Amplified cortical neural responses as animals learn to use novel activity patterns

Bradley Akitake*, Hannah M. Douglas*, Paul K. LaFosse, Ciana E. Deveau, Anna J. Li, Lauren N. Ryan, Samuel P. Duffy, Zhishang Zhou, Yanting Deng, and Mark H. Histed.
* denotes equal contribution

Unit on Neural Computation and Behavior, National Institute of Mental Health Intramural Program, National Institutes of Health, United States

Abstract

Cerebral cortex supports representations of the world in patterns of neural activity, used by the brain to make decisions and guide behavior. Past work has found diverse or limited changes in the primary sensory cortex in response to learning, suggesting the key computations might occur in downstream regions. Alternatively, sensory cortical changes may be central to learning. To study changes in sensory cortical representations, we trained mice to recognize entirely novel, non-sensory patterns of cortical activity in the primary visual cortex (V1) created by direct optogenetic stimulation. As animals learned to use these novel patterns, we found their detection abilities improved by an order of magnitude or more. The behavioral change was accompanied by large increases in V1 neural responses to fixed optogenetic input. Neural response amplification to novel optogenetic inputs had little effect on existing visual sensory responses. Amplification would seem to be desirable to improve decision-making in a detection task, and thus these data suggest adult cortical plasticity plays a significant role in improving the detection of novel sensory inputs during learning.

Introduction

Sensorimotor decision-making involves patterns of neural activity, which propagate through the neural circuits of many brain areas, and are changed by those circuits. The sets of neural computations involved in sensory decision-making have not been fully determined (for reviews, see Chung and Abbott, 2021; Gold and Shadlen, 2007; Niell and Scanziani, 2021; Wang, 2012), but some principles have been identified. One basic neural computation is representation, storing information about the sensory world in patterns of activity. Another is decision, or readout, in which representations are transformed or categorized by circuits into forms suitable for action (Egger and Lisberger, 2022; Wu et al., 2020). This readout process for object identification might, for example, lead to neurons in areas downstream from primary sensory cortex that encode category identity (Freedman et al., 2003), with neurons' firing indicating whether a given sensory stimulus is food or stone; predator or friend.

There is substantial evidence that sensory cortical synapses can be modified by activity (Frégnac and Shulz, 1999; He et al., 2006; Hengen et al., 2013; Malenka and Bear, 2004; Sawtell et al., 2003), but it is less clear whether cortical response changes constitute the computational change that leads to improved behavior with learning. Studies in humans and animals have reported varied effects of learning on visual cortical responses, including increased activity after visual training (Bao et al., 2010; Li et al., 2008; Schoups et al., 2001; Schwartz et al., 2002), selective suppression of activity (Ghose et al., 2002), decreased variability of visual selectivity response properties after training (Goltstein et al., 2013, 2021; Poort et al., 2015), and activity changes that disappeared once early learning has ended (Yotsumoto et al., 2008). Some learning studies have found improvement in primary sensory representations (Goltstein et al., 2021; Henschke et al., 2020; Jurjut et al., 2017; Marshel et al., 2019), along with changes in anticipatory and other signals (Khan et al., 2018; Poort et al., 2015). Other studies have found little task-relevant change in the primary visual cortex, area V1 (but found changes in higher visual areas like V4; Boynton and Finney, 2003; reviewed by Dosher and Lu, 2017; Ghose et al., 2002; Yang and Maunsell, 2004). Thus, it has been unclear whether a major substrate of visual sensory learning is representational improvement in V1, such as increased gain or selectivity, or whether the principal changes are readout changes, perhaps in downstream areas.

One reason it has been difficult to delineate the neural computations underlying sensory decisions and sensory learning is that neurons and brain areas are highly interconnected, and sensory stimuli change activity in many brain areas (e.g. Schmolesky et al., 1998; Steinmetz et al., 2019; Zatka-Haas et al., 2021). Thus, changes in responses to sensory stimuli that are observed in one cortical area may be inherited from input regions, and indeed cognitive factors like attention or arousal can modulate visual activity before it arrives at the cortex (Liang et al., 2020). One way to isolate cortical representations from downstream readout computations is to use stimulation-based behavioral paradigms. Using electrical or optogenetic stimulation methods, entirely novel (non-sensory, or "off-manifold", Jazayeri and Afraz, 2017; Sadtler et al., 2014) activity patterns can be introduced in a chosen brain region. With training, subjects can successfully use stimulation of a wide variety of brain areas to make behavioral decisions (Doty, 1969; Histed et al., 2013; Kesner and Wilburn, 1974; O'Connor et al., 2013). Using such novel

patterns is a way to explore the limits of cortical plasticity, as they are dissimilar from normal sensory patterns — and because optogenetics can induce novel patterns at the level of one cortical area, changes seen in that cortical area are unlikely to be entirely inherited from changes that occur in input regions.

Studying learning is a powerful way to understand the computations behind decision-making, because the neural computations used to perform a task must change to create improved performance on that task. Here, to isolate representational changes that occur as animals improve on a task, we insert a new cortical representation into mouse V1 using optogenetic stimulation. We train mice to use that signal for behavior, and we study how responses change during learning. We first find that animals show dramatic improvements in behavior as they learn this task. Detection thresholds computed via psychophysical curves can change over several orders of magnitude during weeks or months of learning, implying animals require much lower stimulation power, and thus smaller neural activation, as they learn. We then examine how neural responses to the optogenetic stimuli change during learning and find that cortical neurons produce larger responses to the same optogenetic input as learning progresses. This amplification effect does not occur for similar optogenetic stimuli when they are delivered outside the context of the behavioral task. Thus, learning enables a fixed input to produce an increasingly large response in the V1 network, presumably by some adjustment of local, recurrent circuitry (Goldman, 2009; Hennequin et al., 2012; Murphy and Miller, 2009). In sum, we find that learning causes local changes in representations by increasing amplification in V1.

Results

We trained animals to perform an all-optogenetic detection task and measured cortical responses with two-photon imaging. To gain optical access to the primary visual cortex, we implanted a 3mm optical glass window over V1. We expressed a calcium indicator (jGCaMP 7s or 8s) via a virus that expresses in cortical neurons (AAV9-hSyn; Dana et al., 2019; Kügler et al., 2003; Zhang et al., 2020), and expressed the opsin ChrimsonR (targeted to cells' soma; stChrimsonR, Pégard et al., 2017) in layer 2/3 excitatory cells via a second virus (Cre-dependent, FLIP/DIO, Emx1-Cre mouse line; Cardin et al., 2010; Sohal et al., 2001). In this way, we could image visual cortex neurons with a two-photon microscope (Figure 1a) while delivering optogenetic stimulation to activate stChrimsonR-expressing neurons. We stimulated opsin expressing neurons with light pulses delivered at the cortical surface through the window (~500 µm diameter spot; Methods; Histed and Maunsell, 2014).

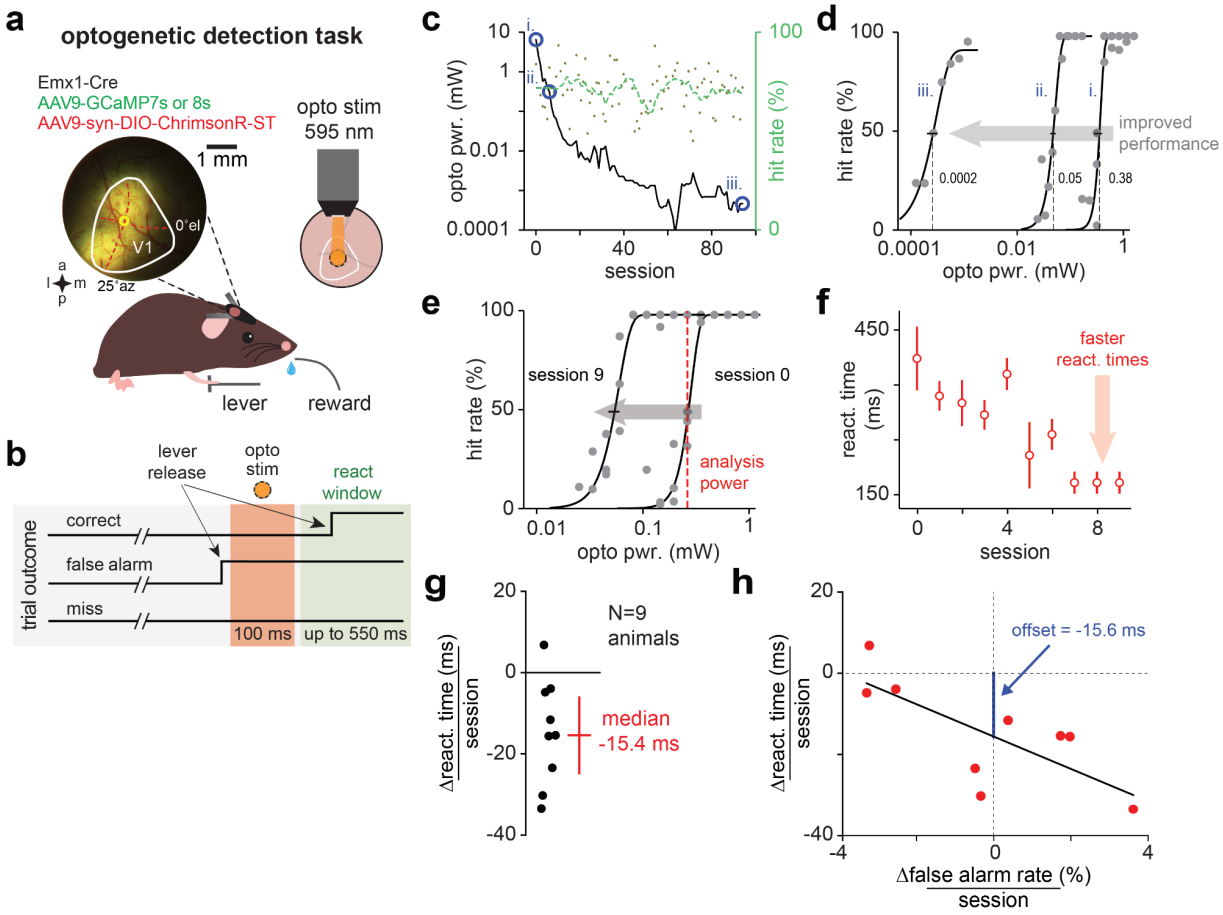

Figure 1 - Optogenetic learning: Mice use direct optogenetic stimulation of V1 neurons and improve their performance in an optogenetic-only perceptual detection task. (**a, b**) Schematic of the optogenetic detection task with simultaneous *in vivo* 2-photon imaging. Animals were trained to release a lever when they detected an optogenetic stimulus (soma-targeted opsin stChrimsonR in excitatory cells; Emx1-Cre animals; 595 nm light) delivered through the imaging objective. Lever releases that occurred during a reaction time window (> 50 ms and up to 550 ms post stim) were scored as correct responses. Lever releases before or after the reaction time window were scored as false alarms and misses respectively. For optogenetic behavior without imaging, the stimulus was delivered through a fixed cemented fiber optic cannula (see Methods). (**c**) Example of long-term optogenetic learning showing an initial fast drop in the threshold of stimulation power (i and ii, blue circles, sessions 0-9), followed by a longer phase of continued behavioral improvement (points from ii to iii, sessions 10-90). We adjusted stimulation power as animals improved with learning, using lower stimulation power to do the task as average hit rate stayed roughly constant (overall hit rate > 70%). (**d**) Psychometric curves from points i, ii, and iii (blue circles in 1c), showing a decrease in the threshold power in the first 10 sessions (0.38 and 0.05 mW, curves i and ii). The power threshold of the final session was three orders of magnitude below the first session threshold (0.38 and 0.0002 mW, curves i and iii, respectively). (**e**) Psychometric curves for the initial phase of optogenetic learning. Leftward shift in the curves indicated improved performance (gray arrow). Red dotted line indicates a common power across sessions used to compare reaction times. (**f**) Distribution of the mean reaction times for the first 10 sessions of optogenetic learning in one animal shows a clear trend to faster reaction times. (**g**) Analysis of reaction times across N = 9 animals shows a decrease over sessions 3 to 9, where common powers were available (median reaction time change = -15.4 ms, IQR = 18.6, $p < 0.01$, N = 9). (**h**) The change in reaction time was negative even when the small false alarm rate changes (x-axis, ±4% false alarm rate change) were controlled for. A linear fit showed a negative vertical offset, indicating changes in reaction time arise from learning and not merely from changes in animals' strategy or perceptual criterion.

