bioRχiv
THE PREPRINT SERVER FOR BIOLOGY

March 2021

Dynamic distortion of the orientation representational space after learning in the mouse primary visual cortex

Julien Corbo[1×], John P. McClure Jr.[1×], O. Batuhan Erkat[1], and Pierre-Olivier Polack[1*]

1. Center for Molecular and Behavioral Neuroscience. Rutgers University – Newark. 197 University Avenue. Newark, NJ 07102. USA.

[×] These authors contributed equally to this work.

* Correspondence should be addressed to Dr. Pierre-Olivier Polack, 197 University Avenue. Newark, NJ 07102. USA. E-mail: polack.po@rutgers.edu

Learning is an essential cognitive mechanism that supports behavioral adaptation through neural processing adjustments. Learning was shown to modify sensory integration, yet the nature of those modifications and the computational advantages they confer remain unclear. By comparing the responses of primary visual cortex (V1) neurons evoked by oriented stimuli in naïve mice and mice performing an orientation discrimination task, we found that the representations of rewarded and non-rewarded cues were sparser, more accurate and more stable in trained mice. This improved representation was associated with a distortion of the V1 orientation space such that stimuli close to the task cues were represented as the task stimuli themselves. This distortion was context-dependent, as it was absent in trained mice passively viewing the cues. Hence, visual processing in V1 was dynamically adapted to enhance the reliability of the representation of the learned cues and favor stimulus generalization in the task-relevant computational space.

Introduction

Learning, by allowing the association of sensory inputs, actions, and outcomes, is crucial for adapting to new situations or environments. Although several cellular mechanisms associated with learning have been identified [1,2], understanding their involvement and importance for the various neuronal assemblies involved in any learned task is complex. Indeed, learning not only induces changes in the neural activity of brain regions participating to the emotional and behavioral engagement or motor skills, but also affects sensory processing [3]. In the primary visual cortex (V1), learning has been linked to various modulations of visual processing. In some studies, the responses of neurons selective for the rewarded stimulus were found to be enhanced [4,5]. In others, it was the number of V1 neurons tuned to the rewarded stimulus that increased [5,6], but this number was found to be stable [7,8] or reduced [4] in other experimental paradigms. The selectivity of individual neurons was also shown to improve [5,6,8]. Yet, a reduced selectivity near the orientation of the rewarded stimulus increasing the number of responsive neurons was also reported [4,7]. Finally, a reduction in trial-to-trial variability at the unit level was also described [5,8]. A corollary of all those reports would be that the evoked neuronal activity in V1 increases after learning. Yet, such an increase would stand against studies in other sensory cortices and/or species showing sparser evoked responses after learning [9,10]. Moreover, it would refute the sparse coding theory stipulating that learning should lead to representations supported by a small amount of reliably active responsive neurons to promote the formation of stronger associations while improving energy efficiency and maintaining overall synaptic homeostasis [11,12]. Despite all the discrepancies, the above-mentioned studies agree that learning should improve the saliency of the representation of behaviorally relevant stimuli and therefore improve the discriminative capabilities of sensory cortices [5,6,8].

To examine how learning impacts the representation of behaviorally relevant stimuli at the population level in V1, we performed calcium imaging in naïve mice passively viewing two drifting gratings of orthogonal orientations, and in trained mice performing a simple go/no go discrimination task using those same stimuli. Using the orientation tuning properties of V1 neurons, we mapped the evoked neuronal activities in the orientation space to obtain an explicit reading of the information on the stimulus orientation captured by the V1 population. We found a sparser and more accurate representation of learned cues in trained mice compared to naïve mice. Moreover, in trained mice, the representation of visual stimuli in the orientation space was distorted, emphasizing the representation of behaviorally relevant stimuli. This distortion was context dependent as it could not

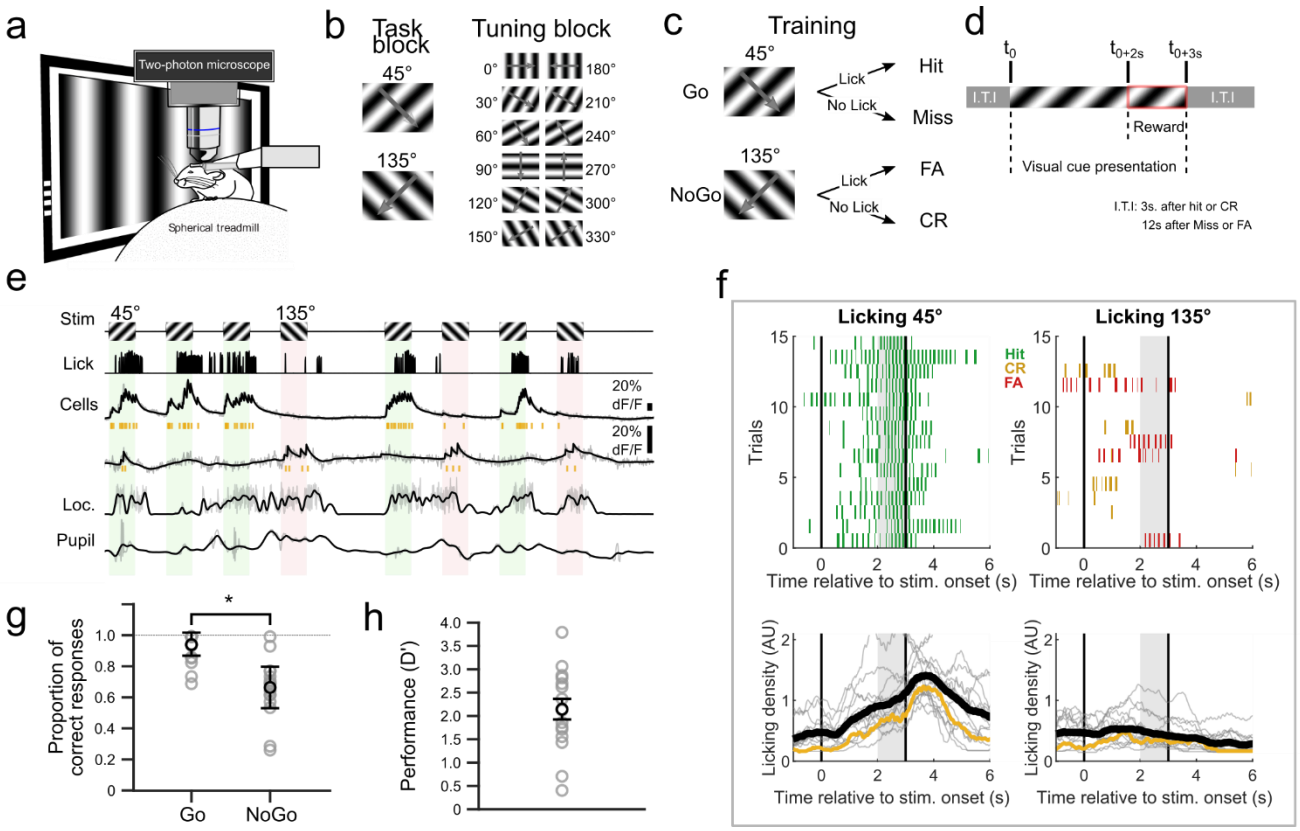

Figure 1. Experimental setup and behavioral training. a. Imaging setup. **b.** Stimuli used. *Left*. drifting gratings of orientations 45° or 135° presented during the task block. *Right*. drifting gratings of 12 orientations evenly spaced in the orientation domain presented during the tuning block. **c.** Schematic description of the Go/No-Go task. FA = False Alarm; CR = Correct Rejection. **d.** Schematic description of the temporal structure of a single trial. I.T.I.: intertrial interval. **e.** Example of recording session, from top to bottom: visual stimulus, licking, two segmented cells ΔF/F (grey) with the deconvoluted spikes (yellow) and the reconstructed signal (black), locomotion, and pupil size. **f.** *Top*: example of licking responses during the presentation of a 45° oriented grating (left) and 135° oriented grating (right). From the same session as depicted in e. *Bottom*: Mean licking density within sessions (grey line) and across sessions (black line). The exemplified session is highlighted in yellow. **g.** Proportion of correct responses to the presentation of the Go and No-Go cues for the 'trained mice' group. * p = 0.0001; $\chi 2$ = 15.2; n = 16 sessions in 10 mice. **h.** Performance (D') of the trained group for each recording session.

be found in trained mice when they did not participate to the task.

Results

Sparser visual responses after orientation discrimination training

To evaluate changes in V1 response saliency and orientation representation induced by learning, we presented the same set of drifting gratings to naïve mice (14 animals and 21 sessions) and to mice trained to use those cues to perform an orientation discrimination task (trained mice, 10 animals and 16 sessions). V1 neuronal activity was recorded by calcium imaging in mice placed head-fixed on a spherical treadmill [13] (**Fig. 1**a). The calcium signal was then deconvolved into Action Potential related Events (APrE) using a benchmarked algorithm [14]. Recording sessions consisted of a *task block* during which two orthogonal drifting gratings were presented, and a *tuning block* during which twelve different orientations/directions, distinct from the two orientations of the task block were displayed in random order (**Fig. 1**b). Trained mice were taught to discriminate between the 45° drifting grating (Go signal) and the 135° drifting grating (No-Go signal; **Fig. 1**c). A task trial consisted in the presentation of one of the two task stimuli for 3 seconds. The response to the Go/No-Go question was requested during the third second of the trial (**Fig. 1**d). As locomotor activity and arousal modulate V1 neuronal activity [15,16], treadmill movement and pupil size were recorded simultaneously with the calcium signal (**Fig. 1**e). After learning, trained mice were successfully licking in the reward period for Go trials and refraining from licking during No-Go trials (**Fig. 1**f). Trained mice performed the task with a hit rate of 94 ± 8% and a correct rejection rate of 67 ± 13% (median ± m.a.d.; p=9.8 x 10^{-5}; Kruskal-Wallis Test: χ^2 = 15.18; **Fig. 1**g). Altogether the mean sensitivity index D'

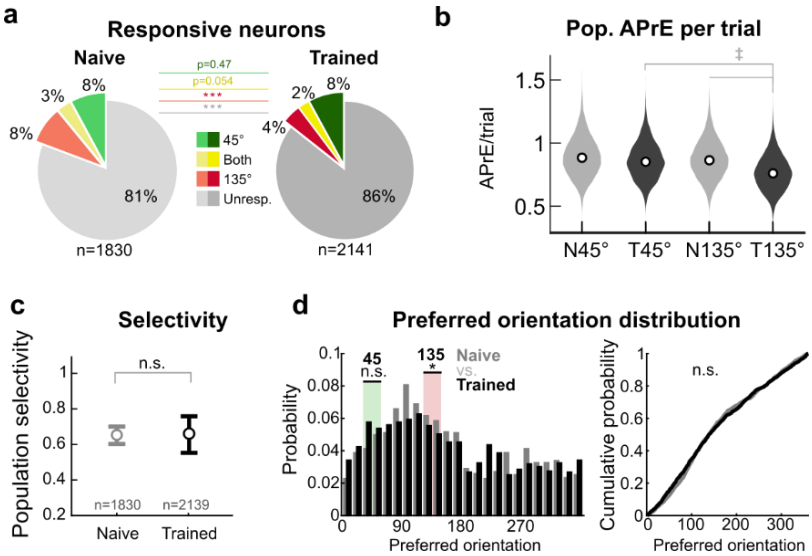

Figure 2. Saliency of the visually evoked responses of V1 neurons before and after training. **a.** Proportion of neurons responsive to the 45° and 135° gratings (binomial test, response vs. baseline, p<0.01). **b.** APrE per trial and per neuron evoked by the 45° and 135° gratings in the naïve (N) and trained (T) groups. **c.** Population selectivity of naïve and trained mice. Observed value and 95% confidence interval obtained by bootstrapping (n= 10,000). **d.** Distribution of preferred orientations in naïve and trained mice. Left: distribution histogram using of 15° width bins. Colored areas indicate the significance of the difference between the proportion of neurons observed in naïve and trained groups for those bins using a permutation test. Right: cumulative probability without binning. Anderson-Darling nonparametric test, corrected for ties, p = 0.69. Effect size (Hedge's g): ‡ = g>0.6. Permutation test: * = p<0.05, **= p<0.01, ***= p<0.001.

across all the sessions was 2.1 ± 0.2 (16 sessions in 10 mice, **Fig. 1**h).

