded from analysis. We estimated the likelihood of observing a false hit by calculating the conditional probability that the animal would release the lever in each 100 ms bin given that a stimulus had yet to occur in that or earlier bins. We then averaged these values across all bins to obtain the average false hit rate across all time bins. The false hit rate represents the probability that the mouse would get a trial correct because of spontaneous lever releases independent of detecting stimulus changes and is thus a lower bound on performance. The false hit rate was low (Emx median 5.3%; range, 3.8%–6.3%; PV median 5.3%; range, 2.8%–8.3%), demonstrating that mice were relying on stimulus changes to guide their responses. To correct for false hits, we subtracted a randomly selected fraction of correct trials from each contrast level in proportion to their frequency. This correction was applied independently to each behavioral session.

When performance data were fit to psychometric functions, we first corrected for false hits as described above. This correction was typically small (median hits removed 6.8%; range, 4.5%–12.9% for 38 sessions from four mice). Corrected performance data were then fit with a Weibull cumulative distribution function using nonlinear least squares and variance weighting of each mean. The two psychometric functions (with and without ChR2 stimulation) were fit simultaneously using four parameters: individual thresholds ($\alpha_{unstimulated}$, $\alpha_{stimulated}$), a common lapse rate (γ), and a common slope (β) such that:

$$\text{Proportion Correct} = (1 - \gamma) \times \left(1 - e^{-\left(\frac{contrast}{\alpha}\right)^{\beta}}\right)$$

Threshold confidence intervals were estimated using a bootstrap (1000 repetitions, $p < 0.05$, one-tailed).

To compare how optogenetic increments and decrements affected lever responses relative to the trial time-matched false rate, we used the stimulus onset times for 0% contrast change trials in which mice correctly responded to optogenetic stimuli. We restricted our analyses to stimulus onset times that occurred before the final 400 ms of all possible onset times as the small number of observations for the longest trial times made this calculation unstable over the response window. Using only trials in which stimuli had yet to occur by the stimulus onset time, we expressed the lever release time for false alarms relative to the stimulus onset time. Thus, the false alarm distribution aligned to the onset of stimuli served as a time-in-trial matched measure for the lever release probability over time for static optogenetic input. This process was repeated within each session and the probability of lever releases were compared for optogenetic increments, decrements, and time-matched false alarms. The same procedure was used to align lever responses with the onset of PV interneuron activation for trials where contrast changes were omitted.

To link RNN unit responses with network decisions to respond or withhold responses to contrast changes, we calculated the mean normalized receiver operating characteristic (ROC) value for each network. First, we calculated the ROC for each excitatory unit by comparing its firing rate distributions when the network decided to release versus continue holding the lever following contrast changes. The ROC values for contrast increases and decreases were then averaged for each unit.

Next, we summed a rescaled version of each unit's ROC value, normalized by the absolute value of this metric, such that:

$$\Sigma(ROC(i) - 0.5) \div \Sigma|(ROC(i) - 0.5)|$$

Where i is the index of all units in the network. Here, a value of -1 is the extreme case where all units in the network have lower firing rates when the network releases than when it withholds, whereas a value of 1 indicates all units had higher firing rates when the network responds compared with trials in which it does not respond.

Data availability. Data that support the findings of this study are available from the corresponding author on reasonable request.

Code availability. Custom code required for experimental control is available from the corresponding author on reasonable request.

Results

We used transgenic mouse lines that expressed Cre-recombinase selectively in one of three major subclasses of cortical neurons: excitatory neurons (*Emx1*, Emx), *Pvalb*-expressing inter neurons, PV, and *Sst*-expressing inter neurons, SST (Gorski et al., 2002; Hippenmeyer et al., 2005; Taniguchi et al., 2011). These strains allow selective targeting of excitatory opsins to the neurons of interest with >95% specificity (Madisen et al., 2012; Pfeffer et al., 2013).

Changes in visual contrast evoke increases and decreases in V1 spiking

We first performed electrophysiological recordings ($n = 8$; 4 Emx, 1 female; 4 PV, 3 females) to characterize how V1 units respond to changes in visual contrast. Electrophysiological recordings done outside of behavioral sessions using passively viewing, head-fixed mice. Two of the PV mice were first used in the behavioral experiments that are described later. The others were prepared only for electrophysiological recordings.

A full-screen, 50% contrast sinusoidal grating stimulus (0.1 cycles/degree, static, vertically oriented) was always present on the visual display, except during contrast changes. We presented a range of randomly interleaved increases and decreases in visual contrasts (500 ms). Spikes were sorted offline, and responses to different presentations of each contrast change were averaged. For each contrast change, we calculated the change in spike rate (Δspikes/s) by subtracting a prestimulus firing rate (50–250 ms before stimulus) from the firing rate during a stimulus epoch (50–250 ms following stimulus onset).

V1 units exhibited diverse responses that could be used to support contrast change detection ($n = 250$ units; 8 mice; Fig. 1). Some units were excited by either increases (doubling) or decreases (halving) in contrast (Fig. 1*A*), while others were inhibited by contrast changes (Fig. 1*B*). Across the population, we classified units as excited or inhibited if the average firing rate increased or decreased by 10% relative to baseline when the visual contrast either halved or doubled. Figure 1*C* depicts the proportions of units that were modulated by contrast changes. Many units were excited either by decreases (117 of 250, 47%) or increases in contrast (109 of 250, 44%). Smaller numbers were inhibited by decreases in contrast (62 of 250, 25%) or by increases (81 of 250, 32%). The mixture of selectivity for increases and decreases in visual contrast is expected given interspersed ON and OFF responses in V1. Overall, the population responded to both increases and decreases in contrast with higher spike rates (Fig. 1*D*). The net change was small both because different neurons had responses of different sign, and because we presented gratings of a single orientation and spatial frequency that were suboptimal for most V1 neurons recorded.

Using optogenetic approaches to perturb visual processing during behavior

To compare how increments and decrements in V1 spike rates influence perceptual detection, we used optogenetic methods to perturb V1 spiking. Mice were surgically implanted with a headpost and a cranial window to give stable optical access to V1 (Goldey et al., 2014). Following window implantation, we mapped retinotopy in V1 using intrinsic signal imaging (Fig. 2*A*). Imaging data were used to target injections of Cre-dependent viruses containing ChR2-tdTomato (Nagel et al., 2003) to monocular V1 (Fig. 2*B*,*C*). All electrophysiological or behavioral

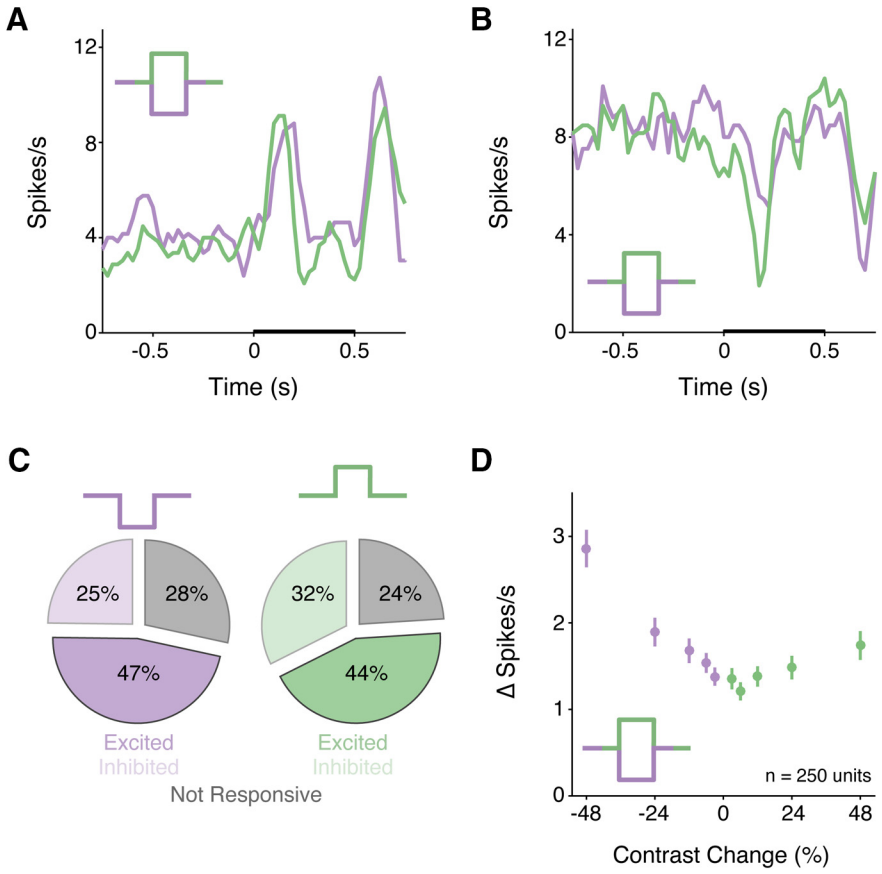

Figure 1. The V1 population exhibits diverse responses to increments and decrements in contrast. ***A***, Average contrast change responses from a representative unit that was excited by halving (decrements, purple) and doubling (increments; green) the stimulus contrast. Legend represents the contrast profile. Visual stimulus duration is indicated by the thickening of the *x* axis in all peristimulus time histograms (PSTHs). Bin size = 25 ms, smoothed. ***B***, Same as in ***A***, but for a unit that was inhibited by contrast changes. ***C***, Pie charts represent the proportions of each response profile to contrast changes. Conventions are the same as in ***A***. Gray represents unresponsive units. Solid color represents excited units. Transparent represents inhibited units. ***D***, Evoked change in firing rate (mean ± SEM) relative to baseline across the population ($n = 250$ units) for halving (purple) or doubling (green) in contrast.

experiments were conducted following stable ChR2 expression (≥1 month post-injection).

In our behavioral experiments, mice were trained to perform a contrast change detection task while head fixed. In this task (Fig. 2*D*), the mouse faced a video display containing static, achromatic, and vertically oriented 50% contrast Gabor (centered at 20°–25° azimuth, −15° to 15° elevation, 5°–7° SD; 0.1 cycles/degree; odd-symmetric during the experimental sessions) on a mid-level gray background. To start a trial, the mouse depressed and held a lever through a randomly varying delay period (600–3000 ms) after which the contrast of the Gabor changed (700 ms). The mouse had to release the lever within a 700 ms response window to receive a reward. We randomly varied the sign and magnitude of the contrast change between trials using a range that spanned behavioral detection threshold. An optical fiber was attached to the headpost to deliver optogenetic stimulation to a consistent cortical location each day (Fig. 2*B*). The optical fiber was aligned with the retinotopic location of the visual stimulus representation in V1 by comparing images of virus expression and intrinsic signal imaging data (Fig. 2*A*,*B*). Light from the optogenetic stimulus was prevented from cueing the mice to respond using a black fabric shield that attached to the headpost.

Optogenetic stimulation of principal versus PV neurons produces opposing effects on V1 unit responses and contrast change detection

During our electrophysiological recordings presented in Figure 1, we also delivered optogenetic stimulation to V1 on a random half of the presentations of 12% (increase in contrast) and −12% (decrease in contrast) stimuli. We used a moderate optogenetic stimulation power (0.3 mW) and visual and optogenetic stimuli were presented for 500 ms. The optogenetic stimulus was delayed by 35 ms relative to the visual stimulus to account for the neuronal response latencies in V1. To assess the effects of optogenetic stimulation on visual responses, we calculated change in the average firing rate (Δspikes/s) for each unit using the 50–500 ms before stimulus onset compared with 50–500 ms after stimulus onset to capture the full duration of optogenetic stimulation. Optogenetic effects were measured by comparing the stimulus evoked firing rates for 12% and −12% contrast changes with and without optogenetic stimulation. To better understand how the effects of optogenetic stimulation compared with the normal physiological response range, we also calculated the change in spike rate for maximal contrast changes (±48% contrast) presented without optogenetic stimulation.

A total of 119 single and multiple units were recorded across 12 V1 sites in 4 Emx mice (3 males). Large decreases in contrast elevated overall firing rates (−48%, black, Fig. 3*A*,*B*: mean 1.6 spikes/s, 0.2 SEM). Moderate decreases in contrast elicited a weaker response (−12%, gray, Fig. 3*A*,*B*: mean 0.7 spikes/s, 0.1 SEM). When −12% decreases in contrast were paired with optogenetic activation of pyramidal neurons, the average response was comparable in magnitude to the response for to the −48% contrast (−12%+Emx, aqua, Fig. 3*A*,*B*: mean 1.9 spikes/s, 0.2 SEM). Average responses differed significantly between stimulus conditions (Fig. 3*B*; all comparisons at least $p < 0.05$; Friedman's test with Dunn–Sidak correction). Thus, optogenetic excitation of V1 Emx-positive neurons significantly elevated spiking when paired with moderate contrast decrements.