We trained animals (N = 17; Methods) to perform an optogenetic detection task. Training occurred in two phases (Figure S1a,b). First, animals learned to respond to the appearance of a monocular Gabor visual stimulus (14˚ FWHM), of varying contrast, until they performed the task with a stable psychometric threshold for three sessions (animals used for imaging: training time 15-29 days, 23.6 ± 6.2 days, mean ± SEM, N = 3 animals). Animals were then trained to report the optogenetic stimulus (Figure 1; Figure S1). We moved the visual stimulus to the retinotopic location of the imaging and stimulation site (see Methods and Goldbach et al., 2021 for retinotopic mapping), and added an optogenetic stimulus (0.5 mW at 595 nm), delivered at the same time as the visual stimulus. Over the course of several sessions, we removed the visual stimulus gradually. We reduced visual stimulus contrast manually (cf. Dalgleish et al., 2020), making it more difficult to perform the task using the visual stimulus, but keeping performance at approximately the same level as animals began to rely on the optogenetic stimulus (Figure 1a,b; Figure S1a,b). The first trial where contrast of the visual stimulus was zero, when animals relied entirely on the optogenetic stimulus (2.3 ± 0.9 days after first optogenetic stimulus, mean ± SEM, animals used for imaging, N = 3), we labeled session 0. We confirmed that animals responded only to the trained optogenetic-evoked neural activity by moving the optogenetic spot during behavior to non-training locations within V1 which resulted in no behavioral responses (Figure S2a,b).

While the optogenetic stimuli we use produce a different pattern of responses than visual inputs, which activate specific cells based on their receptive field properties, our optogenetic stimulation modulates firing rates just as visual inputs do. Cortical neurons *in vivo* generally fire in an asynchronous irregular way, due to large amounts of recurrent input that lead to highly fluctuating membrane potentials (Ahmadian and Miller, 2021; Brunel, 2000; Sanzeni et al.,

2022). Inputs then modulate the firing rate (Destexhe and Paré, 1999; Histed and Maunsell, 2014; Sanzeni et al., 2020) of the neurons — whose individual spike times are determined by the network-driven membrane potential fluctuations (Mainen and Sejnowski, 1995). The upshot is that our optogenetic inputs, which we apply to many cortical neurons simultaneously for 100 ms, increase the firing rates of the cells over that period. However, the cells fire according to near-Poisson statistics, and individual spike times are not synchronized (confirmed with this opsin via physiological recordings *in vivo*, O'Rawe et al., 2022).

Optogenetic learning in a detection task

We found that animals dramatically increase their ability to detect the optogenetic stimulus – that is, their ability to detect the activation of V1 neurons — with practice. We collected psychometric curves during training sessions to track changes in animals' perceptual sensitivity to the optogenetic stimulus (example from 1 animal in Figure 1c). Over the course of long-term training (~90 sessions), we found animals were able to consistently detect the stimulus well (Figure 1c, green dots indicate hit rate), and with practice animals' perceptual thresholds dropped dramatically over sessions (Figure 1c, black line indicates psychometric thresholds). That is, the animals got better at detecting the stimulus: they needed less-strong stimulation over time to achieve the same level of performance. The observed rate of threshold change could be roughly separated into two phases, a phase that occurred within the initial ~10 sessions of training after acquisition of the optogenetic task (example in Figure 1c,d: i and ii, sessions between 2 and 8) and a slower phase over many additional sessions (example in Figure 1c,d: ii and iii, sessions between 8 and 92). Below, we focus on the first six days of this initial learning phase for our experiments examining neural activity changes. In this initial phase, the changes in the thresholds of the psychometric curves were large (example animal in Figure 1e, gray arrow shows change in threshold, Δ thresh. pwr. = -0.28 mW, first threshold: 0.35 mW, 95% CI [0.31 - 0.37] mW, second threshold: 0.058 mW, 95% CI [0.052 - 0.063] mW). The threshold changes were accompanied by decreases in the mean reaction times. We compared reaction times for fixed stimulation powers across days (Figure 1f,g, median = -15.4 ms, IQR = 18.6, $p < 0.01$, over a subset of animals, N = 9, which had common stimulation powers from sessions 3 to 9). The changes in reaction times could not be accounted for by changes in animals' false alarm rates (Figure 1h); while reaction times did change with false alarm rates, as expected if underlying perceptual criterion was fluctuating, reaction time changes remained after regressing out false alarm rate (see negative offset at x = 0 in Figure 1h). Therefore, we found that optogenetic learning results in an improvement of perceptual sensitivity, which allows the detection of lower stimulation powers over days.

Responses of V1 to optogenetic stimulation are amplified by learning

We studied cortical neural responses during learning by measuring activity with 2-photon imaging while simultaneously stimulating. We measured neural responses in layer 2/3 during the first 6 optogenetic learning sessions, where learning is rapid (Figure 1c-f and Figure 2a). During this period animals' showed a greater than 50% drop in their optogenetic detection thresholds (Figure 2a, average change between session 0 and 5, -62 ± 10%, mean ± SEM, N = 4; thresholds extracted from psychometric functions collected in each session, as in Figure 1d).

2-photon imaging revealed strong and even co-expression of GCaMP- and stChrimsonR-bearing neurons in our dual-virus injection preparation (Figure 2b, green and red stained cells respectively, see Methods).

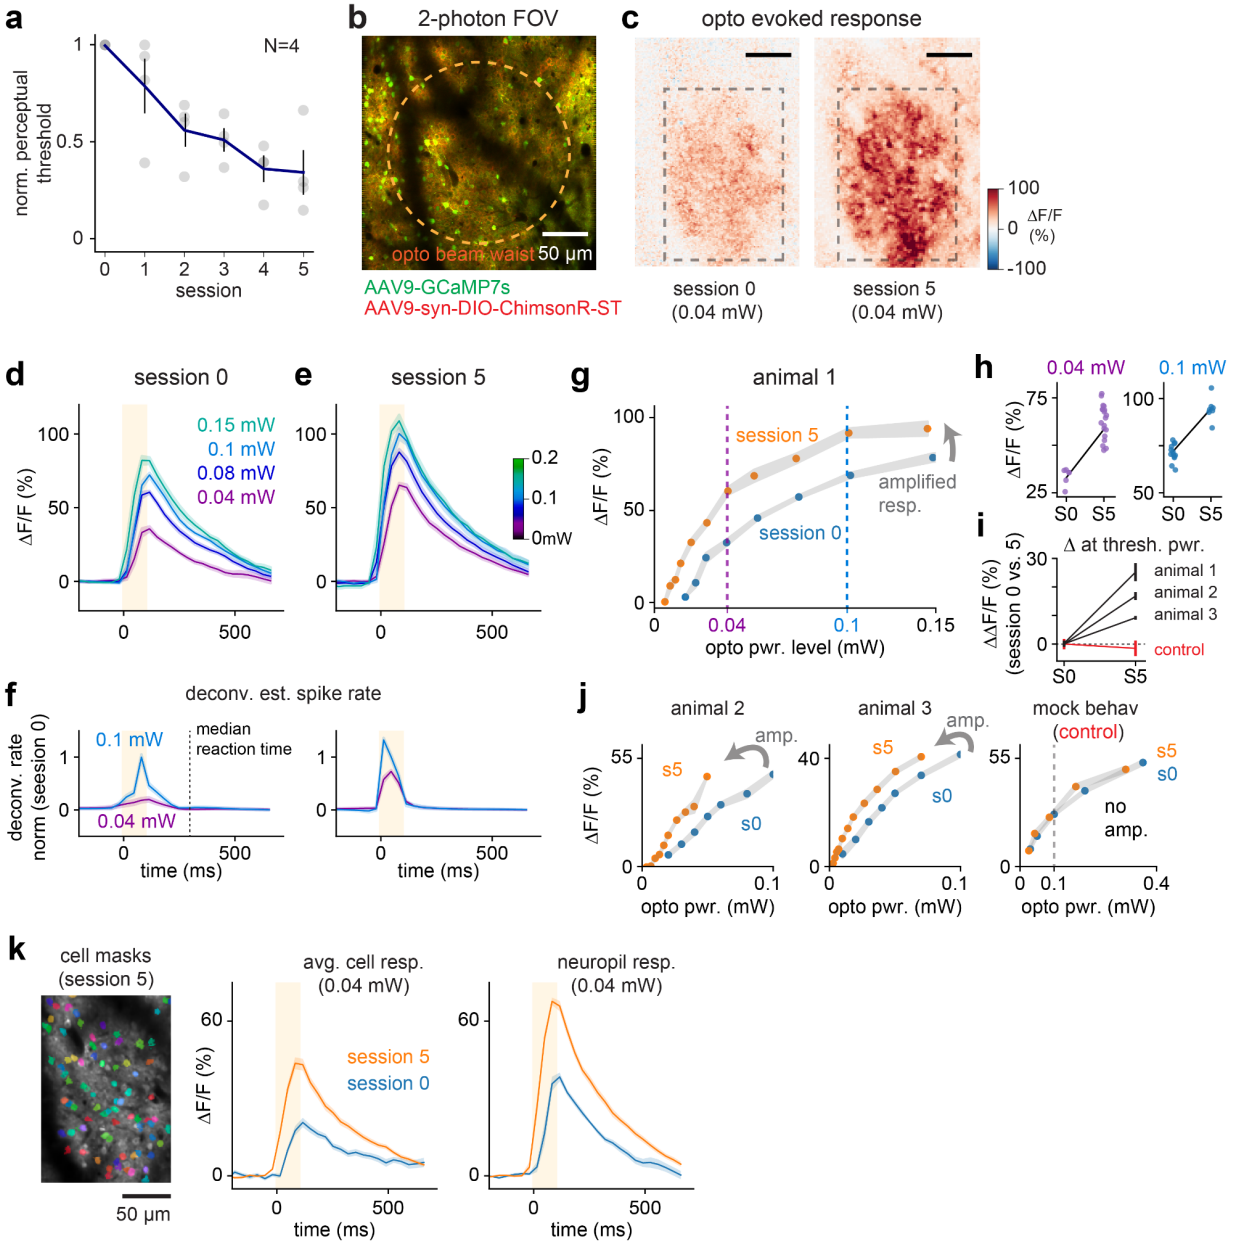

Figure 2 - Responses of V1 to optogenetic stimulation are amplified by learning. (**a**) Optogenetic power thresholds over session across the population of animals imaged during learning (normalized to session 0, N = 4) shows improved task performance. (**b**) Example imaging field of view (animal 1, see also Figure S2) where we performed simultaneous *in vivo* 2-photon imaging and optogenetic stimulation in layer 2/3 neurons. Soma-targeted opsin stChrimsonR and GCaMP7s expression are shown. Approximate beam waist of the optogenetic stimulation path is shown (~200 μm, orange dotted circle; see Methods). (**c**) Response amplification in animal 1 before and after optogenetic learning at a stimulation power near psychometric detection threshold (0.04 mW). Grey dashed box shows the region of interest (ROI) in the frame used for ΔF/F analysis. (**d, e**) ΔF/F time courses show GCaMP response before and after learning (session 0 and session 5). (**f**) Deconvolved spiking shows amplification before and after learning (session 0 and session 5). (**g**) Average ROI ΔF/F response across all tested power levels in animal 1 before and after learning. Shift in the response curve indicates amplification (gray arrow). (**h**) Plots of all ΔF/F responses for near-threshold and above-threshold stimulation powers (animal 1, dotted lines shown in panel g, 0.04 and 0.1 mW respectively) during session 0 and session 5. (**i**) Positive response changes (ΔΔF/F, before versus after learning, mean ± SEM) at near-thresh power for all animals, compared to mock behavioral control (shown in red). (**j**) Responses across all tested power levels for animals 2 and 3 (N = 3 total, one animal in panel (a) not imaged), and the mock behavioral control animal. Leftward shift in the response curve indicates amplification (gray arrow). (**k**) Suite2p-extracted cell masks (left panel, animal 2, session 5). Comparison of the change in average cells versus neuropil responses before and after optogenetic learning.

Imaging of the stimulated region during optogenetic detection behavior revealed clear stimulus-evoked responses that were strongly amplified over the course of training (Figure 2c-k). This amplification could not be explained by shifts in the imaging plane or by changes in virus expression over sessions (Figure S3). It also could not be explained by tissue growth under the window or other optical degradation over time, or because of stimulation, as the effect we measured was in the opposite direction: an increase in responses to stimulation. However, to verify that optical changes did not account for the effects, we measured the effects of stimulation within each preparation at the imaging plane while not imaging cells, and used it to adjust stimulation power, finding that the amplification effects remained with and without this adjustment (Figure S4). As another check to rule out effects of imaging properties or expression contributing to this effect, we stimulated in a control animal (see also Figure 3 for similar control experiments using even higher powers) using matched mock training sessions, with the same imaging, stimulation, reward, optical window, and injection parameters as during training. In these mock behavioral sessions, with timing and stimulation power statistics matched to the learning sessions, optogenetic stimuli were delivered with no lever responses recorded, and reward was given on a randomized set of trials (Figure 2i,j, red labeled control; Methods). We found no amplification in this closely matched control (Figure 2i-j), arguing that the amplification we saw was indeed an increase in neural responses as a function of learning.

