

Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex

Highlights

- V1 neurons increasingly discriminate task-relevant stimuli with learning
- Chronic imaging reveals single cell changes underlying this population effect
- Learning-related changes are reduced when animals ignore task-relevant stimuli
- Anticipatory and behavioral choice-related signals emerge in reward-predicting cells

Authors

Jasper Poort, Adil G. Khan, Marius Pachitariu, ..., Georg B. Keller, Thomas D. Mrsic-Flogel, Sonja B. Hofer

Correspondence

sonja.hofer@unibas.ch

In Brief

By tracking the same visual cortex neurons across days, Poort et al. demonstrate how learning a visual task leads to increasingly distinguishable representations of relevant stimuli. These changes parallel the emergence of diverse non-sensory signals in specific neuronal subsets.



Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex

Jasper Poort,^{2,5} Adil G. Khan,^{1,2,5} Marius Pachitariu,³ Abdellatif Nemri,^{1,2} Ivana Orsolich,¹ Julija Krupic,² Marius Bauza,² Maneesh Sahani,³ Georg B. Keller,⁴ Thomas D. Mrsic-Flogel,^{1,2} and Sonja B. Hofer^{1,2,*}

¹Biozentrum, University of Basel, Klingelbergstrasse 50/70, 4056 Basel, Switzerland

²University College London, 21 University Street, London WC1E 6DE, UK

³Gatsby Computational Neuroscience Unit, University College London, 17 Queen Square, London WC1N 3AR, UK

⁴Friedrich Miescher Institute for Biomedical Research, Maulbeerstrasse 66, 4058 Basel, Switzerland

⁵Co-first author

*Correspondence: sonja.hofer@unibas.ch

<http://dx.doi.org/10.1016/j.neuron.2015.05.037>

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

SUMMARY

We determined how learning modifies neural representations in primary visual cortex (V1) during acquisition of a visually guided behavioral task. We imaged the activity of the same layer 2/3 neuronal populations as mice learned to discriminate two visual patterns while running through a virtual corridor, where one pattern was rewarded. Improvements in behavioral performance were closely associated with increasingly distinguishable population-level representations of task-relevant stimuli, as a result of stabilization of existing and recruitment of new neurons selective for these stimuli. These effects correlated with the appearance of multiple task-dependent signals during learning: those that increased neuronal selectivity across the population when expert animals engaged in the task, and those reflecting anticipation or behavioral choices specifically in neuronal subsets preferring the rewarded stimulus. Therefore, learning engages diverse mechanisms that modify sensory and non-sensory representations in V1 to adjust its processing to task requirements and the behavioral relevance of visual stimuli.

INTRODUCTION

Primary areas of the sensory neocortex are thought to faithfully represent the identity of stimuli in the external environment. Yet as animals learn the association between a sensory stimulus and its behavioral relevance, or improve their perceptual capabilities with training, stimulus representations in sensory cortical areas can change (Schoups et al., 2001; Yang and Maunsell, 2004; Rutkowski and Weinberger, 2005; Blake et al., 2006; Li et al., 2008; Wiest et al., 2010; Gdalyahu et al., 2012; Goltstein et al., 2013; Yan et al., 2014). Such changes may lead to enhanced and more distinct representations of task-relevant stimuli, and therefore improve the salience of information relayed to downstream areas.

The nature and effect sizes of learning-related changes to neural representations vary strongly between different studies, potentially depending on modality, sensory cortical area, and the behavioral task (Schoups et al., 2001; Yang and Maunsell, 2004; Rutkowski and Weinberger, 2005; Li et al., 2008; Ghose et al., 2002; Law and Gold, 2008). The repeated association between a stimulus and reward can lead to lasting, task-independent changes in cortical representations of that stimulus (Schoups et al., 2001; Rutkowski and Weinberger, 2005; Goltstein et al., 2013). Alternatively, the expression of learning-related changes to sensory responses can also depend on the animals being engaged in the task (Li et al., 2004, 2008; Polley et al., 2006), consistent with observations that even in primary sensory cortex neuronal responses can be influenced by non-sensory, task-dependent signals reflecting the animal's attentive state, expectations, or behavior (see, for example, Ress and Heeger, 2003; Shuler and Bear, 2006; Li et al., 2008; Niell and Stryker, 2010; Keller et al., 2012; David et al., 2012; Stanišor et al., 2013; Nienborg and Cumming, 2014). Therefore, the strategies by which learning can modify cortical sensory processing are diverse but remain poorly understood. Specifically, how do individual neurons change their response properties as stimuli acquire behavioral relevance? To what extent do these changes persist when the animals are not engaged in the task? How do learning-induced response changes relate to the appearance of non-sensory, task-dependent signals? Do these non-sensory signals act globally, or do they target specific neuronal subsets encoding behaviorally relevant sensory features?

To address these questions, it is crucial to track the activity of the same cells over the course of learning. We therefore used chronic two-photon calcium imaging (Huber et al., 2012; Chen et al., 2013) in mouse V1 while the animals learned to perform a visual discrimination task in virtual reality. We observed a robust and progressive population-wide increase in neural selectivity in cortical layer 2/3 (L2/3) during learning—an effect related to greater day-to-day stability of single cell response preferences as well as to an increase in the number of cells selective for task-relevant stimuli. Improvements in V1 selectivity were reduced when animals disengaged from the task. Task acquisition additionally led to the appearance of both anticipatory and behavioral choice-related signals in a specific subpopulation of neurons whose firing predicted the reward. Therefore,

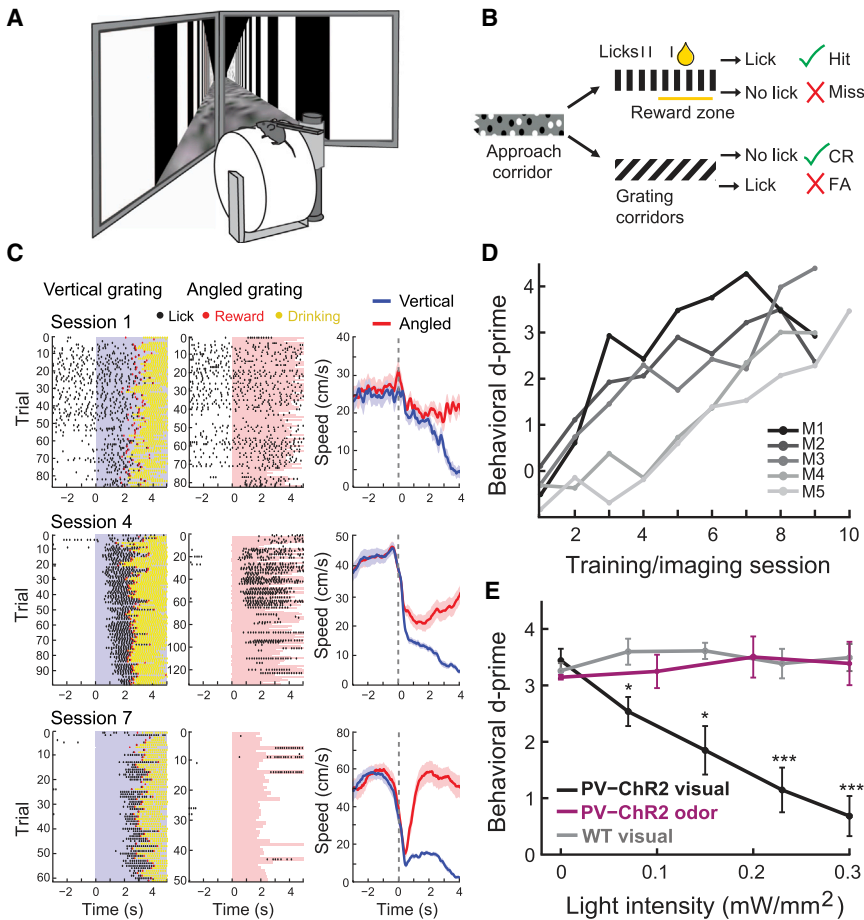


Figure 1. Rapid Learning of a V1-Dependent Visual Discrimination Task in Virtual Reality

(A) Schematic of the virtual reality setup.

(B) Task schematic with virtual corridor wall patterns. CR, correct rejection, FA, false alarm.

(C) Changes in licking over learning in an example mouse. Licks (dots) aligned to grating onset in vertical grating (left, blue shading) and angled grating (middle, pink shading) trials. Red dots, reward delivery; yellow dots, licking after reward delivery. Right, average running speed for session shown on left, aligned to grating onset for vertical (blue) and angled (red) trials. Shading, SEM.

(D) Behavioral performance (behavioral d-prime; see [Experimental Procedures](#)) of five mice imaged on consecutive training sessions. See also [Figure S1](#).

(E) Behavioral performance in the visual and an equivalent odor discrimination task (see [Experimental Procedures](#), average across sessions) as a function of light intensity during bilateral optogenetic silencing of visual cortex. PV-ChR2, transgenic mice expressing Channelrhodopsin-2 in parvalbumin-positive interneurons ($n = 4$ mice, 10 visual and 4 odor discrimination sessions). WT, wild-type mice ($n = 3$ mice, 7 sessions). * $p < 0.05$, *** $p < 0.001$ after Bonferroni correction, Wilcoxon rank-sum test comparing PV-ChR2 to WT in the visual task.

learning the relationship between visual cues and their behavioral relevance leads to concerted changes in the representation of both sensory and non-sensory task-related information in a primary sensory cortical area.

RESULTS

Behavioral Task

Mice can perform complex visually guided behaviors, but they often require weeks of training to achieve high performance levels when head restrained ([Andermann et al., 2010](#); [Glickfeld et al., 2013](#); [Pinto et al., 2013](#)). Virtual reality environments offer an advantage for training head-fixed animals because they allow active engagement with the sensory world, for example, when the animal's locomotion on a treadmill is directly coupled to optic flow changes in the visual scene ([Hölscher et al., 2005](#); [Dombeck et al., 2007](#)). We hypothesized that this type of active visuomotor engagement approximates ethological situations when mice encounter behaviorally relevant stimuli during navigation, exploration, or foraging. Indeed, we found that this enabled rapid visually guided learning (see below).