As previous studies reported various learning-induced enhancements of the neuronal response saliency, we started by measuring those modulations in our experimental paradigm. We therefore quantified the proportion of neurons responsive to the visual stimuli presented (binomial test p<0.01, see **Methods**) in naïve and trained mice [4,5,7,8]. In naïve mice, the proportion of neurons responsive to the 45° grating (8.1% or 149/1830) was similar to that in the trained mice (8.1% or 174/2141; permutation test p=0.47, **Fig. 2**a). A lower proportion of neurons were responsive to the unrewarded stimulus (135° drifting grating) in trained mice (4.2% or n=90/2141) than in naïve mice (8.6% or 157/1830, p<0.001; **Fig. 2**a), In both naïve and trained mice, only a small proportion of neurons responded to both stimuli (naïve: 3.2%, n=58/1830; trained: 2.4%, n=51/2141, p=0.053**).** Overall, the proportion of responsive neurons was lower in trained than naïve mice (naïve: 16.7%, n=306/1830; trained: 12.3%, n=264/2141; p<0.001). We then tested if this decrease in the number of activated neurons in trained animals affected the total amount of action potentials generated by the V1 population in response to the task-relevant cues (**Fig. 2**b). We found a lower amount of APrE per neuron and per trial in the naïve than in the trained group (mean ± standard deviation; naïve: 0.93±0.16; trained: 0.86±0.15; effect size: Hedge's g = 0.48). For the 45° stimulus, that metric was marginally lower in trained mice (naïve: 0.95 ± 0.16; trained: 0.90 ± 0.15; effect size: Hedge's g = 0.29). For the 135° stimulus, the decrease was more pronounced (naïve: 0.92 ± 0.15; trained: 0.81 ± 0.14; effect size: Hedge's g = 0.77). This differential decrease led to a significant difference between the population activity evoked by the presentation of 45° and 135° gratings in trained mice (Hedge's g = 0.61). The decrease in global population activity in trained mice was therefore associated with a specific reduction in the evoked activity of neurons responsive to the task stimuli (**Fig. S1**a,b).

Although neurons responsive to the task stimuli were less active after training, this did not exclude the possibility of an improvement of the population *selectivity* [6], a metric capturing the separability of the responses evoked by the two task cues (see **Methods**). However, the population selectivity of the neurons recorded in trained mice did not show such an improvement (**Fig. 2**e and **Fig. S1**c,d; selectivity in naïve mice: 0.65, bootstrapped 95% confidence interval: [0.61 - 0.70]; trained: 0.77, bootstrapped 95% confidence interval: [0.61 - 0.96]; permutation test: p = 0.16).

Finally, we investigated whether learning recruited untuned neurons to respond to the task cues by testing if the distribution of preferred orientations was different in the naïve and trained neuronal populations (**Fig. 2**f). The proportion of neurons with a preferred orientation ± 15° around 45° was not significantly different between the trained and naïve groups (naïve: 9.1%; trained: 11.1%; permutation test: p=0.09), while the proportion of neurons with a preferred orientation ± 15° around 135° was slightly lower in the trained group (naïve: 12.0%; trained: 10.5%; p=0.02). These relatively small differences did not significantly impact the overall distribution of preferred orientations in the V1 population (Anderson-Darling non-parametric test: p=0.70).

Sharpened representations in the orientation space after training.

Our results so far indicated that after training, the population response in V1 was sparser with fewer responsive neurons and/or a decrease in the total amount of activity evoked by the visual cues. To determine how those changes modified the visual information present in V1, we estimated

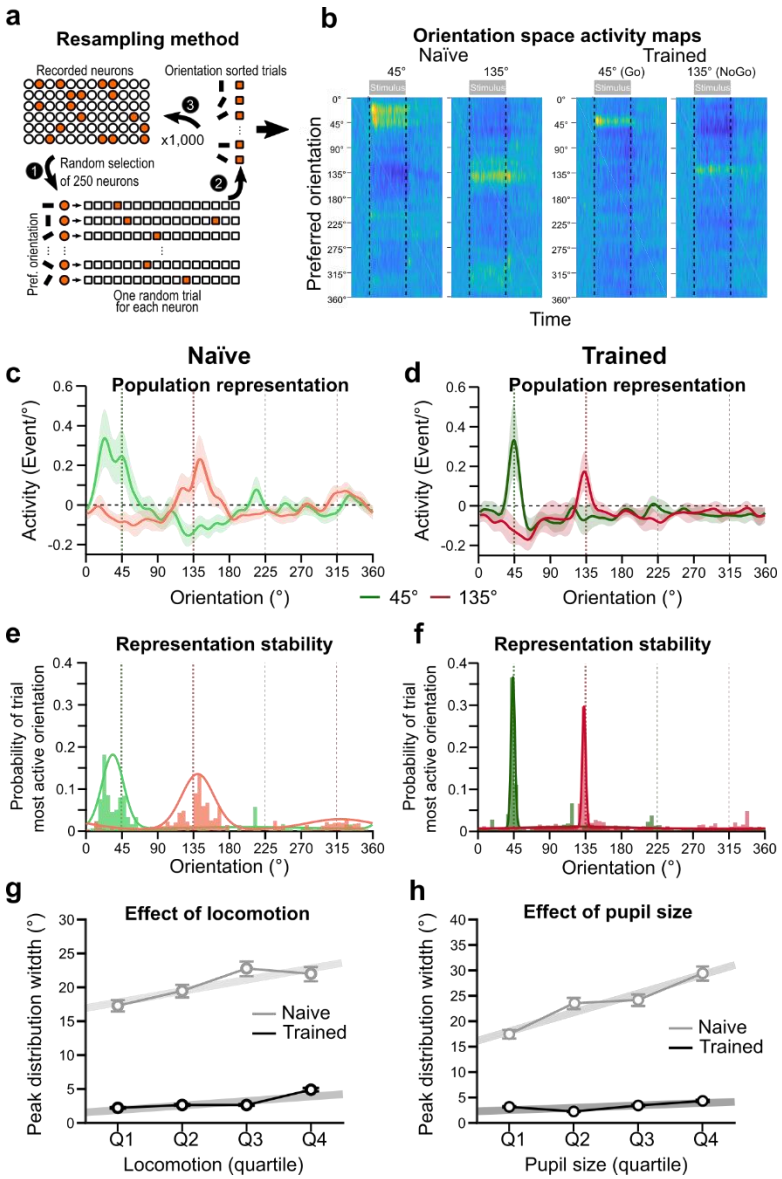

Figure 3. Representations of the orientations of repeated stimuli in the V1 of naïve and trained mice. a. Schematic representation of the method used to generate the orientation space activity maps. **b.** Orientation space activity map of the V1 responses to the presentation of the 45° and 135° task stimuli in naïve (left panels) and trained mice (right panels). Dark blue hues indicate activity lower than baseline. Yellow hues indicate activity greater than baseline. **c.** Average of the orientation space activity map shown in b over the first second of the presentation of the 45° (green trace) and 135° (red) task stimuli in naïve mice. Shaded areas indicate the standard deviation across time. **d.** Same representation as in c. in trained mice. **e.** Probability distribution of the location in the orientation space of the peak of the activity profile during trials when 45° or 135° stimuli were presented to naïve mice. The distributions were fitted with a two component von Mises mixture model. **f.** Same as in e. for trained mice. **g.** Width of the peak distribution, a measure of population representation stability, as a function of the quantity of locomotion occurring during the first second after stimulus onset. All trials were sorted according to the amount of locomotion observed during the first second, with quartiles cutoff values computed from the trials in the trained group. $\sigma \pm C.I._{95\%}$. Lines indicate the best linear fit on the four points in each group. **h.** Same as in g. but for pupil size.

how the orientations of the task block stimuli were represented in V1 using a resampling method (**Fig. 3**a, see **Methods**).

Randomly selected individual trial responses to the presentation of the same stimulus were arranged according to the preferred orientations of the selected neurons, creating a resampled single trial map of the V1 neuronal activity in the orientation space (**Fig. 3**a). This procedure was repeated 1,000 times, to generate an *orientation space activity map* allowing us to capture the trial-to-trial variability of the APrE discharge in the orientation space (**Fig. 3**b). We restricted our analysis to the activity evoked during the first second of the visual stimulus presentation, before trained mice initiated a licking