We characterized how overall V1 responses changed over learning by plotting the timecourses of evoked activity over a large analysis region (or region of interest, ROI, shown for one animal as dashed box in Figure 2c). We found that after learning, optogenetic-evoked responses were amplified at all optogenetic power levels, with strong effects both near the psychometric threshold and also at above-threshold optogenetic stimulation powers, where animals attain their highest behavioral performance (example for animal 1, Figure 2g,h, threshold power dotted purple, above threshold dotted blue; Figure 2i,j summary of N = 3 animals). Though the magnitude of these changes varied somewhat across animals, in all animals we measured we

found positive between-session amplification (thresholds from all sessions, all animals shown in Figure S5).

Since the 1-photon optogenetic stimulus activates neurons both above and below the imaging plane we sought to determine the relative contribution of cells and neuropil responses to the amplification that we measured. We extracted cell masks using suite2p (Pachitariu et al., 2017) and found that both cells and neuropil showed substantial amplification with learning (Figure 2k). The neuropil signal likely reflects both signals from local imaged neurons and signals from out-of-plane cells that increase their responses to the optogenetic stimulation. Therefore, while the cells we imaged were in layer 2/3, the neuropil changes reflect that cells in deeper layers may also be stimulated and participate in the amplification effects.

Overall, the changes we observed in neural activity were smaller than the improvements seen in perception. Animals' perceptual detection performance improved, and thresholds decreased, by a factor of approximately 2.7x after 6 sessions (i.e. power threshold was 37 ± 11%, mean ± SEM of session 0 levels; Figure 2a). However, the changes in ΔF/F over the course of 6 sessions, measured at a fixed stimulation power, were smaller (a 1.7x increase in ΔF/F over the large ROIs: session 0, 25.6 ± 7.4%, mean ± SEM across animals, session 5, 42.9 ± 10.9%, Figure 2i; and a 2.1x increase in cell peak ΔF/F, Figure 2k, 20.7% to 43.7%). Several caveats apply: the readout mechanism presumably sums across large numbers of neurons and thus may not be limited by the change in cortical responses we measure, and opsin saturation at high power may lead to greater changes in power than activity. However, the fact that behavior changes by a larger factor than cortical responses could potentially indicate that there is an improvement in the readout mechanism, occurring along with the amplification changes we see.

The largest neural response changes happened from one day to the next, not within-session

While we observed significant increases in ΔF/F responses across experimental days, we found no evidence of increases within-session. In fact, we found a small decrease in responses to stimulation over the course of each experimental day (average ΔF/F change over 100 trials: -1.2% ΔF/F, 95% CI [-0.9 to 1.6]% ΔF/F, coeff less than zero at $p < 10^{-13}$, via linear regression over trials numbered within day, estimated across animals and sessions, N = 3; see Methods for details). Thus, it appears that optogenetic learning-related changes do not happen within the behavioral day, i.e. from one trial to the next. Instead, these data support that the major changes to neural responses occur outside of training, and may be driven by consolidation: changes in the brain in the hours between the experimental sessions.

No amplification occurs with stimulation outside of the behavioral learning context

To determine if cortical amplification is dependent on learning, or might arise from repeated optogenetic stimulus alone, we performed a stimulation control, in a mock behavioral context, and found no amplification (Figure 2i,j, 'control'). That mock experiment was conducted with stimulation powers matched to those used during optogenetic learning (up to 0.5 mW). To determine if we could drive changes using stronger optogenetic stimulation, we increased

stimulation power levels up to twice that used for behavior. In water-restricted but non-behaving animals, we provided repeated optogenetic stimulation using a range of powers up to 1 mW (100 ms stimulation with ~6 s interpulse interval, 1200 and 1500 repetitions, N = 2; Methods). Even with higher stimulation powers we observed no changes in the optogenetic sensitivity of cells in the stimulated regions (Figure 3a,b).

This result suggests that amplification in response to these novel “off-manifold” stimuli requires an associative context, and that the amplification we observe arises from adult neuronal plasticity due to optogenetic learning.

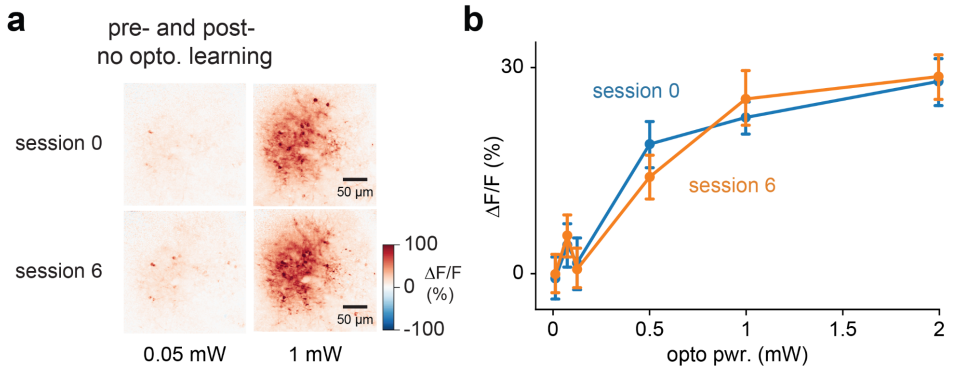

Figure 3 - No amplification occurs for optogenetic stimulation delivered to V1 outside the learning context. (**a**) ΔF/F responses in an example animal to 0.05 and 1 mW of optogenetic stimulation delivered outside the behavioral context; animal was awake and alert but any motor responses were not reinforced (see Methods). 0.5 mW is near the average post-learning threshold power for optogenetic learning animals. 1 mW is a power level three times higher than the maximum used in training optogenetic learning animals. (**b**) No amplification occurs at any power level over seven sessions of optogenetic stimulation (example animal, session 0, blue, to session 6, orange, mean ± SEM). There was no significant change in response across N = 2 animals (session 0 vs. 6, via ANOVA/linear regression; t = 1.1, p = 0.27, also neither animal reaches significance alone, and treating power as a continuous or log-continuous variable did not change the results; see Methods for regression details).

Statistics of visual responses are unchanged after optogenetic learning at both the training and control sites

Previous studies suggest that learning in visual perceptual tasks can lead to changes in the tuning properties of responsive neurons in mouse V1 (Goltstein et al., 2021; Khan et al., 2018). However, it remains unresolved if these changes arise from plasticity in the local cortical networks or if changes may be inherited from thalamic input pathways that could in principle adjust input strength, state, or synchrony (Cano et al., 2006; Hubel and Wiesel, 1962; Kelly et al., 2014; Sadagopan and Ferster, 2012) to change cortical responses. Since optogenetic stimulation bypasses feedforward input from the thalamus, we wished to determine if the visual response properties of neurons would change with optogenetic learning.

To test how optogenetic learning and its associated amplification in V1 affects visual response properties, we performed 2-photon imaging of V1 as mice were shown a series of visual stimuli before and after optogenetic learning. We used drifting gratings, either full-field or monocular

Gabor patches (15˚ FWHM; Figure 4a; Methods). We collected the responses of neurons at both the optogenetic training location (an FOV to which the visual stimulus was retinotopically matched), and an adjacent control location in V1 where stimuli were not delivered for optogenetic learning (Figure 4a).

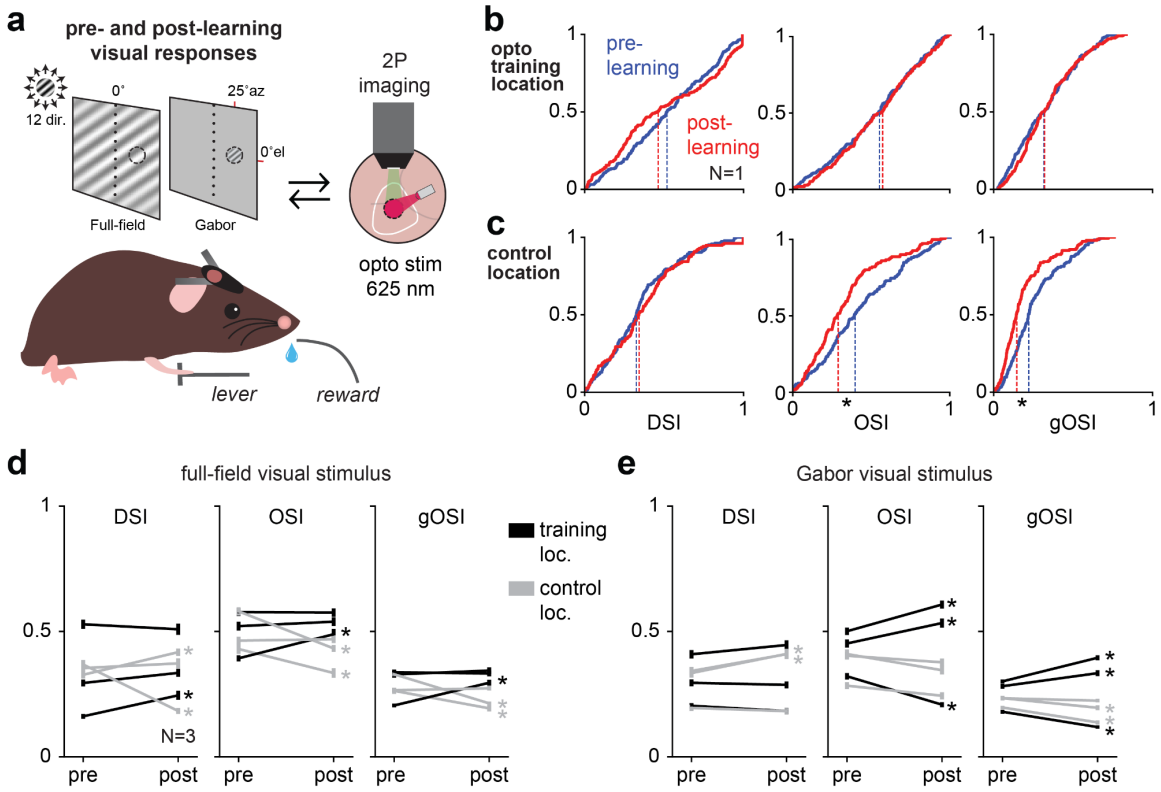

Figure 4 - Visual response properties are unchanged after optogenetic learning at both a training and a control site. (**a**) Schematic of testing visual response properties by measuring visual responses before and after learning (N = 3 animals, different animals than shown in Figure 2). Visual responses to 12-direction full-field drifting gratings or small Gabor patches (drifting gratings masked by a 2-d gaussian, FWHM 12˚) were tested before (pre-) and after (post-) optogenetic learning. (**b,c**) Example distributions of indices for direction selectivity (DSI), orientation selectivity (OSI), and global orientation selectivity (gOSI; Methods) for full-field stimulus, at training and control locations, before and after learning for one animal. (* distributions significantly different via Kolmogorov-Smirnoff 2-sample test: p-values: training location, DSI: 0.06, OSI: 0.90, gOSI: 0.21; control location, DSI: 0.27, OSI: 7.8 x 10^{-4}, gOSI: 5.7 x 10^{-6}). (**d,e**) Summary of all visual response properties for full-field or gabor stimuli, pre- and post- optogenetic learning (mean ± SEM). Training locations shown in black, control locations shown in gray (N = 3 animals). (*: $p < 0.05$, Mann-Whitney U with Holm-Šídák correction for multiple comparisons; number of significant changes either full-field or Gabor: training locations (black), DSI: 1, OSI: 4, gOSI: 4; control locations (gray), DSI: 3, OSI: 3, gOSI: 4).

We found some small but statistically significant changes in the distributions of direction selectivity, orientation selectivity, and global orientation selectivity (DSI, OSI, and gOSI, respectively) at both the optogenetic training and control locations before and after learning (CDFs from one example animal, Figure 4b,c; statistical tests measure differences in medians, see Figure 4 legend). These changes, however, were inconsistent across animals and comparable in size between the training and control locations (Figure 4d,e), perhaps due to

329 learning-related global changes in the cortical network, or perhaps due to representational drift
330 over time (Deitch et al., 2021; Marks and Goard, 2021). These results therefore suggest that,
331 while learning seems to allow the local cortical network to amplify the artificial stimulus,
332 underlying visual response distributions and the overall structure of existing sensory
333 representations remain intact.

Discussion

In this work we examine the capacity of adult mouse primary visual cortex (V1) to undergo plastic changes in response to novel optogenetic stimuli, over a few days of learning. We found clear evidence that neural responses to novel stimuli — optogenetic inputs applied directly to many cells — are amplified in V1, but only if those stimuli are made behaviorally-relevant. The changes in neurons' responses over learning sessions mirrored the animals' perceptual improvements. Responses to visual stimuli, which were not relevant for learning, did not show systematic changes, suggesting that the layer 2/3 cortical network was able to selectively amplify the input pattern created by optogenetic stimulation. Taken together, our results provide evidence for substantial plastic changes in the primary visual cortex of the adult mouse brain that are linked to perceptual learning of a completely novel stimulus.

Amplification is a desirable representational change for a perceptual detection task

In an optogenetic detection task, the principal neural computation that must be performed is a comparison between the activity evoked by optogenetic stimulation and spontaneous, ongoing activity. Therefore, the amplification of the optogenetic signal we found, an increasingly large spiking response to fixed input, seems to be the optimal way for the V1 recurrent network to adjust to improve task performance (assuming no major changes in the noise or variability in the population, Moreno-Bote et al., 2014). Amplification leads to increased separation, in population activity space, of spontaneous activity and the stimulation-evoked response. This increase in separation can also enhance downstream readout or classification, and thus would lead to improved behavioral performance.

