

Reviewed Preprint
v1 • August 3, 2026
Not revised

✉ For correspondence:

anthony.renard@epfl.ch

georgios.foustoukos@epfl.ch

sylvain.crochet@epfl.ch

carl.petersen@epfl.ch

* These authors contributed equally

† These authors contributed equally

Competing interests: CCHP is a Reviewing Editor at eLife. No other competing interests declared.

Funding: See [page 21](#)

Reviewing editor: Andrew J King, University of Oxford, United Kingdom

© 2026, Renard et al. This article is distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use and redistribution provided that the original author and source are credited.

Rapid cortical reorganization tracks goal-directed sensorimotor learning in real time

Anthony Renard^{*} ✉, Georgios Foustoukos^{*} ✉, Maya Iuga, Pol Bech, Axel Bisi, Robin F Dard, Sylvain Crochet[†] ✉, Carl CH Petersen[†] ✉

Laboratory of Sensory Processing, Brain Mind Institute, School of Life Sciences, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

eLife Assessment

This **valuable** study aimed to explore whether sensory representations in the cortex reorganize on the same timescale over which behavioral changes first emerge.

Convincing evidence is presented to show that reward-dependent plasticity takes place in the mouse barrel cortex within a single recording session as mice learn a whisker detection task. However, the evidence supporting the authors' proposal that this rapid reorganization may involve spontaneous reactivation of neurons that gain stimulus responsiveness during training is **incomplete**. The work will be of interest to systems neuroscientists studying cortical circuits and learning.

<https://doi.org/10.7554/eLife.111818.1.sa3>

Abstract

Sensorimotor associations are typically thought to require days of training to consolidate in sensory cortex, yet adaptive behavior can emerge within minutes. Here, we developed a barrel cortex-dependent whisker-based detection task in which mice learned to associate a novel tactile whisker stimulus with reward within a single behavioral session. Longitudinal two-photon calcium imaging of layer 2/3 barrel cortex neurons revealed that reward-driven learning rapidly reorganized the neuronal representation of the whisker deflection within a single session. Population decoding tracked this transition trial-by-trial during learning with neuronal trajectories mirroring behavior. Critically, neurons that gained stimulus responsiveness across training preferentially took part in spontaneous reactivation events during learning, suggesting that online reactivations could act as a potential upstream selection mechanism. Our results suggest that reward-based learning evokes rapid sensory cortical reorganization on the timescale of minutes, which could be mediated by a concurrent reactivation-based mechanism driving plasticity.

Introduction

Adaptive behavior requires the brain to rapidly associate sensory stimuli with appropriate motor responses. While learning-related changes across sensory modalities have been extensively investigated, most studies rely on training paradigms extending over days or weeks, leaving the dynamics of the formation of the initial association unresolved (Kato et al., 2015 [↗](#); Poort et al., 2015 [↗](#); Sachidhanandam et al., 2013 [↗](#)). It remains unclear whether sensory representations in cortex reorganize on the timescale at which behavioral changes first emerge or whether cortical plasticity reflects a slower process requiring days of offline consolidation (Drieu et al., 2025 [↗](#); McClelland et al., 1995 [↗](#)).

A candidate mechanism for rapid representational change is the spontaneous reactivation of stimulus-evoked ensemble patterns during inter-trial periods. In sensory cortex, such reactivations have been shown to predict the direction of change of future representations during passive stimulus viewing, suggesting that they act as an upstream selection process rather than a faithful replay of past activity (Miller et al., 2014 [↗](#); Nguyen et al., 2024 [↗](#)). Moreover, reactivations during learning have also been linked to bidirectional changes in network connectivity (Sugden et al., 2020 [↗](#)). However, whether reactivations operate on the rapid timescale of the formation of the initial association and whether reward gates which neurons they engage remain untested.

Here, we developed a whisker-based detection task in which mice form a novel sensorimotor association within a single behavioral session. Combining longitudinal two-photon calcium imaging of layer 2/3 neurons in the primary whisker-related somatosensory barrel cortex (wS1) (Woolsey and Van der Loos, 1970 [↗](#); Diamond et al., 2008 [↗](#); Petersen, 2019 [↗](#)) with trial-by-trial population decoding, we show that reward drives rapid reorganization of population representations tracking behavior, and that neurons recruited by learning are preferentially engaged in spontaneous reactivations, pointing to an online, reward-gated selection mechanism operating on the timescale of minutes.

Results

Mice rapidly learn to detect a novel tactile stimulus within a single session

To study rapid sensorimotor learning, head-fixed, water-restricted mice were trained on a whisker detection task designed to capture the formation of a novel sensorimotor association in the course of a single behavioral session (Figure 1A [↗](#)). Mice first went through an initial auditory pre-training phase during which they were rewarded with a drop of water for licking a spout in response to a brief auditory tone. Auditory pre-training continued for 5–7 days until stable expert performance was reached for at least two days (Days -2 and -1). On Day 0, a brief deflection of the C2 whisker was introduced as a novel stimulus, and whisker training continued for two additional days (Days +1 and +2). To establish that behavioral and neuronal changes were driven by reward, mice were randomly assigned to one of two groups: rewarded mice (R+, $n = 19$ mice, during functional imaging, described later) for which licking in response to whisker stimulation was rewarded, and non-rewarded mice (R-, $n = 16$ mice, during functional imaging, described later) for which licking in response to whisker stimulation was not rewarded. Auditory trials were interleaved throughout training and remained rewarded in both groups, providing an internal control for task engagement. Catch trials without sensory stimulation were pseudo-randomly interleaved to assess spontaneous lick rates and licking in catch trials was never rewarded.

Whisker-rewarded mice, but not whisker-unrewarded mice, learned to lick in response to whisker stimulation within a single session (Figures 1B [↗](#) and 1C [↗](#)), with the first whisker hit occurring on average within 5 whisker trials (Figure 1 – figure supplement 1 [↗](#)). The Day 0 whisker performance was significantly higher in R+ than R- mice (Mann-Whitney U test, $p = 2 \times 10^{-5}$) and continued to diverge over the following two days (Figure 1C [↗](#)). Auditory performance remained stable across days in both groups, ruling out general task disengagement as an explanation for the divergent whisker learning trajectories. Within Day 0, whisker hit rate progressively increased in R+ mice, while R- mice rapidly learned to suppress licking to the non-rewarded whisker stimulus, with the two groups significantly diverging after 22 whisker trials, 14 ± 4 min into the session (Mann-Whitney U test with Benjamini-Hochberg correction) (Figure 1D [↗](#) and Figure 1 – figure supplement 1 [↗](#)). Together, these results demonstrate that mice can rapidly acquire a novel sensorimotor association driven by reward within a single behavioral session, providing a tractable framework for investigating the neural mechanisms underlying fast goal-directed learning.




Figure 1. Mice can learn to lick for reward in response to a novel tactile whisker stimulus in a single session.

(A) Water-restricted mice were first trained to lick a water reward spout in response to a brief auditory pure tone (Days -2 and -1 correspond to the final two days of the auditory pre-training phase). On Day 0, the whisker stimulus (brief deflection of the C2 whisker) was introduced for the first time as an interleaved active (rewarded / unrewarded) trial type, and whisker training continued for two additional days (Day +1 and Day +2). Rewarded mice (R+, green) received a water reward if they licked in response to the whisker stimulus (whisker hit trials), whereas non-rewarded mice (R-, magenta) did not. Auditory hits (blue) were rewarded for both groups, and catch trials (black, no stimulus) were always unrewarded. At the end of each active behavioral session, a block of 50 consecutive passive whisker trials was presented. (B) Example probability of licking in blocks of 20 trials for one whisker rewarded (R+) and one whisker non-rewarded (R-) mouse across the five training days. Blue data points show licking probability on auditory trials; green (R+) and magenta (R-) data points show licking probability on whisker trials; and black data points show licking probability in no-stimulus trials. (C) Behavioral performance across days (n = 19 R+ mice and n = 16 R- mice, during functional imaging described later; dark blue, auditory lick probability in R+ mice; light blue, auditory lick probability in R- mice; green, whisker lick probability in R+ mice; magenta, whisker lick probability in R- mice; black, catch trial lick probability in R+ mice; grey, catch trial lick probability in R- mice). Thin lines indicate single-mouse trajectories and thick lines indicate averages across reward groups (left, mean ± 95% confidence interval). Bar plot quantifying whisker performance across days, comparing the two reward groups (right, mean ± 95% confidence interval; p-values from Mann-Whitney U test). (D) Whisker performance dynamics during Day 0. The top shaded bar indicates the p-value obtained from Mann-Whitney U tests comparing reward groups at each whisker trial, corrected for multiple testing with the Benjamini-Hochberg procedure. Data points show the mean Bayesian lick probability per trial ± 95% confidence interval. Divergence between the two groups first became significant after 22 whisker trials.

Pharmacological and optogenetic inactivation of barrel cortex prevents learning

Having established that mice rapidly acquire the whisker association within a single session, we next asked whether barrel cortex activity causally contributes to learning. We therefore carried out inactivation experiments in two different cohorts of whisker-rewarded mice, which were not functionally-imaged.

We first carried out pharmacological inactivation experiments. Muscimol, a GABA-A receptor agonist, was injected into the C2 column of the contralateral barrel cortex (wS1) or, as a control, into the forepaw primary somatosensory cortex (fpS1, located ~1.5 mm anterior to wS1) (Figure 2A). Inactivation was performed over the first three days of whisker-rewarded learning (Days 0, +1, +2), followed by three recovery days without injection (Days +3, +4, +5). Learning was abolished during muscimol inactivation in wS1, but mice acquired the task in the subsequent recovery days ($n = 13$ mice) (Figure 2B). In contrast, fpS1 inactivation did not prevent learning ($n = 8$ mice) (Figure 2B). Whisker-evoked licking probabilities on Day 0 were significantly higher in mice with fpS1 inactivation compared to mice with wS1 inactivation (Mann-Whitney U test, $p = 0.012$, Day 0) (Figure 2C).

To further test these findings with a spatiotemporally more precise approach, optogenetic inactivation was performed on Day 0 in VGAT-ChR2 mice (Zhao et al., 2011), in which blue light illumination of the cortical surface activates ChR2-expressing GABAergic interneurons to locally silence excitatory neurons (Guo et al., 2014) (Figure 2D). Optogenetic inactivation was carried out in all auditory, whisker and catch trials, beginning 0.1 s before stimulus onset and lasting throughout the 1-s reporting period. Licking in auditory and whisker trials was rewarded, but not in catch trials. Consistent with the pharmacological results, optogenetic suppression of wS1 activity on Day 0 prevented learning, whereas mice learned to lick in response to whisker stimulation when fpS1 was optogenetically inactivated ($n = 6$ R+ wS1-inactivated mice and $n = 6$ R+ fpS1-inactivated mice; Mann-Whitney U test, $p = 2 \times 10^{-3}$, Day 0) (Figures 2E and 2F). Together, these results demonstrate that barrel cortex activity contributes to the rapid acquisition of the association.