V1 responses to the largest increases in contrast were weaker than those for the largest decreases (Fig. 3*C*,*D*; −48% vs 48%, $p < 0.05$, signed rank test). Large increases in contrast evoked an overall increment in firing across the V1 population (48%, black, Fig. 3*C*,*D*: mean 1.1 spikes/s, 0.2 SEM), while the moderate increases in contrast evoked a weaker net positive response (12%, gray, Fig. 3*C*,*D*: mean 0.7 spikes/s, 0.1 SEM). As before, when moderate contrast changes were paired with optogenetic activation of pyramidal neurons, the population response was significantly enhanced (12%+Emx, gold, Fig. 3*C*,*D*: mean 1.8 spikes/s, 0.2 SEM; $p < 0.0001$ for 12% change with vs without optogenetic stimulation; Friedman's test with Dunn–Sidak

correction). As expected, optogenetic activation of pyramidal neurons significantly affected the firing rate of many units (31%; 37 of 119, $p < 0.05$; Wilcoxon's signed rank test), with the vast majority (95%, 35 of 37) having higher firing rates on trials with optogenetic stimulation compared with trials without stimulation. Thus, optogenetic excitation of V1 Emx neurons enhanced the population response to positive changes in contrast.

We examined how well optogenetic stimulation penetrated the cortical depth by measuring responses to 12% contrast changes of both signs from neurons at different depths. We compared data from superficial (upper half of each shank) versus deep (lower half) electrode contacts and found that modulations were substantially stronger for superficial compared with deep contacts (effect of optogenetic stimulation: superficial median 3.7 spikes/s, interquartile range [IQR] 1.8–4.9 spikes/s, deep: median 2.0 spikes/s, IQR 0.5–3.3 spikes/s, $p < 0.05$, rank sum test). Nevertheless, responses of deep neurons were significantly modulated ($p < 10^{-4}$, signed rank test). Thus, the optogenetic stimulus had effects throughout the cortical thickness, although as expected they were somewhat stronger for more superficial units.

We next asked how optogenetic excitation of V1 during visual stimulation affected behavioral detection of contrast changes. Before experimental sessions, all mice were trained to work reliably across a range of contrast changes (increases and decreases, interleaved) that spanned detection threshold for each change type. We chose to study both increases and decreases in contrast as animals may differentially weight increments or decrements in spiking depending on their reliability for each change type, as one might expect from an ideal observer (Geisler, 2011). During experimental sessions, we delivered optogenetic stimulation on a random half of presentations of a single, near threshold, increase (15%) or decrease (−15%) in contrast. For these measurements, the optogenetic stimulus never exceeded 0.25 mW. As with electrophysiological recordings, the optogenetic stimulus was delayed by 35 ms relative to the onset of the visual stimulus to account for the neuronal response latencies in V1. We restricted optogenetic perturbations to the stimulus epoch so as to only affect V1 spiking responses evoked by the visual stimulus. Both the visual and optogenetic stimulus remained on until the end of the trial so as to prevent stimulus offsets from producing an additional signal that could drive behavioral responses.

Emx mice ($n = 3$, 2 males) detected ±15% contrast changes comparably well when those visual stimuli were presented without optogenetic stimulation (median percent correct, decrements 44%, 3% SEM, vs increments 42%, 2% SEM; 21 sessions in 3 mice; $p > 0.05$, Wilcoxon signed-rank test). Optogenetic excitation of pyramidal neurons significantly increased the proportion of trials in which mice detected increases or decreases in contrast. Figure 3*E* shows data from a representative session in which pyramidal neurons were activated during some trials on which the contrast increased or decreased by 15%. In either case detection was enhanced.

Improvements in detecting contrast changes were seen in virtually every session (Fig. 3*F*; 21 sessions in 3 mice; decreases median stimulated = 66% [range: 46%–92%], unstimulated 44% [range: 16%−81%], $p < 10^{-4}$; increases median stimulated = 68% [range: 44%–97%] vs unstimulated 42% [range: 12%−58%], $p < 10^{-4}$; Wilcoxon signed-rank tests). Activation of excitatory neurons significantly increased the proportion of hits in many individual sessions (contrast decreases: 9 of 21 sessions; contrast increases: 5/21 sessions; $p < 0.05$, both Fisher's exact test). Optogenetic activation of excitatory neurons also shortened reaction times on trials in which the animal correctly detected the stimulus compared with trials with no optogenetic stimulation (decreases in contrast: unstimulated median = 347 ms, 296–414 IQR; stimulated median = 316 ms, 273-388 IQR; $p < 10^{-4}$; increases: unstimulated median = 356 ms, 295–424 IQR; stimulated median = 326 ms, 276–409 IQR; $p < 0.01$; Komolgorov-Smirnov tests). In summary, optogenetic stimulation of excitatory neurons enhanced the ability of mice to detect changes in contrast regardless of the sign of the contrast change.

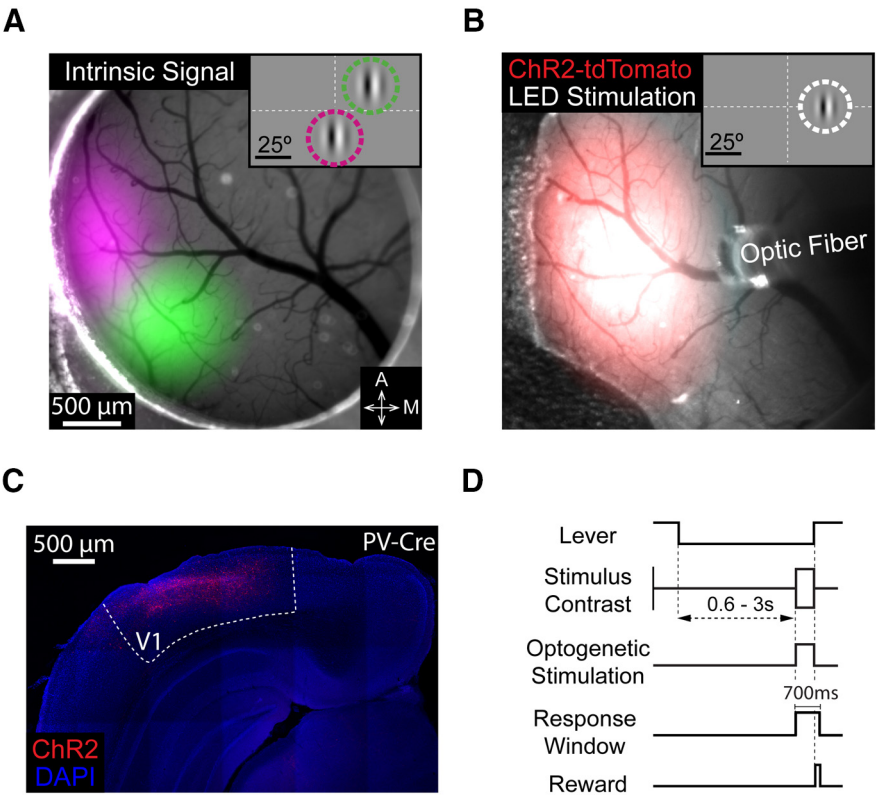

Figure 2. Targeting ChR2 to retinotopically defined areas of visual cortex. ***A***, Pseudo-colored intrinsic autofluorescence responses to visual stimuli presented in two locations in a PV-Cre mouse. Magenta and green features represent 2D-Gaussian fits of responses to stimuli at visual field locations depicted in the inset (magenta: 0° azimuth, −20° elevation; green: 25° azimuth, 20° elevation; Gabor SD = 10°). Dashed lines indicate horizontal and vertical meridians. A, Anterior; M, medial. ***B***, ChR2-tdTomato fluorescence (2D-Gaussian fit) from the same cortical region shown in ***A***. Area of LED illumination (2D-Gaussian fit) through the optic fiber positioned above ChR2-expressing V1. The retinotopic location corresponding to maximal expression was used in all behavioral sessions (shown in inset; 25° azimuth, 0° elevation; Gabor SD = 6.75°). ***C***, Representative confocal image of ChR2-tdTomato expression in the visual cortex of a different PV-Cre mouse. ***D***, Trial schematic of the contrast change detection task. A contrast change of each direction (increase or decrease) was selected for stimulation (±15% contrast change for all mice).

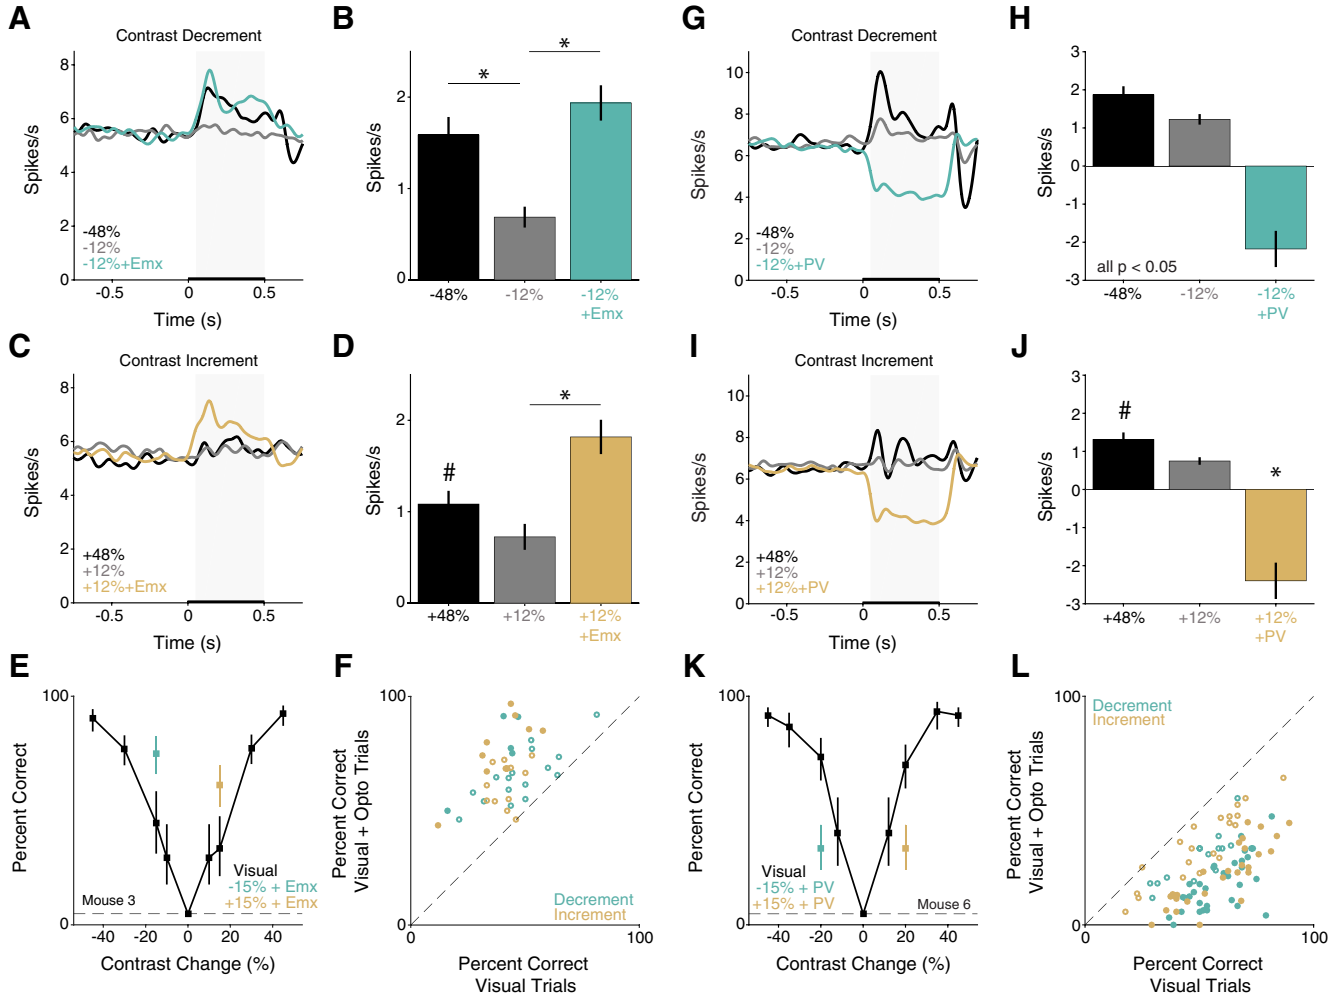

Figure 3. Optogenetic stimulation of principal versus PV neurons produces opposing effects on V1 population responses and contrast change detection. ***A***, Population Gaussian-filtered (σ = 25 ms) PSTH in response to large (black) and moderate decreases in contrast without (gray) or with (aqua) optogenetic stimulation of pyramidal neurons in passively viewing Emx mice ($n = 119$ units). Thickening of the *x* axis represents duration of visual and optogenetic stimuli. Gray box represents analysis window (50-500 ms) used for spike rate quantification in ***B***. ***B***, Average change in spike rate (± SEM) compared with the time-matched baseline period. *$p < 0.0001$, relative to −12% change (Friedman's test with Dunn–Sidak correction). ***C***, Same as in ***A***, but for contrast increments. ***D***, Quantification of spike rate changes evoked by contrast increments with and without optogenetic stimulation. #$p < 0.10$, 48% compared with both 12% contrast change without (gray) and with (gold) optogenetic stimulation. *$p < 0.0001$, 12% change with versus without optogenetic stimulation (Friedman's test with Dunn–Sidak correction). ***E***, Representative behavioral performance from a single session in an Emx mouse. Points and lines represent the percent correct ± 67% CI for trials without (black) optogenetic stimulation or decreases (aqua) and increases in contrast (gold) paired with optogenetic of excitatory neurons. Dashed line indicates false alarm rate. ***F***, Summary of stimulation effects in Emx mice. Circles represent the percent correct in individual behavioral sessions (3 mice, 21 sessions) with (*y* axis) and without (*x* axis) optogenetic stimulation, separately for increases (gold) and decreases (aqua) in contrast. Filled circles represent significant change in detection performance (14 of 42 observations, $p < 0.05$, Fisher's exact test). ***G***, Same as in ***A***, but for decreases in contrast presented with and without optogenetic stimulation of PV interneurons in PV mice ($n = 131$ units). ***H***, Average change in spike rate with and without stimulation of PV interneurons compared with the time-matched baseline period (all comparisons at least $p < 0.05$; Friedman's test with Dunn–Sidak correction). ***I***, Same as in ***G***, but for contrast increments with and without stimulation of PV interneurons. ***J***, Average change in spike rate evoked by contrast increments with and without optogenetic stimulation of PV interneurons. #$p < 0.10$, 48% compared with 12% contrast change without (gray) optogenetic stimulation. *$p < 10^{-8}$, 12% change with optogenetic stimulation (gold) compared with large (48%) and moderate (12%) contrast increments. ***K***, Same as in ***E***, but for a single session in a PV mouse. ***L***, Summary of stimulation effects in PV mice (6 mice, 47 sessions). Conventions are the same as in ***F***. Filled circles represent significant change in detection performance (57 of 92 observations, both increases and decreases; $p < 0.05$, Fisher's exact test).