A few other studies have found evidence for learning-related changes in the sensory cortex with optogenetic-stimulation tasks. Using a discrimination task and stimulating neurons in the somatosensory cortex (S1), Pancholi et al. (2021) found no evidence for amplification but did see other changes, including increases in response sparsity. In the visual cortex, Marshel et al. (2019) trained animals to report activation of specific neural ensembles activated with two-photon holographic stimulation. They found evidence for amplification in two different subnetworks (defined by intrinsic visual responses), as animals learned to discriminate optogenetic activation of those two subnetworks. Amplification appears to be the optimal neural change for those animals to improve task performance: the animals were asked to compare two different patterns of activity, and so amplifying each would increase separation in population activity space, improving readout. The different effects seen in Pancholi et al. might be due to structural differences between V1 and S1 cortical circuits, or may be related to differences in task-specific computations, as their subjects were asked to discriminate stimulation intensity, rather than discriminate patterns of activity. In sum, our work adds to the evidence that cortical areas can, over the course of several days of practice, adjust their input-output transformations to create representations that allow more sensitive perceptual performance.

Readout changes and representational changes

Our results appear to help resolve a contradiction in recent optogenetic stimulation studies. Dalgleish et al. (2020) found that animals detected as few as 40-100 randomly-selected

somatosensory (S1) cortical neurons in an optogenetic stimulation task. Mice can also detect similarly-sized ensembles of stimulated neurons in the olfactory bulb (Gill et al., 2020). Work by Marshel et al. (2019), on the other hand, found that most animals (6 of 7 trained) did not effectively report activation of similarly-sized groups of randomly-selected neurons, only correctly reporting ensembles already defined by visual co-activation. While a possible explanation may be differences between S1 and the olfactory bulb vs V1, our data suggest a different explanation: that detection of randomly-selected ensembles of ~100 neurons requires initial learning with stronger stimulation. Both Gill et al. and Dalgleish et al. initially trained animals using one-photon (widefield) optogenetics, the same approach we use here to produce strong activation at the initial powers used for learning. Both studies note, as we show quantitatively here, that practice with these one-photon optogenetic stimuli leads to less and less stimulation being required for the task. Thus, these optogenetic results, along with electrical stimulation studies (Doron and Brecht, 2015; Doty, 1969; Histed et al., 2009, 2013; Ni and Maunsell, 2010) imply that animals can use completely novel, randomly-chosen patterns of neural stimulation — but to do so, learning must first be induced by strong stimulation of hundreds of neurons or more.

While we found significant changes in cortical representations during learning, it is possible that the readout mechanism improves as well. Our data might suggest there are changes in readout, beyond V1 changes in amplification, as we found larger improvements in behavioral performance than in cortical responses (percent changes in stimulation power needed to do the task vs. percent changes in neural responses; Figures 1 and 2), though interpretation is difficult due to potential opsin saturation and potential nonlinear or variability-dependent readout (Kohn et al., 2016; Moreno-Bote et al., 2014). Dalgleish et al. also provide evidence that readout changes occur in optogenetic-learning tasks: they found that high detection performance generalized across different stimulated patterns of cortical neurons. That is, after learning, animals did well at detecting the activation of not just a single trained subset of up to 100 neurons, but many different sets of up to 100 neurons. On the other hand, Marshel et al., who also stimulated randomly selected groups of up to approximately 100 neurons, found little generalization from one randomly-selected pattern to the next (their Figure 4I). Some possible differences amongst these three studies that could affect the results include differences in cortical area: visual vs. somatosensory cortex, and difference in behavioral task: single-pattern detection vs. two-pattern discrimination. Taken together, these studies leave open the possibility that the decoding mechanism can change during optogenetic learning.

The learning that we observed here seems likely to be a change in sensory properties and not related to changes in movements. Our animals were pre-trained on a visual detection task before introducing the optogenetic stimulus (Figure S1a,b). Thus, the task demands, and motor responses were fixed, and the only learning step needed was for animals to gain the ability to perceive and report the novel optogenetic activity induced in the cortex.

The amplification shows signs of consolidation, with the largest changes outside sessions

Because we measured neural responses during task performance, we were able to determine whether amplification happened within the training sessions or developed from one day to the next. We found that within-session, there were small or negative changes in neural responses to a fixed stimulus (Results section, “The largest neural response changes happened from one day to the next…”), though there were consistent changes from one learning session to the next (Figure 2; all data shown in Figure S5). While some decreases in response within-session could, in principle, be due to bleaching of opsin or indicator, the changes from one session to the next suggests that the major cortical network changes were happening outside sessions, perhaps as animals rested or slept. This is reminiscent of the consolidation that happens in motor learning, where a significant component of the motor improvement also appears to occur outside of the actual learning or practice repetitions (Brashers-Krug et al., 1996; Krakauer and Shadmehr, 2006).

Our physiological recordings found learning-related neural changes over the initial few days of optogenetic learning (5-6 days), consistent with previous reports (Dalgleish et al., 2020, their Figure 2, Supplementary Figure 3; Marshel et al., 2019, with holographic stimulation; and Pancholi et al., 2021, in S1). However, we also measured continued improvement in optogenetic detection performance (without neural imaging) over many weeks to months of training (Figure 1). It seems possible that additional cortical amplification happens during this longer phase as well. This is supported by past studies of long-term deafferentation, which have demonstrated that cortical responses can change over months or years to accommodate input changes (Gilbert and Li, 2012; Pons et al., 1991).

Pattern amplification in cortex due to recurrent connectivity

We found that cortical responses to a fixed optogenetic stimulus are amplified over several days as animals’ behavioral performance improves. We also found that optogenetic learning produced little change in the visual response properties of targeted neurons (Figure 4). In principle, the observed increase in cortical responses to the optogenetic stimulus could have arisen from changes outside the local cortical network that would not be due to modification of recurrent connections. These outside sources might be changes in top-down, higher-order thalamic (e.g. from the lateral posterior nucleus, LP / pulvinar) or neuromodulatory input that change the gain of V1 neurons. However, were top-down input changes or neuromodulatory effects the dominant players, we might expect effects on visual responses as well. Theoretical work also shows that response amplification to a fixed input can be created in recurrent networks by adjusting the synaptic connectivity within the network (Goldman, 2009; Murphy and Miller, 2009; Sadeh and Clopath, 2020). Along with the timecourse of the changes we saw — over several days of practice and with changes occurring outside the session — these observations suggest that changes in local recurrent cortical synapses are a likely mechanism for the learning-related neural changes we observed.

Given this, what circuit mechanism might gate, or enable, cortical recurrent plasticity, to allow changes during behavior but not for inputs presented outside a behavioral context? There is substantial evidence that inhibitory modulation is involved when such cortical network changes occur (Carcea and Froemke, 2013; Fagiolini and Hensch, 2000; Fagiolini et al., 2004; He et al., 2006; Heimel et al., 2011; Keck et al., 2013; Swanson and Maffei, 2019; Trachtenberg, 2015) and alternation of perineuronal networks, which surround many inhibitory neurons, participate in these synaptic changes (Balmer et al., 2009; Banerjee et al., 2017; Gu et al., 2016; Hylin et al., 2013; Le Naour et al., 2001; Reichelt et al., 2019; Sorg et al., 2016). Since the response changes we observed are dependent on animals performing a rewarded behavioral task, a compelling possibility is that task context or reward prediction signals trigger inhibitory neurons, which enables plasticity to begin.

Conclusion

How the cerebral cortex builds sensory representations for use in behavior is key to understanding brain function. Though the adult visual cortex is less plastic than the developing cortex (Desai et al., 2002; Hensch, 2005; Katz and Crowley, 2002), our results – cortical amplification in response to completely novel artificial patterns of optogenetic input – provide key insights into how brains can adapt to behaviorally-relevant sensory information throughout our lifetimes.

Author contributions

B.A., H.D., P.K.L, L.R., and S.D. collected behavior and imaging data, with the help of Y.D. and A.L.. B.A., H.D., P.K.L., C.D., and M.H. performed data analysis. H.D., A.L., and Z.Z. prepared optical windows and did virus injections. B.A., H.D., P.K.L, and M.H. designed the experiments. B.A., H.D., P.K.L, and M.H. wrote the manuscript.

Acknowledgements

We thank Victoria Scott for assistance with breeding and husbandry, and A. Afraz and members of the Histed laboratory for comments and discussion. This work was supported by the NIH Intramural program (ZIAMH002956) and NIH BRAIN Initiative (U19NS107464 and U01NS108683).

Competing Interests

The authors report no competing interests.

Key resources table

Reagent type (species) or resource	Designation	Source or reference	Identifiers	Additional information
Genetic reagent (*M. musculus*)	Emx1-Cre	The Jackson Laboratory	RRID:IMSR_JAX:005628	17 total animals
Recombinant DNA reagent	AAV9-hSyn-FLEX-GCaMP6f	Addgene	#100833	
Recombinant DNA reagent	AAV9-hSyn-jGCaMP7s	Addgene	#104487	
Recombinant DNA reagent	AAV9-hSyn-jGCaMP8s	Addgene	#162377	
Recombinant DNA reagent	AAV1-hsyn-FLEX-ChrimsonR-tdTomato	Addgene	#62723	
Recombinant DNA reagent	AAV9-hSyn-DIO-stChRimsonR-mRuby2	Addgene	#105448	

Methods

Animals

All experimental procedures were approved by the NIH Institutional Animal Care and Use Committee (IACUC) and complied with Public Health Service policy on the humane care and use of laboratory animals. Emx1-Cre mice, (Cre-recombinase targeted at the Emx1 locus, Gorski et al., 2002, Jax stock no. 005628, N = 17), were used for all experiments. N = 13 animals were used for optogenetic behavior without imaging (Figure 1c-h, Figure 2a), N = 3 for optogenetic behavior plus simultaneous 2-photon imaging (Figure 2), N = 3 for mock behavior with optogenetic stimulation only (N = 1, Figure 2j; N = 2, Figure 3), and N = 3 for visual stimulation before and after optogenetic behavior (Figure 4). Animals were housed on a reverse light/dark cycle.

Cranial window implantation and viral injection

Mice were given intraperitoneal dexamethasone (3.2 mg/kg) and anesthetized with isoflurane (1–3% in 100% O_2 at 1 L/min). Using aseptic technique, a titanium headpost was affixed using C & B Metabond (Parkell) and a 3 mm diameter craniotomy was made, centered over V1 (−3.1 mm ML, +1.5 mm AP from lambda).

Mice were injected with a pre-mixed combination of two adenovirus-mediated (AAV9) vectors for expression in the cortex, a functional calcium indicator (AAV9-hSyn-jGCaMP7s or -jGCaMP8s, viral titers 3.0 x 10^{13} and 4.1 x 10^{13} GC/ml respectively, final dilution 1:10) construct and a photoactivatable soma-targeted opsin construct (AAV9-hSyn-stChrimsonR-mRuby2, viral titer 3.2 x 10^{13} GC/ml, final dilution 1:8). Injections were made 150-250 µm below the surface of the brain for expression in layer 2/3 neurons. Multiple 300 nL injections were done at 150 nL/min to achieve widespread coverage across the 3 mm window.

A 3 mm optical window was then cemented into the craniotomy, providing chronic access to the visual cortex. Post-surgery, mice were given subcutaneous 72 hr slow-release buprenorphine (0.5 mg/kg) and recovered on a heating pad. Virus expression was monitored over the course of 3 weeks. We selected animals with good window clarity and high levels of virus co-expression (GCaMP and stChrimsonR) for behavior and imaging experiments.

Retinotopic mapping

We determined the location of V1 in the cranial window using a hemodynamic intrinsic imaging protocol previously described in (Goldbach et al., 2021). Briefly, we delivered small visual stimuli to head-fixed animals at different retinotopic positions and measured hemodynamic-related changes in absorption by measuring reflected 530 nm light. Imaging light was delivered with a 530 nm fiber-coupled LED (M350F2, Thorlabs). Images were collected through a green long-pass emission filter onto a Retiga R3 CCD camera (QImaging Inc., captured at 2 Hz with 4 × 4 binning). The hemodynamic response to each stimulus was calculated as the change in reflectance of the cortical surface between the baseline period and a response window starting

2–3 s after stimulus onset. We fit an average visual area map to the cortex based on the centroids of each stimulus' V1 hemodynamic response.

These retinotopic maps were used during behavioral training to overlap the visual stimulus position in the right monocular hemifield with the imaging/optogenetic stimulation location in the V1. We found that the transition period between visual detection and optogenetic detection was facilitated by a strong overlap.