Learning induces a rapid reorganization of barrel cortex population activity

To examine how barrel cortex representations evolve across learning, we performed longitudinal two-photon calcium imaging of the same layer 2/3 neurons expressing GCaMP6f from Day -2 to Day +2 (R+ mice: 3,210 neurons across 19 mice, 169 ± 49 mean neurons per mouse $\pm$ SD; R- mice: 2,846 neurons across 16 mice, 177 ± 48 mean neurons per mouse $\pm$ SD) (Figure 3 and Figure 3 – figure supplement 1). At the end of each session, including auditory pre-training days (i.e. Days -2 and -1), a block of 50 passive whisker stimulations was delivered after task disengagement, to map sensory responses independently from reward and motor output. A striking divergence in the grand-average response emerged where the whisker-evoked responses progressively increased across days in R+ mice but decreased in R- mice (Figures 3A-C). Comparing response amplitudes between pre-learning (Days -2 and -1) and post-learning (Days +1 and +2) periods confirmed a significant enhancement of whisker-evoked responses in R+ mice ($n = 19$ mice, Wilcoxon signed-rank test, $p = 2 \times 10^{-3}$) (Figure 3D) and a significant suppression of whisker-evoked responses in R- mice ($n = 16$ mice, Wilcoxon signed-rank test, $p = 1.6 \times 10^{-3}$) (Figure 3E), revealing a reward-dependent bidirectional reshaping of the C2 whisker representation.




Figure 2. Pharmacological and optogenetic inactivation of barrel cortex prevent learning.

(A) Schematic illustrating muscimol inactivation of the barrel cortex (wS1) or the forepaw primary somatosensory cortex (fpS1). (B) Behavioral performance across days during wS1 muscimol inactivation (left; n = 13 R+ mice) and fpS1 muscimol inactivation (right; n = 8 R+ mice). Inactivation days, Days 0, +1, +2, are indicated in red; recovery on Days +3, +4, +5. (C) Statistical comparison of whisker performance during muscimol inactivation for wS1 and fpS1 (Mann-Whitney U test). (D) Schematic illustrating spatiotemporally-specific inactivation through optogenetic activation of local inhibitory neurons. (E) Behavioral performance across days during optogenetic inactivation of wS1 (left; n = 6 R+ mice) and fpS1 (right; n = 6 R+ mice) on Day 0 and recovery on Day +1. (F) Statistical comparison of whisker performance during optogenetic inactivation for wS1 and fpS1 (Mann-Whitney U test).




Figure 3. Learning induces a rapid reorganization of barrel cortex population activity.

(A) Two-photon calcium imaging of layer 2/3 neurons expressing GCaMP6f in the barrel cortex of Rasgrf2-Cre $\times$ Ai148 mice. Top: example field of view with individual cell ROIs color-coded by their learning modulation index (LMI: red, positive; blue, negative). Bottom: calcium traces in response to passive whisker trials for example cells. Orange vertical bars indicate whisker stimulus onset. (B) Heatmap of the trial-averaged whisker-evoked responses in the post-behavior passive epoch across the five imaging days (Day -2 to Day +2) for cells significantly modulated by learning (same example mouse as panel A) ordered by their LMI. (C) Population PSTHs over passive whisker trials across the five imaging days for R+ (green) and R- (magenta) mice. Thick lines show grand averages across mice ($n =$

19 R+ mice; $n = 16$ R- mice) and shading indicates 95% confidence interval. **(D)** Comparison of population whisker response amplitude ($\Delta F/F_0$, averaged over 0-300 ms from stimulus onset) before (Days -2 and -1; pre) and after (Days +1 and +2; post) learning in R+ mice ($n = 19$ mice). Left: population PSTHs for pre- (grey) and post- (green) learning. Right: bar plot quantification of individual mouse averages over the same time window (Wilcoxon signed-rank test). **(E)** As in panel D, but for R- mice ($n = 16$ mice). **(F)** Distribution of LMI across all imaged neurons for R+ mice (green; $n = 3,210$ neurons) and R- mice (magenta; $n = 2,846$ neurons) (Kolmogorov-Smirnov test). **(G)** Proportion of neurons with a significant positive LMI (left) and significant negative LMI (right), compared between R+ and R- mice (Mann-Whitney U test). **(H)** Trial-by-trial population similarity matrices averaged across R+ (left) and R- (right) mice. Each entry represents the mean cosine similarity between neural population activity vectors from the indicated pair of days, computed using responses to passive whisker trials averaged over 0-300 ms from stimulus onset. **(I)** Average cosine similarity between trials from the same day compared between R+ and R- mice (Mann-Whitney U test). Thin lines indicate individual mice, thick lines average $\pm$ 95% confidence interval. **(J)** Network reorganization index quantifying the difference in cosine similarity between pre-learning (Days -2 and -1) and post-learning (Days +1 and +2) population representations for R+ and R- mice. Points indicate individual mice, bar graph shows mean $\pm$ 95% confidence interval. **(K)** Schematic of the linear decoding approach. A logistic regression classifier was trained to discriminate pre-learning (Days -2, -1) from post-learning (Days +1, +2) passive whisker trial population activity. **(L)** Pre- vs post-learning decoding accuracy for R+ and R- mice, with 10-fold stratified cross-validation. Dashed line indicates chance level (50%). Points indicate individual mice, bar graph shows mean $\pm$ 95% confidence interval (Mann-Whitney U test). **(M)** Pairwise day decoding accuracy matrices using the fixed pre- vs post-learning decoder, for R+ (left) and R- (right) mice. Each entry shows the mean decoding accuracy when the classifier is applied to classify trials from the indicated pair of days. **(N)** Decoding accuracy comparing pre-learning days (Days -2, -1) vs Day 0 (Pre) and post-learning days (Days +1, +2) vs Day 0 (Post) for R+ and R- mice. Dashed line indicates chance level (50%). Points indicate individual mice, bar graph shows mean $\pm$ 95% confidence interval (Wilcoxon signed-rank test).

We then identified neurons responsible for this population effect by computing a learning modulation index (LMI) for each cell using ROC analysis between pre-learning (Days -2 and -1) and post-learning days (Days +1 and +2). Rather than driven by few outliers, the entire LMI distribution was shifted toward positive values in R+ mice relative to R- mice (two-sample Kolmogorov-Smirnov test, $p = 4 \times 10^{-54}$) (Figure 3F [↗](#)), and a significantly larger fraction of neurons acquired a positive LMI in R+ mice across learning ($n = 19$ mice, Mann-Whitney U test, $p = 8 \times 10^{-6}$) (Figure 3G [↗](#)), while a larger fraction acquired a negative LMI in R- mice ($n = 16$ mice, Mann-Whitney U test, $p = 0.02$) (Figure 3G [↗](#)). These modulations were heterogeneous at the single-cell level spanning recruitment, enhancement of pre-existing responses, and suppression (Figure 3 – figure supplement 1B [↗](#)). In a subset of mice, we retrogradely labelled wS1 neurons projecting to wM1 or wS2, finding that whisker-evoked responses were enhanced in wS2-projecting neurons of R+ mice but decreased in R- mice, whereas wM1-projecting neurons showed no overall significant change (Figure 3 – figure supplement 2 [↗](#)).

We next examined how the population representation changed upon learning. We computed trial-by-trial cosine similarity matrices between population activity vectors across all pairs of passive whisker stimulations (Figure 3H [↗](#)). In R+ mice, the within-day similarity increased, reflecting a progressive stabilization of the representation, while it tended to decrease in R- mice (Mann-Whitney U test R+ vs R-mice, $p = 0.08$ Day 0, $p = 3 \times 10^{-4}$ Day +1 and $p = 6 \times 10^{-4}$ Day +2) (Figure 3I [↗](#)). Across days, R- mice showed a uniform pattern with low similarity, whereas for R+ mice, the average cosine similarity matrix increased with learning. We quantified a network reorganization index, computed as the difference between within-period and between-period cosine similarities, which was significantly larger in R+ than R- mice (Mann-Whitney U test, $p = 1.2 \times 10^{-3}$) (Figure 3J [↗](#)). To test whether this reorganization carried information at the single trial level, we trained a logistic regression classifier to discriminate pre-learning (Days -2 and -1) from post-learning (Days +1 and +2) population activity patterns (Figure 3K [↗](#)). Cross-validated decoding accuracy was significantly higher in R+ than R- mice (Mann-Whitney U test, $p = 1.6 \times 10^{-3}$) (Figure 3L [↗](#)). Critically, this decoding approach allowed us to ask whether the transition had already begun on Day 0 by testing the unseen trials from that day with the fixed weights of the pre- versus post-learning decoder (Figure 3M [↗](#)). In both groups, Day 0 could be significantly distinguished from pre-

learning and from post-learning days (Wilcoxon signed-rank test, $p = 2 \times 10^{-4}$ for R+ Day 0 vs pre-learning, $p = 1 \times 10^{-4}$ for R+ Day 0 vs post-learning, $p = 3 \times 10^{-5}$ for R- Day 0 vs pre-learning, $p = 8 \times 10^{-4}$ for R- Day 0 vs post-learning). Moreover, decoding tended to be more accurate at distinguishing Day 0 from pre-learning than from post-learning (Wilcoxon signed-rank test, $p = 0.2$ for R+, $p = 0.016$ for R-; [Figure 3N](#)). Together, these results establish that reward-driven learning is rapidly accompanied by a new representational state that transitions in the first whisker learning session.

Neuronal representation shift correlates with learning and is accompanied by online reactivations

Having established that population representations are already partially shifting on Day 0, we next asked whether this transition could be tracked online trial-by-trial during active behavior. To do so, we projected whisker-evoked population activity from Day 0 onto the decoding axis of the pre- versus post-learning decoder, trained exclusively on passive whisker trials, excluding Day 0, yielding a continuous readout of the shift from naïve to expert representational states ([Figure 4A](#)). R+ mice exhibited a progressive increase in both lick probability and projection on this learning dimension across Day 0 trials, whereas R- mice showed decreasing lick probability and a flat projection trajectory ([Figures 4B](#) and [4C](#)). The linear slope of the projection on the learning dimension over Day 0 trials was significantly positive in R+ mice (Wilcoxon one-sample test against zero, $p = 0.03$) but not in R- mice (Wilcoxon one-sample test, $p = 0.33$) ([Figure 4D](#)). Furthermore, there was significant correlation between the trajectory on the learning dimension and the behavioral learning curve in R+ mice (Wilcoxon one-sample test, $p = 0.02$) but not in R- mice (Wilcoxon one-sample test, $p = 0.9$) ([Figure 4E](#)). These results demonstrate that barrel cortex population activity progressively realigns from naïve to expert-like representations during learning, tracking behavioral performance trial-by-trial.