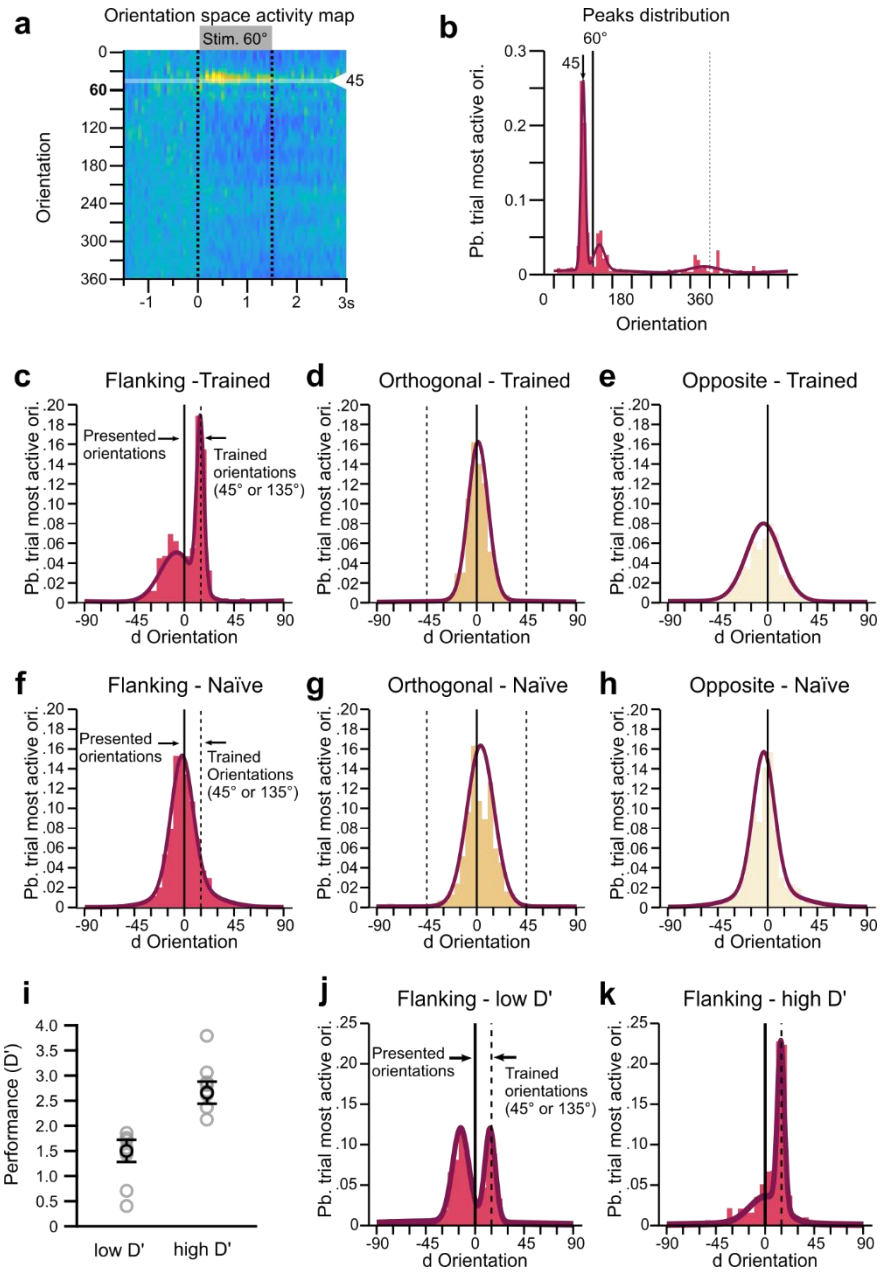

Figure 4. Population representations of the orientations of the tuning block stimuli. a. Example of orientation space activity map for the 60° stimulus, during the tuning block in trained mice. Dark blue hues indicate activity lower than baseline. Yellow hues indicate activity greater than baseline. **b.** Probability distribution of the location in the orientation space of the peak of activity for this stimulus. Purple curve: three component von Mises mixture model fit. **c.** Combination of the histograms obtained for the stimuli of the tuning curve block flanking the orientation of the task stimuli (30°, 60° 120°, and 150°). 60° and 150° histograms were mirrored such as the difference in orientation (d orientation) between the stimulus orientation and the orientation of the flanking cue of the task block was +15°. Purple curve: three component von Mises mixture model fit. **d.** Same representation as in c for the orientations of the tuning block orthogonal to the orientation of the task stimuli (0°, 90° 180°, and 270°). **e.** Same representation as in c for the orientations of the tuning block opposite to the orientation of the task stimuli (210°, 240° 300°, and 330°). **f- h**. Same as c-e for naïve mice. **i.** Individual D' of sessions sorted in low or high D' group. Black markers indicate mean and s.e.m. **j,k.** Same as c, but for the low and high D' groups respectively.

response (Go response time 1.8 ± 0.1 s, No-Go response time 1.9 ± 0.1 s, 16 sessions in 10 mice). The profile of activation in the orientation space was broader in naïve (**Fig. 3**c) than in trained mice (**Fig. 3**d). To test whether this difference resulted from more consistent evoked responses across trials, we computed the distribution of the orientations at which the activity peaked for each individual resampled trial. We used the width of this distribution as an indicator of representation stability. Those distributions were best fitted by von Mises distributions that were significantly broader in naïve mice ($\sigma_{45°}$ = 13.2°; C.I.$_{95\%}$: [12.5 – 13.8]; $\sigma_{135°}$ = 19.2°, C.I.$_{95\%}$: [18.2 – 20.1]; **Fig. 3**e) than in trained mice ($\sigma_{45°}$ = 2.5°, C.I.$_{95\%}$: [2.4 – 2.7]; $\sigma_{135°}$ = 2.5°, C.I.$_{95\%}$: [2.4 – 2.7]; **Fig. 3**f). Thus, the representations of task-relevant cues were more stable and accurate in the orientation space in trained animals, with neurons preferring the orientation of the cues being reliably the most active across trials.

Because behavioral state can modulate the activity of V1 neurons [15-17], we compared the amount of locomotion and the pupil size between naïve and trained mice (**Fig. S2**). Trained animals displayed a greater amount of locomotion during baseline and increased their locomotor activity during the presentation of the stimuli. Locomotor activity was the largest for rewarded trials (**Fig. S2**a,b). Pupil size showed a similar pattern, with a greater pupil dilatation overall in trained mice (**Fig. S2**c,d). To determine whether differences in arousal and locomotion could explain the improved accuracy and stability of the representations observed in the trained animals, we repeated the measure of trial-to-trial representation stability (width of the peak distribution) after sorting the trials according to the amount of locomotion or pupil dilatation measured during the first second post stimulus onset (**Fig. S2**e-h). The representation stability stayed well separated between naïve and trained mice regardless of the amount of locomotion or pupil dilatation (**Fig. 3**g,h). Even when there was no locomotion (Quartile 1, mean locomotion value = 0; **Fig. S2**f), the

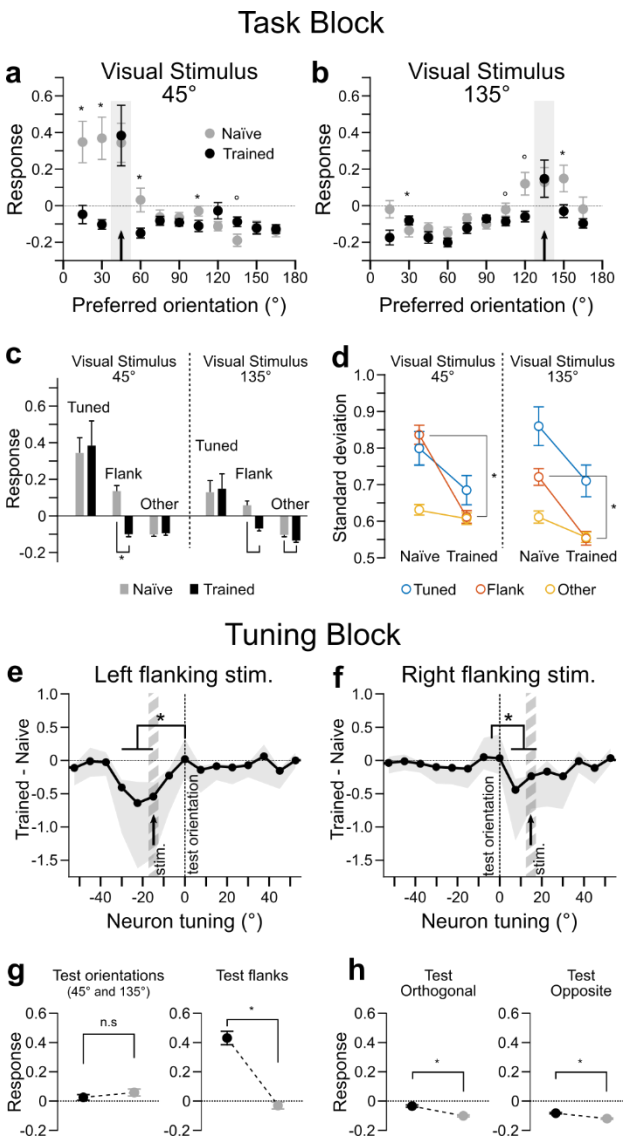

Figure 5. Activity of individual V1 neurons evoked by the task block stimuli as a function of their preferred orientation. a. Average across V1 neurons (mean ± s.e.m.) of the mean response relative to baseline (in APrE.s^{-1}) in response to the presentation of the 45° task stimulus in naïve (grey) and trained (black) mice. Neurons are grouped by preferred orientation (x-axis value ± 7.5°). * $p < 0.05$, º $p<0.01$. Two-sample Kolmogorov-Smirnov test. **b.** Same representation as in *a* for the neuronal responses to the presentation of the 135° task stimulus. **c.** Comparison of the responses of the neurons tuned for the task stimuli (Tuned, left section, preferred orientation range: 45°±7.5°; Tuned, right section, preferred orientation range: 135°±7.5°), neurons tuned for orientations flanking the task stimulus orientation (Flank for 45°: 30°±7.5° and 60°±7.5°; Flank for 135°: 120°±7.5° and 150°±7.5°), and the other neurons having the same direction selectivity as the task stimuli (Other: between 0° and 180°) to the presentation of the two task stimuli between the naïve (blue) and trained (green) mice. *: $p < 0.05$, two-sample Kolmogorov-Smirnov test. **d.** Same representation as in c for the standard deviation of the neuronal response across trials. *: $p < 0.05$, two-sample Kolmogorov-Smirnov test. **e.** Difference between the mean responses to left-flanking orientations during the tuning block in naïve and trained mice, as a function of preferred orientation. Data from the 30° and 120° stimuli were pooled. Shaded area indicates the 95% confidence interval obtained by bootstrapping (1,000 repetitions) **f.** Same representation as in e. for the right-flanking orientations, i.e. 60° and 150°. **g.** Comparison between the naïve (grey) and trained (black) mice of the response of the neurons tuned for the orientations of the task stimuli (Task orientations: 45°±7.5° and 135°±7.5°; left panel) and the neurons tuned for the orientations flanking the task orientations (30°±7.5°, 60°±7.5°, 120°±7.5° and 150°±7.5°; right panel) to the presentation of the orientations flanking the task visual stimuli during the tuning block (30°, 60°, 120° and 150°).*: $p < 0.05$, two-sample Kolmogorov-Smirnov test. **h.** Same representation as in c for neurons with preferred orientations orthogonal (352.5° to 22.5°, 67.5° to 112.5°, and 157.5° to 187.5°; left panel), or opposite (187.5° to 352.5°; right panel) to the task orientations when cue having orientations flanking the task visual stimuli (30°, 60°, 120° and 150°) were presented during the tuning block.

representation stability in the trained group was greater than in the naïve group. Therefore, the difference in locomotion and arousal did not explain the representation changes observed between the two groups.

Distortion of the orientation space in trained mice

We next tested whether the greater stability and precision of orientation representation in trained mice was specific to the context and/or stimuli of the task block. To address this question, we searched for similar changes in the responses evoked during the *tuning block* presented to the trained mice immediately after the task. However, we found that the orientation space activity maps (**Fig. 4**a, **Fig. S3**a) presented unexpected distortions for orientations closest to the Go and NoGo stimuli (flanking orientations: 30° and 60°, 120° and 150°, **Fig. S3**a). For those orientations, the peak of activity was consistently centered on the orientations of the task stimuli (see example in **Fig. 4**b, **Fig. S3**a). For example, when a 60° stimulus was presented, neurons preferring 45° were the most active as if the population categorized close orientations as Go and NoGo instead of representing their orientation faithfully. To characterize this deformation of the orientation space, we compiled the distributions of activity peaks across trials for the flanking orientations. The distributions of 60° and 150° stimuli were mirrored such that the difference between the orientation of the visual stimulus presented and the orientation of the nearest task stimulus (hereafter referred to as relative angle) was +15° (**Fig. 4**c; see Methods). The resulting distribution was better fitted by a mixed model of von Mises distributions whose best component (most variance explained) was tightly centered around the orientation of the task stimulus (peak