For measuring visual response properties, we further refined the visual position by measuring cellular responses in layer 2/3 with 2-photon imaging. Small, oriented noise visual stimuli (14˚ FWHM) were presented at 9 locations (spaced by ±15˚ azimuth and ±10˚ elevation) in the right visual hemifield. The visual stimulus position that evoked the greatest response in the FOV was chosen for characterizing visual responses. We found that the strongest response was typically the center location, selected using the widefield hemodynamic map above.

Behavioral task

Mice were head-fixed and trained first to hold a lever and release in response to a visual stimulus (Gabor patch; 14° FWHM, spatial frequency 0.1 cycle/degree), that increased contrast relative to a gray screen (Goldbach et al., 2021; Histed et al., 2012), and then to an optogenetic stimulus that directly activated layer 2/3 neurons in V1. Mice initiated behavioral trials by pressing and holding a lever for 400-4000 ms (according to a geometric distribution, to reduce variation in the stimulus appearance time hazard function, see Goldbach et al., 2021, their Figure 2), and then the stimulus appeared for 100 ms in the animal's right monocular hemifield. Animals had up to 550 ms to report the stimulus by releasing the lever. Because some minimum time is required to process the stimulus, we counted as false alarm trials those releases that occurred within 50-100 ms of the stimulus onset. Correct detection responses resulted in delivery of a 1-3 µL liquid reward.

All behavioral animals were first trained on a visual detection task (see task schematic, Figure S1, and Goldbach et al., 2021). Once animals were performing well on the visual task and produced stable psychometric curves with low lapses for three consecutive sessions, we transitioned the animal to using the optogenetic stimulus by pairing each visual stimulus appearance with a fixed power (0.5 mW) optogenetic stimulation. During these transition sessions we lowered the contrast of the visual stimulus until animals could perform the task without the visual stimulus. The session where animals started behaving exclusively on the optogenetic stimulus was denoted session 0. During session 0 we generated the first psychometric curve for optogenetic stimulation. Analysis of data from session 0 came only from the part of trials where the animal was exclusively on the optogenetic stimulus. Subsequent behavioral sessions were started and conducted with only optogenetic stimuli.

Optogenetic stimulation

For optogenetic behavior experiments without simultaneous 2-photon imaging we delivered light through a fiber aimed at the cortical surface (Goldbach et al., 2021). A fiber-coupled LED light source (M625F2, Thorlabs, peak wavelength 625 nm) was coupled via a fiber patch cable to a

fiber optic cannula (400 µm core diameter, 0.39 NA, Thorlabs CFMLC14L02) cemented above V1. This method was used for long-term learning and control experiments with increased optogenetic stimulation outside of behavior (powers up to 1mW with 6.3 ± 1.7s between simulations, mean ± SD, N = 2).

For optogenetic behavior experiments conducted with simultaneous 2-photon imaging we activated stChrimsonR expressing neurons by passing 595 nm light (CoolLED pE4000) through the imaging objective to the surface of the brain. The illumination power was measured through the objective at the beginning of each session using a light meter (Newport 1918-C with a 918D-SL-OD3R detector) with a maximum of ~0.5 mW.

Analysis of behavioral data

Analyses were conducted in Matlab and Python. Optogenetic learning effects were characterized by analyzing data collected during animal behavior on the optogenetic stimulation detection task.

Reaction times were averaged across trials for each laser power group and for each training session. Linear fits were calculated for these data points across the start and end sessions in which each laser power group was present during the task. The slope of the linear fit indicated the change in reaction time per session for each laser power group. A mean change in reaction time per training session was then calculated across all laser powers for each animal. Changes in optogenetic detection sensitivity were analyzed by fitting cumulative Weibull functions to data from individual training sessions to estimate detection performance (hit rate) as a function of laser power. Quantifying thresholds with d' (sensitivity) produces similar results to using hit rate in this task, as false alarm rates are nearly constant over time (false alarm hazard rate is near constant, see Goldbach et al., 2021). Threshold was the 50% point of the Weibull functions.

2-photon calcium imaging

2-photon calcium imaging was conducted using a custom microscope based on MIMMS (Modular In vivo Multiphoton Microscopy System, e.g. Kerlin et al., 2019) components (Sutter Instruments, Novato, CA) with a Chameleon Discovery NX tunable femtosecond laser (Coherent, Inc.; Santa Clara, CA). Imaging was performed using a 16X water dipping objective (Nikon; Tokyo, Japan). A small volume of clear ultrasound gel (~1 mL) was used to immerse the lens. Images of calcium responses (~150-200 µm from the surface of the pia, layer 2/3) were acquired at 30 Hz using ≤ 50 mW laser power for static imaging, and ≤ 15 mW for behavior at 920 nm.

Analysis of imaging data

Raw 2-photon image stacks were downsized (512 rows to 256 rows) to facilitate handling of large datasets. For each behavioral session, frames were motion corrected using CaImAn (Giovannucci et al., 2019). Each imaging data set was baseline corrected to an estimated minimum pixel intensity, calculated as the minimum value in the average projection image

across all frames from all trials prior to stimulus presentation (F_{min}, a scalar). The minimum pixel intensity was subtracted from all pixels and all resulting negative values were set to 0.

For quantitative analyses we computed $\Delta F/F$ as $(F-F_0)/F_0$ at each pixel. F_0 was taken over the 10 frames before each stimulus onset, and F_0 did not systematically change over days (see also Figure S3). For statistical analyses F was taken as the frame 120 ms after the stimulus onset (frame 3 post-stimulation, near the peak response). For visual display of responses in entire frames, as in Figure 2c, F was taken over 0-270 ms after stimulus onset (frames 0-9 post-stimulation), and we computed $\Delta F/F$ as $(F-F_0)/F_{div}$, where F_{div} is F_0 smoothed with a gaussian filter (sigma = 20 pix). Using a smoothed divisor image averages overall intensity in small regions of the image, yielding a form of local contrast adaptation. Image ROI fluorescent (F) activity traces were measured by calculating the average pixel intensity within a user-defined ROI, prior to computing $\Delta F/F$ for an ROI. Deconvolved calcium responses to estimate spiking activity for an ROI were calculated using the OASIS method with an autoregressive constant of 1 (Friedrich et al., 2017).

Segmented cell masks were identified using either Suite2p (for Figure 2; Pachitariu et al., 2017) or CaImAn (for Figure 4; Giovannucci et al., 2019) and their resulting calcium responses (F) were extracted. To quantify neuropil activity, we constructed a single mask such that all pixels belonging to an identified cell mask were excluded, and all remaining pixels were included.

Linear regression model for testing for effects of change within experimental day was OLS regression, using all trials on which the stimulus was successfully detected. Data was from N = 3 animals, N = 6 sessions for each animal, and 2633 total number of stimulation trials (all animals and sessions are shown in Supp. Figure 5, leftmost 3 columns). Regression model equation: $\Delta F/F$ ~ C(animal) * C(session) + stimulation_power_mw + trial_number + constant, where C(x) signifies a categorical or dummy variable. Full details of the model definition are in https://patsy.readthedocs.io/en/latest/. We also tested for significant change in $\Delta F/F$ within-session by running the same model over each animals' data, and found all three animals showed a negative change (trial number coefficient: -1.5, -1.1, -0.2% $\Delta F/F$) though only two were significantly different from zero ($p < 1 \times 10^{-12}$, $< 1 \times 10^{-6}$, = 0.6, respectively).

Linear regression model for testing effects of optogenetic stimulation outside of behavior (results in Figure 3) was OLS regression from N = 2 animals, session 0 (S0) vs. session 6 (S6) via ANOVA. Regression model equation: $\Delta F/F$ ~ C(power) + C(S0 v. S6), where C(x) indicates a categorical or dummy variable.

Confirming optogenetic stimulation power between sessions

We measured the power of the stimulation LED light path immediately before each behavioral session. We also measured relative laser excitation power across days by measuring light collected by the PMTs during stimulation. The optogenetic blanking circuit operates the LED illuminator during the flyback phase of scanning image acquisition, and the refractory time of the blanking circuit leaves an up to ~20 pixel artifact at the edges of the raw image stacks that scales with stimulation intensity. We used the mean pixel intensity change for this artifact to scale attenuated sessions and normalize stimulation powers across days (Figure S4), and our

results were unchanged with and without this scaling, confirming we accurately measured stimulation power.

Analysis of visual response properties

2-photon calcium imaging was performed directly before and after optogenetic learning to assess V1 neural responses at both training and control locations (an area with stable expression at least 200 μm away from training location). Visual stimuli were presented on a monitor positioned in front of the head-fixed animal at a 45˚ angle on the animal's right side. The visual stimulus was either a full-field or Gabor patch (15˚ FWHM) drifting grating stimulus at 100% contrast presented in 12 different directions (30˚ increments). Stimuli were presented for 3 second durations (with 4 seconds between presentations) and were delivered in random order for a total of 25 repetitions of each stimulus direction. Gabor patch stimuli were displayed on the monitor at the visual field location corresponding to the retinotopic map at the training and control locations.

To assess potential changes in visual response selectivity, direction and orientation selectivity indices were calculated for each identified cell (Kondo and Ohki, 2016; Swindale, 1998). First, tuning curves for each cell were calculated by averaging ΔF/F responses across the 3 second stimulus period across all repetitions for each of the 12 drifting grating directions. Direction selectivity indices (DSI) were measured as $(R_{pref} - R_{oppo})/(R_{pref} + R_{oppo})$, where R_{pref} is the peak average response across the 12 directions and R_{oppo} is the average response at the opposite direction 180˚ away from the preferred direction. Orientation selectivity indices (OSI) were measured by first averaging responses from opposite pairs of directions (e.g. 0˚ and 180˚, 45˚ and 225˚) and calculating $(R_{pref} - R_{ortho})/(R_{pref} + R_{ortho})$, where R_{pref} is the peak average response across the 6 orientations, and R_{ortho} is the average response of the orthogonal orientation 90˚ away from the preferred orientation. Last, a global OSI (gOSI) metric was calculated as 1 - CV (tuning curve) for each cell, where CV is the circular variance.

Data availability

The datasets generated during the current study are available from the corresponding author on reasonable request. Data with plotting code are available at: https://github.com/histedlab/

Supplemental Figures

a

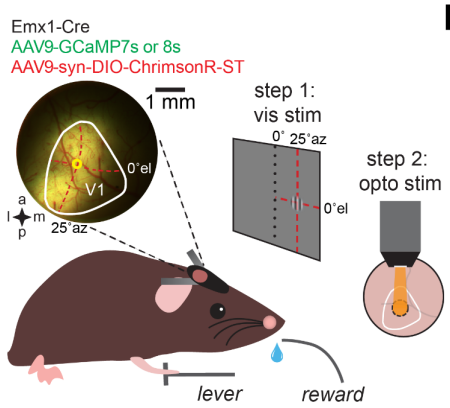

b

training schedule (~21 days post injection)

step	days	stim type	task
1	1 - 20†	visual	train visual detection task
transition*	21-23	visual + opto	paired visual + opto stim (last paired day is session 0)
2*	24-30	opto	opto stim only (sessions 1 to 5)

† stable psychometric threshold for three consecutive sessions
* sessions with 2-photon imaging

Figure S1 - Training timeline for the optogenetic detection task, Related to Figures 1 and 2. (**a**) Schematic of 2-step protocol for behavioral training first on visual stimulus (step 1) then on optogenetic stimulation (step 2). The optogenetic stimulation location was aligned to the retinotopic location of the visual stimulation in V1. (**b**) Typical behavioral training schedule outlining the length of time for visual detection task proficiency and the steps to transition animals from the visual to the optogenetic stimulus (other statistics in Results). Visual detection proficiency was determined by animals achieving a stable psychometric threshold for three consecutive sessions (†). 2-photon imaging was conducted during the transition and step 2 sessions (*).

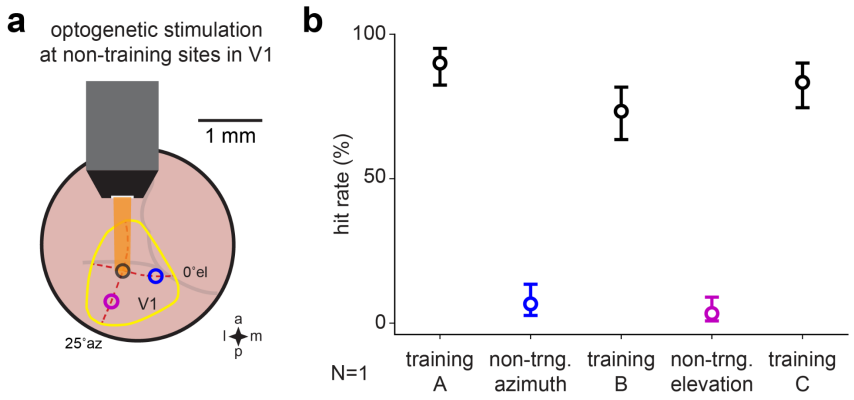

Figure S2 - Animals detect and use the optogenetic-induced cortical activity; they do not detect stray light with their retinas, Related to Figure 1 and 2. (**a**) Schematic of experiment where we moved the stimulation light spot a small amount and found dramatic changes in behavior. This implies that animals' behavior depends on cortical neural optogenetic activation. Black circle indicates optogenetic training location in V1 (yellow outline). After collecting a psychometric curve at the training location, we moved the optogenetic stimulation ~500 μm along the cortex, both in the visual-map-defined azimuth and elevation meridians (red dotted lines). At each of the shifted locations, blue and magenta circles, behavioral performance dropped and was recovered when we moved the stimulation back to the training location. (**b**) Detection hit rates in a trained animal during a session where the optogenetic stimulation location was sequentially moved for 30 trials each to and from non-trained locations in V1 (black, training A: 90.0 CI [82.4 - 95.1]%, training B: 73.3 [63.5 - 81.65]%, training C: 83.3 CI [74.5 - 90.1]%, blue, non-training azimuth change, 6.7 CI [2.7 - 13.4]%, magenta, non-training elevation change, 3.3 CI [0.7 - 9.0]%, hit rate ± Wald CI, N = 1).