Spontaneous cortical reactivations of stimulus-evoked ensemble patterns have been shown to occur during inter-trial periods and to predict future changes in sensory representations rather than merely replaying past activity ([Miller et al., 2014](#); [Nguyen et al., 2024](#); [Sugden et al., 2020](#)). Such reactivations could therefore play a role in selecting which neurons undergo experience-dependent plasticity across learning. To address this, we examined spontaneous reactivations of the whisker-evoked activity pattern during catch trials ([Figures 4F](#) and [4G](#)). Comparing the reactivation rates of R+ and R- mice, we found that reward triggered significantly higher reactivation rates across days (Mann–Whitney U test, $p = 0.014$, Day 0) ([Figure 4H](#)). We next computed how much each cell took part in these reactivations to test whether learning-modulated neurons preferentially contributed. In R+ mice, a significant positive correlation was observed between the learning modulation index (LMI) and the participation rate during Day 0, showing that neurons gaining stimulus responsiveness through learning were more likely to participate in reactivations, whereas in R- mice the correlation was weakly negative (R+: $r = 0.239$, $p = 1 \times 10^{-41}$; R-: $r = -0.067$, $p = 3 \times 10^{-4}$) ([Figure 4I](#)). Across training days, and already on Day 0, significant LMI-positive neurons in R+ mice showed a progressive increase in participation rate (Kruskal–Wallis test, $p = 5 \times 10^{-4}$), while significant LMI-negative neurons did not (Kruskal–Wallis test, $p = 0.05$) ([Figure 4J](#)). Neither LMI-positive nor LMI-negative neurons showed significant changes in participation across days in R- mice (Kruskal–Wallis tests, $p = 0.13$ and $p = 0.2$). This relationship still held after controlling for baseline spontaneous activity levels ([Figure 4 – figure supplement 1](#)). Together, these results indicate that neurons whose sensory responses were strengthened by reward-driven learning were preferentially engaged in spontaneous reactivation events, consistent with online reactivations potentially acting as a mechanism for inducing learning-related plasticity.

Discussion

Here, we developed a task in which mice formed a novel sensorimotor association within a single behavioral session ([Figure 1](#)). Pharmacological and optogenetic inactivation suggested a causal role for sensory cortex in task learning ([Figure 2](#)). Neuronal activity in sensory cortex exhibited a rapid reorganization of sensory representations ([Figure 3](#)) concomitant with single session




Figure 4. Progressive realignment of barrel cortex representations and online reactivations.

(A) Schematic of the Day 0 decoding approach. The pre- vs post-learning classifier (trained on passive whisker trials from Days -2 and -1 vs Days +1 and +2) was applied to active whisker trials during Day 0 to compute a trial-by-trial projection on the learning axis. (B) Example Day 0 trajectories for one R+ mouse (left, green) and one R- mouse (right, magenta). Top: whisker hit probability, estimated by a Bayesian state-space model fitted to binary hit/miss outcomes. Bottom: projection on learning axis applied to the same whisker trials. (C) Average Day 0 trajectories across all mice for R+ (left, green) and R- (right, magenta) groups. Top: licking performance across whisker trials. Bottom: projection on the learning axis across whisker trials. Thick lines report averages across mice with shading indicating the 95% confidence interval. (D) Linear slope of the projection on the learning axis over Day 0 whisker trials, between R+ and R- mice. Small data points show individual mice, and big points show mean $\pm$ 95% confidence interval (Wilcoxon one-sample test). (E) Pearson correlation between the projection on the learning axis and whisker performance across Day 0, between R+ and R- mice. Small data points show individual mice, and big points show mean $\pm$ 95% confidence interval (Wilcoxon one-sample test). (F) Example reactivation heatmap for a R+ mouse on Day 0. Left strips: Learning modulation index (LMI); average response to passive whisker trials (reactivation template); and reactivation participation rate, for the top 20 cells ranked by participation rate. Right: neural activity heatmap of 180 seconds of concatenated catch trial data. Top: template correlation trace with detection threshold of 0.25 (horizontal dashed line); red vertical lines indicate detected reactivation events. (G) Example template correlation traces across all five training days for the same R+ mouse as in panel F. Red vertical lines indicate detected reactivation events and grey dotted lines indicate 0 correlation. (H) Reactivation frequency across training days for R+ (green) and R- (magenta) mice. Bar graph shows mean $\pm$ 95% confidence interval (Mann-Whitney U test). (I) Scatter plots of Day 0 reactivation participation rate versus LMI for individual neurons, shown separately for R+ (left, green) and R- (right, magenta) mice. Each point represents one neuron. Linear regression line is overlaid (Pearson correlation). (J) Participation rate across training days for LMI-positive (red) and LMI-negative (blue) neurons, shown separately for R+ (left) and R- (right) mice. Bar graph shows mean $\pm$ 95% confidence interval (Kruskal-Wallis test).

task learning (Figures 4A-D). Finally, we found that neurons whose stimulus responsiveness was strengthened by learning were preferentially engaged in spontaneous reactivation events, suggesting that online reactivation might drive learning-induced plasticity (Figures 4F-J).

By designing a task that could be learned in a single session, we could track the representational transition as it unfolded trial-by-trial. This goes beyond prior work correlating neural selectivity with performance across days (Banerjee et al., 2020; Chéreau et al., 2020; Goltstein et al., 2013; Henschke et al., 2020; Kato et al., 2015; Makino and Komiyama, 2015; Peron et al., 2015; Poort et al., 2015; Sachidhanandam et al., 2013) and shows that representational changes are concurrent with behavioral acquisition, rather than a delayed consequence. Importantly, the non-rewarded control group rules out passive representational drift driven by repeated stimulus exposure as an explanation (Ahmed et al., 2024; Bauer et al., 2024; Marks and Goard, 2021; Wang et al., 2022). A recent study in the auditory cortex also found that sensory cortex drives rapid acquisition of task contingencies within tens of trials and becomes dispensable at expert performance (Drieu et al., 2025). However, learning in that study did not enhance stimulus representations, which instead habituated as in passively exposed animals. In contrast, our passive whisker trials revealed that stimulus representations in barrel cortex are restructured by reward-driven learning. This divergence likely reflects a difference in sensory demand: unlike the pure tones easily discriminable at the perceptual level used by Drieu et al. (Drieu et al., 2025), near perceptual threshold whisker deflection might require the enhancement of cortical representations. Interestingly, this reorganization was heterogeneous across projection subtypes: neurons projecting to secondary somatosensory cortex showed large learning-related changes while those projecting to motor cortex were unaffected (Figure 3 – figure supplement 2), consistent with the preferential routing of task-relevant information through the wS1-to-wS2 pathway reported previously in expert mice performing closely-related whisker detection tasks (Chen et al., 2015; Vaviladi et al., 2020; Yamashita and Petersen, 2016).

Strikingly, neurons gaining stimulus responsiveness through learning were preferentially engaged in spontaneous reactivations already on Day 0 (Figures 4I-J). Recent work established that cortical reactivations during passive viewing predict the direction of future plastic changes (Nguyen et al., 2024), suggesting an upstream selection mechanism, and that reactivations during reward-based learning are linked to bidirectional network changes (Lensjø et al., 2025; Sugden et al., 2020). Two aspects of our findings elaborate on this framework. First, reward gated both the rate and cellular identity of reactivations: R+ mice showed higher reactivation frequencies than R- mice (Figure 4H), and a significant positive correlation between LMI and participation rate was only found in R+ mice (Figures 4I and Figure 4 – figure supplement 1). Second, the process operated on the timescale of minutes, concurrent with the initial formation of the association, rather than during post-session offline consolidation as described in prior work. Together, these observations are consistent with online reactivations acting as a reward-gated selection mechanism that biases which neurons undergo experience-dependent plasticity.

One limitation is that our current data do not establish a causal role for reactivations. Closed-loop optogenetic experiments disrupting reactivation events during learning could be used to test necessity, which we predict would impair representational enhancement as well as behavioral performance during Day 0. Our study is also restricted to layer 2/3 excitatory neurons in a single cortical area. Deeper layer neurons (Lacefield et al., 2019; Oryshchuk et al., 2024; Takahashi et al., 2020), inhibitory neurons (Deister et al., 2024; Park et al., 2025; Ramamurthy et al., 2023; Sachidhanandam et al., 2016) and neuromodulators (Constantinople and Bruno, 2011; Muñoz et al., 2017; Gasselín et al., 2021) likely play important roles in shaping the network reorganization we observed. Finally, extending this approach to downstream regions such as the secondary somatosensory cortex (Han and Helmchen, 2025; Kwon et al., 2016) and the striatum (Sippy et al., 2021, 2015) would help clarify how rapid cortical reorganization is propagated to drive behavioral change.

Methods

Animals

All procedures were approved by the Swiss Federal Veterinary Office (License numbers VD-1628 and VD-3753) and were conducted in accordance with the Swiss guidelines for the use of research animals. We used both male and female adult C57BL/6 mice of at least 5 weeks at the time of surgery. For calcium imaging experiments we used Rasgrf2-dCre mice [B6;129S-Rasgrf2<tm1(cre/foIA)Hze>/J, JAX: 022864] (Harris et al., 2014 [↗](#)) crossed with Ai148 mice [B6;Ai148(TIT2L-GC6f-ICL-tTA2)-D, JAX: 030328] (Daigle et al., 2018 [↗](#)), which express the green fluorescent calcium indicator GCaMP6f in layer 2/3 neurons. For optogenetic stimulation of inhibitory neurons, we used VGAT-ChR2 mice [B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP)8Gfng/J, JAX: 014548] (Zhao et al., 2011 [↗](#)). For pharmacological inactivation experiments we used wild-type mice.

Head plate surgery and intrinsic optical imaging

Mice were first implanted with a metal head-post under anesthesia (1.5-2.5 % isoflurane). Before the start of surgery, carprofen was injected (7.5 mg/kg at 1.5 mg/ml, subcutaneous) for general analgesia, and a mix of lidocaine and bupivacaine below the scalp as local analgesic. An ocular ointment (Vita-Pos, Pharma Medica AG, Switzerland) applied over the eyes prevented drying during the surgery. A solution of povidone-iodine (Betadine, Mundipharma Medical Company, Bermuda) was applied for skin disinfection before incision. We monitored body temperature throughout the surgery and it was maintained at 37°C with a heating pad (ThermoStar, Intellibio, France). After excising the scalp with surgical scissors, the exposed skull was cleaned with cotton buds and a scalpel blade to remove the periosteal tissue. After further disinfection with Betadine, the skull was dried with cotton buds and made optically clear using a thin layer of super glue applied over the dorsal part of the skull (Loctite super glue 401, Henkel, Germany). We then glued a custom-made head fixation plate above the right hemisphere, and further secured it with self-curing denture acrylic (Paladur, Kulzer, Germany or Ortho-Jet, Lang, USA). Finally, a second thin layer of the glue was applied homogeneously on the left hemisphere to ensure a smooth and equally transparent field of view. This transparent skull preparation was used to perform intrinsic optical signal (IOS) imaging. Mice were returned to their home cages and ibuprofen (Algifor Dolo Junior, Verfora SA, Switzerland) was added to the drinking water for three days after surgery.