We trained head-fixed mice to discriminate two grating patterns of different orientations in a virtual reality environment in which the animals' running controlled their position in a corridor ([Figures 1A and 1B](#); see also [Movie S1](#) available online).

relative to vertical) gratings on both walls. The abrupt appearance of the grating corridors provided precise control of stimulus timing. Mice were rewarded for licking in response to the vertical grating corridor with a drop of soya milk delivered through a reward spout (hit trial; reward was given if a lick was detected in a region a short distance into the grating corridor, referred to as the reward zone). No punishment was given for licking in response to the non-rewarded, angled grating corridor (false-alarm trial). Most mice progressed rapidly from indiscriminate licking (example lick raster plots in [Figure 1C](#), top) to licking only within the grating corridors in response to both gratings ([Figure 1C](#), middle), and finally to nearly exclusive licking in response to the rewarded, vertical grating ([Figure 1C](#), bottom) and withholding licking in the non-rewarded angled grating corridor (correct rejection trials). Mice typically slowed down while licking in the rewarded grating corridor and learned to accelerate upon seeing the non-rewarded grating ([Figure 1C](#), right panels). We quantified task performance by calculating the behavioral d-prime for each training session, which is a measure of the difference in the proportions of hit and false-alarm trials ([Figure 1D](#); see [Experimental Procedures](#)). Mice usually learnt the task within 3–6 days ([Figure 1D](#)) and eventually reached high behavioral accuracies (behavioral d-prime in last session 3.2 ± 0.7 , corresponding to $89\% \pm 8\%$ correct responses, mean \pm SD; see [Figure S1](#)).

After running through a virtual approach corridor (walls with black/white circles) from random starting points, mice were abruptly presented with a corridor containing either vertical or angled (40° relative to vertical) gratings on both walls.

We tested whether V1 activity was required for visual discrimination in this task by optogenetically silencing V1 in both hemispheres of fully trained animals in a random subset of trials. We silenced the cortex during grating corridor presentation by photostimulation of parvalbumin-positive inhibitory interneurons expressing Channelrhodopsin-2 in transgenic mice (Boyden et al., 2005; Lien and Scanziani, 2013; Glickfeld et al., 2013). Visual discrimination performance decreased progressively when increasing the intensity of blue light directed to V1 in transgenic mice, (Figure 1E; Friedman test, $\chi^2[4] = 32.44$, $p < 10^{-5}$), but not in wild-type control mice (Figure 1E; Friedman test, $\chi^2[4] = 5.76$, $p = 0.22$). The same transgenic mice were additionally trained in an analogous odor discrimination task in the same virtual corridor (see Experimental Procedures), which they continued to perform normally even when illuminating V1 with high light intensities (Figure 1E; Friedman test, $\chi^2[3] = 0.20$, $p = 0.98$), demonstrating that only visual processing was affected by this optogenetic manipulation.

Response Dynamics Underlying Increase in Neuronal Selectivity during Learning

Having established the necessity of V1 for this visual discrimination task (see also Glickfeld et al., 2013), we examined how the activity of neuronal populations in V1 changed during learning. For this purpose, we expressed the calcium indicator GCaMP6 (Chen et al., 2013) in V1 using AAV vectors, and chronically recorded calcium signals (32 Hz frame rate) in L2/3 using two-photon microscopy (Denk et al., 1990) while the animals performed the task (Figures 2A and 2B; on average 199 trials per session, range 31–342 trials). We imaged the same populations of neurons (75 ± 27 cells per mouse; mean \pm SD) either in each training session over the entire time course of learning (five mice, Figure 1D), before and after learning (three mice), or only after learning (three mice). Neurons exhibited diverse response profiles during the task (Figures 2A, 2B, and S2). While some neurons responded to features in the approach corridor (Figure 2A, cell 1; Figure S7), many cells responded to both the vertical and angled grating corridors, and their responses were often stronger to one grating than the other (Figures 2B and S2). In other neurons, the calcium signal decreased during grating presentation (Figure 2B, cell 8; Figure S2). Despite variability in response amplitudes and in the degree of response selectivity from session to session (see below and Figures 2E–2G), the majority of neurons maintained their response profiles over time (Figures 2B, S3A, and S3B).

To quantify how the preference and selectivity of individual neurons for the two grating corridors changed during learning, we derived an index of neuronal selectivity for each neuron in each training session (defined as the difference between the average responses to vertical and angled gratings in a time window 0–1 s after grating onset, normalized by the pooled standard deviation of responses across trials). By binning sessions with similar behavioral performance (Figure 2C), we observed a gradual broadening of the distribution of neuronal selectivity over learning, resulting in both more positive values (higher preference for the rewarded, vertical grating) and more negative values (higher preference for the non-rewarded, angled grating). Consequently, the fraction of selective neurons rose significantly

over learning (Figure 2D), including an increase in the percentage of cells preferring the non-rewarded grating corridor (12% to 19%, $p = 0.02$, bootstrap test), and a larger increase in the percentage of cells preferring the rewarded grating corridor (12% to 32%, $p < 10^{-4}$, Figure 2D; see Figure S4 for individual mice). Restricting the analysis only to neurons with a significant response increase after grating corridor onset ($p < 0.01$, Wilcoxon signed-rank test) yielded similar results (Figure S5).

The increase in neuronal selectivity was caused by an increase in reliability of responses (mean standard deviation of responses within 0–1 s window from grating onset, pre learning = 0.088, post learning = 0.063, $p = 0.001$, bootstrap test, 27 sessions before and 52 after learning), as well as an increased difference in response amplitude to the two gratings with learning (mean absolute response difference; pre learning = 0.017, post learning = 0.024, $p = 0.016$, bootstrap test). However, there was no consistent strategy by which individual neurons changed their response amplitudes to the two gratings (Figure S3C).

We next determined whether the increase in selectivity for grating stimuli was restricted to neurons with specific response properties. Neurons preferentially responding to either grating before learning were no more likely to increase their selectivity during learning than non-selective neurons ($R = -0.06$, $p = 0.20$; Figure S3E). Moreover, individual cells showed relatively large variability in how they changed their selectivity over learning (Figures S3D and S3F). The increase in selectivity, therefore, involved diverse modes of response change distributed over many neurons across the L2/3 population in V1.

Neuronal responses can show considerable variability from one day to the next (Huber et al., 2012; Peters et al., 2014; Ziv et al., 2013). We quantified day-to-day fluctuations of stimulus preferences of individual cells and how they changed during learning (Figures 2E–2G). We computed the likelihood of neurons maintaining their grating selectivity from one day to the next (persistence of response preference, Figure 2F) within different stages of learning: before animals showed improvements in their behavioral performance (pre learning), during learning, and after the behavioral performance had stabilized (post learning; see Experimental Procedures). While neurons were relatively more likely to lose their selectivity from one day to the next before learning, it was rare for neurons to completely switch from preferring one grating to the other (on average 3% before learning). Over learning, the persistence of selective responses increased, and cells preferring either the rewarded or the non-rewarded stimulus became more stable in their stimulus preference (Figure 2F; rewarded-grating-preferring cells, pre = 49% to post = 70%, $p < 10^{-3}$; non-rewarded-grating-preferring cells, pre = 17% to post = 55%, $p < 10^{-4}$, bootstrap tests). We additionally determined the probability of non-selective neurons becoming selective from one day to the next (Figure 2G). As learning progressed, non-selective neurons became more likely to acquire a preference for the rewarded, vertical grating, but not for the non-rewarded, angled grating (Figure 2G, rewarded-stimulus-preferring cells, $p < 10^{-4}$; non-rewarded-stimulus-preferring cells, $p = 0.29$, bootstrap tests). Therefore, the increasing preference for task-relevant stimuli in L2/3 of V1 during learning was a result of a stabilization of response selectivity to both gratings as well as an increased

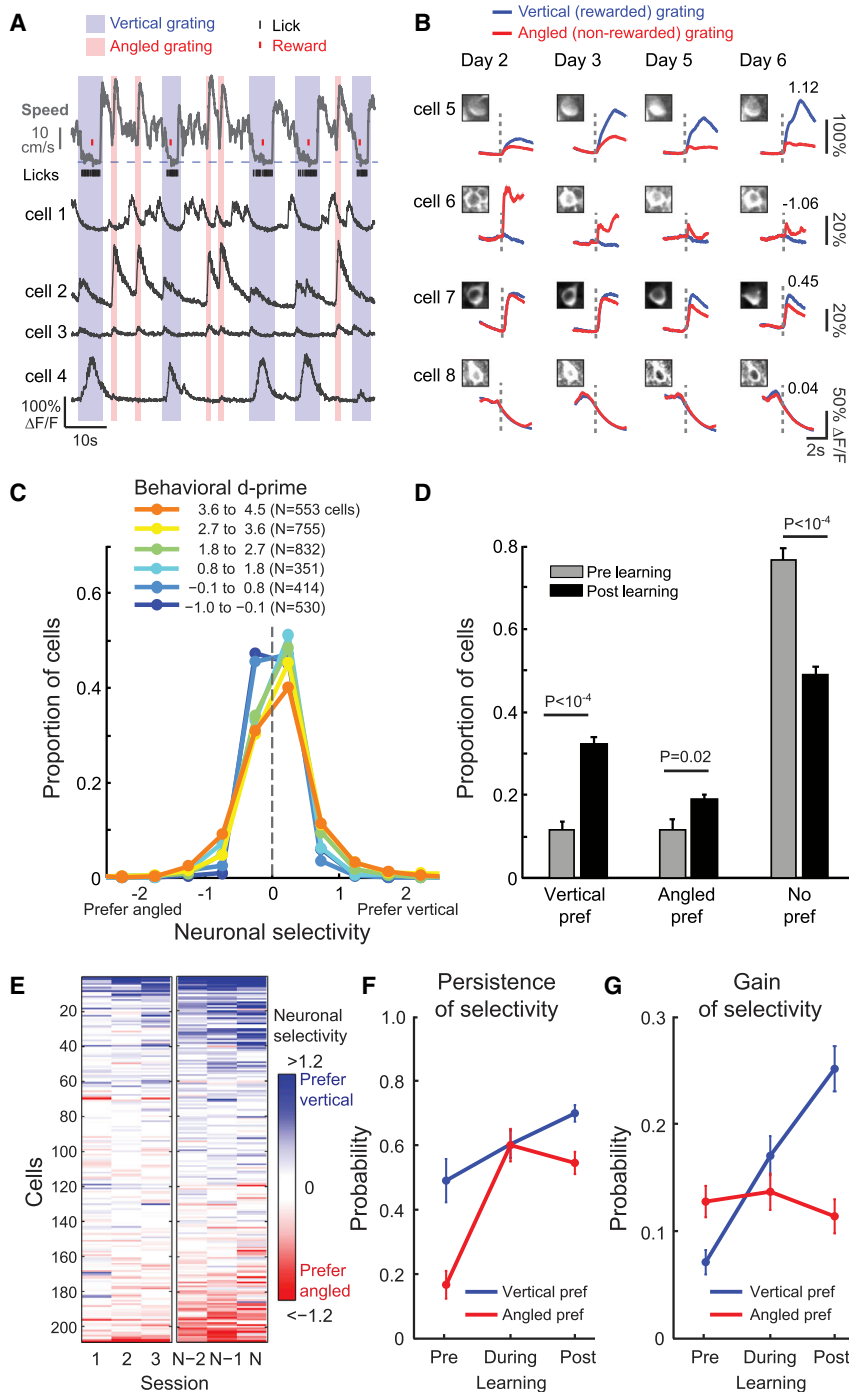


Figure 2. Chronic Two-Photon Imaging of Single Cells across Learning

(A) Example calcium traces of four V1 neurons during the task in an expert mouse, aligned to running speed (gray trace on top), licking (black lines), and reward delivery (red lines). Blue and red shading indicate time spent in the vertical and angled grating corridor, respectively.