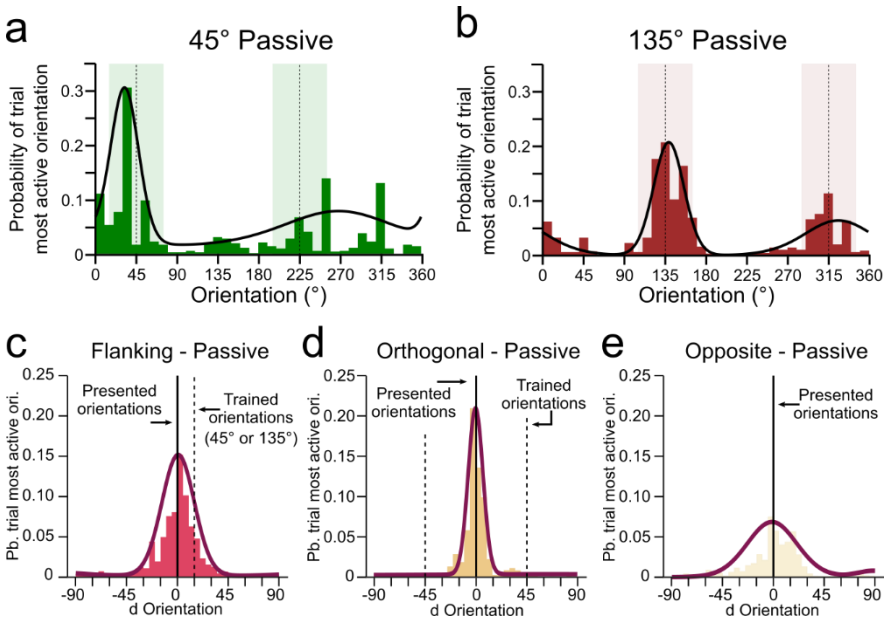

Figure 6. Population representations of the orientations of the task block and tuning block stimuli in trained but passive mice. a. Probability distribution of the location in the orientation space of the peak of activity during trials when the 45° task stimulus was presented to passive mice. The distribution was fitted with a two component von Mises mixture model. **b.** Same representation as in *a* for the presentation of the 135° task stimulus. **c.** Combination the distribution of the peak of activity in the orientation space for orientations of the tuning curve block flanking the orientation of the task stimuli (30°, 60° 120°, and 150°). 60° and 150° histograms were mirrored so the difference in orientation (d orientation) between the stimulus orientation and the orientation of the flanking cue of the task block was +15°. Purple curve: three component von Mises mixture model fit. **d.** Same representation as in c for the orientations of the tuning block orthogonal to the orientation of the task stimuli (0°, 90° 180°, and 270°). **e.** Same representation as in c for the orientations of the tuning block opposite to the orientation of the task stimuli (210°, 240° 300°, and 330°).

relative angle ± $CI_{95\%, \text{two-tails}}$: +15.5 ± 0.7°; σ = 3.9°, σ C.I. $_{95\%}$: [3.7 – 4.1]) while the second largest proportion was broader (σ = 10.9°, σ C.I. $_{95\%}$: [10.3 – 11.4]) and away from the orientation of the tuning block stimulus (peak relative angle: -9.4 ± 2.0°). This deformation of the orientation representation was not found with stimuli orthogonal to the task stimuli (peak relative angle: + 2.6 ± 1.7°; **Fig. 4**d) even though the width of the fit was significantly smaller ($\sigma_{trained}$ = 9.5°, σ C.I. $_{95\%}$: [9.0 – 10.0] vs. σ_{naive} = 13.8°, σ C.I.$_{5\%}$: [13.1 – 14.5]). Finally, the representation of orientations of tuning block stimuli opposite to task stimuli was centered (peak: -2.3 ± 3.3°; **Fig. 4**e) and slightly broader after training ($\sigma_{trained}$ = 17.4°, σ C.I. $_{95\%}$: [16.6 – 18.3] vs. σ_{naive} = 15.6°, σ C.I. $_{95}$%: [14.9 – 16.4]; **Fig. 4**e).

To determine whether the distortion of the orientation space for the flanking stimuli resulted from training on the task or from the mere repetition of the 45° and 135° stimuli, we performed the same analyses in naïve mice. No significant differences were found between the width of the peak distribution for the flanking (30°, 60° 120°, and 150°; σ = 14.6°, C.I. $_{95\%}$: [13.8 – 15.3]; **Fig. 4**f), orthogonal (0°, 90° 180°, and 270°; σ = 13.8°, C.I. $_{95\%}$: [13.1 – 14.5]; **Fig. 4**g), or opposite stimuli (210°, 240° 300°, and 330°; σ = 15.6°, C.I. $_{95\%}$: [14.9 – 16.4]; **Fig. 4**h). These results suggest that the repeated presentation of the 45° and 135° stimuli does not distort the representation of flanking orientations unless the stimuli acquire behavioral significance through learning. As a result of this distortion, orientations close to the task-relevant cues seem to be represented in trained animals as the task-relevant cues themselves.

We hypothesized that a relationship could exist between the orientation space distortion and the behavioral performance of the animals. To test that, sessions were therefore split into two groups: the low D' group contained sessions with the lower half of the D' values (mean D' ± s.e.m.: 1.4±0.2), and the high D' contained the top half performances (2.9±0.1, **Fig. 4**i). Representation stability in the task block, as measured by the width of the peaks' distribution, was not systematically improved in the high-performance group (σ and [C.I. $_{95\%}$]; 45°: low D' = 2.17 [2.07 – 2.28], high D' = 2.38 [2.26 – 2.50]; 135°: low D' = 3.06 [2.92 – 3.22], high D' = 5.56 [5.30 – 5.85]; **Fig. S3**c-f). However, the distortion of orientation space observed during the tuning block was greater in the high D' sessions (**Fig. 4**j,k). In the low performance sessions, the distribution was bimodal, with peaks at -12.1 ± 1.4° and +14.7 ± 0.9° away from the block stimulus orientation (**Fig. 4**j). In the high D' sessions, the peak activity of most trials was located +15.5 ± 0.6° away from the flanking stimulus orientation i.e., at the location of the task orientation. Therefore, for flanking stimuli of the tuning block, single trial representations consistently resembled those of the task stimuli instead of the truly presented orientations (**Fig. 4**k).

Neuronal activity changes underpinning the distortion of the orientation space in trained mice.

We then investigated the individual neuronal changes underpinning the learning-induced transformation of orientation representation in V1. The distortion of the population activity could be the result of an increased response of neurons that preferred the Go/NoGo stimuli and/or a suppression of neurons that preferred flanking orientations. To determine which alternative was at work, we measured the mean APrE evoked during the first second of the presentation of the 45° (**Fig. 5**a) and 135° (**Fig. 5**b) stimuli for each cell.

Neurons were then pooled by preferred orientation (bin size 15°). We did not find a significant difference between the activity of neurons tuned to the orientation of the task stimuli of naïve and trained mice (**Fig. 5**a,b grey box; **Fig. 5**c Tuned neurons). However, the responses of neurons whose preferred orientation flanked the orientation of the task block cue was significantly lower in trained mice relative to naïve mice (**Fig. 5**a-c). This lower activity was associated with a significant decrease of the response variance across trials (**Fig. 5**d).

The presence of a distortion of the orientation representations during the tuning block suggested that a filter was applied to the orientation space. To better visualize it, we pooled the data from the responses to the left-flanking (30° and 120°, **Fig. 5**e) or right-flanking (60° and 150°, **Fig. 5**f) tuning block stimuli and subtracted the bootstrapped mean activity of naive and trained mice neurons for each preferred orientation bin. The responses of the neurons tuned for the orientation of the presented stimulus were significantly lower in trained mice (**Fig. 5**e,f, grey box) while the responses of the neurons tuned for the neighboring task stimulus (45° or 135°) were similar in naïve and trained mice (**Fig. 5**e,f labelled arrow; **Fig. 5**g left panel). Altogether, the responses of neurons flanking the orientations of the task stimuli (**Fig. 5**g, left panel) but also, at a smaller magnitude, the responses of neurons tuned for orientations orthogonal or opposite to those orientations (**Fig. 5**h) were significantly lower in trained mice. This suppression of the activity of neurons flanking the orientation of the task stimuli moved *de facto* the peak of activity toward the orientation of neurons responding preferentially to the orientation of the task stimuli, generating a distorted representation of orientations in V1.

No distortions of the orientation space in trained but passive mice.

The distorted orientation space found in trained mice occurred outside of the task context, yet immediately after task performance. Hence, two phenomena could contribute to this distortion: a transient contextual modulation of V1 processing via top-down influence, and/or a long-lasting local modification of V1 functional connectivity originating from synaptic plasticity [6]. We hypothesized that if the distortion resulted from long-term plasticity and was not context-dependent, we should be able to observe it in trained animals not motivated to perform the task. We therefore performed the same experiments in mice trained to perform the task but satiated for water (passive group, 7 sessions in 4 animals). These sessions were performed 2 to 6 days (mean ± sem: 4.4±1.2 days) after the last task recording. Mice were not involved in the task, yet their locomotor activity and pupil diameter indicated a greater arousal than naïve mice, but lower than during task performance and not stimulus-specific (**Fig. S2**). In these conditions, the width of the distribution of the peak of activity evoked by the two task stimuli was comparable to that found in naïve mice ($\sigma_{45°}$= 14.2°, C.I.$_{95\%}$: [13.5 – 14.9]; $\sigma_{135°}$= 16.9°, C.I.$_{95\%}$: [16.0 – 17.7]; compare **Fig. 6**a,b with **Fig. 3**e) and significantly wider than the distribution found during task performance (compare to **Fig. 3**f). We then tested if the distortion of the orientation space was present in passive mice during the tuning block. We found that the larger component of the mixed model of von Mises distributions fitting the distribution of the peak of activity across trials was centered around the orientation of the displayed stimulus for the orientations flanking (peak: +3.7 ± 2.4°; σ = 13.1°, C.I. $_{95\%}$: [12.5 – 13.8]; **Fig. 6**c), orthogonal (peak: + 1.3 ± 1.2°; σ = 6.7°, C.I. $_{95\%}$: [6.3 – 7.0]; **Fig. 6**d) and opposite (peak:+4.4 ± 2.3°; σ = 12.7°, C.I. $_{95\%}$: [12.1 – 13.3]; **Fig. 6**e) to the task block stimulus orientations. Those results therefore suggest that the distorted representation of orientation observed in trained mice is context dependent.

Discussion

In this study, we compared the population response of layer 2/3 V1 neurons before and after orientation discrimination training to determine whether learning [6,9,18-24] modified the representation of orientations. We found that: [1] Learning is not always paired with enhanced responses to task-relevant cues. [2] The representation of orientations of drifting gratings was sparser, more accurate, and more stable in trained than in naïve mice. [3] The representation of drifting gratings with orientations flanking the orientations of the task cues was distorted such as those orientations were represented as the task cues themselves. [4] The deformation of the orientation space in trained mice was a consequence of a suppression of the responses of neurons whose preferred orientations flanked the orientations associated to the task. [5] The improved representation of the cues in trained mice and the associated deformation of the orientation space was context dependent as it was absent in mice trained to perform the task but passively viewing the cues.

Our findings do not support the hypothesis that learning is systematically associated with an increase of the task cues' representation saliency that would be caused by an increase in the number of recruited neurons [5,6], an increase of the response amplitude [4,5,7], or an increase in the difference between the responses to the rewarded and non-rewarded cues [6,8]. In our study, the only observed effect that contributed to a relative enhancement of the representation of the Go cue compared to the NoGo cue was a decrease of the number of neurons responding to the latter. Contrary to previous reports that showed a specific enhancement of the rewarded cue representation [5-7,25,26], we observed the same post-learning changes of the orientation representations for the rewarded and non-rewarded cues. Discrepancies between studies investigating the impact of learning on sensory processing are common. This variability is most likely due to the existence of differences between the experimental design of those studies [3]. The type of task (operant chamber [7], classical conditioning [8],

or Go/NoGo task in a virtual reality set up where the stimulus presentation is coupled with locomotion [5,6]) and the nature of the reward and/or the punishment likely affect the perceived value of the task stimuli, changing the nature and magnitude of reward-related signaling in the brain [27]. Additionally, the external context (the apparatus in which animals are placed) can change the subset of neurons responding to a given orientation [5], and the coupling or decoupling of locomotion speed and optic flow generates different contextual signals in V1 [28]. Those experimental differences generate *de facto* a differential engagement of the cortical microcircuits preventing the extraction of overarching rules from these different situations. Moreover, the intrinsic properties of the visual stimuli used for the task affect the susceptibility of V1 neurons to present task-related modulations. Indeed, the size of the stimulus is known to influence neuronal responses in V1: full-field drifting gratings trigger surround inhibition that enhances the sparseness of the V1 population response compared to more localized stimuli [29,30]. The contrast of visual stimuli is also affecting context dependent modulations [31,32].