Following head plate implantation, IOS imaging was performed on the left hemisphere in order to localize the right C2 whisker representation in wS1 and wS2, as well as the location of fpS1. Animals were lightly anesthetized (isoflurane at 1%) and head-fixed under a CMOS camera coupled to a stereomicroscope (Leica MZ9.5). We acquired a reference image of the surface vasculature under green light (525 nm, Thorlabs LED). Subsequently, under red light (630 nm, Thorlabs LED), the right C2 whisker or the right paw was stimulated by a capillary glass glued to a piezoelectric actuator (PICMA, PI Ceramic). Each trial consisted of 4 s of rectangular deflections at a 10 Hz frequency. We interleaved stimulation with no-stimulation trials and the average activation map was computed using a custom MATLAB-based software. The resulting map was then overlaid on the anatomical image to locate the region of interest relative to surrounding blood vessels.

Cranial window surgery for two-photon imaging and retrograde labelling

For mice used for two-photon calcium imaging, a circular craniotomy with a 4 mm diameter was performed over the C2 column of the primary whisker somatosensory cortex wS1 identified with intrinsic optical signal imaging. After the bone had been removed, the surface of the cortex was rinsed with Ringer solution and the dura was removed using a 27G sharp needle bent as a hook. To retrogradely label wS1 neurons projecting to wS2 and wM1, cholera toxin subunit-B (CTB) conjugated with two different Alexa-Fluor dyes (Molecular Probes, Invitrogen; 100 nL, 0.5%, wt/vol) were injected in these areas. A thin glass pipette (PCR Micropipettes 1–10 mL, Drummond

Scientific Company, USA) was first pulled and then the tip was broken using a tissue to obtain a 21–27 μm inner tip diameter. The pipette was filled with mineral oil and then tip-filled with the CTB solution. Two injections in the cortex at depth 300 and 500 μm below the pia were performed of 50 nL at each depth for wS2 and 100 nL for wM1 (at stereotaxic coordinates 1 mm anterior and 1 mm lateral from bregma). Injections were performed using a single-axis oil hydraulic micromanipulator (Narishige, Japan) at 20 nL per minute with 5 min of waiting time after injection. A 4 mm glass window was then fixed over wS1 using super glue and self-curing denture acrylic.

Behavioral task and training

We trained head-fixed, water-restricted mice to perform a sensory detection task. During the behavioral experiments, all right whiskers were trimmed except for the C2 whisker, and the mice were water restricted to 0.9–1.2 mL of water/day (depending on initial body weight). Mice were trained daily with one session per day; their weight and general health status were carefully monitored using score sheets. Mice were habituated to be head-fixed to a metal post with their head held at a 30-degree angle. During the first session, mice were also habituated to licking from a water spout positioned on their right side which delivered 5 μL of water in response to some licks. Next, we trained them in an auditory detection task for 5–7 days during which licking in a 1-s response window after the auditory cue (10 ms, 10 kHz tone of 74 dB embedded in the continuous background white noise of 80 dB) was rewarded. Inter-trial interval was uniformly sampled between 6 s and 10 s including a final 2 s to 5 s no-lick window during which mice were required not to lick for the trial to be initiated. When two consecutive days of stable expert performance (hit rate $\leq 80\%$ and false alarm rate $\geq 30\%$) in the auditory detection task was reached (Days -2 and -1), whisker detection training was introduced for the next three days (Days 0, +1 and +2). Whisker deflection was achieved with a 30 mT magnetic pulse driven by a 3 ms cosine impulse sent to an electromagnetic coil positioned under the mouse and acting on a small iron particle placed on the C2 whisker. Trials in auditory and whisker training days were presented pseudo-randomly in blocks of 20 trials. Blocks consisted of 10 auditory and 10 catch (no stimulus) trials on auditory training days and 7 whisker, 3 auditory and 10 catch on whisker training days. If mice licked the reward spout within the reward time window of 1 s from stimulus onset, it was considered a hit trial, and they obtained a 5 μL drop of water. If they did not lick within the reward window, it was considered a miss trial, and no reward was delivered. Each mouse was randomly assigned to the rewarded (R+) or non-rewarded (R-) group at the start of the training protocol. While auditory trials were rewarded for both groups throughout training, only R+ mice obtained a water reward in whisker trials, whereas R- mice did not obtain a water reward if they licked in response to whisker stimulation. Catch trials were never rewarded. In every two-photon imaging session (Days -2 to +2), we additionally presented 50 passive whisker stimulations at the end of the session after the mouse was disengaged from the task (i.e. stopped licking in all trial types). No licking occurred during passive stimulation.

Optogenetic inactivation

Optogenetic inactivation of cortical excitatory neurons was achieved indirectly by activating ChR2-expressing GABAergic interneurons in VGAT-ChR2 mice (Zhao et al., 2011 [DOI](#)). Blue light (470 nm; MBL-F473/200 mW, Changchun New Industries Optoelectronics Technology, China) was delivered through a 200 μm , 0.22 NA patch cable (M84L01, Thorlabs, USA) fitted with a ceramic ferrule cannula (CFMXA05, Thorlabs, USA). The target area, wS1 or fpS1, was identified by IOS imaging and the skull was thinned above the site to facilitate light penetration. The remainder of the exposed skull was covered with Kwik-Cast to prevent off-target illumination of adjacent cortical areas. A blue masking light (M490L4, Thorlabs, USA) was placed in the animal's visual field throughout inactivation sessions to prevent the light stimulus from serving as a visual cue. Light stimulation consisted of a 100 Hz pulse train with a 50% duty cycle, delivered on 100% of trials (whisker, auditory and catch). Stimulation onset was 100 ms before stimulus presentation and lasted 1 s to cover the full reward window, followed by a 100 ms linear ramp-down to prevent

rebound excitation. Light power was measured at 10 mW at the fiber tip. Inactivation was performed on Day 0, during initial whisker learning, and mice were trained for one additional session without light stimulation on Day +1.

Pharmacological inactivation

Cortical inactivation was achieved by injection of the GABA-A receptor agonist muscimol (5 mM in Ringer solution; Biotrend, USA). A small craniotomy was drilled above the C2 barrel column of wS1 or above fpS1 identified by IOS imaging. Muscimol was injected through a glass pipette of 21–27 μ m in diameter at a rate of 20 nL/min, to a total volume of 400 nL distributed across four depths ranging from 200 to 800 μ m below the pia. Injections were performed 45–60 min before the start of the behavioral session. Inactivation was performed during the three days of whisker training (Days 0, +1, +2); mice were then trained for three additional recovery days without injection (Days +3, +4, +5).

Two-photon imaging

A custom made two-photon microscope was used to perform calcium imaging experiments described in this paper. The microscope was equipped with a galvo-resonant mirror pair (8 kHz CRS, Cambridge Technology, USA), allowing a frame rate of 30 Hz for a resolution of 512×512 pixels. A femtosecond tunable infrared laser (Mai Tai or InSight DeepSee, Spectra Physics—Newport, USA) was fed into the light path at a wavelength of 940 nm to excite the genetically-encoded calcium indicator GCaMP6f. Light emission was detected with a GaAsP photosensor module (H10770PA-40, Hamamatsu, Japan), and signal acquisition was performed with National Instruments hardware (NI PXIe-1073, NI PXIe-6341, National Instruments, USA). The microscope head was movable and controlled in three dimensions by motors (Luigs and Neumann, Germany). A 16x immersion objective (16x Nikon CFI LWD, Japan) was used for all the imaging. The system was operated by the Matlab-based software ScanImage (Pologruto et al., 2003) (Vidrio Technologies, USA). For each mouse, the same field of view was imaged for five consecutive days (Days -2 to +2).

Quantification and data analysis

Behavioral performance

To obtain a continuous trial-by-trial estimate of behavioral performance rather than averaging binary hit/miss outcomes across mice, lick probability was estimated trial-by-trial using a Bayesian state-space model implemented with PyMC (Smith et al., 2004). For each session and stimulus type, binary hit/miss outcomes were modelled as Bernoulli observations with a latent lick probability that evolved as a Gaussian random walk in logit space. The logit-transformed lick probability at each trial t was modelled as $\text{logit}(p_t) = x_t$, where x_t followed a Gaussian random walk. Posterior inference was performed by MCMC sampling, the posterior mean of p_t was taken as the estimated lick probability at each trial, and the 80% credible interval was used to characterize uncertainty.

Two-photon calcium imaging preprocessing

To extract time-varying somatic calcium signals, we used the Suite2p toolbox (Pachitariu et al., 2017). Neuropil contamination was corrected by subtracting the fluorescent signal from a surrounding ring $F_{\text{neuropil}}(t)$ from somatic fluorescence: $F(t) = F_{\text{soma}}(t) - 0.7 \times F_{\text{neuropil}}(t)$. Neuropil-corrected fluorescence signals $F(t)$ were then normalized to the moment-to-moment $F_0(t)$ computed from a maximin filter (first lowpass filter at 1 Hz, then applying a rolling minimum followed by a rolling maximum with a rolling window of 2 min) to give $\Delta F/F_0(t)$. Labelling of the wS2- and wM1-projecting neurons were performed by registering the mean red and far red CTB images with the mean GCaMP6f image with the Python-based SimpleElastix toolbox (<https://simpleelastix.github.io/>). Every ROI was then classified by visual inspection of its average fluorescence into wS2-projection, wM1-projection or non-labelled with a custom-based Python GUI.

Learning modulation index

To quantify learning-induced changes in whisker-evoked responses at the single-cell level, a learning modulation index (LMI) was computed for each neuron using receiver operating characteristic (ROC) analysis. For each cell, the whisker-evoked response amplitude ($\Delta F/F_0$) was averaged over 0–300 ms from stimulus onset for each passive whisker trial. ROC analysis was then performed by treating pre-learning trials (Days -2 and -1 pooled) as class 0 and post-learning trials (Days +1 and +2 pooled) as class 1, with the response amplitude serving as the predictor variable. The area under the ROC curve (AUC) was computed and transformed to the LMI using $LMI = (AUC - 0.5) \times 2$. To assess the statistical significance of each LMI, a permutation test was performed by shuffling the class labels (pre- versus post-learning) 1,000 times and recomputing the AUC for each shuffle. The p-value for each cell was defined as the percentile of the observed AUC within the shuffled distribution. Cells with $p \geq 0.975$ were classified as having a significantly positive LMI, while cells with $p \leq 0.025$ were classified as having a significantly negative LMI.

Similarity matrices and network reorganization index

Trial-by-trial similarity matrices were computed for each mouse using passive whisker trials. For each session (Days -2 to +2), the last 40 passive whisker trials were selected and the whisker-evoked response amplitude was averaged over 0–300 ms from stimulus onset for each trial. Cosine similarity was computed between all pairs of trials across the five imaging days, resulting in a 200×200 similarity matrix (5 days $\times$ 40 trials per day) for each mouse. Matrix entries were computed as the normalized dot product between trial pairs: for population activity vectors x and y , $\text{similarity} = (x \cdot y) / (|x| \times |y|)$. A cosine similarity matrix was computed for each mouse and then averaged across mice.