(B) Average responses and corresponding images of four additional example cells in four training sessions aligned to grating onset (dashed vertical line). Values above each trace on day 6 denote neuronal selectivity for grating corridors, computed from responses 0–1 s after grating onset (see [Experimental Procedures](#)).

(C) Histograms of neuronal selectivity (positive values: cells prefer vertical, rewarded gratings; negative values: cells prefer angled, non-rewarded gratings) for different behavioral discrimination performance levels. Colors denote bins of behavioral d-prime from chance performance (blue) to expert performance (orange).

(D) Proportions of neurons significantly preferring the vertical or the angled grating or those without preference, before (sessions with behavioral d-prime < 1) and after learning (behavioral d-prime > 2); session mean \pm SEM computed from responses 0–1 s after grating onset.

(E) Grating selectivity of the same neurons (rows) across sessions (columns) in the first three and last three sessions; cells were ordered based on the selectivity averaged across the middle four sessions; $n = 8$ mice.

(F) Persistence of response selectivity across consecutive training sessions during different stages of learning. Values are the probability of a neuron with a grating preference on one day to maintain this preference on the next day within each learning stage (response 0–1 s after grating onset; vertical grating, $N_{pre} = 51$, $N_{dur} = 121$, $N_{pst} = 279$; angled grating, $N_{pre} = 90$, $N_{dur} = 95$, $N_{pst} = 200$ cells). Error bars depict SEM (determined by bootstrapping with replacement). Pre learning, behavioral d-prime (d') of both sessions < 1, and $\Delta d' < 0.5$ (14 session pairs); during learning: d' first session < 2, d' second session > 0.5, $\Delta d' > 0.5$ (14 session pairs); after learning: d' both sessions > 2, absolute $\Delta d' < 0.5$ (19 session pairs).

(G) The fraction of non-selective cells becoming selective for task-relevant stimuli across consecutive training sessions during different stages of learning (as in F). Values are the probability cells non-selective on one day ($N_{pre} = 549$, $N_{dur} = 417$, $N_{pst} = 422$) to develop a preference for one of the two gratings the next day within each learning stage. $n = 11$ mice for all panels, except where indicated. See also [Figures S2–S5](#).

conversion of unselective neurons into those more selective for the rewarded grating.

Progressive Increase of Population-wide Stimulus Discriminability in V1 with Learning

We next determined how these learning-related changes in single-neuron selectivity influenced the ability of neuronal populations to discriminate the grating stimuli. As a composite mea-

sure of selectivity in a population with both positive and negative selectivity indices, we computed the root-mean-square of grating selectivity of all neurons imaged simultaneously (population selectivity) over the time course of stimulus presentation (200 ms sliding window; see [Experimental Procedures](#)) for different training sessions, grouped by behavioral performance ([Figure 3A](#)). Neuronal population selectivity increased progressively with improving behavioral performance (pre learning =

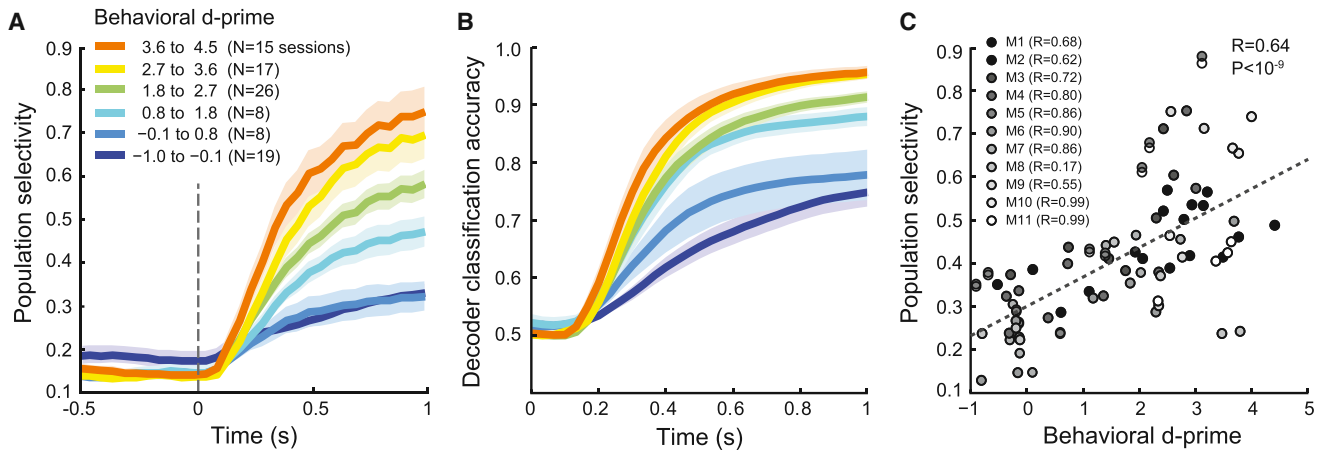


Figure 3. Learning Increases Neuronal Stimulus Selectivity in Populations of V1 Cells

(A) Time course of neuronal population selectivity (see [Experimental Procedures](#)) aligned to grating onset (dashed line; 200 ms sliding time window) for different behavioral performance levels as in [Figure 2C](#). Shading indicates SEM.

(B) Time course of classification accuracy of a linear decoder (probability of correctly identifying vertical versus angled grating corridor trials), based on cumulative neuronal activity of simultaneously imaged cells from grating onset for different behavioral performance levels. Shading indicates SEM.

(C) Relationship between population selectivity (average value 0–1 s after grating onset) and behavioral performance for individual sessions. Shades of gray indicate individual mice. $n = 11$ mice for all panels. See also [Figures S6](#) and [S7](#).

0.26, post learning = 0.46, $p < 10^{-4}$, bootstrap test, comparison within 0–1 s window post grating onset), and rose sharply after grating onset only in well-trained mice.

Additionally, we trained a linear decoder to predict which stimulus the mouse had encountered in each trial (vertical versus angled grating corridor) from calcium responses of all cells imaged simultaneously (see [Experimental Procedures](#)). The ability of the decoder to classify trials correctly increased strongly with improved behavioral performance during learning, such that classification accuracy exceeded 90% in expert mice ([Figure 3B](#)). Therefore, as mice got better at discriminating the two gratings, population-level representations of these task-relevant stimuli became increasingly distinguishable. In individual animals, neuronal population selectivity closely tracked the session-by-session changes in behavioral performance ([Figure S6](#)); there was a high positive correlation between the average population selectivity (0–1 s post grating onset) and the behavioral d-prime for individual sessions ([Figure 3C](#); $R = 0.64$, $p < 10^{-9}$, $n = 78$ sessions).

These results suggest that the increased selectivity of V1 neurons during training is a specific effect of learning the discrimination task. Indeed, neither response amplitude nor response selectivity for stimulus features in the approach corridor increased during learning ($p = 0.38$, pre- versus post learning, Wilcoxon signed-rank test), even though those features did evoke reliable responses in subsets of cells ([Figure S7](#)). Therefore, learning-related changes in V1 activity were specific to task-relevant grating stimuli and were not a consequence of repeated exposure to the same visual environment over multiple sessions ([Frenkel et al., 2006](#)).

Task Dependence of Learning-Induced Increases in Neuronal Selectivity

To what extent did these learning-related changes in V1 representations depend on the animals being engaged in the task? To

address this question, we trained expert mice to switch between blocks of the visual discrimination task and an analogous olfactory discrimination task. Mice learned to lick to obtain a reward in response to one of two different odors while running through the virtual corridor where they occasionally encountered the grating stimuli used in the visual discrimination task (see [Experimental Procedures](#)). Mice learned to switch rapidly between the two tasks within the same session, such that they successfully discriminated the grating stimuli in the visual task but ignored the same grating stimuli (while successfully discriminating odors) during the intervening olfactory blocks ([Figure 4A](#); see [Experimental Procedures](#), [Movie S1](#), and [Figures S8A–S8F](#)). Although the average response amplitudes to the grating stimuli did not change in the olfactory blocks ([Figure S9](#); $p > 0.32$), most neurons became less selective ([Figure 4B](#)), as the fractions of neurons preferring both the rewarded and non-rewarded stimuli decreased ([Figure 4C](#); all p values $< 10^{-4}$, bootstrap test). Consequently, population selectivity for the same grating stimuli decreased significantly in the olfactory blocks compared to the visual blocks ([Figure 4D](#), $p = 0.014$, bootstrap test), but remained above the pre learning level ($p = 10^{-4}$). Moreover, when the same visual stimuli were played back to fully trained but anesthetized mice, the selectivity of V1 populations was further reduced compared to the olfactory blocks ($p = 0.002$) but still higher than before learning ([Figure 4D](#), $p = 0.04$). These results indicate that there may be two causes underlying the learning-related increase in stimulus selectivity in V1: a more lasting, task-independent change in the visual circuits, and a task-dependent modulation that depended on the animals being engaged in visual discrimination. The fact that the selectivity of most neurons increased during visual discrimination ([Figure 4B](#)) suggests that the task-dependent signals mediating these effects have a widespread influence on neuronal populations in V1.