In this study, we chose a task that mice acquire easily, and there should be *a priori* no need to improve the representations of the stimuli in V1 to perform it. The angular distance between the Go and NoGo cues is maximal (90°), we use high-contrast (75%) full-field stimuli, and mice are not required to navigate in a virtual reality environment. Hence, the representation of the 45° and 135° stimuli are already highly distinct in naïve mice V1. They do not overlap in the orientation space (**Fig. 3**c), and we observed similar values of population selectivity in naïve and trained mice (**Fig 2**e). Interestingly, the selectivity measured in naïve and trained animals is comparable to what was reported for expert mice in other studies [6]. This suggest that the V1 representations of the two stimuli were sufficiently separable for naïve mice to discriminate them reliably. Therefore, the change in representation described here is more likely to establish a robust and efficient response to the expected and/or valuable stimuli than to improve the orientation discrimination ability of V1.

One of the main findings of this study was that the response to the visual cues was sparser in trained mice. This increased sparseness was characterized by a reduction of the number of neurons responsive for the Go and NoGo cues and an overall decrease in the number of action potentials elicited by the presentation of the visual stimuli. Similar increases in sparseness have been reported in the mouse barrel cortex after associative fear learning [9] or in the primate visual cortex after orientation discrimination training [10]. More importantly, we showed that this increase in sparseness was associated with an increase in accuracy and trial-to-trial stability. Non-specific visual processing is known to present some redundancy [33], and action potentials elicited in sensory cortices by unattended sensory stimuli outnumber by far the amount of action potentials needed for sensory detection [34,35]. This redundancy in the naïve context is likely to preserve the information diversity of the sensory scene made of overlapping complementary attributes (e.g., stimulus orientation, contrast, spatial frequency, temporal frequency, etc.). However, redundancy has a high metabolic cost and promotes trial-to-trial variability [11]. We found that learning limits the evoked responses to a small number of neurons reliably representing a meaningful attribute of the learned cues. Therefore, it could decrease the metabolic burden of the sensory processing during the task [30,36,37], while bringing other computational advantages of sparse coding. i.e. an enhanced fidelity of the neuronal response across trials [38], and the promotion of associative learning mechanisms such as Hebbian plasticity [11].

In our experimental design, the improved representations of the task cues after learning did not result from a signal increase [39] but rather from a reduction of noise [9,40]. The suppression of neurons tuned for flanking orientations distorted the orientation space during and immediately after the execution of the task. This suppression favored neurons tuned for the trained orientations when close orientations were presented. As a result, the population activity was funneled toward two attractors of the orientation space: the orientations of the Go and No-Go cues. This caused the orientations flanking the trained orientations to be encoded as the trained orientations themselves. This funneling mechanism might be supporting the capability of the visual system to generalize [41]. It is therefore possible that the purpose of the visual cortex in the context of our task might not be to determine with precision the orientation of the visual stimulus but rather to determine whether the stimulus presented is Go-like, No-Go-like, or neither. The suppression of the activity representing orientations surrounding the trained orientations is reminiscent, in the orientation space, of the spatial modulation encountered in V1 after learning in the context of a figure-ground detection task [24] or of spatial attention [17,42-48]. It remains debated whether this surround suppression should be attributed to an attentional mechanism facilitating the selection of what was learned, or to a transient computational step required for learning [41,49,50]. However, the correlation found between performance and distortion of the orientation space suggests that this mechanism is important for the animal when it performs the task.

The large decrease in activity observed in trained mice suggests that the cellular mechanisms underpinning the distorted representation of orientation depends on interneuron-mediated inhibition. Inhibitory interneurons have been shown to facilitate the plasticity of cortical circuits during learning [51-56]. Somatostatin neurons play an important role in gating response plasticity in V1 [57]. Those neurons are, at least in part, under the control of top-down signal arising from the orbitofrontal cortex [58], the cingulate cortex [59] and/or the retrosplenial cortex [55]. Yet, parvalbumin-expressing interneurons were also shown to present some modulation of their tuning properties after training [57]. Furthermore, several mechanisms involving mesencephalic dopaminergic neuromodulation [48,60,61], and cholinergic neuromodulation [62]

were also associated to perceptual learning in the context of attention and reward. It was recently proposed that cholinergic signals promote cortical plasticity during associative learning via the activation of vasoactive intestinal polypeptide-expressing (VIP) interneurons [63]. In our results, the changes in orientation representation were context-dependent, suggesting a pivotal role of top-down modulation. However, we cannot exclude the participation of local synaptic plasticity [6]. The respective role of top-down modulation, local synaptic plasticity and neuromodulation in the distortion of the orientation space described in this study will need to be further investigated in the future.

To conclude, our results provide new evidence that, during the performance of a learned perceptual task, visual processing is biased toward the cues of the task. This bias effectively structures the neuronal responses around attractors in the encoding space. It is possible that this mechanism facilitates learning by allowing the generalization of the properties of the cue of interest to a larger proportion of the encoding space. Because we cannot find those attractors in passive mice, our results also suggest that the deformation of orientation representational space that results from perceptual learning is an *ad-hoc* modulation. The function of this filter imposed transiently to the sensory cortices (possibly by top-down signals), would be to facilitate the extraction of task-relevant information.

Conflict of interest statement: The authors declare no competing financial interests.

Acknowledgments: This work was funded by the Whitehall Foundation (grant 2015-08-69), the Charles and Johanna Busch Biomedical Grant Program, and the National Institutes of Health – National Eye Institute (Grant #R01 EY030860). J.C. was supported by a Fyssen Foundation postdoctoral fellowship. The authors are grateful to Denis Paré, Bart Krekelberg and the members of the laboratory for their comments on the manuscript.

Methods

All the procedures described below have been approved by the Institutional Animal Care and Use Committee (IACUC) of Rutgers University-Newark, in agreement with the Guide for the Care and Use of Laboratory Animals (National Research Council, 2011).

Surgery.

Head-bar implants. 10 minutes after systemic injection of an analgesic (carprofen, 5 mg per kg of body weight), adult (3-6 months old) male and female Gad2-IRES-Cre (Jackson stock #019022) x Ai9 (Jackson stock #007909) mice were anesthetized with isoflurane (5% induction, 1.2% maintenance) and placed in a stereotaxic frame. Body temperature was kept at 37°C using a feedback-controlled heating pad. Pressure points and incision sites were injected with lidocaine (2%). Eyes were protected from desiccation with artificial tear ointment (Dechra). Next, the skin covering the skull was incised and a custom-made lightweight metal head-bar was glued to the skull using Vetbond (3M). In addition, a large recording chamber capable of retaining the water necessary for using a water-immersion objective was built using dental cement (Ortho-Jet, Lang). Mice recovered from surgery for 5 days, during which amoxicillin was administered in drinking water (0.25 mg/mL).

AAV virus injection. After recovery from the head-bar surgery, mice were anesthetized using isoflurane as described above. A circular craniotomy (diameter = 3 mm) was performed above V1. The AAV vector AAV1.eSyn.GCaMP6f.WPRE.SV40 (UPenn Vector Core) carrying the gene of the fluorescent calcium sensor GCaMP6f was injected at three sites 500 µm apart around the center of V1 (stereotaxic coordinates: -4.0 mm AP, +2.2 mm ML from bregma) using a MicroSyringe Pump Controller Micro 4 (World Precision Instruments, WPI) at a rate of 30 nl/min. Injections started at a depth of 550 µm below the pial surface and the tip of the pipette was raised in steps of 100 µm during the injection, up to a depth of 200 µm below the dura surface. The total volume injected across all depths was 0.7 µl. After removal of the injection pipette, a 3-mm-diameter coverslip was placed over the dura, such that the coverslip fits entirely in the craniotomy and was flush with the skull surface. The coverslip was kept in place using Vetbond and dental cement. Mice were left to recover from the surgery for at least 3 weeks to obtain a satisfactory gene expression.

Imaging.

Calcium imaging setup. During the last week of recovery, mice were trained to stay on a spherical treadmill consisting of a ball floating on a small cushion of air that allowed for full 2D movement [15]. During three daily twenty-minute sessions, the mouse head-bar was fixed to a post holding the mouse on the apex of the spherical treadmill. Ball motion was tracked by an IR camera taking pictures of the ball at 30 Hz. Eye motion was monitored at 15 Hz using a second IR camera imaging the reflection of the eye on an infrared dichroic mirror. Functional imaging was performed at 15 frames per second using a resonant scanning two-photon microscope (Neurolabware) powered by a Ti-Sapphire Ultra-2 laser (Coherent) set at 910 nm. The microscope scanning mirrors were hermetically sealed in a chamber to bring the scanning hum below the room ambient noise (< 59 dBA). The laser beam was focused 200 microns below the cortical surface using a 16×, 0.8 NA Nikon water-immersion objective. The objective was tilted 30 degrees such that the objective lens was parallel to the dura surface. Laser power was kept below 70 mW. Frames (512x796 pixels) were acquired using the software Scanbox developed by Neurolabware.

Naïve imaging session. Mice were placed head fixed in the functional imaging rig in front of a screen placed such as it

covers the visual field of the right eye, contralateral to the craniotomy. Visual stimuli of the task block consisted of the presentation of one of two vertical sinewave gratings that drifted toward the right and were rotated clockwise by 45° and 135° (temporal frequency = 2 Hz, spatial frequency = 0.04 cycle per degree, contrast = 75%; duration: 3 s; intertrial interval: 3 s). Visual cues were presented in a pseudorandom order, such as the same stimulus could not be presented more than three times in a row. Some recording sessions contained blocks during which the two visual cues were paired with pure tones. Those blocks were not included in the analysis. At the end of the imaging session, after a break of at least five minutes, we assessed the orientation tuning of the imaged neurons by presenting an orientation tuning block that consisted of the presentation of a series of drifting sinewave gratings (12 orientations evenly spaced by 30° and randomly permuted). The spatiotemporal parameters of the orientation tuning stimuli were identical to those for the task block except for their duration (temporal frequency = 2 Hz, spatial frequency = 0.04 cycle per degree, contrast = 75%; duration: 1.5 s; intertrial interval: 3 s). As scanning was not synced to the stimuli, a photodiode located at the top left corner of the screen was used to detect the exact timing of the visual stimulus onset and offset. The photodiode signal was acquired along with the following signals: 1) a signal provided by the two-photon microscope, which indicated the onset of each frame, and 2) two analog signals encoding the orientation of the drifting grating. These signals were digitized (NiDAQ, National Instruments) and recorded with the software WinEDR (John Dempster, University of Strathclyde). Imaging sessions started by recording one thousand frames with the green and red channels. The red channel was used to exclude GABAergic neurons from the analysis.

Behavioral training.

Behavioral training. After the naïve recording session, mice were water-deprived to 90% of their body weight and acclimated to head fixation on a spherical treadmill in custom-built, soundproof training rigs. Each rig was equipped with a monitor (Dell), water dispenser with a built-in lickometer (to monitor licking, infrared beam break). Data acquisition boards (National Instruments and Arduino) were used to actuate water delivery and vacuum reward retrieval as well as monitor animal licking. The monitor and data acquisition boards were connected to a computer that ran the custom-made training program scripted in MATLAB. Once animals reached the target weight and were acclimated to the training setup, they were trained to perform the orientation discrimination task. In this task, drifting sine-wave gratings oriented 45° below the vertical were paired with a water reward, and the animal was expected to lick (Go). Drifting gratings orthogonal to the Go signal (orientation 135°) signaled the absence of reward, and the animal was expected to withhold licking (No-Go) during those trials. When the stimulus instructed the animal to lick, the water delivery had to be triggered by the mouse licking during the third second of the stimulus presentation. No water was dispensed in the no-lick condition or if the mouse failed to trigger water delivery in the lick condition. If the animal responded correctly (Hit or Correct Rejection (CR)), the intertrial interval was 3 s. If the animal responded incorrectly (Miss or False Alarm (FA)), the intertrial interval was increased to 9.5 s as negative reinforcement. Performance was measured using the D' statistic (D'=norminv(Hit rate) - norminv(False alarm rate), norminv = inverse of the normal cumulative distribution function). Animals are considered experts if their performance during training sessions is greater than 1.7 (probability of chance behavior < 0.1%, Monte Carlo simulation).