To quantify the decorrelation between pre- and post-learning representations, a network reorganization index was defined for each mouse as the difference between the average within-period similarity and the average between-period similarity. Within-period similarity was calculated as the average similarity among all pre-learning trial pairs (Days -2 and -1) and the average similarity among all post-learning trial pairs (Days +1 and +2). Between-period similarity was calculated as the average similarity between all pre-learning and post-learning trial pairs. The network reorganization index was computed as: $[(\text{Within-pre-similarity} + \text{Within-post-similarity}) / 2] - \text{Between-pre-post-similarity}$.

Decoding of pre- and post-learning and projection on learning dimension

To test whether population activity patterns could discriminate pre- from post-learning representations, a logistic regression classifier was trained on passive whisker trials for each mouse. The whisker-evoked response amplitude ($\Delta F/F_0$) was averaged over 0–300 ms from stimulus onset for each trial. Pre-learning trials (Days -2 and -1 pooled) were labelled as class 0 while post-learning trials (Days +1 and +2 pooled) were labelled as class 1. Population activity vectors (neurons $\times$ trials) were transposed to form a feature matrix (trials $\times$ neurons). Features were z-scored across trials and a logistic regression classifier with L2 regularization was trained using 10-fold stratified cross-validation. Decoding accuracy was computed as the mean accuracy across folds.

To evaluate whether Day 0 activity patterns resembled pre- or post-learning representations, the trained pre- versus post-learning decoder was applied to all pairwise combinations of days (Days -2, -1, 0, +1, +2) using held-out test folds. For each mouse, data were split into 10 folds per day. At each fold iteration, a decoder was trained exclusively on pre-learning (Days -2, -1) versus post-learning (Days +1, +2) trials from the training folds, then tested on all pairwise day combinations using the corresponding test folds. For each day pair (day_i, day_j), trials from day_i were assigned to class 0 and trials from day_j were assigned to class 1, and classification accuracy was computed.

To track the progressive realignment of population representations during learning, whisker-evoked activity from active Day 0 trials was projected onto the learning axis defined by the pre- versus post-learning decoder. For each mouse, a fixed logistic regression decoder was trained on passive whisker trials from Days -2 and -1 (pooled) versus Days +1 and +2 (pooled), excluding Day

0. The decoder weights and scaler were saved and subsequently applied to population activity from Day 0 active whisker trials. For each Day 0 whisker trial, the response amplitude was averaged over 0–300 ms from stimulus onset, scaled using the fitted scaler, and projected onto the decision axis using the trained classifier's decision function. The resulting trial-by-trial projection values quantified the shift from naïve to expert-like representations during Day 0 learning.

Detection of reactivations and participation rate

To detect spontaneous reactivations of whisker-evoked activity patterns during inter-trial periods, a template-matching approach was applied to catch trial data (trials without auditory or whisker stimuli). For each mouse and day, a whisker response template was first constructed from passive whisker trials. The template was defined as the trial-averaged whisker-evoked population activity vector, computed by averaging the response amplitude of each neuron over 0–300 ms from stimulus onset across the last 40 passive trials of each session. This template captured the day-specific ensemble activity pattern evoked by the whisker stimulus.

For each day, neural activity during catch trials was then continuously compared to the corresponding template by computing the Pearson correlation between the template vector and the population activity vector at each imaging frame. The resulting correlation time series quantified the moment-by-moment similarity between spontaneous activity and the stimulus-evoked template. A Savitzky-Golay filter (window length = 5 frames, polynomial order = 2) was applied to the correlation time series to reduce jitter from frame-to-frame noise.

Reactivation events were identified as peaks in the correlation timeseries exceeding a threshold. To ensure that detected events were unlikely to arise from chance fluctuations, a detection threshold was determined for each mouse by randomly shifting activity in time independently for all cells: for each mouse, catch trial data were circularly shuffled 1,000 times along the time axis, the correlation timeseries was recomputed for each shuffle and the 99th percentile was computed for each timeseries. The detection threshold was set to the median of that percentile distribution. Peak detection was performed using `scipy.signal.find_peaks` with a minimum inter-peak distance of 150 ms and a minimum prominence of 0.15.

To quantify each neuron's engagement in reactivation events, a participation rate was computed for each cell. For each detected reactivation event, the average response amplitude of each neuron was calculated over a ± 150 ms window centered on the event time. A neuron was classified as participating in a given event if its average response within this window exceeded 10% $\Delta F/F_0$. The participation rate for each neuron was defined as the proportion of detected reactivation events in which the neuron participated, computed separately for each day.

To assess whether the relationship between LMI and reactivation participation was robust to baseline spontaneous activity, we computed a binary measure of significant participation for each neuron using a circular-shift control. For each cell, the catch trial data were circularly shuffled 1,000 times along the time axis. For each shuffle, the participation rate was computed, yielding a null distribution of 1,000 participation rate values per cell. A neuron was classified as significantly participating on a given day if its observed participation rate exceeded the 95th percentile of its corresponding null distribution. The proportion of significantly participating neurons was then computed separately for LMI-positive and LMI-negative cells for each mouse and day. The effect of day on the proportion of participating cells was assessed with a Kruskal–Wallis test, applied independently to each combination of reward group (R+, R-) and LMI category.

Statistics

We used non-parametric tests and corrected for multiple testing with the Benjamini-Hochberg procedure. Error bars indicate 95% confidence interval throughout.

Data availability

Imaging data, behavioral data, and calcium traces were combined into a single NWB file for each behavioral session. NWB offers a common format for sharing and analyzing neurophysiology data (Rübel et al., 2022 [DOI](#)). Subsequently, we developed open-source Python scripts to analyze data in

the NWB format. The full dataset will be available via Zenodo.

Code for data acquisition and behavior control is available on Github (https://github.com/LENS-BMI-EPFL/behavior_control). The code used to preprocess and convert the data into NWB format is available on Github (https://github.com/LENS-BMI-EPFL/NWB_converter). The code used for data analysis is available on Github (https://github.com/LENS-BMI-EPFL/renard_foustoukos). The Python code used for analyses will also be available via Zenodo.

Figure supplements




Figure 1 – figure supplement 1. Behavioral quantification of rapid whisker detection learning. (A) Trial number of the first whisker hit on Day 0 for R+ and R- mice. Points indicate individual mice. Bar graph shows mean $\pm$ 95% confidence interval (Mann-Whitney U test). (B) Whisker performance on Day 0. Points indicate individual

mice. Bar graph shows mean $\pm$ 95% confidence interval (Mann-Whitney U test). **(C)** Stimulation particle off control performed on Day +3 for a subset of 14 R+ mice where the particle is removed from the C2 whisker to control for potential stimulus artefacts. First particle-on block (ON₁), particle-off block (OFF), and second particle-on block (ON₂). Left: whisker hit rate across the three epochs. Middle: catch trial false alarm rate across the three epochs. Right: comparison of whisker hit rate and false alarm rate during the OFF epoch alone. Grey lines show individual mice. Bar graph shows mean $\pm$ 95% confidence interval (Mann-Whitney U test). **(D)** Single-trial whisker binary hit/miss outcomes averaged across mice during Day 0 for R+ (green) and R- (magenta) mice. Top grey-scale bar indicates statistical significance comparing R+ and R- for each day (Mann-Whitney U test). **(E)** As in panel D with lick probability estimated by a Bayesian state-space model fitted to binary hit/miss outcomes and with whisker trials realigned to the first whisker hit. **(F)** Lick probability for whisker, auditory, and no-stimulus trial types across time on Day 0, interpolated onto a common time grid. **(G)** Mean reaction time per stimulus type (auditory, whisker, catch) across training days for R+ and R- mice, computed from lick trials. Points indicate individual mice. Bar graphs show mean $\pm$ 95% confidence interval (Mann-Whitney U test). **(H)** Mean reaction time $\pm$ 95% confidence interval across lick trials within Day 0 for each stimulus type, for R+ and R- mice.




Figure 3 – figure supplement 1. Relationship between single-cell learning modulation and population decoding.

(A) Stability of the field of view across imaging days for an example mouse. Images are obtained by averaging the first 1,000 frames of the movement corrected data. (B) Responses to passive whisker stimulations of three example learning modulated cells during the five imaging days: a cell with a strong positive LMI that did not respond before learning (top), a cell with a strong positive LMI with an enhanced response (middle) and a cell with a strong negative LMI with an abolished response (bottom). Individual trials are shown in grey; the trial-averaged trace is shown in black. Orange line indicates whisker stimulus onset. (C) Scatter plot of logistic regression classifier weight versus LMI for all imaged neurons, color-coded by mouse. Black line indicates linear regression line. (D) Pre- vs post-learning decoding accuracy as a function of the proportion of most-modulated cells removed from the population. Cells are ranked by $|LMI|$ in descending order and progressively excluded; at each step, decoding accuracy is computed by 10-fold stratified cross-validated logistic regression.




Figure 3 – figure supplement 2. Pathway-specific reorganization of barrel cortex representations.

(A) Example two-photon calcium imaging field of view (top) and labelling of projecting neuron subtypes (bottom). Green: GCaMP6f-expressing neurons; yellow: CTB retrogradely labelled neurons projecting to the secondary somatosensory cortex (wS2p); blue: CTB retrogradely labelled neurons projecting to the primary motor cortex (wM1p). (B) Population PSTHs over passive whisker trials across the five imaging days for wS2p (top) and wM1p (bottom) neurons, shown separately for R+ mice (green) and R- mice (magenta) (wS2p: $n = 17$ R+ mice, $n = 17$ R- mice; wM1p: $n = 12$ R+ mice, $n = 11$ R- mice). Orange line indicates whisker stimulus onset. (C) Comparison of whisker-evoked response amplitude ($\Delta F/F_0$, averaged over 0–300 ms from stimulus onset) before (Days -2 and -1; Pre) and after (Days +1 and +2; Post) learning in R+ mice, for wS2p neurons ($n = 17$ mice). Left: population PSTHs for pre- (grey) and post- (green) learning. Right: bar plot quantification of individual mouse averages over the same time window. Points indicate individual mice. Bar graphs show mean $\pm$ 95% confidence interval (Wilcoxon signed-rank test). (D) As in panel C, but for R- mice ($n = 17$ mice). (E) Distribution of the LMI across wS2p neurons for R+ mice (green; $n = 328$ neurons from 17 mice with 19 ± 10 neurons per mouse) and R- mice (magenta; $n = 254$ neurons from 17 mice with 21 ± 12 neurons per mouse). (F) Proportion of wS2p neurons with a significant positive LMI (left) and significant negative LMI (right), compared between R+ and R-mice. Points indicate individual mice. Bar graphs show mean $\pm$ 95% confidence interval (Mann-Whitney U test). (G) As in panel C, but for wM1p neurons ($n = 12$ mice). (H) As in panel D, but for wM1p neurons ($n = 11$ mice). (I) As in panel E, but for wM1p neurons (green, R+ mice: $n = 304$ neurons from 12 mice with 18 ± 10 neurons per mouse; magenta, R- mice: $n = 182$ neurons from 11 mice with 17 ± 12 neurons per mouse). (J) As in panel F, but for wM1p neurons. (K) Contribution of projection subtypes to LMI values. Cumulative distribution of LMI values for wS2p (orange) and wM1p (blue) neurons, shown separately for positive and negative LMI populations and for R+ and R-groups. (L) As in panel K, but for logistic regression classifier weights. (M) Pairwise Pearson correlation between wS2p neuron pairs (wS2p-wS2p) and wM1p neuron pairs (wM1p-wM1p) during the 2-second pre-stimulus quiet window, compared between pre- (Days -2, -1) and post- (Days +1, +2) learning for R+ and R-groups. Bar graphs show mean $\pm$ 95% confidence interval (Wilcoxon signed-rank test).