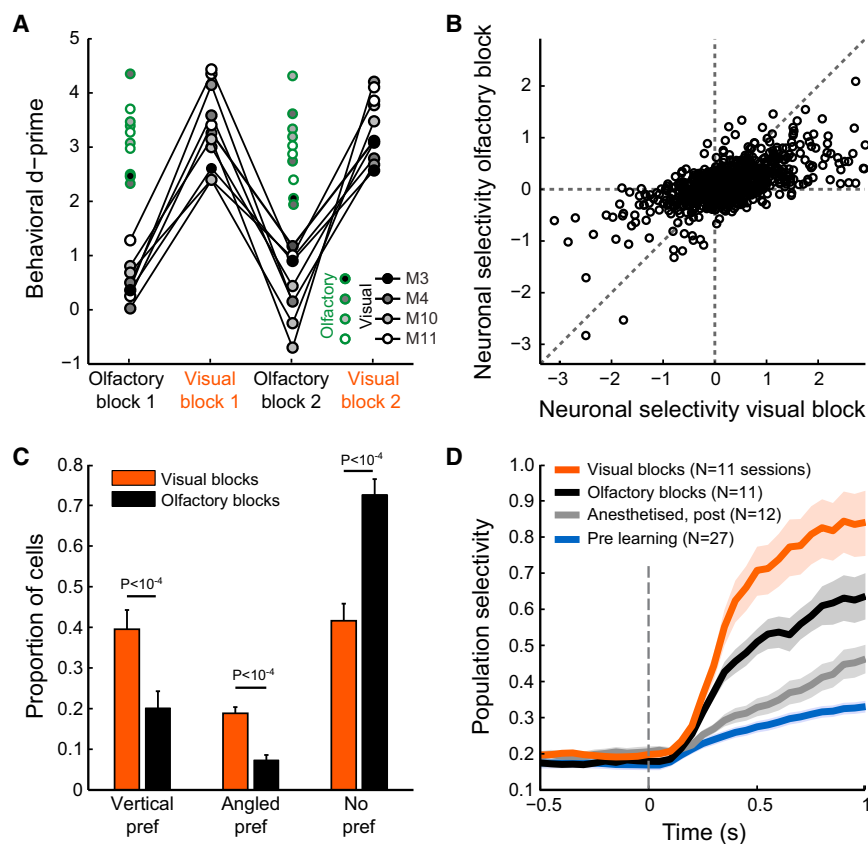


Figure 4. Learning-Related Increase in Neuronal Selectivity Is Partly Task-Dependent

(A) Visual discrimination performance (behavioral d -prime) in interleaved blocks of the visual discrimination task and an analogous olfactory discrimination task in which the same grating stimuli were shown but ignored by the animals (black circles and lines; shades of gray indicate individual mice). Odor discrimination performance in the olfactory blocks is additionally shown (green circles). $n = 4$ mice.

(B) Selectivity for grating stimuli of individual neurons during the visual and the olfactory task. The majority of neurons reduce their selectivity for the same grating stimuli during olfactory discrimination. (C) Proportions of neurons significantly preferring the vertical or the angled grating, or those without preference in the visual and the olfactory task ($n = 11$ sessions).

(D) Time course of visual stimulus population selectivity aligned to grating corridor onset (dashed line) in the visual task (red, gratings relevant), in the olfactory task (black, gratings irrelevant), when the same stimuli are presented during anesthesia (gray, $n = 11$ mice), and before learning (blue, $n = 8$ mice). Error bars and shading are SEM. See also Figures S8 and S9.

Changes in Motor Behavior with Training Cannot Account for the Increase in Neuronal Selectivity

Several possible causes may underlie the task-dependent changes of stimulus selectivity in V1 during learning. Responses to task-relevant stimuli could be specifically modified to give rise to more distinguishable representations at the population level, thus allowing for easier perceptual discrimination. In addition, changes in V1 activity could also reflect signals associated with the behavioral outcome of the task, including signals related to the animals' motor behavior, which are known to modulate the activity of V1 neurons (Niell and Stryker, 2010; Keller et al., 2012; Saleem et al., 2013). Neither the average running speed nor running speed variability at grating onset changed systematically over the course of training (median speed = 45.3 cm/s before, 43.4 cm/s after learning, $p = 0.46$; median SD = 12.2 cm/s before, 13.9 cm/s after learning, $p = 0.84$, Wilcoxon signed-rank test). However, the running profile after the animals had identified the grating did change with training: mice slowed down in response to the rewarded grating and increasingly accelerated when detecting the non-rewarded grating during learning (see Figures 1C and S10 for more examples from different mice).

To determine whether these changes in running behavior or an associated change in optic flow speed could explain the learning-related increase of stimulus discriminability in V1, we carried out several independent controls. First, we trained a separate set of animals in a modified version of the task in which expert mice encountered grating corridors whose optic flow was uncoupled from running speed for 1 s and exactly matched to pre learning

optic flow speed profiles (Figures S11A and S11B; see Supplemental Experimental Procedures). In separate experiments, gratings were presented at a fixed speed profile during the visual and olfactory discrimination task (Figures S11C–S11G). Note that in these tasks, gratings were always preceded by a gray corridor to ensure also that the visual input preceding the task-relevant stimuli was uniform across conditions. Thus, when the speed profiles of task-relevant visual stimuli were identical in all conditions, we again found that V1 neurons increased their grating selectivity over the course of learning, as well as when the animals engaged in the visual compared to the olfactory discrimination task (Figures S11B, and S11F and S11G, respectively).

Second, we tested if locomotion-related response modulation in V1 influenced the learning-related changes in neuronal selectivity. We did not observe any speed-related differences in neuronal population selectivity computed from trials with matched running speed profiles in all conditions (Figures 5A and 5B). Specifically, V1 neurons showed increased selectivity after learning independent of running speed (Figure 5A; population selectivity within 0–0.5 s from grating onset, pre- versus post learning, slow: $p = 0.02$; fast: $p = 0.01$, bootstrap test). Moreover, while some neurons showed a correlation between their calcium signal and running speed, as expected from previous studies (Niell and Stryker, 2010; Keller et al., 2012; Saleem et al., 2013), there was no positive relationship between how strongly cells were modulated by running and/or optic flow speed and their change in grating selectivity over learning (Figures S12A and S12D; see Experimental Procedures). Indeed, the exclusion of neurons whose responses were modulated by running did not alter the increase in V1 population selectivity over learning (Figures S12B and S12C, and S12E and S12F).

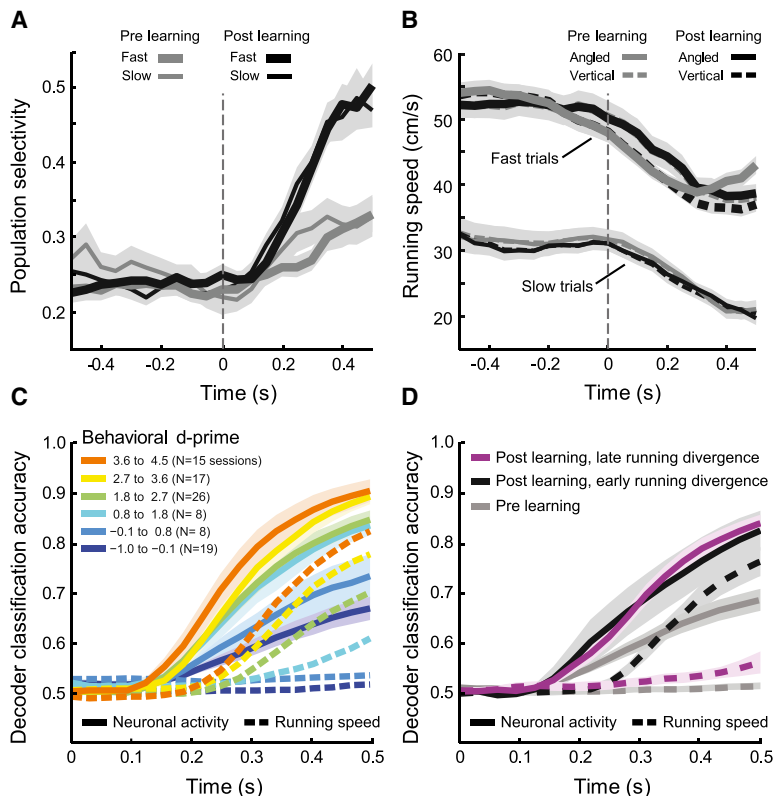


Figure 5. Neuronal Changes during Learning Cannot Be Explained by Changes in Running Behavior

(A) Population selectivity for different running speeds (thick traces, fast running trials; thin traces, slow running trials) matched across sessions before (behavioral d-prime < 1, gray traces) and after learning (behavioral d-prime > 2, black traces). Data are from 27 sessions pre learning and 52 sessions post learning. See also Figures S10–S13.

(B) Average running speeds corresponding to conditions in (A). Solid lines indicate vertical grating trials and dashed lines angled grating trials. There was no difference in running speeds within the same speed bin across stimuli, nor before/after learning (all comparisons $p > 0.08$).

(C) Time course of decoding performance from grating onset (probability of correct classification of grating corridor type, vertical versus angled). Decoder was based either on cumulative neuronal activity (solid lines) or cumulative running speed (dashed lines) for different behavioral discrimination performance levels during learning. $n = 11$ mice for (A)–(C).

(D) Decoding performance as in (C), before learning (behavioral d-prime < 1, gray lines, 23 sessions, 11 mice), and after learning (behavioral d-prime > 2) for sessions with delayed divergence of running behavior in vertical and angled grating trials (see Experimental Procedures; purple, 8 sessions, 7 mice), and sessions with matched behavioral d-prime but early divergence of running behavior (black lines, 8 sessions, 6 mice). In all panels, 0 s = grating corridor onset.

Third, while there was some modulation of V1 activity by signals related to the animals' licking, excluding neurons modulated by licking did not change the learning effect (Figures S12G–S12L). Fourth, we found that any signals related to eye position, eye movements, and pupil size could not account for the increased neuronal selectivity after learning (see Supplemental Information and Figures S13A–S13F). Furthermore, we conducted similar analyses to control for any differences in motor behavior during the visual and olfactory discrimination task (Figures S8G–S8J and S13G–S13J), and found that variations in locomotion, licking, eye movements, or pupil size could not explain the task-dependent improvements of neuronal selectivity in V1.