The population response to moderate increases in contrast paired with pyramidal neuron stimulation (12%+Emx, gold, Fig. 3*C*,*D*) was larger than the response to large contrast increases (48%, black, Fig. 3*C*,*D*). During behavioral sessions, animals typically detected large increases in contrast with greater frequency than trials with optogenetic stimulation (Fig. 3*E*,*F*). This difference is likely because of our recordings being conducted in different animals outside of the behavioral task, using stimulation powers at the upper limit of those used during behavior.

These data suggest that detection of changes in visual contrast benefits from increases in V1 spike rates. Previously, we showed that PV neuron stimulation impairs contrast change detection (Cone et al., 2019), but decreases in contrast were not tested and no neurophysiological measurements were made. Given that many (~30%; compare Fig. 1) of V1 neurons were inhibited by contrast changes, we next sought to test whether reductions in V1 spiking could support contrast change detection. To do so, we reduced V1 spike rates with optogenetic activation of PV inhibitory interneurons.

Prior work has shown that activation of PV interneurons in V1 inhibits visually evoked neuronal responses (Atallah et al., 2012; Wilson et al., 2012; Glickfeld et al., 2013). To confirm that PV stimulation indeed inhibited population spiking, we recorded from 131 units (including multiple units) from 12 sites in 4

awake, passively viewing mice (3 females) with ChR2 expressed in V1 PV interneurons. As with Emx mice, we compared the change in firing rate evoked by large contrast changes (±48%), moderate contrast changes (±12%), and moderate contrast changes with concurrent PV stimulation (±12% + PV; Fig. 3*G*, *H*). Large decreases in contrast evoked a robust increase in firing rate across the population (black, Fig. 3*G*,*H*; −48%: mean 1.9 spikes/s, 0.2 SEM), whereas moderate decreases in contrast evoked a weaker response (gray, Fig. 3*G*,*H*; −12%: mean 1.2 spikes/s, 0.1 SEM). Pairing moderate reductions in contrast with optogenetic activation of PV interneurons robustly decremented V1 output compared with prestimulus firing rates (aqua, Fig. 3*G*,*H*; −12% + PV: mean −2.2 spikes/s, 0.5 SEM). Average responses differed significantly between stimulus conditions (Fig. 3*H*; all comparisons $p < 0.05$; Friedman's test with Dunn–Sidak correction). Thus, PV stimulation produced a change in V1 output that was comparable in magnitude, but opposite in sign, to the response to large contrast changes.

Responses to increases in contrast were weaker than decreases, although still above baseline firing rates ($p < 0.05$ for both −48% vs 48% and −12% vs 12%, signed rank tests). Large increases in contrast-enhanced V1 output (black, Fig. 3*I*,*J*; 48%: mean 1.3 spikes/s, 0.2 SEM) and a moderate change evoked a smaller response (gray, Fig. 3*I*,*J*; 12%: mean 0.7 spikes/s, 0.1 SEM). Pairing a moderate contrast increment with optogenetic activation of PV interneurons strongly and significantly suppressed V1 output compared with prestimulus firing rates (gold, Fig. 3*I*,*J*; 12% + PV: mean −2.4 spikes/s, 0.5 SEM; 12% change with optogenetic stimulation vs other conditions; both $p < 10^{-8}$; Friedman's test with Dunn–Sidak correction).

Consistent with the strong effect of PV neurons on population responses, most recorded units (65%; 86 of 131) were significantly modulated by optogenetic stimulation ($p < 0.05$; Wilcoxon's signed rank test comparing evoked responses with and without optogenetic stimulation). Of these units, almost all (93%, 80 of 86) had lower firing rates on trials with optogenetic stimulation compared with trials without stimulation, as expected for activation of inhibitory interneurons. As was done in Emx mice, we measured whether the effects of optogenetic stimulation in PV mice varied across the thickness of cortex. Among units that were significantly modulated by optogenetic stimulation, optogenetic effects were similar for superficial versus deep contacts (median optogenetic modulation: superficial: −2.6 spikes/s, IQR −6.5 to −0.4 spikes/s, deep: −3.3 spikes/s, IQR −7.7 to −1.6 spikes/s, $p > 0.05$, rank sum test).

To examine how reductions in V1 spiking affected contrast change detection, we prepared and trained PV-Cre mice ($n = 6$, 2 females) to detect interleaved, bidirectional changes in contrast as above. We delivered optogenetic stimulation concurrently with a single, near threshold, increase, or decrease in contrast. In the absence of optogenetic stimulation, animals had comparable levels of detection performance for ±15% contrast changes without optogenetic stimulation (median across 47 sessions, decreases 54% correct, 2% SEM vs increases 59%, 3% SEM; 47 sessions in 6 mice; $p > 0.05$, Wilcoxon signed-rank test). Optogenetically stimulating PV interneurons produced behavioral effects that were opposite to those observed in Emx mice. Figure 3*K* shows data from a representative session in which PV interneurons were activated during some trials on which the contrast increased or decreased by 15%. In either case, detection was impaired.

Impairments in detecting contrast changes were seen in virtually every session (Fig. 3*L*; decreases median stimulated percent correct = 24% [range: 3%–56%] vs unstimulated 54% [range: 29%–82%], $p < 10^{-8}$; increases median stimulated = 29% [range: 0%–66.7%] vs unstimulated 59% [range: 18%–90%], $p < 10^{-8}$; both Wilcoxon signed-rank test; Fig. 3*K*,*L*). PV interneuron stimulation significantly reduced the proportion of hits in most individual sessions (decreases: 30 of 47 sessions; increases: 27 of 47 sessions; $p < 0.05$, both Fisher's exact test; Fig. 3*L*). PV stimulation also affected reaction times, generally slowing responses (decreases in contrast: unstimulated median = 338 ms, 290–406 IQR; stimulated median = 373 ms, 293–474 IQR; $p < 0.01$; increases: unstimulated median = 326 ms, 284–396 IQR; stimulated median = 330 ms, 278–404 IQR; $p = 0.45$; both Komolgorov-Smirnov tests). These data show that reductions in V1 spiking impair detection of contrast changes regardless of sign. Moreover, that we observed opposing results in Emx and PV mice argue against the possibility that the facilitation of detection observed in Emx mice was because of scattered light from the optogenetic stimulus cueing animals to respond.

Optogenetic stimulation of PV interneurons impairs detection of contrast increments in dynamic stimuli

In our main experiment, a static 50% Gabor stimulus was always present on the video display, which can attenuate visual responses because of adaptation (Blakemore and Campbell, 1969; Maffei et al., 1973; Movshon and Lennie, 1979). We wanted to compare those effects with detection of contrast changes in counterphasing stimuli, which are made up of ongoing positive and negative changes in contrast and produce a sustained, elevated spike rate in V1.

In a subset of the recordings from PV mice ($n = 3$) shown in Figure 3, we also presented contrast changes on the background of a 20% contrast counterphasing Gabor stimulus. The V1 population responded to contrast increments (Fig. 4*A*). Few units were significantly inhibited by counterphase modulated contrast increments (6.5%, 6 of 93 units, $p < 0.05$ relative to baseline, Wilcoxon signed-rank test), whereas many units were excited (43%, 40 of 93 units). On average, moderate increases in contrast evoked an increase in spiking (gray, Fig. 4*A*; 24%: mean 1.5 spikes/s, 0.2 SEM). However, when increases in contrast were paired with optogenetic stimulation of PV neurons, the average population signal changed sign (gold, Fig. 4*A*; 24% + PV: mean −2.2 spikes/s, 0.5 SEM). The population responses were significantly different between conditions (Fig. 4*A*; $p < 10^{-9}$; Wilcoxon signed rank test).

We next wanted to assess how decrementing V1 spike rates on the background of a counterphase modulated Gabor would affect contrast change detection performance. We trained PV mice ($n = 2$; both female; 1 was retrained after completing the experiments presented in Fig. 3) to detect contrast increments of a Gabor stimulus that was counterphase modulated (Fig. 4*B*; 2 or 4 Hz). We optogenetically activated PV neurons on a random half of trials during a moderate contrast increment (30%). As before, optogenetic stimulation was delivered starting at stimulus onset through to the end of the trial. Figure 4*C* shows data from a representative session in which PV interneurons were activated during some trials in which the contrast increased 30%. Consistent with our primary findings, optogenetic stimulation of PV interneurons greatly impaired task performance. Across all sessions (2 mice, 28 sessions; for 2 Hz [$n = 18$ sessions] and 4 Hz [$n = 10$ sessions]), the proportion of trials in which mice successfully detected a 30% contrast increment was greatly reduced on trials with optogenetic stimulation of PV neurons (Fig. 4*D*; median = 57.1% vs 26.9%; $p < 10^{-5}$, Wilcoxon signed-rank test).

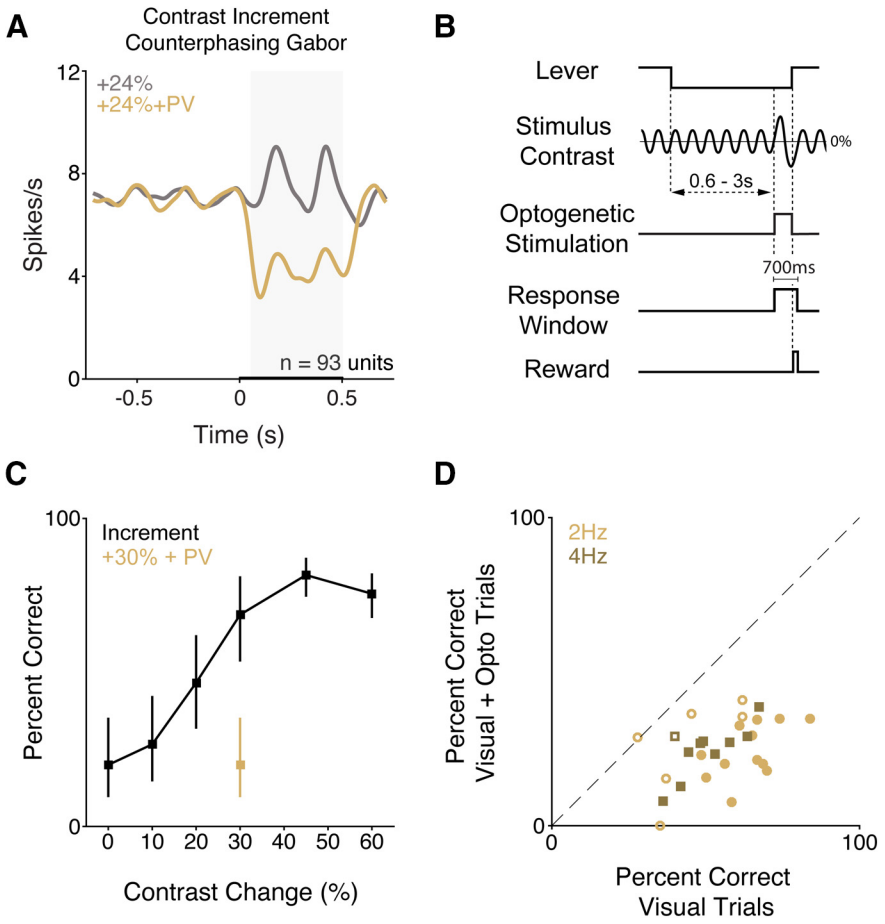

Figure 4. Optogenetic stimulation of PV interneurons impairs detection of contrast increments presented on a counterphase modulated Gabor. ***A***, Population Gaussian-filtered (σ = 25 ms) PSTH in response to moderate increases in contrast presented on a counterphase modulated Gabor without (gray) or with (gold) optogenetic stimulation of PV neurons in passively viewing mice (n = 93 units). Thickening of the x axis represents duration of visual and optogenetic stimuli. Gray box represents analysis window (50–500 ms) used for spike rate quantification. The population responses were significantly different across conditions ($p < 10^{-9}$, Wilcoxon signed-rank test). ***B***, Trial schematic for behavioral experiments. The average contrast of the Gabor stimulus was held at 20%. Contrast changes were synchronized with the zero crossings of the phase modulation and were randomly selected from values spanning threshold. ***C***, Representative single-session performance (hit rate ± 67% CI) from a mouse trained to detect contrast increments of counterphase modulated stimuli. ***D***, Symbols represent the percent correct in individual behavioral sessions (2 mice, 28 total sessions) with (y axis) and without (x axis) optogenetic stimulation, separately for sessions in which the counterphase modulation frequency was 2 Hz (gold circles; 18 sessions) or 4 Hz (brown squares; 10 sessions). Filled circles represent a significant change in detection performance (21 of 28 sessions; 2 Hz: 12 of 18 sessions $p < 0.05$; 4 Hz: 9 of 10 sessions $p < 0.05$; Fisher's exact test).