Figure 4 – figure supplement 1. Spontaneous activity does not account for the LMI-reactivation relationship.

(A) Scatter plot of reactivation participation rate versus spontaneous calcium transient frequency on Day 0 for individual neurons, shown separately for R+ (left) and R- (right) mice. Each point represents one neuron, color-coded by LMI value. The black lines indicate the Pearson correlation. (B) Correlation between residuals of LMI and reactivation participation rate obtained from regressing out spontaneous frequency, shown separately for R+ mice (left) and R- mice (right). This analysis tests whether the LMI-participation relationship holds independently of baseline spontaneous activity. Dots indicate individual neurons. The black lines indicate the Pearson correlation. (C) Proportion of neurons classified as significantly participating in reactivation events across training days for LMI-positive (red) and LMI-negative (blue) neurons, shown separately for R+ mice (left) and R- mice (right). Each cell was classified as significantly participating if its observed participation rate exceeded the 95th percentile of a null distribution built from 1,000 circular shifts of the neuronal data. The effect of days on the proportion of reactivations was tested with a Kruskal-Wallis test. This analysis tests whether the LMI-participation relationship holds for a binary measure of significant reactivation independent of the magnitude of the participation rate and accounts for cells that would take part in reactivation by chance due to high spontaneous activity.

Acknowledgements

We thank James Priestley and the members of the Petersen laboratory for helpful discussions. This work was supported by grants 310030_219343 and TMAG-3_209271 from the Swiss National Science Foundation.

Additional information

Author contributions

A.R., G.F., S.C. and C.C.H.P. conceptualized the study; A.R., G.F., R.F.D., A.B. and P.B. wrote the data acquisition codes. A.R., G.F. and M.I. performed all animal experiments and constructed the database. A.R. and G.F. analyzed the data. A.R., S.C. and C.C.H.P. wrote and edited the manuscript, with comments from all authors; S.C. and C.C.H.P. acquired funding and provided overall supervision.

Funding

Funder	Grant reference number	Author
Schweizerischer Nationalfonds zur Förderung der Wissenschaftlichen Forschung (SNF)	310030_219343	Carl CH Petersen
Schweizerischer Nationalfonds zur Förderung der Wissenschaftlichen Forschung (SNF)	TMAG-3_209271	Carl CH Petersen

Author ORCID iDs

Robin F Dard:  <https://orcid.org/0009-0000-6221-0407>

Sylvain Crochet:  <https://orcid.org/0000-0002-2965-8899>

Carl CH Petersen:  <https://orcid.org/0000-0003-3344-4495>

References

- Ahmed A**, Voelcker B, Peron S (2024) Representational drift in barrel cortex is receptive field dependent. *Current Biology* **34**:5623-5634.e4. <https://doi.org/10.1016/j.cub.2024.10.021> | PubMed
- Banerjee A**, Parente G, Teutsch J, Lewis C, Voigt FF, Helmchen F (2020) Value-guided remapping of sensory cortex by lateral orbitofrontal cortex. *Nature* **585**:245-250 <https://doi.org/10.1038/s41586-020-2704-z> | PubMed
- Bauer J**, Lewin U, Herbert E, Gjorgjieva J, Schoonover CE, Fink AJP, Rose T, Bonhoeffer T, Hübener M (2024) Sensory experience steers representational drift in mouse visual cortex. *Nature Communications* **15**:9153 <https://doi.org/10.1038/s41467-024-53326-x> | PubMed
- Chen JL**, Margolis DJ, Stankov A, Sumanovski LT, Schneider BL, Helmchen F (2015) Pathway-specific reorganization of projection neurons in somatosensory cortex during learning. *Nature Neuroscience* **18**:1101-1108 <https://doi.org/10.1038/nn.4046> | PubMed
- Chéreau R**, Bawa T, Fodouliau L, Carleton A, Pagès S, Holtmaat A (2020) Dynamic perceptual feature selectivity in primary somatosensory cortex upon reversal learning. *Nature Communications* **11**:3245 <https://doi.org/10.1038/s41467-020-17005-x> | PubMed
- Constantinople CM**, Bruno RM (2011) Effects and mechanisms of wakefulness on local cortical networks. *Neuron* **69**:1061-1068 <https://doi.org/10.1016/j.neuron.2011.02.040> | PubMed
- Daigle TL**, Madisen L, Hage TA, Valley MT, Knoblich U, Larsen RS, Takeno MM, Huang L, Gu H, Larsen R, et al. (2018) A Suite of Transgenic Driver and Reporter Mouse Lines with Enhanced Brain-Cell-Type Targeting and Functionality. *Cell* **174**:465-480.e22 <https://doi.org/10.1016/j.cell.2018.06.035> | PubMed

- Deister CA**, Moore AI, Voigts J, Bechek S, Lichtin R, Brown TC, Moore CI (2024) Neocortical inhibitory imbalance predicts successful sensory detection. *Cell Reports* **43**:114233 <https://doi.org/10.1016/j.celrep.2024.114233> | PubMed
- Diamond ME**, von Heimendahl M, Knutsen PM, Kleinfeld D, Ahissar E (2008) “Where” and “what” in the whisker sensorimotor system. *Nature Reviews. Neuroscience* **9**:601-612 <https://doi.org/10.1038/nrn2411> | PubMed
- Drieu C**, Zhu Z, Wang Z, Fuller K, Wang A, Elnozahy S, Kuchibhotla K (2025) Rapid emergence of latent knowledge in the sensory cortex drives learning. *Nature* **641**:960-970 <https://doi.org/10.1038/s41586-025-08730-8> | PubMed
- Gasselin C**, Hohl B, Vernet A, Crochet S, Petersen CCH (2021) Cell-type-specific nicotinic input disinhibits mouse barrel cortex during active sensing. *Neuron* **109**:778-787.e3. <https://doi.org/10.1016/j.neuron.2020.12.018> | PubMed
- Goltstein PM**, Coffey EBJ, Roelfsema PR, Pennartz CMA (2013) In Vivo Two-Photon Ca²⁺ Imaging Reveals Selective Reward Effects on Stimulus-Specific Assemblies in Mouse Visual Cortex. *Journal of Neuroscience* **33**:11540-11555 <https://doi.org/10.1523/JNEUROSCI.1341-12.2013> | PubMed
- Guo ZV**, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K (2014) Flow of cortical activity underlying a tactile decision in mice. *Neuron* **81**:179-194 <https://doi.org/10.1016/j.neuron.2013.10.020> | PubMed
- Han S**, Helmchen F (2025) Coordinated multi-level adaptations across neocortical areas during task learning. *Nature Communications* **16**:7719 <https://doi.org/10.1038/s41467-025-62949-7> | PubMed
- Harris JA**, Hirokawa KE, Sorensen SA, Gu H, Mills M, Ng LL, Bohn P, Mortrud M, Ouellette B, Kidney J, *et al.* (2014) Anatomical characterization of Cre driver mice for neural circuit mapping and manipulation. *Frontiers in Neural Circuits* **8**:76 <https://doi.org/10.3389/fncir.2014.00076> | PubMed
- Henschke JU**, Dylida E, Katsanevaki D, Dupuy N, Currie SP, Amvrosiadis T, Pakan JMP, Rochefort NL (2020) Reward Association Enhances Stimulus-Specific Representations in Primary Visual Cortex. *Current Biology* **30**:1866-1880.e5. <https://doi.org/10.1016/j.cub.2020.03.018> | PubMed
- Kato HK**, Gillet SN, Isaacson JS (2015) Flexible Sensory Representations in Auditory Cortex Driven by Behavioral Relevance. *Neuron* **88**:1027-1039 <https://doi.org/10.1016/j.neuron.2015.10.024> | PubMed
- Kwon SE**, Yang H, Minamisawa G, O'Connor DH (2016) Sensory and decision-related activity propagate in a cortical feedback loop during touch perception. *Nature Neuroscience* **19**:1243-1249 <https://doi.org/10.1038/nn.4356> | PubMed
- Lacefield CO**, Pnevmatikakis EA, Paninski L, Bruno RM (2019) Reinforcement Learning Recruits Somata and Apical Dendrites across Layers of Primary Sensory Cortex. *Cell Reports* **26**:2000-2008.e2. <https://doi.org/10.1016/j.celrep.2019.01.093> | PubMed
- Lensjø KK**, Nystuen IN, Rogge FS, Tøndel K, Sugden A, Shurnayte I, Grødem S, Malthe-Sørensen A, Hafting T, Andermann ML, *et al.* (2025) Local inhibitory circuits mediate cortical reactivations and memory consolidation. *Science Advances* **11**:eadu9800 <https://doi.org/10.1126/sciadv.adu9800> | PubMed
- Makino H**, Komiyama T (2015) Learning enhances the relative impact of top-down processing in the visual cortex. *Nature Neuroscience* **18**:1116-1122 <https://doi.org/10.1038/nn.4061> | PubMed
- Marks TD**, Goard MJ (2021) Stimulus-dependent representational drift in primary visual cortex. *Nature Communications* **12**:5169 <https://doi.org/10.1038/s41467-021-25436-3> | PubMed
- McClelland JL**, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review* **102**:419-457 <https://doi.org/10.1037/0033-295X.102.3.419> | PubMed
- Miller JK**, Ayzenshtat I, Carrillo-Reid L, Yuste R (2014) Visual stimuli recruit intrinsically generated cortical ensembles. *Proceedings of the National Academy of Sciences* **111**:E4053-E4061 <https://doi.org/10.1073/pnas.1406077111> | PubMed