Finally, we trained the linear decoder introduced above on either the population activity of V1 neurons or the running speed of the mouse to predict trial type (vertical versus angled grating corridor; see Experimental Procedures; Figure 5C). Due to the systematic divergence of running speed after mice had entered the grating corridors (see above), the ability of the decoder to classify trials correctly based on running speed strongly improved over learning. However, the decoder trained on V1 activity allowed for earlier classification of the stimulus than the decoder trained on running speed (Figure 5C, top behavioral d-prime bin V1 activity versus running speed at 150 ms, $p < 10^{-4}$, bootstrap test). Indeed, even the short-latency V1 activity before running speed divergence (typical divergence > 220 ms after stimulus onset) allowed for a significant improvement in grating classification during learning (bottom versus top behavioral d-prime bin at 220 ms, $p = 0.001$, bootstrap test). Importantly, in post learning sessions (behavioral d-prime > 2), during which mice showed a delayed divergence in

their running speeds in response to the rewarded and non-rewarded gratings (running divergence > 400 ms after grating onset), neuronal activity allowed for an equally early and accurate classification of the grating stimuli compared to sessions with matched behavioral d-prime but with earlier running speed divergence (neuronal decoding performance early versus late running divergence: $p > 0.1$ for all time bins 0 – 0.5 s from grating onset; Figure 5D). Therefore, learning led to improvements in the ability of V1 populations to discriminate task-relevant stimuli before the animal acted on its decision either to slow down and lick for reward, or to speed up and suppress licking. Taken together, the increase of neuronal selectivity in V1 with training cannot be explained by the modulation of V1 activity by any of the measured motor parameters (running, licking, eye movements, pupil dilation) nor by any differences in optic flow before and after learning.

The Emergence of Signals Reflecting Behavioral Outcome during Learning

The information related to the animal's own action is not the only non-sensory signal that can influence V1 activity. Other task-related signals relaying information about the attentional state, expectations, or behavioral choice have also been observed in visual cortical areas (Moran and Desimone, 1985; Britten et al., 1996; Shuler and Bear, 2006; Stănişor et al., 2013; Nienborg and Cumming, 2014). To identify such signals in V1 activity, we compared responses to the non-rewarded, angled grating during correct rejection trials (CR, mouse withheld licking and accelerated) and false-alarm trials (FA, mouse incorrectly licked and slowed down). Because the visual stimulus identity during CR and FA

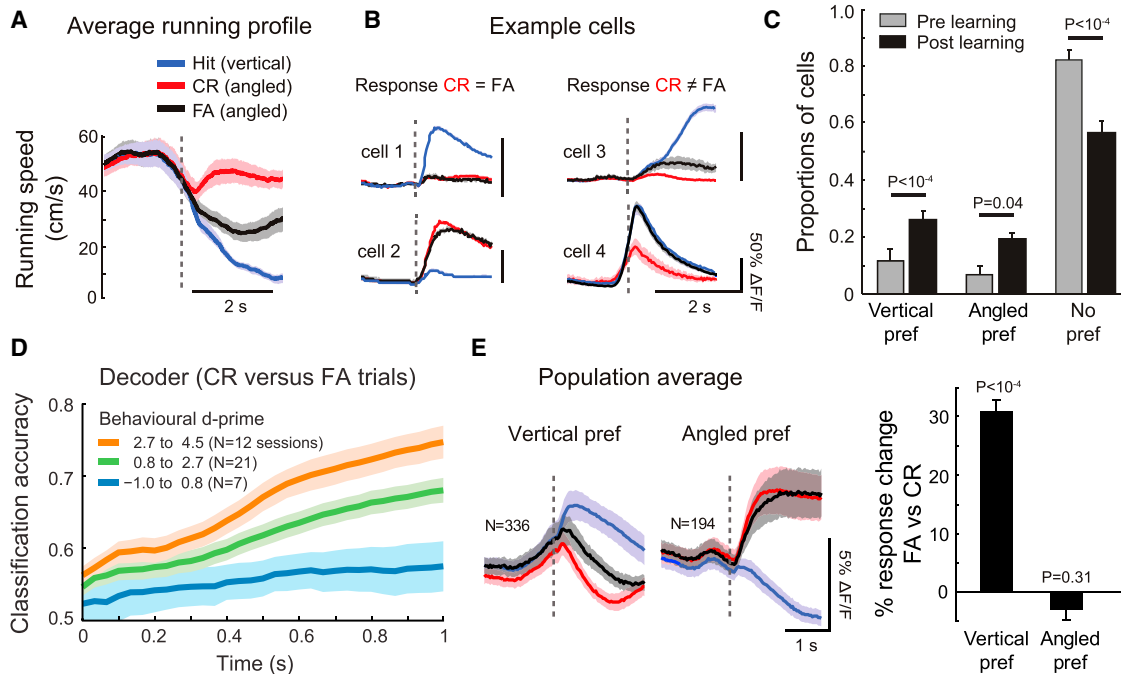


Figure 6. Emergence of Signals Related to Behavioral Trial Outcome during Learning

(A) Average running speed profile in different trial types: hit, correct rejection (CR) and false alarm (FA) (17 sessions). Dashed line indicates grating corridor onset. (B) Average grating responses of four example cells classified as behaviorally-modulated (cells 3 and 4; responses are different in FA and CR trials) and behaviorally not modulated (cells 1 and 2; responses are similar in FA and CR trials). (C) Proportions of neurons not modulated by behavior (only neurons with similar responses in FA and CR trials) which significantly preferred the vertical or the angled grating and without preference, before (behavioral d-prime < 1) and after learning (behavioral d-prime > 2); session mean ± SEM, 0–1 s after grating onset. (D) Time course of decoding performance, i.e., the probability of correct classification of CR and FA trials based on cumulative neuronal activity for three different behavioral discrimination performance levels during learning. (E) Average responses of cells preferring the vertical, rewarded grating (left) or angled, non-rewarded grating (middle) on Hit (blue), CR (red) and False alarm (black) trials. Right, change in response strength 0–1 s after angled grating onset during FA trials relative to CR trials. n = 11 mice for all panels.

trials was the same but the behavior of the animal was different (i.e., stopping and licking versus running; see also Figure 6A), we could identify neurons whose responses were not behaviorally modulated (no significant response difference between CR and FA trials despite a strong difference in behavior; Figure 6B) and those that were (significantly different responses between CR and FA trials; Figure 6B). When we excluded all behaviorally modulated cells from the analysis, we still found that the proportion of neurons selective for the rewarded and non-rewarded gratings significantly increased over learning (Figure 6C; all p values < 0.04, bootstrap test), similar to the effects for the entire population (Figure 2D). These results again demonstrate that the improvement in V1 selectivity for both task-relevant stimuli after learning is not caused by signals related to the change in the animals' behavior during learning, associated changes in optic flow speed, or task-related signals such as reward expectation.

Importantly, however, visually evoked activity of many cells was modulated by the behavioral response (up to 40% of selective neurons; Figure 6B). This difference was apparent at the population level because a decoder trained on predicting the behavioral choice in response to the non-rewarded grating (CR versus FA trials) from neuronal activity of all cells performed above chance and improved with learning (Figure 6D; highest versus lowest behavioral d-prime, p = 0.01, bootstrap test). Inter-

estingly, on average, neurons preferentially responding to the rewarded grating showed significantly different responses between CR and FA trials, while neurons preferring the non-rewarded grating did not (Figure 6E; rewarded-stimulus-preferring cells, p < 10⁻⁴, n = 336; non-rewarded-stimulus-preferring cells, p = 0.31, n = 194, Wilcoxon rank-sum test). Therefore, signals related to the behavioral outcome developed over learning and mainly influenced a specific subgroup of neurons preferring the rewarded stimulus.

The Emergence of Anticipatory Signals during Learning

Analysis of neuronal activity just before the onset of the grating corridors revealed another task-dependent signal that developed during training, presumably related to the animals' anticipation. While mice started each new trial at a different, random position in the approach corridor, the abrupt onset of the grating corridors was always preceded by the same pattern of black and white circles on the corridor walls (see Figure S7A). Some neurons increased their activity just before grating onset with learning (Figure 7A), suggesting that they had developed anticipatory signals (Jaramillo and Zador, 2011; Totah et al., 2013), which might reflect the animals' ability to eventually predict and anticipate the time point of appearance (but not the identity) of the grating corridors from the preceding corridor wall pattern.

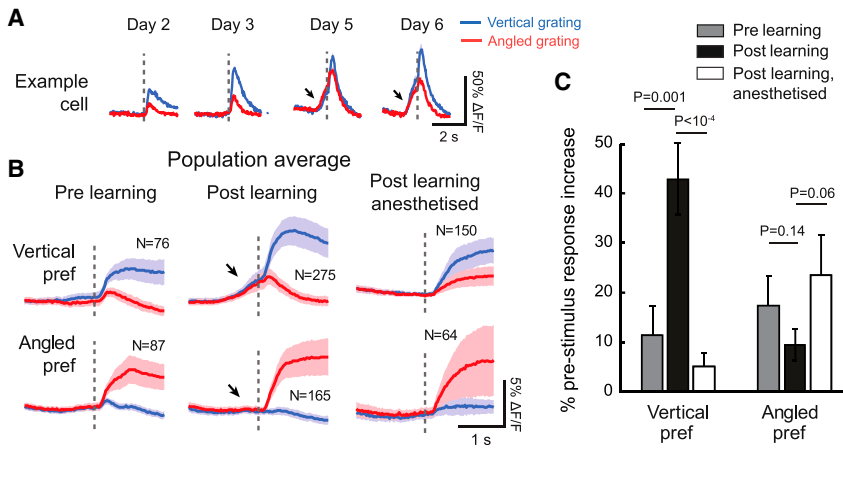


Figure 7. Emergence of Anticipatory Signals during Learning

(A) Increase in ramping activity prior to grating corridor onset (dashed line) in an example cell over the course of learning.

(B) Average population responses to the vertical (blue) and angled grating (red) for vertical- (top) and angled-preferring (bottom) cells in the first training session (pre learning), in the first session after learning (post learning, behavioral d -prime > 2), and during anesthesia post learning.