Trained and passive imaging sessions. Imaging sessions in trained mice (active or passive) were similar to the imaging sessions performed in naïve mice. However, in trained animals, imaging was performed while mice were performing the Go/No-Go discrimination task. In passive sessions, imaging was performed in trained mice at least 48 hours after water restriction was terminated. In this case, mice were not performing task despite the water dispenser apparatus being present in front of the animal.

Timing of the response. The timing of the response (lick onset during Go trials and lick offset during NoGo trials) was quantified by averaging for each animal the lick traces recorded during the presentation of the Go and No-Go stimuli respectively. The mean trace across correct trials (hit and correct rejection respectively) was fitted by a sigmoid function. Lick onset (Go trials) and offset (NoGo trials) were defined as the inflection points of the rising (Go trials) or falling sigmoid fit (No-Go trials). The measure of lick cessation during No-Go trial was possible because mice presented a background licking activity that stopped shortly before the decision windows (3rd second of the stimulus presentation) during correct rejection trials.

Data analysis

All the analyses detailed below were performed using custom MATLAB routines.

Imaging data pre-processing. Calcium imaging frames were realigned offline to remove movement artifacts using the Scanbox algorithm (Neurolabware). A region of interest (ROI) was determined for each neuron using a semi-automatic segmentation routine. For every frame, the fluorescence level was averaged across the pixels of the ROI. Potential contamination of the soma fluorescence by the local neuropil was removed by subtracting the mean fluorescence of a 2-5 µm ring surrounding the neuron's ROI, excluding the soma of neighboring neurons, and then adding the median value across time of the subtracted background. We then computed the fractional fluorescence from the background-subtracted fluorescence data. The fractional fluorescence ($dF/F = (F - F_0) / F_0$), was calculated with F_0 defined as the median of the raw

fluorescence measured during every inter-trial interval. The fractional fluorescence was then deconvolved into Action Potential related Events (APrE) using a benchmarked deconvolution method [14,63]. Only neurons whose correlation between the inferred fractional fluorescence (convolved back from the inferred APrE) and the measured fractional fluorescence was greater than 0.8 were used for the analysis. This inclusion criteria removed neurons for which the calcium signal was too poor to allow the algorithm to accurately detect APrE. Trials were isolated and stored on a SQL database. The orientation tuning curve of each neuron was computed from the APrE responses recorded during the tuning curve blocks using a resampling-based Bayesian method [13,64]. The preferred orientation was defined as the peak of the orientation tuning curve. When a neuron was not direction selective (i.e. responding equally to the same oriented stimulus moving in opposite directions), the preferred orientation was defined as the orientation included in the range [0 – 180].

Responsive neurons. Responsive neurons were identified using a binomial test that compared the probability of having at least one APrE during the stimulus presentation, to the probability of having at least one APrE during the pre-stimulus time. We first estimated the probability of each neuron to fire APrE during the last second of the intertrial interval (binofit function (MATLAB, MathWorks); Clopper–Pearson method). Then, we compared the probability of observing activity during a one second window sliding across the presentation of the stimulus to the upper bound of the 99% confidence interval of the probability of activity calculated for the intertrial interval. A neuron was considered responsive if probability of activity in at least one of the sliding windows was greater than the confidence interval intertrial activity probability. The use of a binomial model to detect neural responses was motivated by the sparseness of the observed activity. Indeed, 84% of the intertrial interval windows did not contain any APrE; only 8% contained a single APrE and another 8% two or more. Hence, assessing the presence of any amount of APrE captured most of the variability of the neuronal activity. The use of a Wilcoxon test approach to compare either the responsiveness (1s pre-stimulus vs post-stimulus) or the selectivity (1s post-stimulus 45° vs 135°; see Poort et al., 2015) led to similar results and conclusions (not shown).

Population selectivity. A selectivity index was computed for each neuron using the ratio of the difference in the mean response during the first second of 45° and 135° stimuli presentation over the pooled standard deviation of the two variables [6]. The population selectivity was computed by taking the square root of the sum of the squared individual selectivity indices values divided by the total number of cells.

$$\left(\sqrt{\sum_i^{Ncell} SelectivityIndex_i^2 / Ncell}\right)$$

The trained neuronal population contained a small number of neurons with outlier large modulation index values (see **Fig. S1**c). The two largest outliers in trained mice (0.999 quantile, 2/2141 neurons) were removed.

Representation of the orientation space in V1. Neurons were separated into three pools of neurons based on the behavioral context of the imaging session (naïve, trained, passive). The preferred orientation of every neuron was first assessed (see above). A relatively small contingent of neurons did not show significant orientation or direction tuning (naïve: 11.3%, n = 206/1830; trained: 16.7%, n = 357/2141). 90.8% and 90.6% of those untuned neurons were unresponsive in naïve and passive mice respectively, and their activity was significantly lower than that of tuned neurons (see **Fig. S1** for details). These neurons were not considered further. For each stimulus of each group (e.g. 45° presented to trained mice), we generated a raster of 1,000 resampled trial using the following procedure. For each resampled trial, 250 neurons of the pool of interest (e.g. trained mice pool) were randomly selected, and for each of those neuron, one trial during which the stimulus of interest was presented (e.g. 45°) was selected at random. The activity of the neuron during the selected trial was included in an orientation/time matrix in which each line (binned every degree from 0 to 360 degrees) corresponded to the time course of the activity of neurons having this preferred orientation. The resulting trial matrix was baseline-subtracted and smoothed across the orientation space by a gaussian weighted mean (std= 6°). The 1000 trial matrices were averaged to obtain the representation of the orientation of the stimulus by the V1 population. To determine *the distribution across trials of the peak of activity*, the peak of activity in the orientation space was determined for each trial, and the distribution of the 1,000 resulting orientations was fit by a mixture of von Mises distributions (two to four mixture components) using mvmdist (https://github.com/chrschy/mvmdist), a Matlab package for probabilistic modeling of circular data.

Locomotion and pupil size. Locomotion (ball motion) and pupil size traces were lowpass filtered (0.25Hz) and then z-scored for each session. For pupil size, a small number of trials were discarded because either the mouse was blinking or the position the eye was not allowing us to reliably assess the size of the pupil. To reflect the resampling procedure used to determine the representation of the orientation of the visual stimuli by the V1 neuronal population (see above), we performed the same resampling procedure, replacing the activity of the neurons by the amount of locomotion or pupil size data. To evaluate the impact of the differences of locomotion and pupil size between groups, we sorted all trials in 4 categories according to the amount of locomotion/pupil size. The cut-off values were based on the quartile values of the trained group, the group displaying the largest amplitude of locomotion and greater pupil size values compared to the other groups.

Statistics

Permutation tests. To determine if the mean across trials computed from two different pools was significantly different, we compared the value obtained to the distribution of 1,000 or 10,000 differences obtained when the pool labels were shuffled. The two-tailed confidence intervals of the null hypothesis at the alpha level 0.05 were defined as the 2.5 and 97.5 percentile of the distribution obtain from the permutations. The difference between the observed means was considered significant if located outside the confidence interval of the null distribution.

Comparison of distributions. Comparison of the preferred orientation distribution was achieved with an Anderson-Darling nonparametric test, adjusted for ties [8].

Circular statistics. Circular statistics were computed with the Circular Statistics Toolbox for MATLAB (Berens, 2009). The difference between groups was considered significant if the values of the confidence interval boundaries were of same sign.

Neuronal activity. The significance of the difference in activity naïve and trained mice of neurons belonging to a specific neuronal group defined by their orientation preference was assessed by a two-sample Kolmogorov-Smirnov test with a significance threshold of $p < 0.05$.

Effect size. Because large number of observations can lead to conclude on significant difference of even irrelevantly small magnitude, we complimented or replace tests in certain cases with an effect size measure, Hedges'g, using a MATLAB toolbox [66]. Traditionally, values under 0.2 correspond to the absence of effect, values between 0.2 and 0.4 are considered to reflect a small effect, values between 0.4 and 0.6 correspond to a medium effect, and values larger than 0.6 correspond to a large effect size. In the figures of this article, we only represent medium (>0.4) and higher values.

References

1 Bi, G.-q. & Poo, M.-m. Synaptic Modification by Correlated Activity: Hebb's Postulate Revisited. *Annual Review of Neuroscience* **24**, 139-166, doi:10.1146/annurev.neuro.24.1.139 (2001).

2 Lisman, J., Cooper, K., Sehgal, M. & Silva, A. J. Memory formation depends on both synapse-specific modifications of synaptic strength and cell-specific increases in excitability. *Nat Neurosci* **21**, 309-314, doi:10.1038/s41593-018-0076-6 (2018).

3 Weinberger, N. M. & Bieszczad, K. M. in *Neurobiology of Sensation and Reward Frontiers in Neuroscience* (ed J. A. Gottfried) (2011).

4 Goltstein, P. M., Meijer, G. T. & Pennartz, C. M. Conditioning sharpens the spatial representation of rewarded stimuli in mouse primary visual cortex. *Elife* **7**, doi:10.7554/eLife.37683 (2018).

5 Henschke, J. U. *et al.* Reward Association Enhances Stimulus-Specific Representations in Primary Visual Cortex. *Curr Biol* **30**, 1866-1880 e1865, doi:10.1016/j.cub.2020.03.018 (2020).

6 Poort, J. *et al.* Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. *Neuron* **86**, 1478-1490, doi:10.1016/j.neuron.2015.05.037 (2015).

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

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

9 Gdalyahu, A. *et al.* Associative fear learning enhances sparse network coding in primary sensory cortex. *Neuron* **75**, 121-132, doi:10.1016/j.neuron.2012.04.035 (2012).

10 Ghose, G. M., Yang, T. & Maunsell, J. H. Physiological correlates of perceptual learning in monkey V1 and V2. *J Neurophysiol* **87**, 1867-1888, doi:10.1152/jn.00690.2001 (2002).

11 Olshausen, B. A. & Field, D. J. Sparse coding of sensory inputs. *Curr Opin Neurobiol* **14**, 481-487, doi:10.1016/j.conb.2004.07.007 (2004).

12 Perrinet, L. U. Role of homeostasis in learning sparse representations. *Neural Comput* **22**, 1812-1836, doi:10.1162/neco.2010.05-08-795 (2010).

13 McClure, J. P. & Polack, P.-O. Pure tones modulate the representation of orientation and direction in the primary visual cortex. *bioRxiv*, 531194, doi:10.1101/531194 (2019).

14 Deneux, T. *et al.* Accurate spike estimation from noisy calcium signals for ultrafast three-dimensional imaging of large neuronal populations in vivo. *Nat Commun* **7**, 12190, doi:10.1038/ncomms12190 (2016).

15 Polack, P. O., Friedman, J. & Golshani, P. Cellular mechanisms of brain state-dependent gain modulation in visual cortex. *Nat Neurosci* **16**, 1331-1339, doi:10.1038/nn.3464 (2013).