Performance was impaired for both counterphase modulation rates (2 Hz: median = 61.5% vs 25.7, $p < 10^{-3}$; 4 Hz: median = 47.8% vs 26.7 $p < 0.01$; both Wilcoxon signed-rank test) and both mice (2 Hz: both mice at least $p < 10^{-11}$; 4 Hz: both mice at least $p < 10^{-10}$; all Fisher's exact test) individually. Thus, decreasing V1 spike output on the background of a dynamic visual stimulus impairs contrast change detection. These data suggest that our primary findings do not depend on the adaptation of sustained neuronal responses to the static Gabor.

Different behavioral consequences of optogenetic increments and decrements of excitatory neuron spiking

The above data show that mice preferentially respond to increments in V1 spiking. However, this may be because spiking increments and decrements were not tested on equal footing. We wondered whether mice might exploit spiking decrements if asked to detect reductions relative to a sustained, elevated baseline level of V1 spiking. This approach would also explicitly reward mice for responding to decrements in V1 spiking. Moreover, our prior experiments perturbed different cellular and synaptic elements (excitatory vs inhibitory neurons), which could contribute to the differences in behavioral sensitivity to increments or decrements in overall spiking.

To measure the effects of increments and decrements in optogenetic stimulation on V1 spiking, we recorded from 50 units (including multiple units) from five sites in 3 Emx mice (2 males) that expressed ChR2. Here, mice passively viewed a stimulus monitor that was held at a neutral gray, and we recorded V1 neuronal responses to optogenetic increments and decrements presented on a background of optogenetic excitation. Most V1 units were sensitive to changes in optogenetic input (Fig. 5*A-C*; 62%, 31 of 50; $p < 0.05$; Wilcoxon signed-rank test on firing rates during the optogenetic step relative to baseline).

Prolonged optogenetic manipulations can produce paradoxical effects in cortical networks (Tsodyks et al., 1997; Sadeh et al., 2017), but we found no evidence of instability in our data, likely owing to the low powers used for stimulation. For the subpopulation of optogenetically modulated units, firing rates closely followed the optogenetic stimulation (Fig. 5*A*,*B*). We quantified the evoked change in spike rate during the change in optogenetic input relative to baseline. The firing rate change evoked by decrements and increments in optogenetic input followed the step size (Fig. 5*C*; both $p < 10^{-4}$; Kruskal–Wallis test). Thus, V1 principal neurons can faithfully follow ramps and steps in optogenetic input at modest input powers.

We then trained a new cohort of Emx mice on a modified visual detection task (Fig. 5*D*; 3 mice, 1 female; 50 sessions) that incorporated optogenetic stimulation that persisted through the random delay period. Mice were first trained to detect interleaved contrast changes and injected with opsins. Once proficient, we shaped mice to perform the task with optogenetic stimulation. When the visual stimulus changed contrast, the optogenetic input either incremented or decremented to potentiate or reduce its input into V1. In addition to examining changes in optogenetic stimulation from a sustained baseline, this design allowed us to examine sensitivity to V1 spiking increments and decrements in the same animals using perturbation of the same circuit (as opposed to effects in different Cre lines).

In all cases, the sustained optogenetic stimulus was modest (0.10–0.25 mW) and stepped up or down by an amount comparable with the baseline power. Our electrophysiological recordings showed that the absolute spike rate changes evoked by decrements and increments within this power range were comparable. The optogenetic stimulus step was brief, 75 ms (19

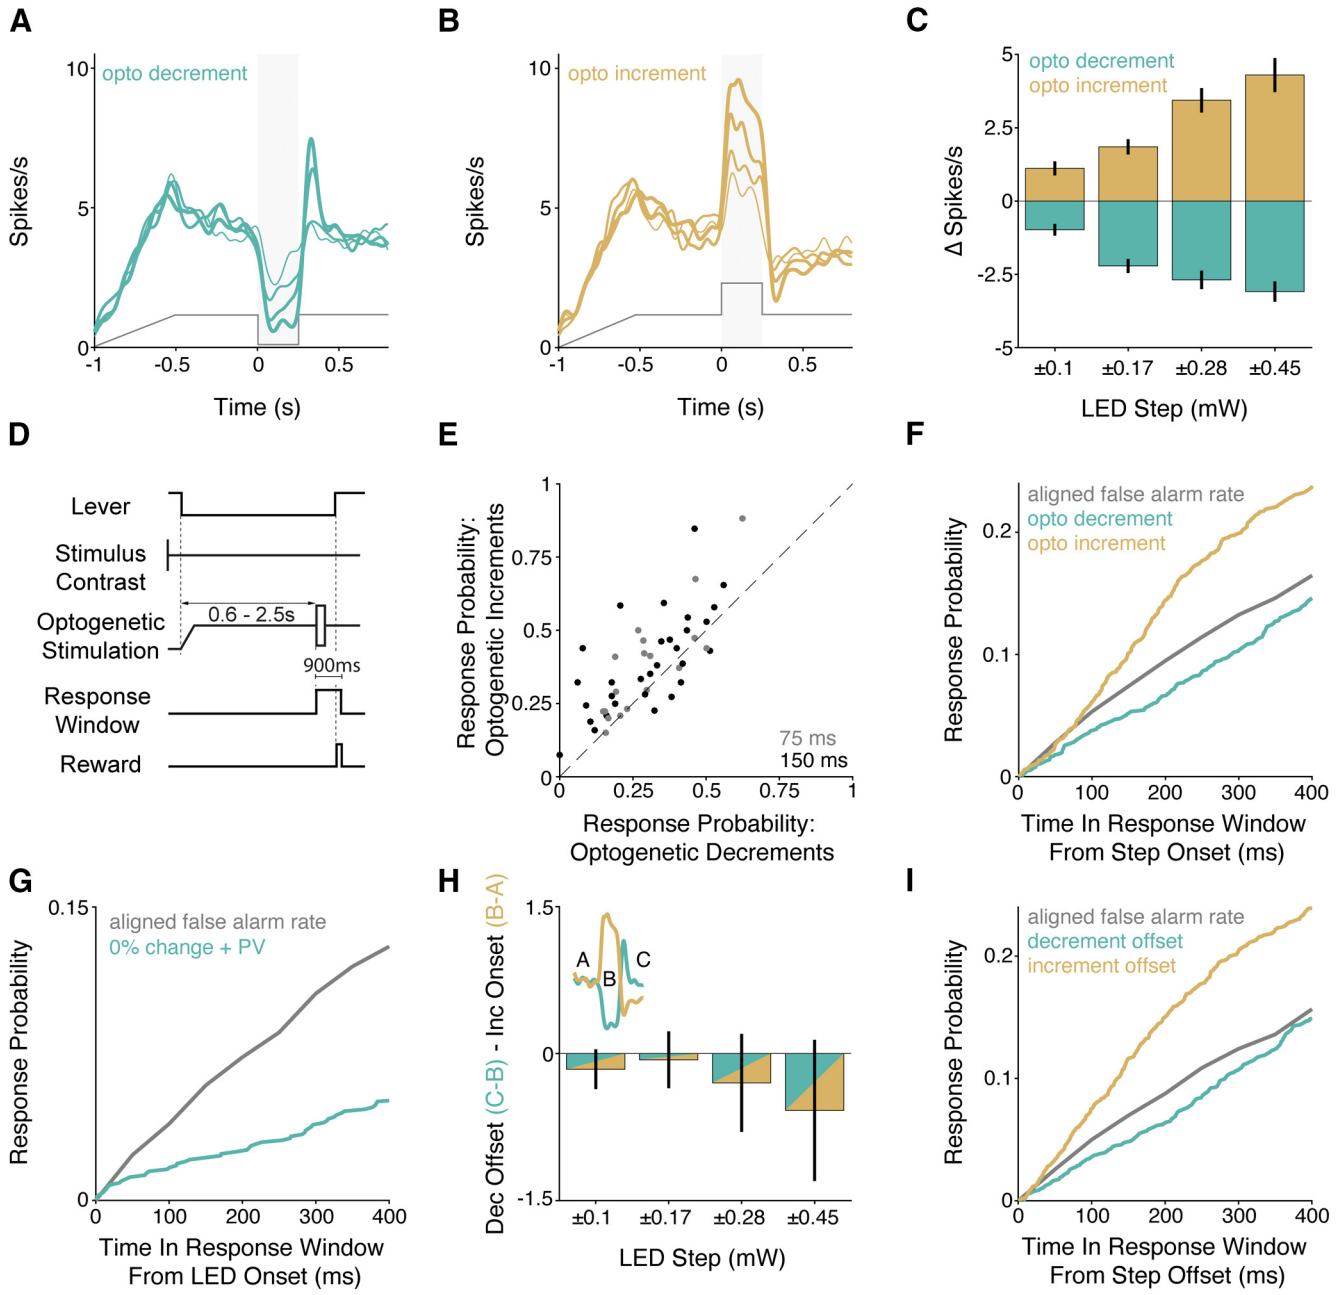

Figure 5. Optogenetically incrementing or decrementing excitatory input into the V1 population asymmetrically affects the probability of behavioral responses. ***A–C***, Electrophysiological recordings confirm V1 spiking follows complex optogenetic input. ***A***, Average PSTHs for optogenetically responsive units ($n = 31$ units) in response to different magnitude decrements in optogenetic input. Increasing line thickness corresponds to the magnitude of optogenetic decrement (step sizes listed in ***C***). Gray trace represents the profile of optogenetic stimulation. ***B***, Same as in ***A***, but for increments in optogenetic input. ***C***, Quantification of V1 spike rate changes in response to decrements (aqua) and increments (gold) in optogenetic input (Mean ± SEM). The magnitude of spike rate change significantly depends on the step size (increments and decrements, both $p < 10^{-4}$; Kruskal–Wallis test). ***D***, Trial schematic. The visual contrast was fixed at 50%. At trial onset, the LED power ramped up and then held at a fixed value. After the random delay, the LED power briefly (75 or 150 ms) stepped up or down. ***E***, Scatter plot depicts % correct for trials with optogenetic decrements (*x* axis) and increments (*y* axis) measured in the same session ($n = 3$ mice; 50 sessions). Gray points represent 75 ms optogenetic step duration (Mean ± SEM). Black represents 150 ms optogenetic step duration. ***F***, Time course of changes in the lever release probability following optogenetic decrements (aqua) or increments (gold) compared with a time-in-trial matched false alarm rate (gray). ***G***, Same as in ***F***, but for optogenetic stimulation of PV interneurons (aqua; $n = 4$ mice; 38 sessions) on trials without contrast changes compared with a time-in-trial matched false alarm rate (gray). ***H***, Quantification of differences between V1 spike rate changes induced by the offset of optogenetic decrements (aqua) versus the onset of increments (gold) across the different optogenetic step sizes. Inset, Epochs of spiking being compared. Error bars indicate SEM. Dec, decrement; Inc, increment. ***I***, Same as in ***F***, but behavioral responses were aligned to the end of the optogenetic steps.

sessions) or 150 ms (31 sessions), while the contrast change persisted until the end of the trial (900 ms). On a subset of trials, the optogenetic stimulus incremented or decremented after the random delay while the contrast of the visual stimulus did not change. On these trials, the animal was rewarded for responding to increments or decrements in the optogenetic stimulus.

The task design encouraged animals to set a liberal response criterion, which elevated the probability of false alarms. We compared the probability of false alarms during the first second of all possible stimulus times for the same 50% contrast baseline Gabor when it was paired with optogenetic excitation relative to earlier experiments when the baseline period was unstimulated

(compare Fig. 3). Baseline optogenetic excitation produced a small but significant increase in the probability of lever releases (median with baseline optogenetic excitation = 6.7%, IQR 5.7%–7.4% [50 sessions]; median without = 5.2%, IQR 4.7%–6.0% [68 sessions] from PV and Emx mice shown in Fig. 3), $p < 10^{-7}$, Wilcoxon rank-sum test).