(C) Relative increase in pre-stimulus activity within the second preceding grating onset (see [Experimental Procedures](#)) for vertical and angled grating-preferring cells in the first training session, the first session post learning, and when virtual reality stimuli were played back to trained animals during anesthesia. p values from bootstrap test; error bars and shading are SEM. $n = 11$ mice.

Importantly, only the neurons preferring the rewarded stimulus, and not the neurons preferring the non-rewarded stimulus, developed this pre-stimulus activity increase during learning (Figures 7B and 7C, pre- versus post learning, rewarded-stimulus-preferring cells: $p = 0.001$; non-rewarded-stimulus-preferring cells: $p = 0.14$, bootstrap test). The existence of these specific, putative anticipation signals was supported by a significant decrease in pre-stimulus activity during anesthesia after learning only in cells preferring the rewarded grating (Figure 7C, rewarded-stimulus-preferring cells: $p < 10^{-4}$; non-rewarded-stimulus-preferring cells: $p = 0.06$, trend in the opposite direction, bootstrap test). Taken together, non-sensory signals, both before and after appearance of the task-relevant stimuli, seem to influence primarily a specific ensemble of cells that preferentially responded to the stimulus that predicts the reward.

DISCUSSION

We show that learning leads to concerted changes in how L2/3 neurons in V1 process visual and non-visual signals related to the behavioral task. By tracking individual neurons during learning, we observed a net recruitment and stabilization of neurons selective for task-relevant stimuli, resulting in improved stimulus discriminability at the population level, which closely correlated with the behavioral performance of the animals. These learning-induced enhancements of stimulus representation in V1 diminished substantially when animals did not engage in the visual discrimination task, suggesting that putative top-down signals contribute to increased population-level discriminability. In parallel, we observed the emergence of additional task-dependent signals in a specific subpopulation of cells—neurons preferentially responding to the rewarded stimulus developed anticipatory responses prior to the appearance of task-relevant stimuli and additional activity related to the animal's behavioral choice after stimulus onset.

Learning-Related Changes in Mouse V1

We developed a visually guided task in which head-fixed mice learned to discriminate two grating patterns in a virtual reality environment in which the animals' running controlled their posi-

tion in a corridor (Hölscher et al., 2005; Dombeck et al., 2007). Most mice learned to perform this task with high behavioral accuracy within 1 week (behavioral d -prime > 3 , corresponding to accuracy levels of $> 90\%$). We speculate that task acquisition was facilitated by the fact that mice had active control over their visual environment (locomotion coupled to visual feedback), resulting in a more naturalistic visual experience (Gibson, 1979) that seemed to promote engagement in the task. We showed that task performance was dependent on visual cortex activity and, importantly, that responses of V1 neurons to task-relevant stimuli became progressively more distinguishable, leading to more selective task-relevant information in V1 circuits.

The closed-loop nature of behavioral tasks in virtual reality makes it necessary to separate sensory and motor influences on neuronal responses. Specifically, it was important to control for the changes in running speed and the resultant changes in the optic flow speed over the course of training in relation to the observed changes in V1 activity. The learning-related increase in neuronal selectivity did not decrease (1) when comparing responses only in running speed-matched conditions before and after learning, (2) when comparing responses to identical optic flow before and after learning, (3) when excluding neurons from the analysis whose responses were modulated by running and visual flow speed, and (4) when only including neurons with similar responses to the same grating in FA and CR trials even though the animals' behavior (running speed and licking) and the optic flow differed. Moreover, learning-induced increase in V1 selectivity did not diminish when controlling for licking, eye position, eye movements, and pupil size. Therefore, the improvement in V1 stimulus discriminability during training could not be accounted for by any changes in the animals' motor behavior we could measure or by associated changes in visual input.

Finally, even though somatic GCaMP6 signals the occurrence of spiking with a slight delay (time to peak for one action potential $> \sim 40$ ms; Chen et al., 2013), we found improved discriminability of task-relevant stimuli in V1 within approximately 200 ms after stimulus onset, which preceded the animal's behavioral response and changes in locomotion. This suggests that learning may increase the salience of information relayed to

downstream areas to better inform behavioral decisions. Importantly, these results are comparable with those of a recent study of learning-related changes in V1 of macaque monkeys using multiunit recordings (Yan et al., 2014), suggesting that learning exerts similar effects on a primary sensory cortex in rodents and primates.

Selectivity Changes in Individual Neurons during Learning

Tracking the activity of the same identified cells throughout learning allowed us to investigate which changes in single cells underlie population-wide improvements in stimulus selectivity. Previous studies in visual cortex have shown differences in orientation tuning at or close to task-relevant grating orientations in animals trained in visually guided tasks compared to control conditions (Schoups et al., 2001; Yang and Maunsell, 2004; Goltstein et al., 2013). These results suggest that increases in population selectivity might have been mainly due to an increase in response selectivity of neurons that already had shown some orientation tuning before learning. However, we did not find that learning-related changes are especially pronounced in or even restricted to neurons with particular visual response properties. Specifically, neurons already selectively responding to one of the two task-relevant grating stimuli before learning were not more likely to increase their selectivity than non-selective neurons during learning.

One change in single cell responses that led to increased stimulus discriminability at the population level was a learning-induced decrease in day-to-day fluctuations of selectivity for task-relevant stimuli in individual neurons, akin to response stabilization observed in the motor cortex (Huber et al., 2012; Peters et al., 2014). Neurons preferring either the rewarded or the non-rewarded stimulus became more likely to maintain their response selectivity across consecutive training sessions. In parallel, we found an increased recruitment of previously non-selective neurons to become selective for the rewarded grating stimulus during training, which may explain the larger proportion of neurons selective for this stimulus in expert mice.

Task Engagement Enhances Neural Selectivity in V1

We successfully trained mice to switch between a visual and an olfactory discrimination task several times within the same training session. Mice ignored the grating stimuli during the olfactory discrimination task, and this allowed us to test whether the learning-related enhancement in task-relevant visual stimulus processing was hardwired or task-dependent. Population-level discriminability for grating stimuli was reduced but not decreased to pre learning levels when expert animals were not engaged in the visual discrimination task. Therefore, learning led to both task-independent and task-dependent enhancements in the processing of relevant stimuli in V1. Task-independent changes likely reflect more persistent alterations to visual circuits, akin to those previously observed outside the task or under anesthesia in visual cortex after learning (Schoups et al., 2001; Yang and Maunsell, 2004; Goltstein et al., 2013). The existence of task-dependent changes, however, suggests that non-sensory signals directly contribute to the enhanced processing

of behaviorally relevant stimuli (Li et al., 2004, 2008; Polley et al., 2006). Such modulatory signals, which depend on the animals' behavioral context, could be relayed by excitatory projections of cortical or subcortical origin (Krauzlis et al., 2013; McAlonan et al., 2008; Zhang et al., 2014), or may additionally involve cholinergic input from the basal forebrain (Pinto et al., 2013). Importantly, we found that these signals seem to increase the selectivity of most neurons encoding both the rewarded and non-rewarded stimuli when animals actively engaged in visual discrimination.

Emergence of Task-Specific Anticipatory and Behavioral-Choice-Related Signals in V1

Coinciding with the changes in the representations of task-relevant stimuli, we observed the appearance of two additional types of task-dependent signals during learning. First, neurons preferring the rewarded stimulus developed anticipatory responses prior to the appearance of task-relevant stimuli (Jaramillo and Zador, 2011; Totah et al., 2013). These signals are unlikely to be visually evoked, as they are not visible in neurons preferring the vertical grating before learning or under anesthesia. Instead, they likely arise through the learned association between a specific corridor position and the appearance of a grating stimulus, suggesting that processing in V1 is influenced by stimulus expectation, perhaps to prime activity in those neurons whose firing best predicts a reward. These anticipatory signals may thus reflect reward expectation (the rewarded stimulus will appear with 50% likelihood). For example, they could be the neural signature of a type of "wishful thinking" by the animals—stimulus expectation that preferentially evokes the cortical representation of the rewarded and therefore preferred stimulus.

Over the course of training, some neurons also increasingly exhibited enhanced responses during error trials in which the animals incorrectly sought reward in response to angled gratings, suggesting their activity might be related to the animal's behavioral choice (Britten et al., 1996; Ress and Heeger, 2003; Nienborg and Cumming, 2014), or reward expectation as previously observed in V1 (Shuler and Bear, 2006; Stănișor et al., 2013). Importantly, both the anticipatory and the behavioral choice-related signals emerged predominantly in neurons responding preferentially to the rewarded stimulus. We hypothesize that these signals may arise by strengthening of inputs from areas encoding reward expectation (e.g., orbitofrontal cortex; Tremblay and Schultz, 1999). Activity-dependent Hebbian mechanisms would permit this strengthening to occur specifically on V1 neurons preferring the rewarded stimulus, because these are consistently active before and during the time of reward delivery. With learning, as the animals increasingly develop an expectation of reward (i.e., just before and during the task-relevant stimulus appearance), neurons preferring the rewarded stimulus in V1 would be preferentially activated by projections conveying these putative top-down signals. This mechanism may act in concert with cholinergic signaling that has been proposed to explain reward timing-related plasticity in V1 (Chubykin et al., 2013).

The appearance of non-sensory signals in neuronal ensembles preferring the stimulus associated with a reward contrasts

with the modulation of sensory stimulus responses when mice were engaged in the visual discrimination task, which acted more generally by increasing the selectivity of neurons encoding both the rewarded and the non-rewarded stimuli. Identifying the sources of these diverse task-dependent signals is an important next step for clarifying their role in shaping early sensory processing. The sophisticated genetic tools available in mice will help elucidate the role of the many cortical and subcortical areas providing input to V1 during learned behaviors, as well as specific inhibitory cell types or different neuromodulator systems in the emergence and expression of learning-related changes.

In summary, as a mouse learns the behavioral significance of a visual stimulus, the responses of L2/3 neurons in V1 become more selective for task-relevant stimuli, leading to enhanced stimulus discriminability at the population level. In parallel, multiple task-dependent signals emerge during learning and differentially influence the firing of neurons within the V1 circuit. This demonstrates the remarkable flexibility by which a primary sensory cortex can tailor its processing to the requirements of a task and to the behavioral relevance of sensory stimuli.