- Muñoz W, Tremblay R, Levenstein D, Rudy B (2017) Layer-specific modulation of neocortical dendritic inhibition during active wakefulness. *Science (New York, N.Y.)* **355**:954-959 <https://doi.org/10.1126/science.aag2599> | PubMed
- Nguyen ND, Lutas A, Amsalem O, Fernando J, Ahn AY-E, Hakim R, Vergara J, McMahon J, Dimidschstein J, Sabatini BL, et al. (2024) Cortical reactivations predict future sensory responses. *Nature* **625**:110-118 <https://doi.org/10.1038/s41586-023-06810-1> | PubMed
- Oryshchuk A, Sourmpis C, Weverbergh J, Asri R, Esmaili V, Modirshanechi A, Gerstner W, Petersen CCH, Crochet S (2024) Distributed and specific encoding of sensory, motor, and decision information in the mouse neocortex during goal-directed behavior. *Cell Reports* **43**:113618 <https://doi.org/10.1016/j.celrep.2023.113618> | PubMed
- Pachitariu M, Stringer C, Dipoppa M, Schröder S, Rossi LF, Dalglish H, Carandini M, Harris KD (2017) Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv* 061507 <https://doi.org/10.1101/061507>
- Park E, Kuljis DA, Swindell RA, Ray A, Zhu M, Christian JA, Barth AL (2025) Somatostatin neurons detect stimulus-reward contingencies to reduce neocortical inhibition during learning. *Cell Reports* **44**:115606 <https://doi.org/10.1016/j.celrep.2025.115606> | PubMed
- Peron SP, Freeman J, Iyer V, Guo C, Svoboda K (2015) A Cellular Resolution Map of Barrel Cortex Activity during Tactile Behavior. *Neuron* **86**:783-799 <https://doi.org/10.1016/j.neuron.2015.03.027> | PubMed
- Petersen CCH (2019) Sensorimotor processing in the rodent barrel cortex. *Nature Reviews. Neuroscience* **20**:533-546 <https://doi.org/10.1038/s41583-019-0200-y> | PubMed
- Pologruto TA, Sabatini BL, Svoboda K (2003) ScanImage: flexible software for operating laser scanning microscopes. *Biomedical Engineering Online* **2**:13 <https://doi.org/10.1186/1475-925X-2-13> | PubMed
- Poort J, Khan AG, Pachitariu M, Nemri A, Orsolic I, Krupic J, Bauza M, Sahani M, Keller GB, Mrsic-Flogel TD, et al. (2015) Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. *Neuron* **86**:1478-1490 <https://doi.org/10.1016/j.neuron.2015.05.037> | PubMed
- Ramamurthy DL, Chen A, Zhou J, Park C, Huang PC, Bharghavan P, Krishna G, Liu J, Casale K, Feldman DE (2023) VIP interneurons in sensory cortex encode sensory and action signals but not direct reward signals. *Current biology: CB* **33**:3398-3408.e7. <https://doi.org/10.1016/j.cub.2023.06.086> | PubMed
- Rübel O, Tritt A, Ly R, Dichter BK, Ghosh S, Niu L, Baker P, Soltesz I, Ng L, Svoboda K, et al. (2022) The Neurodata Without Borders ecosystem for neurophysiological data science. *eLife* **11**:e78362 <https://doi.org/10.7554/eLife.78362> | PubMed
- Sachidhanandam S, Sermet BS, Petersen CCH (2016) Parvalbumin-expressing GABAergic neurons in mouse barrel cortex contribute to gating a goal-directed sensorimotor transformation. *Cell Reports* **15**:700-706 <https://doi.org/10.1016/j.celrep.2016.03.063> | PubMed
- Sachidhanandam S, Sreenivasan V, Kyriakatos A, Kremer Y, Petersen CCH (2013) Membrane potential correlates of sensory perception in mouse barrel cortex. *Nature Neuroscience* **16**:1671-1677 <https://doi.org/10.1038/nn.3532> | PubMed
- Sippy T, Chaimowitz C, Crochet S, Petersen CCH (2021) Cell Type-Specific Membrane Potential Changes in Dorsolateral Striatum Accompanying Reward-Based Sensorimotor Learning. *Function* **2**:zqab049 <https://doi.org/10.1093/function/zqab049> | PubMed
- Sippy T, Lapray D, Crochet S, Petersen CCH (2015) Cell-type-specific sensorimotor processing in striatal projection neurons during goal-directed behavior. *Neuron* **88**:298-305 <https://doi.org/10.1016/j.neuron.2015.08.039> | PubMed
- Smith AC, Frank LM, Wirth S, Yanike M, Hu D, Kubota Y, Graybiel AM, Suzuki WA, Brown EN (2004) Dynamic Analysis of Learning in Behavioral Experiments. *Journal of Neuroscience* **24**:447-461 <https://doi.org/10.1523/JNEUROSCI.2908-03.2004> | PubMed

Sugden AU, Zaremba JD, Sugden LA, McGuire KL, Lutas A, Ramesh RN, Alturkistani O, Lensjø KK, Burgess CR, Andermann ML (2020) Cortical reactivations of recent sensory experiences predict bidirectional network changes during learning. *Nature Neuroscience* **23**:981-991 <https://doi.org/10.1038/s41593-020-0651-5> | PubMed

Takahashi N, Ebner C, Sigl-Glöckner J, Moberg S, Nierwetberg S, Larkum ME (2020) Active dendritic currents gate descending cortical outputs in perception. *Nature Neuroscience* **23**:1277-1285 <https://doi.org/10.1038/s41593-020-0677-8> | PubMed

Vavladeli A, Daigle T, Zeng H, Crochet S, Petersen CCH (2020) Projection-specific Activity of Layer 2/3 Neurons Imaged in Mouse Primary Somatosensory Barrel Cortex During a Whisker Detection Task. *Function* **1**:zqaa008 <https://doi.org/10.1093/function/zqaa008> | PubMed

Wang HC, LeMessurier AM, Feldman DE (2022) Tuning instability of non-columnar neurons in the salt-and-pepper whisker map in somatosensory cortex. *Nature Communications* **13**:6611 <https://doi.org/10.1038/s41467-022-34261-1> | PubMed

Woolsey TA, Van der Loos H (1970) The structural organization of layer IV in the somatosensory region (S I) of mouse cerebral cortex: The description of a cortical field composed of discrete cytoarchitectonic units. *Brain Research* **17**:205-242 [https://doi.org/10.1016/0006-8993\(70\)90079-X](https://doi.org/10.1016/0006-8993(70)90079-X) | PubMed

Yamashita T, Petersen CCH (2016) Target-specific membrane potential dynamics of neocortical projection neurons during goal-directed behavior. *eLife* **5**:15798 <https://doi.org/10.7554/eLife.15798> | PubMed

Zhao S, Ting JT, Atallah HE, Qiu L, Tan J, Gloss B, Augustine GJ, Deisseroth K, Luo M, Graybiel AM, *et al.* (2011) Cell type-specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. *Nature Methods* **8**:745-752 <https://doi.org/10.1038/nmeth.1668> | PubMed

Peer reviews

Reviewer #1 (Public review):

Summary:

The paper submitted by Renard et al. seeks to capture the moment when learning occurs and to identify the associated changes in neuronal activity within cortical circuits. Specifically, the study aims to test whether sensory representations in the cortex reorganize on the same timescale over which behavioral changes first emerge.

To address this question, the authors developed a new behavioral paradigm in which mice were first trained on an auditory detection task and then introduced to whisker stimulation, which they learned to associate with reward. This design allowed mice to form a new whisker-reward association within a single behavioral session, enabling the authors to track learning-associated neuronal changes during the course of the experiment.

Using pharmacological and optogenetic interventions, the authors first show that learning depends on the whisker somatosensory cortex. They then combined the task with longitudinal two-photon calcium imaging to examine real-time changes in neuronal representations that accompany improvements in task performance over trials within a session and across days. By applying a range of analytical approaches, they show that learning induces a rapid reorganization of sensory cortical representations over tens of trials, on the timescale of minutes. They further propose that spontaneous reactivation of neurons during the task may contribute to these representational changes during learning.

Strengths:

(1) Overall, the experiments are thoughtfully designed, well controlled, and clearly presented. The conclusions are generally well supported by the data. The manuscript is clearly written,

and the Discussion acknowledges potential caveats while outlining future directions.

(2) A major strength of the study is the design of a new learning paradigm in which head-fixed mice rapidly form a new sensory-motor association within a single session, on the timescale of minutes. This offers a unique opportunity to track real-time changes in neuronal dynamics associated with learning during a single recording experiment.

(3) Taking advantage of this behavioral design, the authors show that learning induces rapid reorganization of sensory cortical representations. They also report an increase in spontaneous reactivation of neurons that gained stimulus responsiveness during training, and propose that these reactivations may contribute to rapid representational reorganization. These findings provide important insights into the neural dynamics associated with learning.

Weaknesses:

(1) The authors propose that spontaneous reactivation mediates rapid reorganization of neuronal representations and thereby supports rapid task learning. However, as they also acknowledge in the Discussion, the present study does not directly test a causal role for these reactivations in facilitating representational changes or behavioral improvement.

(3) The authors show reorganization of neuronal representations even on the first day of training with the new whisker task. However, because there is no explicit control for natural representational drift, it remains unclear to what extent these changes reflect learning-related reorganization rather than spontaneous day-to-day drifts in neuronal responses.

<https://doi.org/10.7554/eLife.111818.1.sa2>

Reviewer #2 (Public review):

Summary:

Renard, Foustoukos and colleagues present a study of rapid sensorimotor learning in the mouse barrel cortex. Head-fixed water-restricted mice already trained on an auditory detection task are introduced to a novel C2 whisker stimulus, and the authors show that reward-paired mice acquire the whisker-lick association within a single behavioral session, with the two groups (rewarded vs non-rewarded) diverging behaviorally within ~22 whisker trials and ~14 minutes. Both pharmacological inactivation of wS1 across Days 0/+1/+2 and optogenetic inactivation on Day 0 impair whisker-guided performance, while fpS1 manipulations do not, establishing that wS1 activity is required for whisker-guided behavior during the initial learning period. Longitudinal two-photon imaging of GCaMP6f-expressing L2/3 neurons across five days (-2 to +2 relative to whisker introduction) reveals a bidirectional, reward-dependent reorganization of population responses to passive whisker stimuli: rewarded mice show enhancement, non-rewarded mice show suppression. The authors use a logistic-regression decoder trained to discriminate pre- vs post-learning passive trials and then project Day 0 active whisker trials onto this learning axis; the projection rises monotonically across Day 0 in R+ mice and is significantly correlated with behavioral performance, with no such trajectory in R- mice. Finally, the authors detect reactivation events during catch trials by template-matching to the average passive whisker response, and show that on Day 0, the neurons most positively modulated by learning (LMI-positive) participate in these reactivations more than LMI-negative neurons in R+ but not R- mice. The authors interpret this as evidence that online, reward-gated reactivations may act as an upstream selection mechanism for which neurons undergo learning-related plasticity, operating on the minutes-timescale of within-session learning. There is much to like in this paper, with some moderate-to-major concerns that could largely be addressed with re-analysis or re-framing.

Strengths:

The single-session learning paradigm is a key aspect of this paper, given the rapid learning observed. Coupled with the R+ and R- design, there's a lot to like with the behavioral approach. The bidirectional response change across these R+ and R- groups (enhancement vs suppression) is also a nice finding.

The causal manipulations demonstrate that the imaged region is used during the task. By doing both pharmacological and optogenetic inactivation, each with a control in the spatially adjacent region (fpS1), the authors make a strong case that wS1 activity is necessary for whisker-guided behavior during the initial learning period (though see below about the limitations of the current approach).