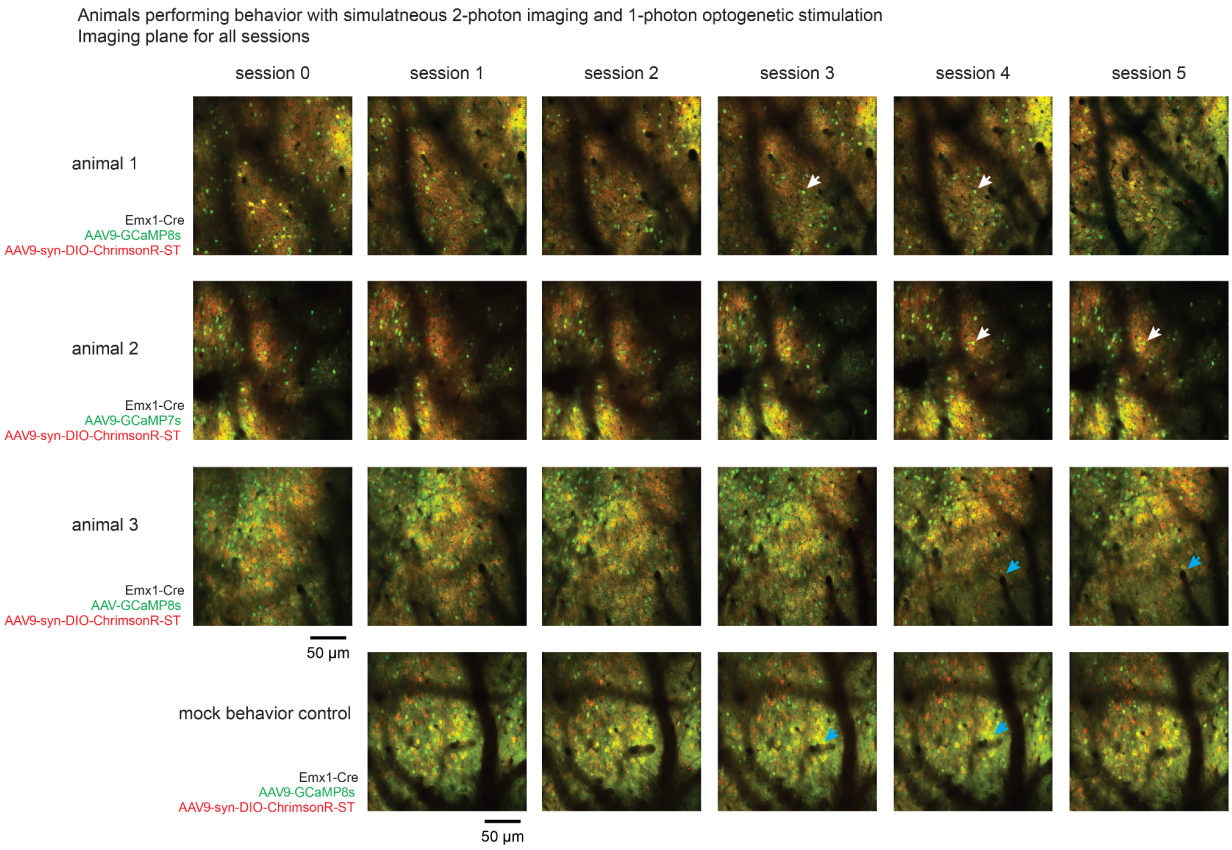

Figure S3 - Imaging plane over sessions for optogenetic learning animals and mock behavioral control, Related to Figure 2. Genotypes and viral injections are listed for each animal tested. Imaging planes were aligned to reference GCaMP expressing cells (examples, white arrows) and vasculature patterns (examples, blue arrows) between sessions. All Red/Green images shown are 300 frame averages acquired with the same amplifier gain settings at 1000 nm excitation (~35-45 mW). While some neurons differ from day to day, many of the same neurons were imaged across days.

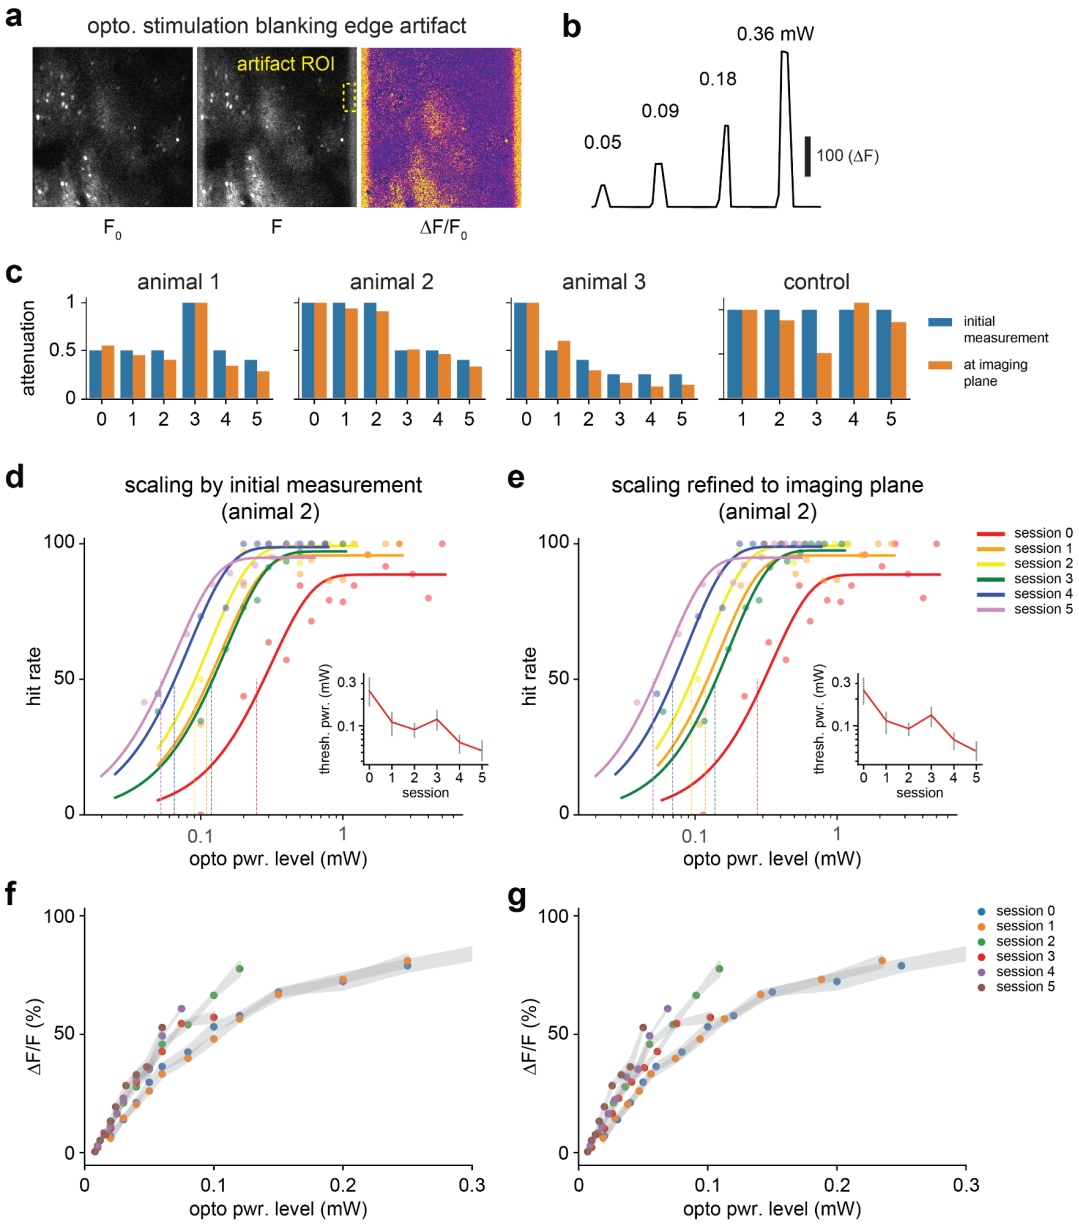

Figure S4 - Optogenetic stimulation blanking artifact allows normalization of optogenetic power at the imaging plane between sessions, Related to Figure 2. (**a**) Optogenetic stimulation produces an ~20 pixel edge artifact that is visible during imaging, as the optogenetic light source offset lasts a few microseconds into each imaging line after horizontal flyback. (**b**) Intensity of the edge artifact scales with applied optogenetic stimulation power. (**c**) Plots of attenuation based on initial measurement of power out of the objective and normalized scaling for all animals and control. (**d,e**) Normalized scaling refines the position of psychometric curves but does not change the order. Normalized scaling does not alter the relationship between threshold powers (insets). (**f,g**) Normalized scaling does not alter the relationship between ΔF/F and power over sessions.

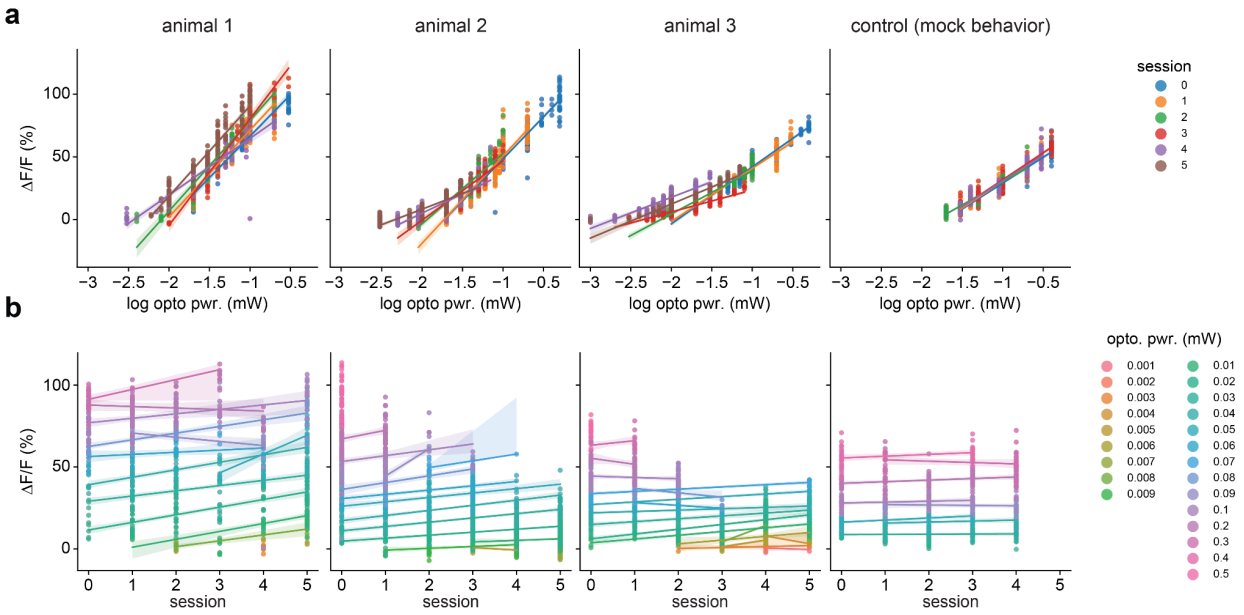

Figure S5 - V1 amplification effect, all sessions, all animals, Related to Figure 2. (**a**) Linear regression model for testing for amplification effects between behavioral sessions. OLS regression using all trials the stimulus was successfully detected ($p = 1.73 \times 10^{-12}$, N = 4 animals, 3209 trials, model: ΔF/F ~ C(animal) * C(session) + stimulation_power_mw + trial_number + constant, where C(x) signifies a categorical or dummy variable). Treating power as a continuous variable did not change the results. In the three training animals, lines fit on each session (colors) moved leftward as learning progressed, signifying amplification. (**b**) Comparison of amplification at each power across all behavioral sessions. Here, at many powers common across sessions (colors, lines), the slope of the corresponding line was positive, signifying amplification.

References

Ahmadian, Y., and Miller, K.D. (2021). What is the dynamical regime of cerebral cortex? Neuron *109*, 3373–3391. https://doi.org/10.1016/j.neuron.2021.07.031.