16 Vinck, M., Batista-Brito, R., Knoblich, U. & Cardin, J. A. Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron* **86**, 740-754, doi:10.1016/j.neuron.2015.03.028 (2015).

17 Eriksen, C. W., Pan, K. & Botella, J. Attentional distribution in visual space. *Psychol Res* **56**, 5-13, doi:10.1007/BF00572128 (1993).

18 Schoups, A., Vogels, R., Qian, N. & Orban, G. Practising orientation identification improves orientation coding in V1 neurons. *Nature* **412**, 549-553, doi:10.1038/35087601 (2001).

19 Yang, T. & Maunsell, J. H. The effect of perceptual learning on neuronal responses in monkey visual area V4.

J Neurosci **24**, 1617-1626, doi:10.1523/JNEUROSCI.4442-03.2004 (2004).

20 Rutkowski, R. G. & Weinberger, N. M. Encoding of learned importance of sound by magnitude of representational area in primary auditory cortex. *Proc Natl Acad Sci U S A* **102**, 13664-13669, doi:10.1073/pnas.0506838102 (2005).

21 Blake, D. T., Heiser, M. A., Caywood, M. & Merzenich, M. M. Experience-dependent adult cortical plasticity requires cognitive association between sensation and reward. *Neuron* **52**, 371-381, doi:10.1016/j.neuron.2006.08.009 (2006).

22 Li, W., Piech, V. & Gilbert, C. D. Learning to link visual contours. *Neuron* **57**, 442-451, doi:10.1016/j.neuron.2007.12.011 (2008).

23 Wiest, M. C., Thomson, E., Pantoja, J. & Nicolelis, M. A. Changes in S1 neural responses during tactile discrimination learning. *J Neurophysiol* **104**, 300-312, doi:10.1152/jn.00194.2010 (2010).

24 Yan, Y. *et al.* Perceptual training continuously refines neuronal population codes in primary visual cortex. *Nat Neurosci* **17**, 1380-1387, doi:10.1038/nn.3805 (2014).

25 Shuler, M. G. & Bear, M. F. Reward timing in the primary visual cortex. *Science* **311**, 1606-1609, doi:10.1126/science.1123513 (2006).

26 Stanisor, L., van der Togt, C., Pennartz, C. M. & Roelfsema, P. R. A unified selection signal for attention and reward in primary visual cortex. *Proc Natl Acad Sci U S A* **110**, 9136-9141, doi:10.1073/pnas.1300117110 (2013).

27 Schultz, W. Neuronal Reward and Decision Signals: From Theories to Data. *Physiol Rev* **95**, 853-951, doi:10.1152/physrev.00023.2014 (2015).

28 Jordan, R. & Keller, G. B. Opposing Influence of Top-down and Bottom-up Input on Excitatory Layer 2/3 Neurons in Mouse Primary Visual Cortex. *Neuron* **108**, 1194-1206 e1195, doi:10.1016/j.neuron.2020.09.024 (2020).

29 Adesnik, H., Bruns, W., Taniguchi, H., Huang, Z. J. & Scanziani, M. A neural circuit for spatial summation in visual cortex. *Nature* **490**, 226-231, doi:10.1038/nature11526 (2012).

30 Vinje, W. E. & Gallant, J. L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**, 1273-1276 (2000).

31 Meijer, G. T., Montijn, J. S., Pennartz, C. M. A. & Lansink, C. S. Audiovisual Modulation in Mouse Primary Visual Cortex Depends on Cross-Modal Stimulus Configuration and Congruency. *J Neurosci* **37**, 8783-8796, doi:10.1523/JNEUROSCI.0468-17.2017 (2017).

32 Ibrahim, L. A. *et al.* Cross-Modality Sharpening of Visual Cortical Processing through Layer-1-Mediated Inhibition and Disinhibition. *Neuron* **89**, 1031-1045, doi:10.1016/j.neuron.2016.01.027 (2016).

33 O'Connor, D. H., Peron, S. P., Huber, D. & Svoboda, K. Neural activity in barrel cortex underlying vibrissa-based object localization in mice. *Neuron* **67**, 1048-1061, doi:10.1016/j.neuron.2010.08.026 (2010).

34 Houweling, A. R. & Brecht, M. Behavioural report of single neuron stimulation in somatosensory cortex. *Nature* **451**, 65-68, doi:10.1038/nature06447 (2008).

35 Huber, D. *et al.* Sparse optical microstimulation in barrel cortex drives learned behaviour in freely moving mice. *Nature* **451**, 61-64, doi:10.1038/nature06445 (2008).

36 Attwell, D. & Laughlin, S. B. An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab* **21**, 1133-1145, doi:10.1097/00004647-200110000-00001 (2001).

37 Lennie, P. The cost of cortical computation. *Curr Biol* **13**, 493-497, doi:10.1016/s0960-9822(03)00135-0 (2003).

38 Willmore, B. D., Mazer, J. A. & Gallant, J. L. Sparse coding in striate and extrastriate visual cortex. *J Neurophysiol* **105**, 2907-2919, doi:10.1152/jn.00594.2010 (2011).

39 Gold, J., Bennett, P. J. & Sekuler, A. B. Signal but not noise changes with perceptual learning. *Nature* **402**, 176-178, doi:10.1038/46027 (1999).

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

41 Ahissar, M. & Hochstein, S. Task difficulty and the specificity of perceptual learning. *Nature* **387**, 401-406, doi:10.1038/387401a0 (1997).

42 Carr, T. H. & Dagenbach, D. Semantic priming and repetition priming from masked words: evidence for a center-surround attentional mechanism in perceptual recognition. *J Exp Psychol Learn Mem Cogn* **16**, 341-350 (1990).

43 Steinman, B. A., Steinman, S. B. & Lehmkuhle, S. Visual attention mechanisms show a center-surround organization. *Vision Res* **35**, 1859-1869, doi:10.1016/0042-6989(94)00276-r (1995).

44 Caputo, G. & Guerra, S. Attentional selection by distractor suppression. *Vision Res* **38**, 669-689, doi:10.1016/s0042-6989(97)00189-2 (1998).

45 Mounts, J. R. Attentional capture by abrupt onsets and feature singletons produces inhibitory surrounds. *Percept Psychophys* **62**, 1485-1493, doi:10.3758/bf03212148 (2000).

46 Slotnick, S. D., Hopfinger, J. B., Klein, S. A. & Sutter, E. E. Darkness beyond the light: attentional inhibition surrounding the classic spotlight. *Neuroreport* **13**, 773-778, doi:10.1097/00001756-200205070-00008 (2002).

47 Cutzu, F. & Tsotsos, J. K. The selective tuning model of attention: psychophysical evidence for a suppressive annulus around an attended item. *Vision Res* **43**, 205-219, doi:10.1016/s0042-6989(02)00491-1 (2003).

48 Muller, C. P. & Huston, J. P. Dopamine activity in the occipital and temporal cortices of rats: dissociating effects of sensory but not pharmacological stimulation. *Synapse* **61**, 254-258, doi:10.1002/syn.20366 (2007).

49 Dolan, R. J. *et al.* How the brain learns to see objects and faces in an impoverished context. *Nature* **389**, 596-599, doi:10.1038/39309 (1997).

50 Gilbert, C. D. & Li, W. Top-down influences on visual processing. *Nat Rev Neurosci* **14**, 350-363, doi:10.1038/nrn3476 (2013).

51 Letzkus, J. J. *et al.* A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature* **480**, 331-335, doi:10.1038/nature10674 (2011).

52 Donato, F., Rompani, S. B. & Caroni, P. Parvalbumin-expressing basket-cell network plasticity induced by experience regulates adult learning. *Nature* **504**, 272-276, doi:10.1038/nature12866 (2013).

53 Chen, S. X., Kim, A. N., Peters, A. J. & Komiyama, T. Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. *Nat Neurosci* **18**, 1109-1115, doi:10.1038/nn.4049 (2015).

54 Kato, H. K., Gillet, S. N. & Isaacson, J. S. Flexible Sensory Representations in Auditory Cortex Driven by Behavioral Relevance. *Neuron* **88**, 1027-1039, doi:10.1016/j.neuron.2015.10.024 (2015).

55 Makino, H. & Komiyama, T. Learning enhances the relative impact of top-down processing in the visual cortex. *Nat Neurosci* **18**, 1116-1122, doi:10.1038/nn.4061 (2015).

56 Sachidhanandam, S., Sermet, B. S. & Petersen, C. C. H. Parvalbumin-Expressing GABAergic Neurons in Mouse Barrel Cortex Contribute to Gating a Goal-Directed Sensorimotor Transformation. *Cell Rep* **15**, 700-706, doi:10.1016/j.celrep.2016.03.063 (2016).

57 Khan, A. G. *et al.* Distinct learning-induced changes in stimulus selectivity and interactions of GABAergic interneuron classes in visual cortex. *Nat Neurosci* **21**, 851-859, doi:10.1038/s41593-018-0143-z (2018).

58 Liu, D. *et al.* Orbitofrontal control of visual cortex gain promotes visual associative learning. *Nat Commun* **11**, 2784, doi:10.1038/s41467-020-16609-7 (2020).

59 Zhang, S. *et al.* Selective attention. Long-range and local circuits for top-down modulation of visual cortex processing. *Science* **345**, 660-665, doi:10.1126/science.1254126 (2014).

60 Bao, S., Chan, V. T. & Merzenich, M. M. Cortical remodelling induced by activity of ventral tegmental dopamine neurons. *Nature* **412**, 79-83, doi:10.1038/35083586 (2001).

61 Rivera, A. *et al.* Cellular localization and distribution of dopamine D(4) receptors in the rat cerebral cortex and their relationship with the cortical dopaminergic and noradrenergic nerve terminal networks. *Neuroscience* **155**, 997-1010, doi:10.1016/j.neuroscience.2008.05.060 (2008).

62 Gavornik, J. P., Shuler, M. G., Loewenstein, Y., Bear, M. F. & Shouval, H. Z. Learning reward timing in cortex through reward dependent expression of synaptic plasticity. *Proc Natl Acad Sci U S A* **106**, 6826-6831, doi:10.1073/pnas.0901835106 (2009).

63 Gasselin, C., Hohl, B., Vernet, A., Crochet, S. & Petersen, C. C. H. Cell-type-specific nicotinic input disinhibits mouse barrel cortex during active sensing. *Neuron* **109**, 778-787 e773, doi:10.1016/j.neuron.2020.12.018 (2021).

64 Berens, P. *et al.* Community-based benchmarking improves spike rate inference from two-photon calcium imaging data. *PLoS Comput Biol* **14**, e1006157, doi:10.1371/journal.pcbi.1006157 (2018).

65 Cronin, B., Stevenson, I. H., Sur, M. & Kording, K. P. Hierarchical Bayesian modeling and Markov chain Monte Carlo sampling for tuning-curve analysis. *J Neurophysiol* **103**, 591-602, doi:10.1152/jn.00379.2009 (2010).

66 Hentschke, H. & Stuttgen, M. C. Computation of measures of effect size for neuroscience data sets. *Eur J Neurosci* **34**, 1887-1894, doi:10.1111/j.1460-9568.2011.07902.x (2011).