EXPERIMENTAL PROCEDURES

Surgical Procedures and Imaging

All experimental procedures were carried out in accordance with the institutional animal welfare guidelines and licensed by the UK Home Office and the Swiss cantonal veterinary office. A virus expressing GCaMP6f or GCaMP6m (AAV2/1-hsyn-GCaMP6-WPRE; Chen et al., 2013) was injected in the primary visual cortex (V1) in the right hemisphere of C57Bl/6J mice (P49–P57). Imaging and behavioral training started approximately 3 weeks after surgery. We imaged GCaMP6-labeled neurons in layer 2/3 in 93 training sessions and 12 recording sessions under isoflurane anesthesia in 11 mice with a custom-built resonant scanning two-photon microscope with a frame rate of 32 Hz. [Supplemental Experimental Procedures](#) contain further details about surgical and imaging procedures.

Behavioral Tasks

Mice were head-fixed and trained to run on a styrofoam cylinder. A reward delivery spout was positioned near the snout of the mouse, and licks were detected using a piezo disc sensor. Mice were then trained in a visual discrimination task in which the running speed on the cylinder was detected with an optical mouse and used to control the speed at which mice moved through a virtual environment presented on two screens in front of them. A trial started when the mouse was positioned at a random starting point in an approach corridor with walls showing black and white circles on a gray background. When the mouse reached a specific point in the corridor, it was randomly teleported to one of two grating corridors with either a vertical or an angled grating on the walls. In the vertical grating corridor, the mouse was rewarded with a drop of soya milk, for licking the spout after it had entered a “reward zone,” a short distance into the grating corridor. No punishment was given for licking in the angled grating corridor.

A subset of mice was trained to switch between blocks of an olfactory and visual discrimination task. In olfactory blocks, mice performed an analogous olfactory go-no go discrimination task in which they were rewarded for licking in response to one of two odors. During this task, mice were also presented with the vertical and angled grating corridor at different positions in the approach corridor. Mice learnt to ignore these irrelevant grating stimuli while accurately discriminating the odors. On switching to the visual block, mice started licking selectively to the rewarded grating as before. See [Supplemental Experimental Procedures](#) for further details about the visual stimulus, behavioral tasks, and training.

Bilateral Optogenetic Silencing of V1 Activity

Bilateral silencing of V1 was carried out in four transgenic mice (three males, one female) expressing channelrhodopsin-2 in parvalbumin-expressing interneurons (Hippenmeyer et al., 2005; Madisen et al., 2012). Additionally, three male wild-type C57Bl/6J mice underwent identical surgical and experimental procedures. Mice were implanted with two cranial windows over both visual cortices. Intrinsic imaging was used to determine the extent of V1, and all regions excluding V1 were covered with black paint. In expert mice (>90% performance levels), V1 was silenced by illuminating both cranial windows with 470 nm light at one of four intensities shortly before and during the grating corridor. In 30% of trials no light stimulation was applied. The same mice were also trained on an olfactory discrimination task as described above (but without grating stimuli). V1 was silenced shortly before and during presentation of the odors. For further details, see [Supplemental Experimental Procedures](#).

Data Analysis

Image stacks were corrected for motion, and regions of interest (ROIs) were selected for each cell in each session. Raw fluorescence time series $F(t)$ were obtained for each cell by averaging across pixels within each ROI. Baseline fluorescence $F_0(t)$ was computed by smoothing $F(t)$ (causal moving average of 0.75 s) and determining for each time point the minimum value in the preceding 60 s time window. The change in fluorescence relative to baseline, $\Delta F/F$, was computed by taking the difference between F and F_0 , and dividing by F_0 .

To analyze responses to the vertical and angled grating corridors, neuronal activity was aligned to the onset of the grating corridor for each trial. A Wilcoxon rank-sum test was used to determine if responses—the average $\Delta F/F$ in a time window of 1 s after grating onset—in the two conditions were significantly different ($p < 0.05$), and the sign of the difference determined the response preference. The persistence of stimulus preference (Figure 2F) was defined as the probability that a cell that significantly preferred one of the two gratings on one day also preferred the same grating on the next day. Recruitment of non-selective cells (Figure 2G) was defined as the probability that a cell with no stimulus preference on one day became selective to one of the two gratings on the next day. We computed these measures for three stages of learning, based on the behavioral d-prime (bDP) of two consecutive sessions: before learning (bDP of both sessions < 1 , and $\Delta bDP < 0.5$, $N_{\text{session}} = 14$), during learning (bDP of first session < 2 , bDP second session > 0.5 , and $\Delta bDP > 0.5$, $N_{\text{session}} = 14$), and after learning (both bDP > 2 and absolute change in bDP < 0.5 , $N_{\text{session}} = 19$). Varying the criteria to define different stages of learning led to similar results (data not shown).

To quantify the selectivity of neural responses we computed a response selectivity index (SI) for individual cells from the difference between the mean response in the first second after grating onset to the vertical and angled grating corridor, divided by the pooled standard deviation of the responses

$$SI = (\bar{R}_V - \bar{R}_A) / s_p^{VA},$$

where

$$s_p^{VA} = \sqrt{\sum_{i=1}^{k=2} (n_i - 1) s_i^2 / \sum_{i=1}^k (n_i - 1)},$$

and n_i is the number of trials in condition i for k conditions. Therefore, positive values indicate a preference for the vertical grating corridor and negative values a preference for the angled grating corridor. Please note that in the manuscript text the term selectivity substitutes for SI. To obtain a combined measure of grating discriminability for simultaneously imaged populations of neurons, population selectivity was computed by taking the average of the squared selectivity index across cells and taking the square root:

$$\sqrt{\left(\sum_i^{N_{\text{cell}}} SI_i^2 \right) / N_{\text{cell}}}.$$

A bootstrap test (Efron, 1979) was used to test for significant differences between conditions that contained both dependent and independent data points. To test whether changes in the proportion of cells preferring the vertical or angled grating, or without preference across two conditions (typically before and after learning), were significant, we first computed for each session the

proportions of cells in each category. Next, we randomly picked the same number of sessions (the minimum across conditions) from both conditions, and repeated this 10,000 times. We then computed in both conditions the average cell proportion across sessions, and we also computed the proportion after randomly assigning sessions to one of the two conditions. The p value was given by the number of bootstraps in which the proportion change in the actual data was greater than the proportion change with randomly assigned condition labels. Similarly, bootstrapping was also used to assign significance to the differences in population selectivity, decoding performance, and pre-stimulus activity increase, by comparing the difference in the original data to the difference with randomly assigned condition labels.

To control for the effect of running speed and optic flow on neural responses and selectivity across learning, grating responses were compared specifically in trials that were matched for running speed across sessions and stimulus conditions (Figures 5A and 5B). First, the average running speed was determined in sliding 200 ms time windows from -0.5 to $+0.5$ s around the onset of the grating corridor (50 ms step size). Then responses in each time window of each trial were assigned to one of three groups, depending on running speed (three bins divided equally from the 2.5% percentile to the 97.5% percentile of the average running speed, across all sessions). Data for each time window were only included if it contained at least ten trials of both grating conditions. In the highest speed bin, not enough matched data were available across learning, thus restricting the analyses to the lowest speed bin (referred to as “slow”) and the intermediate speed bin (“fast”).

To quantify the accuracy with which two conditions (either trials with vertical and angled grating corridors (Figures 3B, 5C, and 5D) or FA and CR trials (Figure 6D) could be classified at time t relative to grating onset, a cumulative decoder was employed. From training data (30 trials of both conditions), the decoder constructed for each neuron n a model of the response using as parameters the mean response to the vertical ($\mu_n^V(t)$) and angled grating corridor ($\mu_n^A(t)$) and the variance of the noise σ_n to maximize the observed log-likelihood of the data under a Gaussian noise model. On test trials (the remaining trials that were not used as training data), the log-likelihood at time t that trial k belongs to condition C (where C was for instance the vertical (V) or angled grating corridor (A) condition) is proportional to

$$L_C(t) = - \sum_n^{N_{\text{cell}}} \sum_0^{T_{\text{start}}} (D_{n,k}(t - T_{\text{start}}) - \mu_n^V(t - T_{\text{start}}))^2 / (2\sigma_n^2),$$

where D indicates deconvolved $\Delta F/F$ (see Supplemental Experimental Procedures). If $L_V > L_A$, the trial was assigned to the vertical condition, otherwise to the angled grating condition. To obtain at each time point t the cumulative likelihood L_C , the summation only included time points starting from T_{start} , which was the time of the grating onset, up until time t . Note that without the temporal accumulation of log-likelihood, the decoder would be equivalent to a linear discriminant analysis. To determine the time point at which there was a detectable divergence of running speed between vertical and angled grating trials, we performed a Wilcoxon rank-sum test on the average speed in nonoverlapping, consecutive 50 ms windows. The time of divergence was defined as the center of the first window with $p < 0.01$ followed by $p < 0.01$ in at least four consecutive windows. For Figure 5D, we defined post learning sessions with delayed divergence as sessions with behavioral d-prime > 2 and time of running speed divergence greater than 400 ms ($n = 8$ sessions in $n = 7$ mice, average d-prime 2.59). We paired each of these sessions to a unique session with the smallest difference in behavioral d-prime, but with time of divergence < 400 ms ($n = 8$ sessions in $n = 6$ mice, average d-prime 2.61).

To analyze responses during FA and CR trials, only sessions with at least 15 FA trials were included in the analysis (Figure 6). These were predominantly sessions at intermediate learning stages, as most expert mice made very few mistakes by the end of training (see Figure S1). Behaviorally modulated cells were defined as cells with significantly different activity for FA and CR trials in the first second after grating corridor onset ($p < 0.05$, Wilcoxon rank-sum tests). To obtain average responses for cells preferring the vertical or the angled grating corridor (Figure 6E), neurons were classified as vertical (or angled) preferring if they significantly preferred the vertical (or the angled) grating corridor in at least one session and never switched preference, and re-

sponses of such cells were averaged across the sessions in which they showed a significant preference.

The relative response increase before grating onset (Figure 7C) was calculated for each cell as the difference in the average $\Delta F/F$ signal between two time windows, -0.25 s to 0 s and -1 s to -0.75 s, divided by the average $\Delta F/F$ signal in the -1 to -0.75 s window, where $t = 0$ s is time of grating onset. To compare pre-stimulus responses before and after learning, responses were averaged on the first day of training and the first day post learning (behavioral d-prime > 2) for each cell. Neurons were classified as vertical (or angled) preferring if they significantly preferred the vertical (or the angled) grating corridor ($p < 0.05$, Wilcoxon rank-sum test).