Balmer, T.S., Carels, V.M., Frisch, J.L., and Nick, T.A. (2009). Modulation of perineuronal nets and parvalbumin with developmental song learning. J. Neurosci. *29*, 12878–12885. https://doi.org/10.1523/JNEUROSCI.2974-09.2009.

Banerjee, S.B., Gutzeit, V.A., Baman, J., Aoued, H.S., Doshi, N.K., Liu, R.C., and Ressler, K.J. (2017). Perineuronal Nets in the Adult Sensory Cortex Are Necessary for Fear Learning. Neuron *95*, 169–179.e3. https://doi.org/10.1016/j.neuron.2017.06.007.

Bao, M., Yang, L., Rios, C., He, B., and Engel, S.A. (2010). Perceptual learning increases the strength of the earliest signals in visual cortex. J. Neurosci. *30*, 15080–15084. https://doi.org/10.1523/JNEUROSCI.5703-09.2010.

Boynton, G.M., and Finney, E.M. (2003). Orientation-Specific Adaptation in Human Visual Cortex. J. Neurosci. *23*, 8781–8787. https://doi.org/10.1523/jneurosci.23-25-08781.2003.

Brashers-Krug, T., Shadmehr, R., and Bizzi, E. (1996). Consolidation in human motor memory. Nature *382*, 252–255. https://doi.org/10.1038/382252a0.

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. J. Comput. Neurosci. *8*, 183–208. https://doi.org/10.1023/a:1008925309027.

Cano, M., Bezdudnaya, T., Swadlow, H.A., and Alonso, J.-M. (2006). Brain state and contrast sensitivity in the awake visual thalamus. Nat. Neurosci. *9*, 1240–1242. https://doi.org/10.1038/nn1760.

Carcea, I., and Froemke, R.C. (2013). Chapter 3 - Cortical Plasticity, Excitatory–Inhibitory Balance, and Sensory Perception. In Progress in Brain Research, M.M. Merzenich, M. Nahum, and T.M. Van Vleet, eds. (Elsevier), pp. 65–90.

Cardin, J.A., Carlén, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L.-H., and Moore, C.I. (2010). Targeted optogenetic stimulation and recording of neurons in vivo using cell-type-specific expression of Channelrhodopsin-2. Nat. Protocols *5*, 247–254. https://doi.org/10.1038/nprot.2009.228.

Chung, S., and Abbott, L.F. (2021). Neural population geometry: An approach for understanding biological and artificial neural networks. Curr. Opin. Neurobiol. *70*, 137–144. https://doi.org/10.1016/j.conb.2021.10.010.

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

749 Dana, H., Sun, Y., Mohar, B., Hulse, B.K., Kerlin, A.M., Hasseman, J.P., Tsegaye, G., Tsang,
750 A., Wong, A., Patel, R., et al. (2019). High-performance calcium sensors for imaging activity in
751 neuronal populations and microcompartments. Nat. Methods *16*, 649–657.
752 https://doi.org/10.1038/s41592-019-0435-6.

753 Deitch, D., Rubin, A., and Ziv, Y. (2021). Representational drift in the mouse visual cortex. Curr.
754 Biol. *31*, 4327–4339.e6. https://doi.org/10.1016/j.cub.2021.07.062.

755 Desai, N.S., Cudmore, R.H., Nelson, S.B., and Turrigiano, G.G. (2002). Critical periods for
756 experience-dependent synaptic scaling in visual cortex. Nat. Neurosci. *5*, 783–789.
757 https://doi.org/10.1038/nn878.

758 Destexhe, A., and Paré, D. (1999). Impact of network activity on the integrative properties of
759 neocortical pyramidal neurons in vivo. J. Neurophysiol. *81*, 1531–1547.
760 https://doi.org/10.1152/jn.1999.81.4.1531.

761 Doron, G., and Brecht, M. (2015). What single-cell stimulation has told us about neural coding.
762 Philos. Trans. R. Soc. Lond. B Biol. Sci. *370*, 20140204. https://doi.org/10.1098/rstb.2014.0204.

763 Dosher, B., and Lu, Z.-L. (2017). Visual Perceptual Learning and Models. Annu. Rev. Vis. Sci.
764 *3*, 343–363. https://doi.org/10.1146/annurev-vision-102016-061249.

765 Doty, R.W. (1969). Electrical stimulation of the brain in behavioral context. Annu. Rev. Psychol.
766 *20*, 289–320. https://doi.org/10.1146/annurev.ps.20.020169.001445.

767 Egger, S.W., and Lisberger, S.G. (2022). Neural structure of a sensory decoder for motor
768 control. Nat. Commun. *13*, 1829. https://doi.org/10.1038/s41467-022-29457-4.

769 Fagiolini, M., and Hensch, T.K. (2000). Inhibitory threshold for critical-period activation in
770 primary visual cortex. Nature *404*, 183–186. https://doi.org/10.1038/35004582.

771 Fagiolini, M., Fritschy, J.-M., Löw, K., Möhler, H., Rudolph, U., and Hensch, T.K. (2004).
772 Specific GABAA circuits for visual cortical plasticity. Science *303*, 1681–1683.
773 https://doi.org/10.1126/science.1091032.

774 Freedman, D.J., Riesenhuber, M., Poggio, T., and Miller, E.K. (2003). A comparison of primate
775 prefrontal and inferior temporal cortices during visual categorization. J. Neurosci. *23*, 5235–
776 5246. https://doi.org/10.1523/JNEUROSCI.23-12-05235.2003.

777 Frégnac, Y., and Shulz, D.E. (1999). Activity-dependent regulation of receptive field properties
778 of cat area 17 by supervised Hebbian learning. J. Neurobiol. *41*, 69–82.
779 https://doi.org/10.1002/(sici)1097-4695(199910)41:1<69::aid-neu10>3.0.co;2-1.

780 Friedrich, J., Zhou, P., and Paninski, L. (2017). Fast online deconvolution of calcium imaging
781 data. PLoS Comput. Biol. *13*, e1005423. https://doi.org/10.1371/journal.pcbi.1005423.

782 Ghose, G.M., Yang, T., and Maunsell, J.H.R. (2002). Physiological correlates of perceptual

783 learning in monkey V1 and V2. J. Neurophysiol. *87*, 1867–1888.
784 https://doi.org/10.1152/jn.00690.2001.

785 Gilbert, C.D., and Li, W. (2012). Adult visual cortical plasticity. Neuron *75*, 250–264.
786 https://doi.org/10.1016/j.neuron.2012.06.030.

787 Gill, J.V., Lerman, G.M., Zhao, H., Stetler, B.J., Rinberg, D., and Shoham, S. (2020). Precise
788 Holographic Manipulation of Olfactory Circuits Reveals Coding Features Determining
789 Perceptual Detection. Neuron *108*, 382–393.e5. https://doi.org/10.1016/j.neuron.2020.07.034.

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

793 Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci.
794 *30*, 535–574. https://doi.org/10.1146/annurev.neuro.29.051605.113038.

795 Goldbach, H.C., Akitake, B., Leedy, C.E., and Histed, M.H. (2021). Performance in even a
796 simple perceptual task depends on mouse secondary visual areas. eLife *10*.
797 https://doi.org/10.7554/eLife.62156.

798 Goldman, M.S. (2009). Memory without feedback in a neural network. Neuron *61*, 621–634.
799 https://doi.org/10.1016/j.neuron.2008.12.012.

800 Goltstein, P.M., Coffey, E.B.J., Roelfsema, P.R., and Pennartz, C.M.A. (2013). In vivo two-
801 photon Ca2+ imaging reveals selective reward effects on stimulus-specific assemblies in mouse
802 visual cortex. J. Neurosci. *33*, 11540–11555. https://doi.org/10.1523/JNEUROSCI.1341-
803 12.2013.

804 Goltstein, P.M., Reinert, S., Bonhoeffer, T., and Hübener, M. (2021). Mouse visual cortex areas
805 represent perceptual and semantic features of learned visual categories. Nat. Neurosci. *24*,
806 1441–1451. https://doi.org/10.1038/s41593-021-00914-5.

807 Gorski, J.A., Talley, T., Qiu, M., Puelles, L., Rubenstein, J.L.R., and Jones, K.R. (2002). Cortical
808 excitatory neurons and glia, but not GABAergic neurons, are produced in the Emx1-expressing
809 lineage. J. Neurosci. *22*, 6309–6314. https://doi.org/20026564.

810 Gu, Y., Tran, T., Murase, S., Borrell, A., Kirkwood, A., and Quinlan, E.M. (2016). Neuregulin-
811 dependent regulation of fast-spiking interneuron excitability controls the timing of the critical
812 period. J. Neurosci. *36*, 10285–10295. https://doi.org/10.1523/JNEUROSCI.4242-15.2016.

813 He, H.-Y., Hodos, W., and Quinlan, E.M. (2006). Visual deprivation reactivates rapid ocular
814 dominance plasticity in adult visual cortex. J. Neurosci. *26*, 2951–2955.
815 https://doi.org/10.1523/JNEUROSCI.5554-05.2006.

816 Heimel, J.A., van Versendaal, D., and Levelt, C.N. (2011). The role of GABAergic inhibition in
817 ocular dominance plasticity. Neural Plast. *2011*, 391763. https://doi.org/10.1155/2011/391763.

Hengen, K.B., Lambo, M.E., Van Hooser, S.D., Katz, D.B., and Turrigiano, G.G. (2013). Firing rate homeostasis in visual cortex of freely behaving rodents. Neuron *80*, 335–342. https://doi.org/10.1016/j.neuron.2013.08.038.

Hennequin, G., Vogels, T.P., and Gerstner, W. (2012). Non-normal amplification in random balanced neuronal networks. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *86*, 011909. https://doi.org/10.1103/PhysRevE.86.011909.

Hensch, T.K. (2005). Critical period plasticity in local cortical circuits. Nat. Rev. Neurosci. *6*, 877–888. https://doi.org/10.1038/nrn1787.

Henschke, J.U., Dylda, E., Katsanevaki, D., Dupuy, N., Currie, S.P., Amvrosiadis, T., Pakan, J.M.P., and Rochefort, N.L. (2020). Reward Association Enhances Stimulus-Specific Representations in Primary Visual Cortex. Curr. Biol. *30*, 1866–1880.e5. https://doi.org/10.1016/j.cub.2020.03.018.

Histed, M.H., and Maunsell, J.H.R. (2014). Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. Proc. Natl. Acad. Sci. U. S. A. *111*, E178–E187. https://doi.org/10.1073/pnas.1318750111.

Histed, M.H., Bonin, V., and Reid, R.C. (2009). Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation. Neuron *63*, 508–522. https://doi.org/10.1016/j.neuron.2009.07.016.

Histed, M.H., Carvalho, L.A., and Maunsell, J.H.R. (2012). Psychophysical measurement of contrast sensitivity in the behaving mouse. J. Neurophysiol. *107*, 758–765. https://doi.org/10.1152/jn.00609.2011.

Histed, M.H., Ni, A.M., and Maunsell, J.H.R. (2013). Insights into cortical mechanisms of behavior from microstimulation experiments. Prog. Neurobiol. *103*, 115–130. https://doi.org/10.1016/j.pneurobio.2012.01.006.

Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J. Physiol. *160*, 106–154. https://doi.org/10.1113/jphysiol.1962.sp006837.

Hylin, M.J., Orsi, S.A., Moore, A.N., and Dash, P.K. (2013). Disruption of the perineuronal net in the hippocampus or medial prefrontal cortex impairs fear conditioning. Learn. Mem. *20*, 267–273. https://doi.org/10.1101/lm.030197.112.

Jazayeri, M., and Afraz, A. (2017). Navigating the Neural Space in Search of the Neural Code. Neuron *93*, 1003–1014. https://doi.org/10.1016/j.neuron.2017.02.019.

Jurjut, O., Georgieva, P., Busse, L., and Katzner, S. (2017). Learning Enhances Sensory Processing in Mouse V1 before Improving Behavior. J. Neurosci. *37*, 6460–6474. https://doi.org/10.1523/JNEUROSCI.3485-16.2017.

853 Katz, L.C., and Crowley, J.C. (2002). Development of cortical circuits: lessons from ocular
854 dominance columns. Nat. Rev. Neurosci. *3*, 34–42. https://doi.org/10.1038/nrn703.

855 Keck, T., Keller, G.B., Jacobsen, R.I., Eysel, U.T., Bonhoeffer, T., and Hübener, M. (2013).
856 Synaptic scaling and homeostatic plasticity in the mouse visual cortex in vivo. Neuron *80*, 327–
857 334. https://doi.org/10.1016/j.neuron.2013.08.018.

858 Kelly, S.T., Kremkow, J., Jin, J., Wang, Y., Wang, Q., Alonso, J.-M., and Stanley, G.B. (2014).
859 The role of thalamic population synchrony in the emergence of cortical feature selectivity. PLoS
860 Comput. Biol. *10*, e1003418. https://doi.org/10.1371/journal.pcbi.1003418.