In general, mice detected contrast changes more often and faster when the optogenetic stimulus incremented than when it decremented. However, we were primarily interested in the behavioral effects on trials where the visual contrast did not change, when behavioral responses were cued by changes in the optogenetic stimulus alone. Mice did not respond reliably to steps in the optogenetic stimulus in the absence of visual stimulus changes but were more likely to respond to optogenetic increments compared with decrements. Figure 5*E* shows the session-average probability of lever releases for optogenetic increments and decrements on trials without visual stimulus changes. Although many releases were false alarms, there were more releases on trials when the optogenetic stimulus increased (increment response rate = 38%, 35–40 95% CI; decrement response rate = 29%, 27–31 95% CI; t statistic = 8.1, $p < 10^{-15}$, binomial logistic regression). Correspondingly, when mice responded to optogenetic stimuli, reaction times were consistently faster for increments than decrements (increments median: 357 ms, 241-585 IQR; decrements median: 467 ms, 257–685 IQR; $p < 0.001$; Kruskal–Wallis test).

Because the animals were operating with an elevated false alarm rate, we could examine how changes in the optogenetic stimulus affected the probability of lever releases over time. Figure 5*F* plots probabilities of lever release following increments (gold) or decrements (aqua) in optogenetic stimulation across all 50 sessions (both 75 and 150 ms step durations), together with trial time-matched false alarm probability (see Materials and Methods). The probability of lever responses increased relative to the false alarm rate when the optogenetic stimulus power incremented (gold) and decreased when the stimulus power decremented (aqua). We repeated this analysis on new data collected during additional sessions in a subset of the PV mice shown in Figure 3 ($n = 4$, 38 sessions). In these additional sessions, a subset of trials featured stimulation of PV interneurons (700 ms) at a random point in the trial and the contrast of the Gabor did not change. Figure 5*G* shows that inhibiting V1 via PV stimulation suppresses the lever release probability relative to the false alarm rate, replicating what we observed in Emx mice following optogenetic decrements. This shows that mice are less likely to report a stimulus change when the rate of V1 spiking is reduced, independent of how that reduction is implemented.

Because the optogenetic perturbations were transient, both polarities of optogenetic steps were associated with increments and decrements in spike rate at the step onset (offset) and offset (onset). To compare the V1 signal available following onsets and offsets of optogenetic stimuli, we calculated the change in spike rate following increment onsets and decrement offsets for each optogenetically responsive unit (Fig. 5*H*, inset). The average difference between these values were slightly negative across the optogenetic step sizes, but in all cases the means were within 1 SEM of zero (Fig. 5*H*; mean difference, ±0.1 mW: −0.16, 0.2 SEM; ±0.17 mW: −0.06, 0.3 SEM; ±0.28 mW: −0.30, 0.5 SEM; ±0.45 mW: −0.58, 0.7 SEM). This indicates that the changes in V1 spiking at the start and end of optogenetic pulses were about the same size, but opposite sign.

If perceptual reports were driven by instantaneous increases in firing rate, the end of the optogenetic decrement might be equally well detected as the as the start of the optogenetic increment. As described above, there was no evidence of a sharp increase in lever releases at any time during or after optogenetic decrements when we aligned behavioral responses to the onset of optogenetic decrements (Fig. 5*F*). Aligning those same data on the end of optogenetic decrements (Fig. 5*I*) similarly did not reveal an increase in the probability of lever releases. Together, these data suggest that mice do not respond to instantaneous increments in spike rate relative to immediately preceding epochs of activity but instead depend on integrating spike counts over periods that likely span hundreds of milliseconds (see Discussion).

Brief activation of PV or SST interneurons impairs contrast change detection

The previous experiment suggests that behavioral responses depended on the integrated spike count and not the rate of change in firing rate (derivative of spike rate) in V1. However, the visual stimulus lasted for 900 ms, which could have encouraged mice to adopt an integration-based strategy. Thus, we did an additional experiment in which both visual and optogenetic stimuli were brief (100 ms). In contrast to our earlier experiments in PV mice where optogenetic stimulation persisted until the end of the response window (700 ms; compare Fig. 3), the 100 ms optogenetic stimulus allowed mice to respond to the termination of inhibition (Fig. 6*A*). We also used higher stimulation powers (0.5–1.0 mW vs 0.20–0.25 mW in Fig. 3) to produce a strong reduction in V1 spiking relative to baseline firing, which in turn generates a rebound in spiking at inhibition offset (see Fig. 8*A*).

In addition to changing the temporal profile of inhibition, optogenetic stimulation occurred on 50% of the trials, including many trials where the visual stimulus was subthreshold, to encourage the mice to learn to exploit these brief changes in spike rate in guiding their responses. We also included animals expressing ChR2 in SST neurons to directly compare optogenetic stimulation of different classes of inhibitory interneurons. PV and SST interneurons each suppress V1 responses (Wilson et al., 2012) but act through distinct synaptic mechanisms (Kubota et al., 2016). We reasoned that, if mice could use instantaneous increases in V1 activity to guide their responses, the offset of robust, brief inhibition, and the resulting rebound in spiking could potentially facilitate detection. Conversely, if mice rely on the integrated change in spiking, then inhibitory neuron stimulation should always impair detection by removing spikes from the population signal.

For these experiments, we collected new data from PV ($n = 2$; 1 female) and SST ($n = 2$, both male) mice that were part of a previously published report (Cone et al., 2019). For the new dataset, mice detected decreases in contrast. A 75% contrast vertically oriented Gabor stimulus (0.1 cycles/deg, centered at 20°–25° azimuth, −15° to 15° elevation, 12° SD, odd-symmetric) was always present on the video display except when its contrast transiently reduced to a lower value (100 ms; Fig. 6*A*). Contrast changes spanned psychophysical threshold, and on a randomly selected half of all trials, were synchronous with 100 ms illumination of ChR2-expressing PV or SST neurons. We then fit psychometric functions separately to performance values for trials with and without stimulation (see Materials and Methods).

Without optogenetic perturbations, there were no detected differences in detection performance between the mouse strains (PV median threshold 11%; SST median threshold 11%; $p = 0.89$; Wilcoxon rank-sum test). PV or SST activation during the visual

stimulus impaired detection, shifting the psychometric functions to the right (Fig. 6*B*,*C*). With the illumination powers used, PV or SST stimulation elevated detection thresholds approximately twofold (PV 2.4-fold, SEM 0.1, range 1.3–4.2; SST 2.0-fold, SEM 0.1, range 1.3–3.0). Across all sessions (4 mice, 38 sessions), contrast change detection thresholds were significantly greater when the visual stimulus was paired with either PV or SST activation (Fig. 6*D*; medians = 23% vs 11%, $p < 10^{-7}$; Wilcoxon signed-rank test). This effect was significant for both genotypes individually (PV medians = 23% vs 11%, $p < 10^{-4}$; SST medians = 21% vs 11%, $p < 10^{-3}$; Wilcoxon signed-rank tests), and in most individual sessions (PV: 20 of 20; SST 17 of 18).

We next asked whether using other optogenetic stimulus powers could facilitate detection. We conducted additional sessions in 1 PV mouse in which we presented visual stimuli as above and delivered optogenetic inhibition on a subset of trials for a single contrast decrement (Fig. 6*E*). We varied the LED power across sessions (0.1-2.0 mW in 8 sessions) and measured the change in performance (Δ percent correct) at the tested contrast level. The tested contrast change level was held constant across sessions. We found that the impairments produced by PV stimulation scaled approximately linearly with power, such that higher powers produced correspondingly larger deficits in detection performance (Fig. 6*F*; maximum likelihood linear regression anchored at the origin; r^2 = 0.67). Thus, although increasing the degree of inhibition should produce more robust rebounds in spiking, mice appear unable to use this signal to guide their responses. That reductions in V1 spiking, even when followed by strong positive rebounds, act to suppress the probability of detection is largely consistent with the idea that detection is mediated by integrating V1 spikes over hundreds of milliseconds. Combining our previously published data with the current results, PV or SST stimulation never facilitated detection performance across 125 behavioral sessions ($n = 8$ mice; 5 PV, three SST). Together, these data suggest that mice cannot learn to respond to decrements in V1 spike rate even with extended practice.

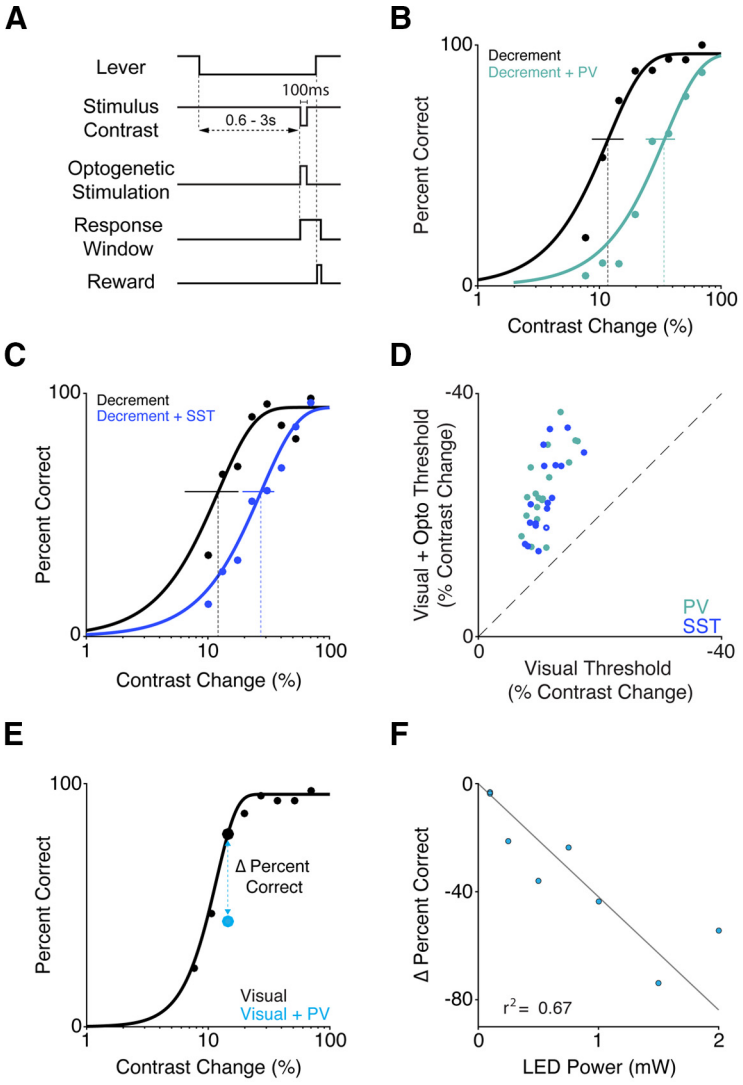

Figure 6. Brief optogenetic stimulation of either PV or SST neurons similarly impairs contrast change detection. ***A***, Trial schematic. The visual stimulus contrast is fixed at 75%, except during a 100 ms reduction. ChR2-expressing interneurons were illuminated with blue light for 100 ms concurrent with reductions in contrast on a randomly selected half of the trials. ***B***, Representative single-session performance for a PV mouse. Dots represent false-alarm corrected performance for trials with (teal) and without (black) activation of PV interneurons. Curves indicate best fitting Weibull functions that were used to determine detection thresholds (dotted vertical lines) and 95% CIs (solid horizontal lines). ***C***, Same as in ***B***, but for an SST mouse. Blue represents trials with SST stimulation. ***D***, Summary of PV and SST stimulation effects. Circles represent the contrast change detection threshold from individual sessions with (*y* axis) and without (*x* axis) PV (teal, 2 mice, 20 sessions) or SST (blue, 2 mice, 18 sessions) stimulation. Filled circles represent sessions with a significant shift in threshold (37 of 38; bootstrapped). ***E***, Representative behavioral performance from a single contrast decrement session where one 100 ms contrast decrement was paired with PV stimulation (LED power was 0.5 mW in this example session). Filled dots represent performance for trials with (blue) and without (black) activation of PV interneurons for different contrast decrements. Solid line indicates psychometric function fit to performance on visual only trials. Blue dashed line connecting dots indicates magnitude of perceptual impairment (Δ Percent Correct) on the selected contrast decrement with (blue) and without (black) optogenetic activation of PV neurons. ***F***, The perceptual impairment induced by PV activation scales with optogenetic stimulation intensity. Individual points depict performance impairment from individual sessions (there are two sessions at 0.1 mW). Gray line indicates a maximum likelihood linear regression anchored at the origin ($r^2 = 0.67$).

Spike costs force RNNs to preferentially rely on increments in spiking to detect changes in contrast

Preferential readout of spike rate increments compared with decrements might result from constraints that impact the robustness of particular decoding strategies (Gjorgjieva et al., 2014, 2019). One potentially strong constraint is the low baseline firing rates of cortical neurons. This limits the dynamic range available for spike rate decrements to render information to downstream areas and how quickly changes can be detected. As the question of how baseline spiking level influences decoding strategies adopted by neural systems is inaccessible in biological networks, we turned to artificial RNNs.

We trained RNNs ($n = 50$) to perform a contrast change detection task similar to that which we used in mice. RNNs were rewarded for detecting brief increases and decreases in contrast that occurred against an otherwise static baseline contrast (Fig.