SUPPLEMENTAL INFORMATION

Supplemental Information includes 13 figures, one movie, and Supplemental Experimental Procedures and can be found with this article at <http://dx.doi.org/10.1016/j.neuron.2015.05.037>.

AUTHOR CONTRIBUTIONS

J.P., A.G.K., T.D.M.-F., and S.B.H. designed the experiments. J.P. and A.G.K. performed the experiments with help from I.O. and A.N. J.P., A.G.K., and M.P. analyzed the data with advice from M.S., T.D.M.-F., and S.B.H. A.G.K. developed the behavioral tasks. A.N. and A.G.K. performed optogenetic silencing. Based on advice and software from G.B.K., J.P. and J.K. built the custom two-photon resonance scanning microscope. M.B. created the visual virtual environment. S.B.H., T.D.M.-F., J.P., and A.G.K. wrote the paper. All authors contributed to discussions and commented on the manuscript.

ACKNOWLEDGMENTS

We thank Petr Znamenskiy, Pieter Roelfsema, Daniel Huber, Jerry Chen, and Fritjof Helmchen for comments on earlier versions of the manuscript, and all members of the lab for comments and discussions. We thank John O’Keefe for support throughout the project, and Dinu Florin Albeanu for help with the intrinsic imaging setup. We thank the GENIE Program and Janelia Farm Research Campus of the Howard Hughes Medical Institute for making GCaMP6 material available, and the UCL Mechanical Workshop for building custom parts for the microscope and experimental setup. This work was supported by the Wellcome Trust (S.B.H., 095853, and T.D.M.-F., 095074, J.K., 50096 and M.B., 50096), the European Research Council (S.B.H., HigherVision 337797, T.D.M.-F., NeuroVision 616509), the Marie Curie Actions of the European Union’s FP7 program (J.P., 332141 and A.G.K., 301742), the Gatsby Charitable Foundation (M.P. and M.S.), and the Biozentrum core funds (University of Basel).

Received: December 8, 2014

Revised: March 7, 2015

Accepted: April 29, 2015

Published: June 4, 2015

REFERENCES

- Andermann, M.L., Kerlin, A.M., and Reid, R.C. (2010). Chronic cellular imaging of mouse visual cortex during operant behavior and passive viewing. *Front. Cell Neurosci.* 4, 3.
- Blake, D.T., Heiser, M.A., Caywood, M., and Merzenich, M.M. (2006). Experience-dependent adult cortical plasticity requires cognitive association between sensation and reward. *Neuron* 52, 371–381.
- Boyden, E.S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nat. Neurosci.* 8, 1263–1268.
- Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* 13, 87–100.

- Chen, T.-W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295–300.
- Chubykin, A.A., Roach, E.B., Bear, M.F., and Shuler, M.G.H. (2013). A cholinergic mechanism for reward timing within primary visual cortex. *Neuron* 77, 723–735.
- David, S.V., Fritz, J.B., and Shamma, S.A. (2012). Task reward structure shapes rapid receptive field plasticity in auditory cortex. *Proc. Natl. Acad. Sci. USA* 109, 2144–2149.
- Denk, W., Strickler, J.H., and Webb, W.W. (1990). Two-photon laser scanning fluorescence microscopy. *Science* 248, 73–76.
- Dombeck, D.A., Khabbaz, A.N., Collman, F., Adelman, T.L., and Tank, D.W. (2007). Imaging large-scale neural activity with cellular resolution in awake, mobile mice. *Neuron* 56, 43–57.
- Efron, B. (1979). Bootstrap methods: another look at the jackknife. *Ann. Stat.* 7, 1–26.
- Frenkel, M.Y., Sawtell, N.B., Diogo, A.C.M., Yoon, B., Neve, R.L., and Bear, M.F. (2006). Instructive effect of visual experience in mouse visual cortex. *Neuron* 51, 339–349.
- Gdalyahu, A., Tring, E., Polack, P.-O., Gruver, R., Golshani, P., Fanselow, M.S., Silva, A.J., and Trachtenberg, J.T. (2012). Associative fear learning enhances sparse network coding in primary sensory cortex. *Neuron* 75, 121–132.
- Ghose, G.M., Yang, T., and Maunsell, J.H.R. (2002). Physiological correlates of perceptual learning in monkey V1 and V2. *J. Neurophysiol.* 87, 1867–1888.
- Gibson, J.J. (1979). *The Ecological Approach to Visual Perception* (Boston: Houghton Mifflin).
- Glickfeld, L.L., Histed, M.H., and Maunsell, J.H.R. (2013). Mouse primary visual cortex is used to detect both orientation and contrast changes. *J. Neurosci.* 33, 19416–19422.
- Goltstein, P.M., Coffey, E.B.J., Roelfsema, P.R., and Pennartz, C.M.A. (2013). In vivo two-photon Ca²⁺ imaging reveals selective reward effects on stimulus-specific assemblies in mouse visual cortex. *J. Neurosci.* 33, 11540–11555.
- Hippenmeyer, S., Vrieseling, E., Sigrist, M., Portmann, T., Laengle, C., Ladle, D.R., and Arber, S. (2005). A developmental switch in the response of DRG neurons to ETS transcription factor signaling. *PLoS Biol.* 3, e159.
- Hölscher, C., Schnee, A., Dahmen, H., Setia, L., and Mallot, H.A. (2005). Rats are able to navigate in virtual environments. *J. Exp. Biol.* 208, 561–569.
- Huber, D., Gutnisky, D.A., Peron, S., O'Connor, D.H., Wiegert, J.S., Tian, L., Oertner, T.G., Looger, L.L., and Svoboda, K. (2012). Multiple dynamic representations in the motor cortex during sensorimotor learning. *Nature* 484, 473–478.
- Jaramillo, S., and Zador, A.M. (2011). The auditory cortex mediates the perceptual effects of acoustic temporal expectation. *Nat. Neurosci.* 14, 246–251.
- Keller, G.B., Bonhoeffer, T., and Hübener, M. (2012). Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. *Neuron* 74, 809–815.
- Krauzlis, R.J., Lovejoy, L.P., and Zénon, A. (2013). Superior colliculus and visual spatial attention. *Annu. Rev. Neurosci.* 36, 165–182.
- Law, C.-T., and Gold, J.I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat. Neurosci.* 11, 505–513.
- Li, W., Piëch, V., and Gilbert, C.D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nat. Neurosci.* 7, 651–657.
- Li, W., Piëch, V., and Gilbert, C.D. (2008). Learning to link visual contours. *Neuron* 57, 442–451.
- Lien, A.D., and Scanziani, M. (2013). Tuned thalamic excitation is amplified by visual cortical circuits. *Nat. Neurosci.* 16, 1315–1323.
- Madisen, L., Mao, T., Koch, H., Zhuo, J.M., Berenyi, A., Fujisawa, S., Hsu, Y.-W.A., Garcia, A.J., 3rd, Gu, X., Zanella, S., et al. (2012). A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nat. Neurosci.* 15, 793–802.
- McAlonan, K., Cavanaugh, J., and Wurtz, R.H. (2008). Guarding the gateway to cortex with attention in visual thalamus. *Nature* 456, 391–394.
- Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science* 229, 782–784.
- Niell, C.M., and Stryker, M.P. (2010). Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* 65, 472–479.
- Nienborg, H., and Cumming, B.G. (2014). Decision-related activity in sensory neurons may depend on the columnar architecture of cerebral cortex. *J. Neurosci.* 34, 3579–3585.
- Peters, A.J., Chen, S.X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature* 510, 263–267.
- Pinto, L., Goard, M.J., Estandian, D., Xu, M., Kwan, A.C., Lee, S.-H., Harrison, T.C., Feng, G., and Dan, Y. (2013). Fast modulation of visual perception by basal forebrain cholinergic neurons. *Nat. Neurosci.* 16, 1857–1863.
- Polley, D.B., Steinberg, E.E., and Merzenich, M.M. (2006). Perceptual learning directs auditory cortical map reorganization through top-down influences. *J. Neurosci.* 26, 4970–4982.
- Ress, D., and Heeger, D.J. (2003). Neuronal correlates of perception in early visual cortex. *Nat. Neurosci.* 6, 414–420.
- Rutkowski, R.G., and Weinberger, N.M. (2005). Encoding of learned importance of sound by magnitude of representational area in primary auditory cortex. *Proc. Natl. Acad. Sci. USA* 102, 13664–13669.
- Saleem, A.B., Ayaz, A., Jeffery, K.J., Harris, K.D., and Carandini, M. (2013). Integration of visual motion and locomotion in mouse visual cortex. *Nat. Neurosci.* 16, 1864–1869.
- Schoups, A., Vogels, R., Qian, N., and Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature* 412, 549–553.
- Shuler, M.G., and Bear, M.F. (2006). Reward timing in the primary visual cortex. *Science* 311, 1606–1609.
- Stănişor, L., van der Togt, C., Pennartz, C.M.A., and Roelfsema, P.R. (2013). A unified selection signal for attention and reward in primary visual cortex. *Proc. Natl. Acad. Sci. USA* 110, 9136–9141.
- Totah, N.K.B., Kim, Y., and Moghaddam, B. (2013). Distinct prestimulus and poststimulus activation of VTA neurons correlates with stimulus detection. *J. Neurophysiol.* 110, 75–85.
- Tremblay, L., and Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature* 398, 704–708.
- Wiest, M.C., Thomson, E., Pantoja, J., and Nicolelis, M.A.L. (2010). Changes in S1 neural responses during tactile discrimination learning. *J. Neurophysiol.* 104, 300–312.
- Yan, Y., Rasch, M.J., Chen, M., Xiang, X., Huang, M., Wu, S., and Li, W. (2014). Perceptual training continuously refines neuronal population codes in primary visual cortex. *Nat. Neurosci.* 17, 1380–1387.
- Yang, T., and Maunsell, J.H. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *J. Neurosci.* 24, 1617–1626.
- Zhang, S., Xu, M., Kamigaki, T., Hoang Do, J.P., Chang, W.-C., Jenvay, S., Miyamichi, K., Luo, L., and Dan, Y. (2014). Selective attention. Long-range and local circuits for top-down modulation of visual cortex processing. *Science* 345, 660–665.
- Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., El Gamal, A., and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* 16, 264–266.