861 Kerlin, A., Mohar, B., Flickinger, D., MacLennan, B.J., Dean, M.B., Davis, C., Spruston, N., and
862 Svoboda, K. (2019). Functional clustering of dendritic activity during decision-making. eLife *8*.
863 https://doi.org/10.7554/eLife.46966.

864 Kesner, R.P., and Wilburn, M.W. (1974). A review of electrical stimulation of the brain in context
865 of learning and retention. Behavioral Biology *10*, 259–293. https://doi.org/10.1016/s0091-
866 6773(74)91894-x.

867 Khan, A.G., Poort, J., Chadwick, A., Blot, A., Sahani, M., Mrsic-Flogel, T.D., and Hofer, S.B.
868 (2018). Distinct learning-induced changes in stimulus selectivity and interactions of GABAergic
869 interneuron classes in visual cortex. Nat. Neurosci. *21*, 851–859.
870 https://doi.org/10.1038/s41593-018-0143-z.

871 Kohn, A., Coen-Cagli, R., Kanitscheider, I., and Pouget, A. (2016). Correlations and Neuronal
872 Population Information. Annu. Rev. Neurosci. https://doi.org/10.1146/annurev-neuro-070815-
873 013851.

874 Kondo, S., and Ohki, K. (2016). Laminar differences in the orientation selectivity of geniculate
875 afferents in mouse primary visual cortex. Nat. Neurosci. *19*, 316–319.
876 https://doi.org/10.1038/nn.4215.

877 Krakauer, J.W., and Shadmehr, R. (2006). Consolidation of motor memory. Trends Neurosci.
878 *29*, 58–64. https://doi.org/10.1016/j.tins.2005.10.003.

879 Kügler, S., Kilic, E., and Bähr, M. (2003). Human synapsin 1 gene promoter confers highly
880 neuron-specific long-term transgene expression from an adenoviral vector in the adult rat brain
881 depending on the transduced area. Gene Therapy *10*, 337–347.
882 https://doi.org/10.1038/sj.gt.3301905.

883 Le Naour, F., Hohenkirk, L., Grolleau, A., Misek, D.E., Lescure, P., Geiger, J.D., Hanash, S.,
884 and Beretta, L. (2001). Profiling changes in gene expression during differentiation and
885 maturation of monocyte-derived dendritic cells using both oligonucleotide microarrays and
886 proteomics. J. Biol. Chem. *276*, 17920–17931. https://doi.org/10.1074/jbc.M100156200.

887 Li, W., Piëch, V., and Gilbert, C.D. (2008). Learning to link visual contours. Neuron *57*, 442–451.
888 https://doi.org/10.1016/j.neuron.2007.12.011.

889 Liang, L., Fratzl, A., Reggiani, J.D.S., El Mansour, O., Chen, C., and Andermann, M.L. (2020).
890 Retinal Inputs to the Thalamus Are Selectively Gated by Arousal. Curr. Biol. *30*, 3923–3934.e9.
891 https://doi.org/10.1016/j.cub.2020.07.065.

892 Mainen, Z.F., and Sejnowski, T.J. (1995). Reliability of spike timing in neocortical neurons.
893 Science *268*, 1503–1506. https://doi.org/10.1126/science.7770778.

894 Malenka, R.C., and Bear, M.F. (2004). LTP and LTD: an embarrassment of riches. Neuron *44*,
895 5–21. https://doi.org/10.1016/j.neuron.2004.09.012.

896 Marks, T.D., and Goard, M.J. (2021). Stimulus-dependent representational drift in primary visual
897 cortex. Nat. Commun. *12*, 5169. https://doi.org/10.1038/s41467-021-25436-3.

898 Marshel, J.H., Kim, Y.S., Machado, T.A., Quirin, S., Benson, B., Kadmon, J., Raja, C.,
899 Chibukhchyan, A., Ramakrishnan, C., Inoue, M., et al. (2019). Cortical layer-specific critical
900 dynamics triggering perception. Science *365*. https://doi.org/10.1126/science.aaw5202.

901 Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., and Pouget, A. (2014).
902 Information-limiting correlations. Nat. Neurosci. *17*, 1410–1417. https://doi.org/10.1038/nn.3807.

903 Murphy, B.K., and Miller, K.D. (2009). Balanced amplification: a new mechanism of selective
904 amplification of neural activity patterns. Neuron *61*, 635–648.
905 https://doi.org/10.1016/j.neuron.2009.02.005.

906 Ni, A.M., and Maunsell, J.H.R. (2010). Microstimulation reveals limits in detecting different
907 signals from a local cortical region. Curr. Biol. *20*, 824–828.
908 https://doi.org/10.1016/j.cub.2010.02.065.

909 Niell, C.M., and Scanziani, M. (2021). How Cortical Circuits Implement Cortical Computations:
910 Mouse Visual Cortex as a Model. Annu. Rev. Neurosci. *44*, 517–546.
911 https://doi.org/10.1146/annurev-neuro-102320-085825.

912 O'Connor, D.H., Hires, S.A., Guo, Z.V., Li, N., Yu, J., Sun, Q.-Q., Huber, D., and Svoboda, K.
913 (2013). Neural coding during active somatosensation revealed using illusory touch. Nat.
914 Neurosci. *16*, 958–965. https://doi.org/10.1038/nn.3419.

915 O'Rawe, J., Zhou, Z., Goldbach, A.L.P.L., and Histed, M. (2022). Recurrent suppression in
916 visual cortex explained by a balanced network with sparse synaptic connections.
917 (Computational and Systems Neuroscience Meeting 2022 (poster)), pp. II – 034.

918 Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Federico Rossi, L., Dalgleish, H.,
919 Carandini, M., and Harris, K.D. (2017). Suite2p: beyond 10,000 neurons with standard two-
920 photon microscopy.

921 Pancholi, R., Ryan, L., and Peron, S. (2021). Sensory cortical dynamics during optical
922 microstimulation training. bioRxiv 2021.12.17.473191.
923 https://doi.org/10.1101/2021.12.17.473191.

924 Pégard, N.C., Mardinly, A.R., Oldenburg, I.A., Sridharan, S., Waller, L., and Adesnik, H. (2017).
925 Three-dimensional scanless holographic optogenetics with temporal focusing (3D-SHOT).
926 Nature Communications *8*. https://doi.org/10.1038/s41467-017-01031-3.

927 Pons, T.P., Garraghty, P.E., Ommaya, A.K., Kaas, J.H., Taub, E., and Mishkin, M. (1991).
928 Massive cortical reorganization after sensory deafferentation in adult macaques. Science *252*,
929 1857–1860. https://doi.org/10.1126/science.1843843.

930 Poort, J., Khan, A.G., Pachitariu, M., Nemri, A., Orsolic, I., Krupic, J., Bauza, M., Sahani, M.,
931 Keller, G.B., Mrsic-Flogel, T.D., et al. (2015). Learning Enhances Sensory and Multiple Non-
932 sensory Representations in Primary Visual Cortex. Neuron *86*, 1478–1490.
933 https://doi.org/10.1016/j.neuron.2015.05.037.

934 Reichelt, A.C., Hare, D.J., Bussey, T.J., and Saksida, L.M. (2019). Perineuronal Nets: Plasticity,
935 Protection, and Therapeutic Potential. Trends Neurosci. *42*, 458–470.
936 https://doi.org/10.1016/j.tins.2019.04.003.

937 Sadagopan, S., and Ferster, D. (2012). Feedforward origins of response variability underlying
938 contrast invariant orientation tuning in cat visual cortex. Neuron *74*, 911–923.
939 https://doi.org/10.1016/j.neuron.2012.05.007.

940 Sadeh, S., and Clopath, C. (2020). Theory of neuronal perturbome in cortical networks. Proc.
941 Natl. Acad. Sci. U. S. A. *117*, 26966–26976. https://doi.org/10.1073/pnas.2004568117.

942 Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M.,
943 and Batista, A.P. (2014). Neural constraints on learning. Nature *512*, 423–426.
944 https://doi.org/10.1038/nature13665.

945 Sanzeni, A., Akitake, B., Goldbach, H.C., Leedy, C.E., Brunel, N., and Histed, M.H. (2020).
946 Inhibition stabilization is a widespread property of cortical networks. eLife *9*.
947 https://doi.org/10.7554/elife.54875.

948 Sanzeni, A., Histed, M.H., and Brunel, N. (2022). Emergence of Irregular Activity in Networks of
949 Strongly Coupled Conductance-Based Neurons. Phys. Rev. X *12*, 011044.
950 https://doi.org/10.1103/PhysRevX.12.011044.

951 Sawtell, N.B., Frenkel, M.Y., Philpot, B.D., Nakazawa, K., Tonegawa, S., and Bear, M.F. (2003).
952 NMDA receptor-dependent ocular dominance plasticity in adult visual cortex. Neuron *38*, 977–
953 985. https://doi.org/10.1016/s0896-6273(03)00323-4.

954 Schmolesky, M.T., Wang, Y., Hanes, D.P., Thompson, K.G., Leutgeb, S., Schall, J.D., and
955 Leventhal, A.G. (1998). Signal timing across the macaque visual system. J. Neurophysiol. *79*,
956 3272–3278. https://doi.org/10.1152/jn.1998.79.6.3272.

957 Schoups, A., Vogels, R., Qian, N., and Orban, G. (2001). Practising orientation identification
958 improves orientation coding in V1 neurons. Nature *412*, 549–553.
959 https://doi.org/10.1038/35087601.

960 Schwartz, S., Maquet, P., and Frith, C. (2002). Neural correlates of perceptual learning: a
961 functional MRI study of visual texture discrimination. Proc. Natl. Acad. Sci. U. S. A. *99*, 17137–
962 17142. https://doi.org/10.1073/pnas.242414599.

963 Sohal, D.S., Nghiem, M., Crackower, M.A., Witt, S.A., Kimball, T.R., Tymitz, K.M., Penninger,
964 J.M., and Molkentin, J.D. (2001). Temporally regulated and tissue-specific gene manipulations
965 in the adult and embryonic heart using a tamoxifen-inducible Cre protein. Circ. Res. *89*, 20–25.
966 https://doi.org/10.1161/hh1301.092687.

967 Sorg, B.A., Berretta, S., Blacktop, J.M., Fawcett, J.W., Kitagawa, H., Kwok, J.C.F., and Miquel,
968 M. (2016). Casting a Wide Net: Role of Perineuronal Nets in Neural Plasticity. J. Neurosci. *36*,
969 11459–11468. https://doi.org/10.1523/JNEUROSCI.2351-16.2016.

970 Steinmetz, N.A., Zatka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of
971 choice, action and engagement across the mouse brain. Nature *576*, 266–273.
972 https://doi.org/10.1038/s41586-019-1787-x.

973 Swanson, O.K., and Maffei, A. (2019). From Hiring to Firing: Activation of Inhibitory Neurons
974 and Their Recruitment in Behavior. Front. Mol. Neurosci. *12*, 168.
975 https://doi.org/10.3389/fnmol.2019.00168.

976 Swindale, N.V. (1998). Orientation tuning curves: empirical description and estimation of
977 parameters. Biol. Cybern. *78*, 45–56. https://doi.org/10.1007/s004220050411.

978 Trachtenberg, J.T. (2015). Competition, inhibition, and critical periods of cortical plasticity. Curr.
979 Opin. Neurobiol. *35*, 44–48. https://doi.org/10.1016/j.conb.2015.06.006.

980 Wang, X.-J. (2012). Neural dynamics and circuit mechanisms of decision-making. Curr. Opin.
981 Neurobiol. *22*, 1039–1046. https://doi.org/10.1016/j.conb.2012.08.006.

982 Wu, Z., Litwin-Kumar, A., Shamash, P., Taylor, A., Axel, R., and Shadlen, M.N. (2020). Context-
983 Dependent Decision Making in a Premotor Circuit. Neuron *106*, 316–328.e6.
984 https://doi.org/10.1016/j.neuron.2020.01.034.

985 Yang, T., and Maunsell, J.H.R. (2004). The effect of perceptual learning on neuronal responses
986 in monkey visual area V4. J. Neurosci. *24*, 1617–1626.
987 https://doi.org/10.1523/JNEUROSCI.4442-03.2004.

988 Yotsumoto, Y., Watanabe, T., and Sasaki, Y. (2008). Different dynamics of performance and
989 brain activation in the time course of perceptual learning. Neuron *57*, 827–833.
990 https://doi.org/10.1016/j.neuron.2008.02.034.

991 Zatka-Haas, P., Steinmetz, N.A., Carandini, M., and Harris, K.D. (2021). Sensory coding and
992 the causal impact of mouse cortex in a visual decision. eLife *10*.
993 https://doi.org/10.7554/eLife.63163.

994 Zhang, Y., Rózsa, M., Bushey, D., Zheng, J., Reep, D., Broussard, G.J., Tsang, A., Tsegaye,

995 G., Patel, R., Narayan, S., et al. (2020). jGCaMP8 fast genetically encoded calcium indicators.
996 Janelia Research Campus *10*, 13148243. .