Supplementary Figures

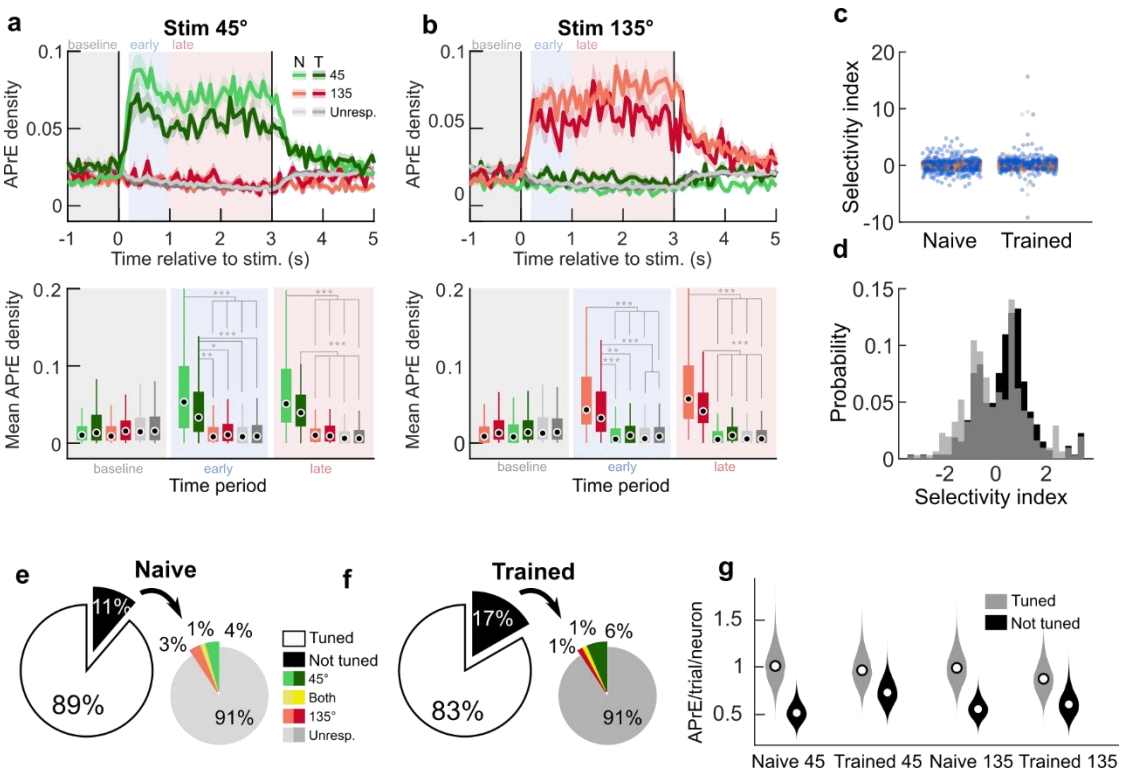

Figure S1. Individual neuronal responses, selectivity indices and untuned neurons a. Time course of the response evoked by the 45° grating for neurons grouped by stimulus responsiveness. Top panel: mean across neurons ± s.e.m. Bottom panel: mean APrE density over baseline (-1s to 0s), early response (0.150s to 1s) and late response (1s to 3s) periods, box plots representing quartiles, dots indicate medians. Effect size (Hedge's g): ★ = g>0.4 ★★ = g>0.6, ★★★ = g>0.8. **b.** Same representation as c. but for the 135° grating presentation. **c.** Individual neuron selectivity indices in naïve and trained mice. Orange dots indicate neurons labelled unresponsive/responsive to both visual cues, blue points stand for neurons labelled responsive to at least one of the repeated visual cues. **d.** Distribution of the selectivity indices for the responsive neurons. **e.** Proportion of untuned neurons in naive animals, and proportion of responsive/unresponsive neurons among those. **f.** Same representation as in c for the neurons recorded in trained animals. **g.** Population APrEs per trial for a resampled (10,000 repetitions) subpopulation of 200 neurons of tuned and untuned neurons.

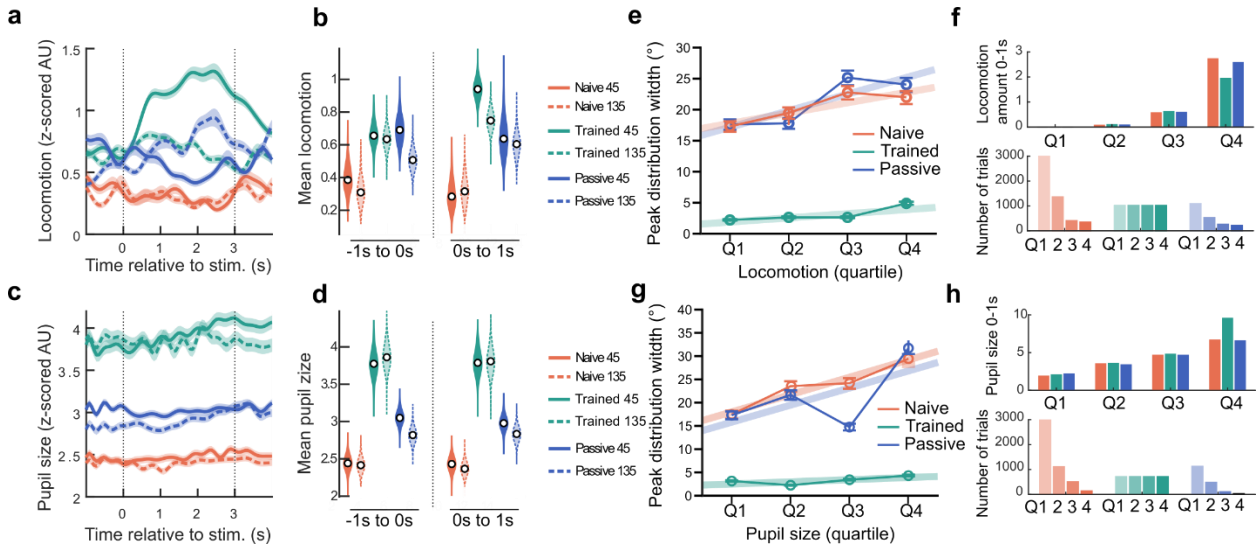

Figure S2. Locomotion and pupil size do not explain the improved population representation stability in trained mice. a. Mean locomotion time course in z-scores during the presentation of 45° and 135° cues in naïve, trained and passive animals. The traces were obtained following the same resampling method as for the neural representations. Mean ± standard deviation. **b.** Mean values distribution for the second preceding the stimulus onset (left group) and the second following the stimulus onset (right group). **c,d.** Same as a and b, but for pupil size. **e.** Peak distribution width, a measure of population representation stability, as a function of the quantity of locomotion recorded during the first second after stimulus onset. All trials were sorted according to quartiles cutoff values computed from the trials of the trained group. σ± $C.I._{95\%}$ Lines indicate the best linear fit on the four points in each group. **f.** Top: mean locomotion values measured during the first second after stimulus onset for each group and each quartile. Bottom: Number of trials sorted in each category. **g,h.** Same as in c and d, but for pupil size.

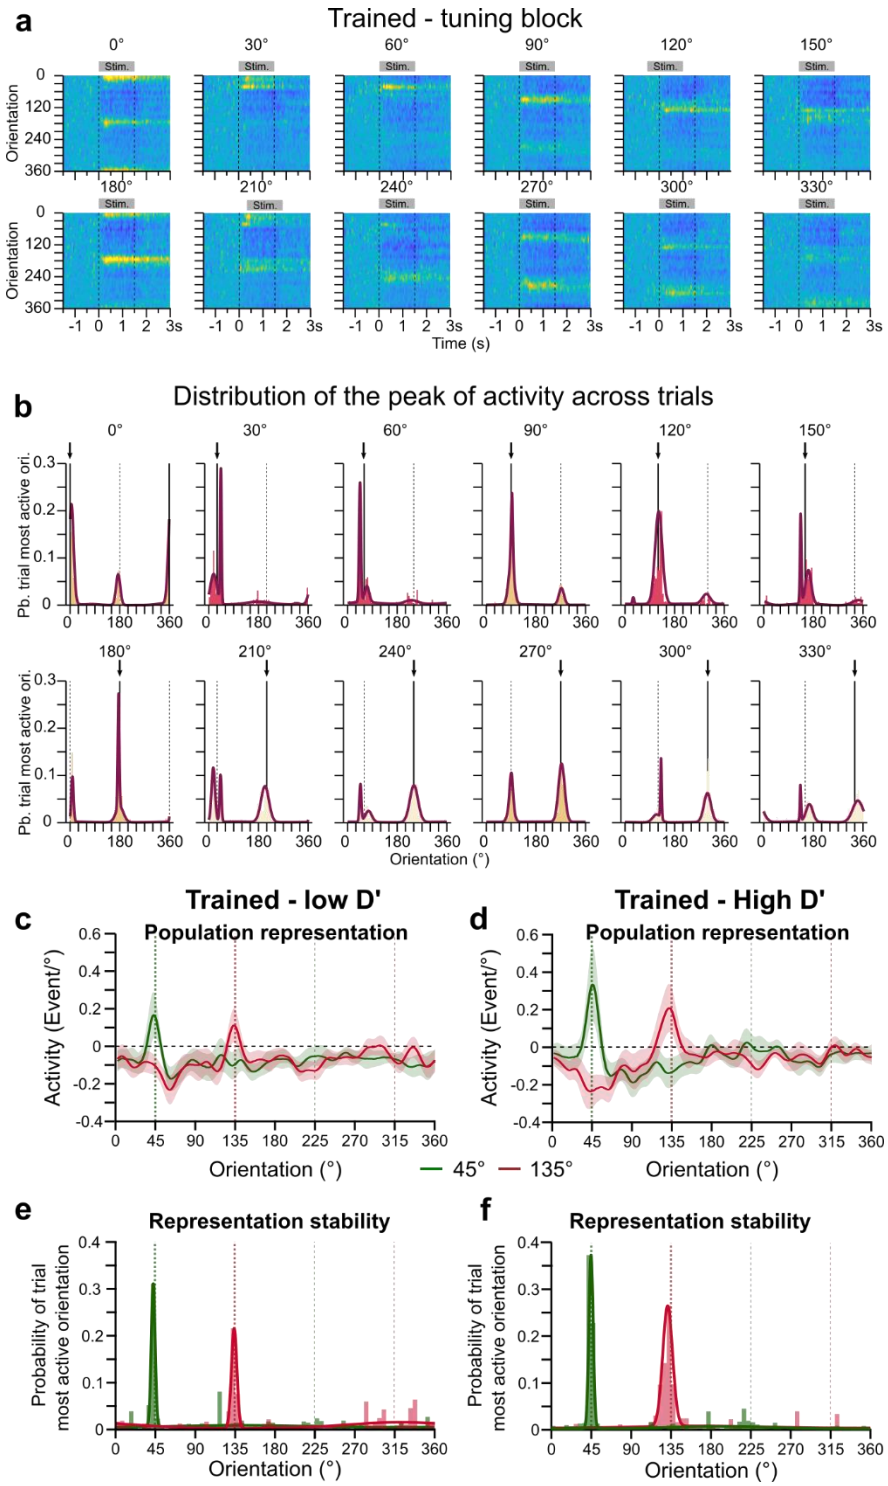

Figure S3. Tuning and task block responses according to the mouse behavioral performance. a. Orientation space activity maps during the presentation of the 12 orientations of the tuning block in trained mice. Dark blue hues indicate activity lower than baseline. Yellow hues indicate activity greater than baseline. **b.** Probability distribution of the location in the orientation space of the peak of activity for each of the 12 oriented stimuli of the tuning block. Purple curves: three component von Mises

mixture model fit. **c.** Average activity profile in the orientation space over the first second of the presentation of the 45° (green trace) and 135° (red) task stimuli in trained mice sessions with a low D' (see Fig. 4i). Shaded areas indicate the standard deviation across time. **d.** Same as c for high D' sessions. **e.** Probability distribution of the peak of the activity profile when 45° or 135° stimuli were presented in low D' sessions. The distributions were fitted with a two component von Mises mixture model. **f.** Same as e for high D' sessions.