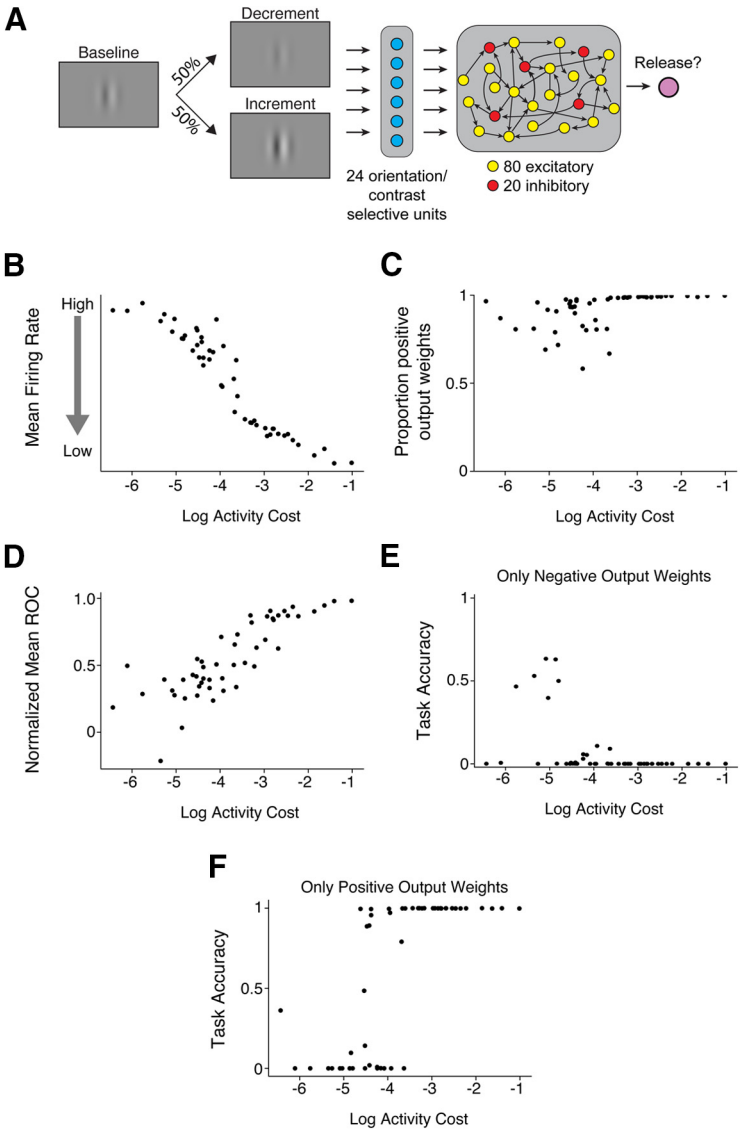

Figure 7. Activity costs force RNNs to adopt low firing rates and ignore firing rate decrements. ***A***, Schematic of a network trained to detect interleaved increases and decreases in contrast. The input is a single, oriented Gabor that changes in contrast, which is passed to 24 units that are tuned to orientation and contrast (12 excited or inhibited by contrast changes). The contrast-responsive layer passes input to a recurrent layer with 100 interconnected units (80 excitatory/20 inhibitory) that project to a single output node that controls lever releases. ***B***, Increasing activity costs (*x* axis) drive recurrent networks to lower basal rates of firing (*y* axis). Points represent the results obtained from different networks trained with different activity costs. ***C***, The proportion of output weights from the recurrent layer that are positive increases as function of activity cost. ***D***, Normalized ROC (mean of all contrast changes). Range extends from −1 (all units decrement in firing) to 1 (all units increment). ***E***, Task accuracy as a function of activity cost when all units with positive weights are removed. ***F***, Task accuracy as a function of activity cost when all units with negative weights are removed.

7*A*, randomly interleaved). The RNN architecture is described in Materials and Methods. Briefly, visual input was passed to a layer of 24 contrast and orientation sensitive units that were either excited or inhibited by changes in visual contrast (50/50 split). The split was chosen to mimic the mixture of excitatory and inhibitory responses we recorded in mouse V1 as well as to give RNNs ample opportunity to use decrements in activity to guide responses. The contrast-responsive layer projected to a recurrent layer of 100 units (80 excitatory, 20 inhibitory) and the excitatory units linearly projected onto an output unit that signaled network responses. Because it is likely that the metabolic cost of neuronal activity contributes to the low baseline spike rates in cortical networks (Attwell and Laughlin, 2001; Lennie, 2003), networks were trained using different activity costs. We used rate-based RNNs rather than spiking networks, as recent modeling work suggests that the average firing rate, rather than temporal patterning of spikes, determines the metabolic cost of neuronal activity (Yi and Grill, 2019). This approach allowed us to explore the relationship between activity cost, overall activity, and the decoding strategies used by RNNs to perform the task.

Activity costs profoundly affected the average firing rate observed in RNNs. As activity costs increased, networks adopted significantly lower average rates of firing (Fig. 7*B*; Spearman $\rho = -0.96$, $p < 10^{-27}$). Consistent with the idea that lower levels of activity impact the dynamic range available for decrements in spiking to signal stimulus changes, the proportion of positive output weights from the recurrent layer depended strongly on activity cost. High activity costs reliably caused RNNs to use exclusively positive weights (Fig. 7*C*; Spearman $\rho = 0.79$, $p = 10^{-10}$). Thus, constraints that affect overall levels of activity can shift networks toward using firing rate increments.

We next used ROC measures (see Materials and Methods) to explore how changes in RNN activity related to behavioral responses. For each network, we calculated a population ROC based on responses to contrast changes. A value of −1 indicates that all units in a network fired less when RNNs correctly detected contrast changes (hit) compared with trials in which the network failed not respond (miss). Conversely, a value of 1 indicates that activity was higher for all units in the network when the network responded. The normalized ROC value converged to 1 as activity costs increased (Fig. 7*D*; Spearman $\rho = 0.85$, $p < 10^{-14}$). This shows that detection of contrast changes were correlated with firing rate increases for high activity costs in our RNNs.

To examine the causal role for increments and decrements in activity rate in guiding responses, we presented new trials to trained networks but turned off outputs from the recurrent layer that were either positive or negative. To align with our neurophysiological data in V1, this meant that decrements in activity still exist and can influence activity within the recurrent layer, but only positive or negative outputs from the recurrent layer can impact RNN responses. Eliminating the contribution of positive output weights obliterated performance (Fig. 7*E*), suggesting that positive weights were critical for detection. However, removing the contribution of negative output weights had almost no impact on performance once activity costs became moderate (Fig. 7*F*; Spearman $\rho = 0.77$, $p < 10^{-10}$). These data show that constraining levels activity in artificial neural networks can produce decoding strategies that are consistent with those observed in the mouse: increments in neuronal activity appear to be preferentially used for detecting changes in visual stimuli.

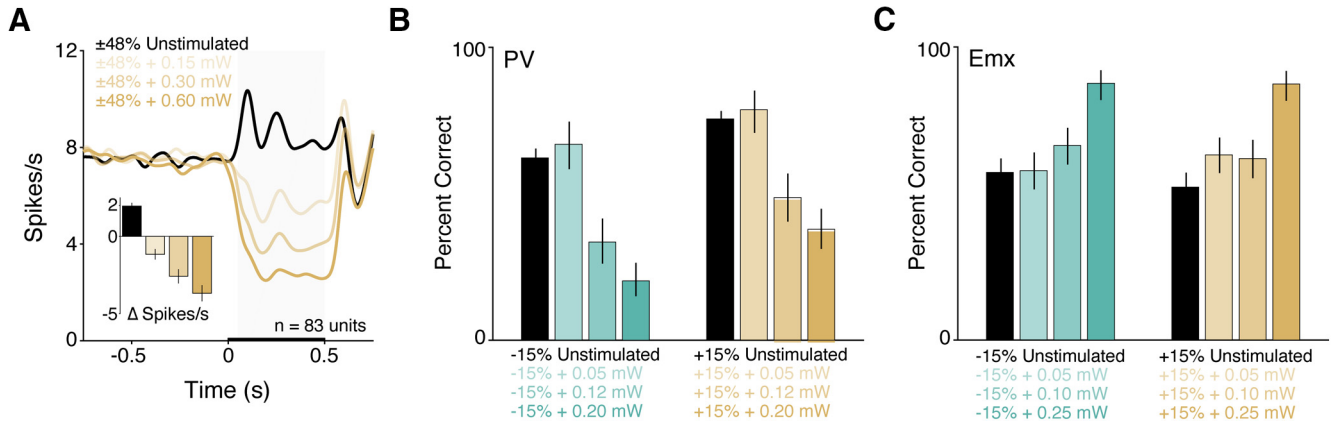

Figure 8. Optogenetic effects on V1 spiking and detection performance scale with the optogenetic stimulus power. ***A***, Population PSTH ($n = 83$ units) combined for both contrast change types (±48%, each neuron contributes two responses to each trace while the color scheme follows the increment convention used in the manuscript). Spikes were convolved with a Gaussian with $\sigma =$ 25 ms. Black line indicates the population response without optogenetic activation of PV interneurons. Line transparency corresponds to the different PV stimulation powers. Inset, Bar plot represents the evoked change in spiking for the different stimulus conditions. The strength of PV-mediated suppression significantly depends on LED power ($p < 10^{-34}$, Friedman's test). ***B***, Summary of behavioral performance for 1 PV mouse tested with different optogenetic stimulus powers. Bar plot represents the cumulative percent correct (±95% CI) for decreases (left, aqua) or increases in contrast (right, gold) presented without (black) or with optogenetic stimulation (colored) of PV neurons. Powers used = 0.05 mW, 7 sessions; 0.12 mW, 7 sessions; 0.20 mW, 8 sessions. The reduction in detection performance scales with power in the PV mouse (contrast decreases, increases both $p < 10^{-16}$, logistic regression). ***C***, Equivalent experiments in 1 Emx mouse revealed that optogenetic improvements in detection performance similarly scale with power (minimum 7 sessions per condition; contrast decreases, increases both $p < 10^{-7}$, logistic regression).

Experimental controls

Our prior work has shown that perceptual impairments produced by activation of V1 PV interneurons depend strongly on retinotopic alignment between visual and optogenetic stimuli (Glickfeld et al., 2013). In a subset of the PV mice shown in Figure 3 ($n = 3$), we collected additional behavioral sessions ($n = 22$) in which we moved the visual stimulus away from the retinotopic location of optogenetic stimulation. In these additional sessions, we offset the stimulus by 15°–20° from the location used for primary data collection ($n = 25$ sessions) while the LED power remained the same. Thus, despite comparable levels of PV neuron activation, offsetting the visual stimulus, should attenuate suppression of stimulus-evoked activity in V1. Indeed, optogenetic perturbations produced larger changes in performance when the visual stimulus was aligned with the stimulated patch of V1. This was true for both decreases (median change in percent correct, aligned: −31.2%, −0.41 to −0.17 IQR; offset: −11%, −0.21 to −0.03 IQR, $p < 10^{-4}$, Wilcoxon rank-sum test) as well as increases in contrast (median change in percent correct, aligned: −22%, −0.35 to −0.13 IQR; offset: −0.02%, −0.09 to 0.05 IQR, $p < 10^{-5}$, Wilcoxon rank-sum test). Thus, the ability of PV stimulation to reduce behavioral responses depends on alignment between the optogenetic stimulus and the visual representation in V1. This argues against the possibility that PV stimulation is somehow impairing performance by distracting the animal or impairing motor planning or execution.

Next, we sought to confirm that our observations were not specific to the power selected for stimulation. Thus, we characterized how the strength of optogenetic stimuli affected V1 spiking, and in separate experiments, perceptual reports.

In a subset of electrophysiology sessions from the PV mice shown in Figure 3 ($n = 3$; 12 recording sites), we presented only large contrast changes (±48% contrast changes from a 50% contrast static Gabor) to test our ability to augment the strongest visual responses. In addition to trials without optogenetic stimulation, we also presented three different stimulation powers (0.15, 0.3, 0.6 mW). As before, visual and optogenetic stimuli were presented for 500 ms and we quantified evoked responses by calculating the change in firing rate from 50 to 500 ms after stimulus onset relative to −500 to −50 ms before stimulus onset. As expected, we found that reductions in V1 spiking evoked by PV neuron activation significantly depended on the LED power (Fig. 8*A*; $p < 10^{-34}$, Friedman's test; *post hoc*: all comparisons $p < 0.01$, Dunn–Sidak correction). This was true when decreases or increases in contrast were tested in isolation (decreases: $p < 10^{-15}$, *post hoc*: 0.15 mW vs 0.3 mW $p > 0.05$, otherwise all other comparisons $p < 0.01$; increases: $p < 10^{-16}$, *post hoc*: 0.3 mW vs 0.6 mW $p > 0.05$, otherwise all other comparisons $p < 0.05$; both Friedman's test with Dunn–Sidak correction).

In other mice (1 Emx, 1 PV), we conducted additional behavioral sessions. These sessions were identical to those described above, except that ±15% contrast changes were paired with different optogenetic stimulation powers in different sessions. Figure 8*B* depicts the cumulative performance data (±95% CI) for the PV mouse for each optogenetic stimulus power (0.05, 0.12, 0.20 mW; min 7 sessions per power). We combined the data across sessions and performed a logistic regression that compared the probability of hits as a function of stimulation power independently for each tested contrast change. The probability of successfully detecting either decreases or increases in contrast significantly depended on the power used for optogenetic stimulation (Fig. 8*B*; decreases hit rate ± 95% CI: 0 mW = 60%, 57%–64%; 0.05 mW = 64%, 55%–72%; 0.12 mW = 33%, 25%–41%; 0.2 mW = 21%, 16%–27%; t statistic = −10.5, $p < 10^{-25}$; Increases hit rate ± 95% CI: 0 mW = 73%, 70%–76%; 0.05 mW = 76%, 68%–83%; 0.12 mW = 47%, 39%–55%; 0.2 mW = 37%, 30%–44%; t statistic = −8.34, $p < 10^{-16}$). We performed similar experiments in an Emx mouse (minimum 7 sessions per power) and found that the effects of Emx stimulation on performance similarly scaled with the optogenetic stimulus power (Fig. 8*C*; decreases hit rate ± 95% CI: 0 mW = 54%, 48%-59%; 0.05 mW = 58%, 51%-63%; 0.15 mW = 66%, 60%-72%; 0.25 mW = 83%, 77%-88%; t statistic = 5.47, $p < 10^{-7}$; Increases hit rate ± 95%: 0 mW = 49%; 43%-55%; 0.05 mW = 61%, 55%–67%; 0.15 mW = 62%, 56%–68%; 0.25 mW = 82%, 77–88%: t statistic = 5.48, $p < 10^{-7}$). In sum, effects of optogenetic stimulation on population spiking and behavioral performance scale with the power used for stimulation suggesting that we would have obtained qualitatively similar results had we used other stimulation powers in our primary experiments.

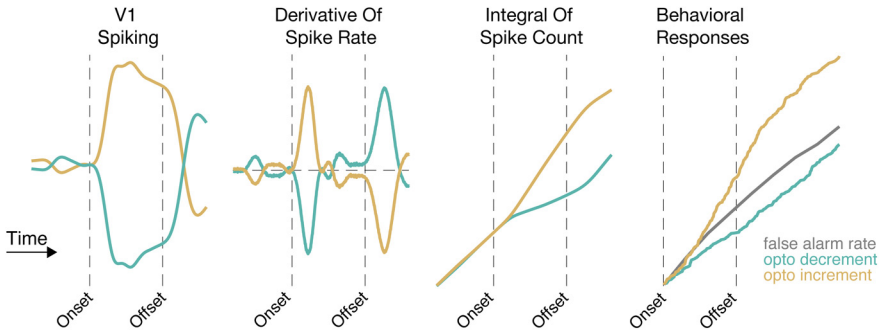

Figure 9. Summary of experimental results. Left, Optogenetic stimuli were associated with both increments and decrements in the V1 spike rate at the step onset (offset) and offset (onset). Vertical dashed lines indicate onsets/offsets of optogenetic stimuli. Second from left, Instantaneous changes in the V1 spike rate following onsets and offsets in optogenetic stimuli produce robust signals that mice could use to guide their responses. Horizontal dashed line indicates 0. Third from left, Alternatively, integration of spikes relative to the prestimulus period generates distinct signals such that any suppression of V1 spiking reduces the integrated spike count. Right, Behavioral responses triggered by changes in optogenetic input reflect the integrated spike counts, as behavioral responses were always suppressed by reductions in spiking regardless of whether reductions were followed by rapid release from inhibition.

Discussion

How do downstream brain regions decode changes in cortical spiking? Our data show that subpopulations of V1 units have firing rates that either increment or decrement in response to changes in visual contrast, likely owing to the ON/OFF juxtaposition of early stage visual receptive fields. This demonstrates that either increments or decrements (or both) in neuronal firing could convey behaviorally relevant information and thus inform decisions about visual stimuli (Fig. 1). We found that optogenetic excitation of primary neurons facilitated contrast change detection regardless of the sign of the change (Figs. 3, 5, 8). Inhibition of V1 neurons, or reductions in V1 spiking, impaired detection of contrast changes and suppressed the probability of responses in the absence of a visual stimulus (Figs. 3-6, 8). Our data suggest that perceptual reports are triggered exclusively by increments in V1 spiking. This finding corroborates a recent study from Jin and Glickfeld (2019) who showed that behavioral responses (hits and false alarms) were elevated by optogenetic excitation and suppressed by optogenetic inhibition. We hypothesize that the brain does not process reductions in cortical spiking because low baseline firing rates in cortex impacts the coding range for spike rate decrements. Indeed, we found that imposing a simple constraint that reduced overall activity levels in RNNs encouraged decoding strategies based exclusively on increases in activity (Fig. 7).

When mice were trained to respond to brief optogenetically induced increments and decrements of V1 spiking in the same neuronal populations, only steps that elevated the rate of firing above baseline were detected, although steps of either sign results in rapid increments and decrements in the rate of firing (Figs. 5, 9). Moreover, mice appeared unable to respond to rapid changes in the spike rate (e.g., the derivative) over short periods, regardless of its sign (Figs. 5, 6, 9). Rather, behavioral responses bore a striking resemblance to the integrated spike count over hundreds of milliseconds (Fig. 9, right panels). The sensory periphery readily converts stimulus onsets and offsets into robust, transient changes in firing, emphasizing spatial and temporal changes. However, it is not clear that transient changes are emphasized in all neuronal processing. The detection of signals in cortical sensory areas may rely on spike count integration to a criterion value, as in the popular drift-diffusion models of decision-making in higher-order sensory areas (Gold and Shadlen, 2002; Ratcliff and McKoon, 2008).

Our goal was to measure behavioral performance at threshold over weeks to months in the same animals and examine how behavioral performance was affected by optogenetic perturbations. For practical reasons, we recorded neuronal activity outside of the behavioral task. While population spiking profiles are certain to differ between passively viewing and behaving mice, differences are most likely to arise from levels of arousal or movement (Fu et al., 2014; Lee et al., 2014). Such state dependencies will affect the magnitude of V1 contrast responses or optogenetic perturbations. However, it is unlikely such differences would invert the sign of optogenetic perturbations on V1 spiking. Given the striking differences in how optogenetic excitation or inhibition, even in the same neuronal populations, differentially affected perceptual reports, our data argue that reductions in V1 output are not decoded by downstream areas.

Why would cortex preferentially decode from increments in neuronal spiking?

Tight coupling between excitation and inhibition keeps baseline firing rates low in cortical circuits (Wehr and Zador, 2003; Okun and Lampl, 2008; Xue et al., 2014; Moore et al., 2018). In the visual system, natural scenes generate sparse firing rates in V1 (Vinje and Gallant, 2000, 2002). Consequently, normal operating regimes limit the dynamic range available for spike rate decrements to encode information, while enhancing the encoding potential for spike rate increments. These features might be related to the metabolic cost of spiking. Energy usage by the brain depends primarily on the rate of neuronal spiking (Attwell and Laughlin, 2001; Lennie, 2003), which likely favors energy-efficient neuronal codes (Levy and Baxter, 1996; Gjorgjieva et al., 2014, 2019). Our modeling experiments support this general idea, as higher activity costs encouraged RNNs to preferentially rely on increments in neuronal activity to detect changes in visual stimuli (Fig. 7).

In the retina, bipolar cells segregate increases and decreases of luminance into distinct ON and OFF processing streams (Schiller, 1992), rather than a single, bidirectional ON (or OFF) channel. This split allows a given number of retinal ganglion cells to convey information more efficiently than a single channel, especially when the costs of spiking are considered (Gjorgjieva et al., 2014, 2019). That pharmacological inhibition of ON cells disrupts perceptual reports for light increments without affecting decrements (Schiller, 1982; Schiller et al., 1986) suggests central structures primarily decode spike increments. Recently, Smeds et al. (2019) used a transgenic mouse with differentially elevated luminance thresholds of the ON and OFF retinal ganglion cells to show that mice performing an absolute luminance detection task depend the ON pathway, even when the OFF pathway could provide greater sensitivity. In cortical visual areas beyond V1, it is unlikely that any sensory stimulus produces an overall reduction in population spiking. Population averages, even to suboptimal visual stimuli, are overwhelmingly positive across extrastriate visual areas, including V4 (Cohen and Maunsell, 2009), MT (Maunsell and Van Essen, 1983), LIP (Meister et al., 2013), and IT

(Jaegle et al., 2019). These results strongly support the idea that decoding by central visual structures depends on increases in spiking.

Implications for signal readout in V1

Primates can learn to report electrical microstimulation of all cortical areas tested (Doty, 1965, 1969; Murphey and Maunsell, 2007, 2008; Murphey et al., 2009), but detection improved markedly with extended practice (Ni and Maunsell, 2010). If mice could decode decrements in V1 output, behavioral responses to PV or SST stimulation might emerge only slowly as mice learn to use the resultant signal to guide responses. However, we never observed an increase in detection probability associated with optogenetic activation of interneurons that reduce cortical spiking (up to 32 sessions, ~4500 stimulation trials; Figs. 3, 4, 6, 8) (Cone et al., 2019). For comparison, mice can learn to report optogenetic activation of vasoactive intestinal peptide-expressing interneurons (Cone et al., 2019), which potentiate cortical spiking (Fu et al., 2014). Moreover, mice readily perceive optogenetic activation of V1 excitatory neurons, and the change in spike rate produced near threshold is low (Histed and Maunsell, 2014) (Δ1.1 spikes/unit). Our PV manipulations produced changes in spiking that were similar in magnitude but opposite in sign (Figs. 3, 6, 8). Together, our data show that no matter which cortical circuit element is used to affect V1 spiking, mice can readily act on increments in V1 spiking, whereas they appear unable to use spiking decrements to guide their responses.

Even when decrements in excitatory input are followed by sharp increases in the rate of V1 spiking, behavioral responses remain suppressed (Fig. 5). These mice were trained using sustained visual stimuli, which may have encouraged them to integrate spikes over long time windows. We encouraged a separate group of mice to use brief stimulus epochs to guide their responses. Here, strong suppression of V1 spiking with PV stimulation followed by release from inhibition consistently reduced the probability of behavioral responses (Fig. 6). Together, these data argue that mice do not decode rapid changes in cortical spiking, but instead guide their behavioral responses by integrating V1 spiking over hundreds of milliseconds. Our observations are consistent with previous work where detection of optogenetic stimulation in V1 was dependent not on the timing of excitation but on the total excitation delivered over short periods (100 ms) (Histed and Maunsell, 2014). Thus, although sensory neurons convert both the spatial and temporal edges of visual stimuli into robust, transient changes in spiking, perceptual reports appear to rely on spike count integration. While accumulation of information closely describes links between neuronal spiking and perceptual decisions in higher order brain areas (Gold and Shadlen, 2007), similar mechanisms may underlie the production of percepts throughout sensory cortex, and possibly beyond. Our data also raise important questions about the timescale of integration and whether it can be flexibly adjusted to cover longer periods. Extrapolating from studies of attention and decision-making, it is likely that windows for cortical integration are flexible (Ghose and Maunsell, 2002; Huk and Shadlen, 2005). Future work should address the periods over which perceptual integration takes place and the degree of its flexibility.

Interactions between visual and optogenetic input

How do optogenetic perturbations interact with visual signals in a V1 neuronal population? We previously showed that the effects of PV or SST stimulation on contrast detection were better explained by a divisive scaling of the stimulus contrast compared with models in which stimulation acted directly on the probability of lever releases (Cone et al., 2019). This argues against the possibility that V1 inhibition impairs performance by disrupting motor planning or by distracting the animal. Optogenetic stimulation of excitatory neurons potentiates spiking across the V1 population, which may enhance stimulus representations or produce a signal that is distinct from the visual signal, or both. Work from others suggests that artificial and natural signals are merged into a common percept (see Histed et al., 2013). Here, we delivered moderate optogenetic excitation alongside weak (or nonexistent) visual stimuli and produced subsaturating enhancements in detection performance, suggesting that the optogenetic stimulus was unlikely to be readily perceptible in isolation. Regardless of the exact mechanism by which potentiating excitation facilitates performance in our experiments, the conclusion is unchanged: adding signal that the brain can decode can increase the probability of detection.

Implications for future work

It remains to be determined how general these results are for cerebral cortex and other brain structures. Increments and decrements in spiking are likely to be equally important for signaling in other brain areas. In the cerebellum and the basal ganglia, aspects of eye position (Lisberger and Fuchs, 1978), head position (Barter et al., 2015), speed and rotation (Muzzu et al., 2018), and action initiation (Krause et al., 2010) all appear to critically rely on decrements in spike rate. Compared with cerebral cortex, these structures have relatively high baseline rates of firing, or signal using inhibitory rather than excitatory neurotransmission. It would be unsurprising if different constraints have shaped diverse sets of information processing strategies across different brain regions.

Determining how changes in spike rate mediate the functions of neural circuits is critical for understanding the brain. Currently, there is very little causal evidence describing how the brain uses changes in neuronal spiking to render information to downstream areas (but see Lerman et al., 2018; Carrillo-Reid et al., 2019; Marshel et al., 2019; Russell et al., 2019). Further work will be required to establish how changes in spiking in neuronal populations contribute to different brain computations. Nevertheless, the data presented here highlight a strong asymmetry in how rapid and sustained increments and decrements in neuronal spiking in V1 relate to perceptual reports.

References

Abadi M, Martín A, Paul B, Jianmin C, Zhifeng C, Andy D, Jeffrey D, Matthieu D, Sanjay G, Geoffrey I, Michael I, Manjunath K, Josh L, Rajat M, Sherry M, Derek GM, Benoit S, Paul T, Vijay V, Pete W, et al. (2016) TensorFlow: a system for large-scale machine learning. In: Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation, pp 265–283. Berkeley, CA: USENIX Association.

Andermann ML, Kerlin AM, Roumis DK, Glickfeld LL, Reid RC (2011) Functional specialization of mouse higher visual cortical areas. Neuron 72:1025–1039.

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

Attwell D, Laughlin SB (2001) An energy budget for signaling in the grey matter of the brain. J Cereb Blood Flow Metab 21:1133–1145.

Barter JW, Li S, Sukharnikova T, Rossi MA, Bartholomew RA, Yin HH (2015) Basal ganglia outputs map instantaneous position coordinates during behavior. J Neurosci 35:2703–2716.

Barto AG, Sutton RS, Anderson CW (1983) Neuronlike adaptive elements that can solve difficult learning control problems. IEEE Transactions on Systems, Man, & Cybernetics 13:834–846.

Blakemore C, Campbell FW (1969) On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. J Physiol 203:237–260.

Carrillo-Reid L, Han S, Yang W, Akrouh A, Yuste R (2019) Controlling visually guided behavior by holographic recalling of cortical ensembles. Cell 178:447–457.e5.

Cohen MR, Maunsell JHR (2009) Attention improves performance primarily by reducing interneuronal correlations. Nat Neurosci 12:1594–1600.

Cone JJ, Scantlen MD, Histed MH, Maunsell JHR (2019) Different inhibitory interneuron cell classes make distinct contributions to visual contrast perception. eNeuro 6:ENEURO.0337-18.2019.

Doty RW (1965) Conditioned reflexes elicited by electrical stimulation of the brain in macaques. J Neurophysiol 28:623–640.

Doty RW (1969) Electrical stimulation of the brain in behavioral context. Annu Rev Psychol 20:289–320.

Ferguson JE, Boldt C, Redish AD (2009) Creating low-impedance tetrodes by electroplating with additives. Sens Actuators A Phys 156:388–393.

Fu Y, Tucciarone JM, Espinosa JS, Sheng N, Darcy DP, Nicoll RA, Huang ZJ, Stryker MP (2014) A cortical circuit for gain control by behavioral state. Cell 156:1139–1152.

Gallant JL, Connor CE, Van Essen DC (1998) Neural activity in areas V1, V2 and V4 during free viewing of natural scenes compared to controlled viewing. Neuroreport 9:2153–2158.

Geisler WS (2011) Contributions of ideal observer theory to vision research. Vision Res 51:771–781.

Ghose GM, Maunsell JHR (2002) Attentional modulation in visual cortex depends on task timing. Nature 419:616–620.

Gjorgjieva J, Sompolinsky H, Meister M (2014) Benefits of pathway splitting in sensory coding. J Neurosci 34:12127–12144.

Gjorgjieva J, Meister M, Sompolinsky H (2019) Functional diversity among sensory neurons from efficient coding principles. PLoS Comput Biol 15: e1007476.

Glickfeld LL, Histed MH, Maunsell JHR (2013) Mouse primary visual cortex is used to detect both orientation and contrast changes. J Neurosci 33:19416–19422.

Gold JI, Shadlen MN (2002) Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. Neuron 36:299–308.

Gold JI, Shadlen MN (2007) The neural basis of decision making. Annu Rev Neurosci 30:535–574.

Goldey GJ, Roumis DK, Glickfeld LL, Kerlin AM, Reid RC, Bonin V, Schafer DP, Andermann ML (2014) Versatile cranial window strategies for long-term two-photon imaging in awake mice. Nat Protoc 9:2515–2538.

Gorski JA, Talley T, Qiu M, Puelles L, Rubenstein JL, Jones KR (2002) Cortical excitatory neurons and glia, but not GABAergic neurons, are produced in the Emx1-expressing lineage. J Neurosci 22:6309–6314.

Grinvald A, Lieke EE, Frostig RD, Hildesheim R (1994) Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. J Neurosci 14:2545–2568.

Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K (2014) Flow of cortical activity underlying a tactile decision in mice. Neuron 81:179–194.

Haefner RM, Gerwinn S, Macke JH, Bethge M (2013) Inferring decoding strategies from choice probabilities in the presence of correlated variability. Nat Neurosci 16:235–242.

Hartline HK, Wagner HG, Ratliff F (1956) Inhibition in the eye of Limulus. J Gen Physiol 39:651–673.

Hippenmeyer S, Vrieseling E, Sigrist M, Portmann T, Laengle C, Ladle DR, Arber S (2005) A developmental switch in the response of DRG neurons to ETS transcription factor signaling. PLoS Biol 3:e159.

Histed MH, Maunsell JHR (2014) Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. Proc Natl Acad Sci USA 111:E178–E187.

Histed MH, Carvalho LA, Maunsell JHR (2012) Psychophysical measurement of contrast sensitivity in the behaving mouse. J Neurophysiol 107:758–765.

Histed MH, Ni AM, Maunsell JHR (2013) Insights into cortical mechanisms of behavior from microstimulation experiments. Prog Neurobiol 103:115–130.

Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J Neurosci 25:10420–10436.

Jaegle A, Mehrpour V, Mohsenzadeh Y, Meyer T, Oliva A, Rust N (2019) Population response magnitude variation in inferotemporal cortex predicts image memorability. Elife 8:e47596.

Jin M, Glickfeld LL (2019) Contribution of sensory encoding to measured bias. J Neurosci 39:5115–5127.

Keefer EW, Botterman BR, Romero MI, Rossi AF, Gross GW (2008) Carbon nanotube coating improves neuronal recordings. Nat Nanotechnol 3:434–439.

Kohn A (2007) Visual adaptation: physiology, mechanisms, and functional benefits. J Neurophysiol 97:3155–3164.

Krause M, German PW, Taha SA, Fields HL (2010) A pause in nucleus accumbens neuron firing is required to initiate and maintain feeding. J Neurosci 30:4746–4756.

Kubota Y, Karube F, Nomura M, Kawaguchi Y (2016) The diversity of cortical inhibitory synapses. Front Neural Circuits 10:1–15.

Lee AM, Hoy JL, Bonci A, Wilbrecht L, Stryker MP, Niell CM (2014) Identification of a brainstem circuit regulating visual cortical state in parallel with locomotion. Neuron 83:455–466.

Lennie P (2003) The cost of cortical computation. Curr Biol 13:493–497.

Lerman GM, Gill JV, Rinberg D, Shoham S (2018) Precise optical probing of perceptual detection. bioRxiv. https://doi.org/10.1101/456764.

Levy WB, Baxter RA (1996) Energy efficient neural codes. Neural Comput 8:531–543.

Lisberger SG, Fuchs AF (1978) Role of primate flocculus during rapid behavioral modification of vestibuloocular reflex: I. Purkinje cell activity during visually guided horizontal smooth-pursuit eye movements and passive head rotation. J Neurophysiol 41:733–763.

Madisen L, Mao T, Koch H, Zhuo JM, Berenyi A, Fujisawa S, Hsu YW, Garcia AJ, Gu X, Zanella S, Kidney J, Gu H, Mao Y, Hooks BM, Boyden ES, Buzsáki G, Ramirez JM, Jones AR, Svoboda K, Han X, et al. (2012) A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. Nat Neurosci 15:793–802.

Maffei L, Fiorentini A, Bisti S (1973) Neural correlate of perceptual adaptation to gratings. Science 182:1036–1038.

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.

Martinez-Conde S, Macknik SL, Hubel DH (2000) Microsaccadic eye movements and firing of single cells in the striate cortex of macaque monkeys. Nat Neurosci 3:251–258.

Masse NY, Yang GR, Song HF, Wang XJ, Freedman DJ (2019) Circuit mechanisms for the maintenance and manipulation of information in working memory. Nat Neurosci 22:1159–1167.

Maunsell JHR, Van Essen DC (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey: II. Binocular interactions and sensitivity to binocular disparity. J Neurophysiol 49:1148–1167.

Meister ML, Hennig JA, Huk AC (2013) Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. J Neurosci 33:2254–2267.

Moore AK, Weible AP, Balmer TS, Trussell LO, Wehr M (2018) Rapid rebalancing of excitation and inhibition by cortical circuitry. Neuron 97:1341–1355.e6.

Movshon JA, Lennie P (1979) Pattern-selective adaptation in visual cortical neurones. Nature 278:850–852.

Müller JR, Metha AB, Krauskopf J, Lennie P (1999) Rapid adaptation in visual cortex to the structure of images. Science 80:1405–1408.

Murphey DK, Maunsell JHR (2007) Behavioral detection of electrical microstimulation in different cortical visual areas. Curr Biol 17:862–867.

Murphey DK, Maunsell JHR (2008) Electrical microstimulation thresholds for behavioral detection and saccades in monkey frontal eye fields. Proc Natl Acad Sci USA 105:7315–7320.

Murphey DK, Maunsell JHR, Beauchamp MS, Yoshor D (2009) Perceiving electrical stimulation of identified human visual areas. Proc Natl Acad Sci USA 106:5389–5393.

Muzzu T, Mitolo S, Gava GP, Schultz SR (2018) Encoding of locomotion kinematics in the mouse cerebellum. PLoS One 13:e0203900.

Nagel G, Szellas T, Huhn W, Kateriya S, Adeishvili N, Berthold P, Ollig D, Hegemann P, Bamberg E (2003) Channelrhodopsin-2, a directly light-gated cation-selective membrane channel. Proc Natl Acad Sci USA 100:13940–13945.

Ni AM, Maunsell JHR (2010) Microstimulation reveals limits in detecting different signals from a local cortical region. Curr Biol 20:824–828.

Okun M, Lampl I (2008) Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. Nat Neurosci 11:535–537.

Pfeffer CK, Xue M, He M, Huang ZJ, Scanziani M (2013) Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. Nat Neurosci 16:1068–1076.

Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput 20:873–922.

Resulaj A, Ruediger S, Olsen SR, Scanziani M (2018) First spikes in visual cortex enable perceptual discrimination. Elife 7:1–22.

Russell LE, Yang Z, Tan PL, Fiş ek M, Packer AM, Dalgleish HW, Chettih S, Harvey CD, Häusser M (2019) The influence of visual cortex on perception is modulated by behavioural state. biorxiv. https://doi.org/10.1101/706010.

Sadeh S, Silver RA, Mrsic-Flogel TD, Muir DR (2017) Assessing the role of inhibition in stabilizing neocortical networks requires large-scale perturbation of the inhibitory population. J Neurosci 37:12050–12067.

Schiller PH (1982) Central connections of the retinal ON and OFF pathways. Nature 297:580–583.

Schiller PH (1992) The ON and OFF channels of the visual system. Trends Neurosci 15:86–92.

Schiller PH, Sandell JHR, Maunsell JHR (1986) Functions of the ON and OFF channels of the visual system. Nature 322:824–825.

Smeds L, Takeshita D, Turunen T, Tiihonen J, Westö J, Martyniuk N, Seppänen A, Ala-Laurila P (2019) Paradoxical rules of spike train decoding revealed at the sensitivity limit of vision. Neuron 104:576–587.e11.

Song HF, Yang GR, Wang XJ (2016) Training excitatory-inhibitory recurrent neural networks for cognitive tasks: a simple and flexible framework. PLoS Comput Biol 12:e1004792.

Taniguchi H, He M, Wu P, Kim S, Paik R, Sugino K, Kvitsiani D, Kvitsani D, Fu Y, Lu J, Lin Y, Miyoshi G, Shima Y, Fishell G, Nelson SB, Huang ZJ (2011) A resource of Cre driver lines for genetic targeting of GABAergic neurons in cerebral cortex. Neuron 71:995–1013.

Tsodyks M, Skaggs WE, Sejnowski TJ, McNaughton BL (1997) Paradoxical effects of inhibitory interneurons. J Neurosci 17:4382–4388.

Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. Science 287:1273–1276.

Vinje WE, Gallant JL (2002) Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1. J Neurosci 22:2904–2915.

Wehr M, Zador AM (2003) Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. Nature 426:442–446.

Wiegert JS, Mahn M, Prigge M, Printz Y, Yizhar O (2017) Silencing neurons: tools, applications, and experimental constraints. Neuron 95:504–529.

Wilson NR, Runyan CA, Wang FL, Sur M (2012) Division and subtraction by distinct cortical inhibitory networks in vivo. Nature 488:343–348.

Xue M, Atallah BV, Scanziani M (2014) Equalizing excitation-inhibition ratios across visual cortical neurons. Nature 511:596–600.

Yi G, Grill WM (2019) Average firing rate rather than temporal pattern determines metabolic cost of activity in thalamocortical relay neurons. Sci Rep 9:6940.

Yizhar O, Fenno LE, Davidson TJ, Mogri M, Deisseroth K (2011) Optogenetics in neural systems. Neuron 71:9–34.