The longitudinal two-photon imaging of the same L2/3 neurons across five days underlies essentially every neural analysis in the paper and enables the single-cell LMI and population-trajectory analyses.

The pathway-specific analysis in Figure 3 - figure supplement 2 is very interesting, but not much time is spent on it (lines 151-155). The dissociation between wS2-projecting neurons (which show learning-related enhancement in R+ and suppression in R-) and wM1-projecting neurons (which do not) is (in my opinion) a nice instance of projection specificity - it also aligns with the known routing of task-relevant whisker information through the wS1 → wS2 pathway. I would encourage the authors to motivate this experiment in the main text rather than leaving it all to the discussion (lines 256-262).

The methods are generally well documented and easy to follow.

Weaknesses:

(1) Conflation of de novo association learning with generalization from auditory pre-training.

All mice have already learned a task structure with the auditory task - "detect the salient sensory cue → lick → reward". Under these conditions, the rapid emergence of licking to the whisker stimulus could reflect either de novo formation of a whisker-specific association or generalization of an instrumental policy to a novel salient cue. The manuscript frames the result as the former ("acquisition of a novel sensorimotor association"), but the experiment cannot distinguish between the two alternatives. This distinction between de novo learning and generalization may have a meaningful impact on the interpretation, though it doesn't impact the specific results. It would be helpful for the authors to discuss the two possibilities and generally consider the contribution of generalization from auditory pre-training to Day 0 performance.

Relatedly, the R- group is introduced (lines 69-74) and later used (lines 244-247) as a passive-exposure control that rules out representational drift. While R- group is an important control for repeated whisker stimulation and task context, it does not appear to be a pure passive-exposure control: Figure 1B shows that on Day 0 the mice lick more to the R- stimulus than with no stimulus and then extinguish that licking by Day 1. Thus, one possibility is that R- mice actively learn to suppress licking to an unrewarded stimulus (whisker) in a context where other stimuli (auditory) remain rewarded. This would be a different cognitive operation (response suppression) from a purely passive exposure condition. The manuscript therefore lacks a true passive-exposure baseline, and several claims that rely on R- as such a baseline (including that bidirectional changes are reward-driven rather than reflecting passive drift, lines 244-247) need to be reframed.

(2) The inactivation experiments establish that wS1 is necessary on Day 0, but they cannot separate detection, acquisition, and expression.

Both the muscimol manipulation (whole session, Days 0/+1/+2) and the optogenetic manipulation (0.1 s before stimulus onset through the 1 s reporting window) silence wS1 during the moments when the whisker stimulus must be detected for a successful trial. Under these conditions, impaired performance could reflect that the animal cannot detect the stimulus, cannot express the learned response on that trial, or cannot acquire the association. These are causally distinct processes, and the manuscript currently treats them as equivalent.

Specifically, on Day +1 of the opto experiment (light off), do mice learn at the same rate as a naive Day 0 cohort (e.g., the R+ imaging mice on Day 0), or is performance already higher than the naive group? If higher than the naïve group, this would suggest that there is learning occurring and would suggest that something that may have been acquired during Day 0 inactivation, even if it could not be expressed.

(3) The interpretation of the LMI-participation correlation is complicated by the peaked LMI distribution and neuron-level pooling.

Two related issues arise from the results shown in Figure 4I. First, the LMI distribution in Figure 3F (and visible in 4I) is sharply peaked near zero. The reported $r = 0.24$ in R+ mice is therefore difficult to interpret biologically because the distribution is dominated by near-zero-LMI neurons and the slope may be disproportionately influenced by neurons in the tails. The key claim is better tested by comparing significantly LMI-positive, LMI-negative, and non-modulated neurons. The authors do address this in Figure 4J - showing that participation rate rises across days for significantly LMI-positive R+ neurons ($p = 5 \times 10^{-4}$) but not for LMI-negative neurons ($p = 0.05$) - but this analysis is not the lead result. To my understanding, Figure 4J is more interpretable and should be the key piece of data supporting their claim.

Second, the p-value of $p = 1 \times 10^{-41}$ in Figure 4I comes from treating thousands of neurons pooled across 19 mice as independent observations. Neurons within an animal are correlated through shared behavioral state, shared imaging session, and circuit-level interactions, so it would be helpful to consider a different statistical unit of comparison (FOV, animal, etc). For example, a linear mixed-effects model with mouse as a random effect could work.

(4) The reactivation-LMI relationship is partially circular, and the framing in the abstract could be more constrained.

The "reactivation template" is the trial-averaged passive whisker-evoked population vector from each session, and reactivations are detected as moments in catch-trial activity that correlate with this template above a shuffled threshold. This approach is reasonable, but it means that the reactivation-LMI relationship is not fully independent of template construction, and the framing in the abstract blurs that line. LMI-positive neurons are defined as neurons whose passive whisker-evoked responses increase from pre- to post-learning. Therefore, neurons with strong whisker responses, or neurons that become stronger components of the whisker-evoked template across learning, may be more likely to contribute to template-matching events by construction. Thus, the LMI-participation relationship could partly reflect template weighting or sensory-response amplitude, rather than showing that reactivation events selectively recruit neurons for future learning-related plasticity. It would be helpful and more reassuring if the authors could control for each neuron's whisker-template weight, baseline whisker responsiveness, and overall calcium event rate when relating LMI to reactivation participation.

A complementary unsupervised approach could also help: rather than starting from the whisker template, one can derive co-activity assemblies directly from spontaneous activity (e.g., via PCA or ICA on the catch-trial population activity), and then ask, separately, whether any of these assemblies overlap with the whisker ensemble. The interesting test is then whether whisker-like assemblies become more frequently expressed across Day 0 in R+ but not R- mice, and whether LMI-positive neurons are preferentially loaded onto these whisker-

like assemblies. This logic inverts the current pipeline and can be complementary to the current analysis. By identifying structure in nominally spontaneous activity first and then comparing to the whisker response, this could help avoid the circularity in which the template both defines the events and contains the cells being tested. The Figure 4 - figure supplement 1B partial-correlation analysis is a step in this direction but addresses only spontaneous firing rate, not template coupling. Without such a complementary approach, the authors may want to clarify that the reactivation detection is anchored to a template defined in part by the same cells whose participation is being tested.

(5) The reactivation-as-selection-mechanism interpretation is not supported by the current data.

The Discussion (lines 278-281) acknowledges that the authors have not shown necessity, but the end of the intro and part of the discussion (Lines 275-277) frame reactivations as a "reward-gated selection mechanism" for plasticity. An equally plausible alternative is that neurons whose synaptic inputs or intrinsic excitability have been potentiated by reward-driven learning will simply co-fire more often during quiet periods - meaning reactivations would be a consequence of plasticity that has already occurred rather than a mechanism that selects which neurons to potentiate. The current data cannot distinguish these.

A separate concern is the use of the term "spontaneous." The authors' usage is defensible in one sense - catch trials are stimulus-free, so the activity is not externally driven. However, "spontaneous" in the reactivation literature typically connotes offline, internally generated activity during quiet wakefulness or sleep, which carries different implications for plasticity than activity during active task engagement. Catch trials in this paradigm occur within the behavioral session, with the animal still engaged in the task, potentially anticipating reward or licking. The authors should either acknowledge this distinction in the text or qualify the term - "within-session" or "inter-trial" reactivations would be more accurate and would avoid borrowing the conceptual weight of the offline-replay literature.

The authors should also clarify whether catch-trial activity around licks (false alarms, anticipatory licks) is excluded from the reactivation analysis, and whether reactivation rates depend on recent reward, recent whisker trial outcome, or behavioral state. Specificity controls - template-matching with shuffled templates and with auditory templates - would help establish that detected events reflect whisker-specific patterns rather than generic high-coactivity moments.

(6) Motor, lick, and behavioral-state confounds in the neural analyses are not fully addressed.

I have two specific concerns. First, for the Day 0 active-trial projection, mean whisker reaction times in Figure 1 - figure supplement 1G are around 350-500 ms, but the distributions extend into the 0-300 ms analysis window. The correlation between the projection trajectory and the behavioral learning curve (Figure 4E, lines 196-198) is the key piece of evidence that the neural shift tracks learning. However, on hit trials the lick may fall within or close to the analysis window, so a motor confound could in principle contribute to the rising projection. The authors could repeat the projection using an earlier/shorter window, exclude trials with early licks, or regress out lick timing. It would be helpful to better understand whether this effect is, in part, driven by licking activity.

Second, the central evidence for representational reorganization (Figure 3) rests on a post-session passive epoch in which 50 whisker stimulations are delivered after "task disengagement" (lines 131, 387-389). The concern is that the brain state during this epoch is unlikely to be matched across groups or across days. R+ mice receive additional water rewards on whisker trials, whereas R- mice receive rewards only on auditory trials. This could lead to systematic differences in satiety, arousal, and disengagement state during the passive block. Because cortical sensory responses are strongly modulated by arousal, some of

the apparent learning-related enhancement (R+) or suppression (R-) of passive whisker responses across days could reflect systematic state differences during the passive epoch rather than plasticity. The disengagement criterion ("stopped licking in all trial types") is also qualitative - no consecutive-miss or time-window threshold is specified - so the epoch may begin at slightly different behavioral states across mice. To resolve this, the authors could (i) specify the disengagement criterion quantitatively and (ii) compare pupil diameter and whisker self-motion (if available) across R+ vs R- and across days during the passive epoch.

<https://doi.org/10.7554/eLife.111818.1.sa1>

Reviewer #3 (Public review):

This is a methodologically sound manuscript and provides reasonably interpretable results. While being appropriate, they do not seem to bring entirely novel concepts; nevertheless, most of my comments concern the calibration of the interpretive claims rather than the quality of the data.

Strengths:

(1) Longitudinal within-subject imaging:

Tracking the same layer 2/3 neurons across learning allows the bidirectional effect (enhancement in R+, suppression in R-) to be measured within identified cells rather than inferred across cohorts.

(2) Appropriate behavioural controls:

The R+/R- design controls for repeated sensory exposure, and maintaining rewarded auditory trials in both groups controls for engagement and arousal, arguing against disengagement as the source of the R- effect.

(3) Convergent causal manipulations:

Muscimol and optogenetic inactivation both abolish acquisition and include an adjacent control region (fpS1); the temporally restricted optogenetic result partially addresses the concern (Hong et al., 2018) that sustained inactivation may destabilise downstream circuits.

(4) Convergent analyses:

Single-cell learning modulation indices, population similarity measures, and a trial-resolved decoder projection onto a naïve-to-expert axis provide consistent evidence that representational change is concurrent with behavioural acquisition.

(5) Projection-specific resolution:

Retrograde labelling shows learning-related changes in wS2-projecting, but not wM1-projecting neurons, consistent with preferential routing of task-relevant signals through the wS1 to wS2 pathway.

(6) Mechanistically motivated reactivation analysis:

Relating rapid, reward-dependent plasticity to spontaneous reactivations on a timescale of minutes is an original use of the single-session paradigm.

Weaknesses and points requiring clarification

(1) The stimulus is not strictly novel: Passive whisker stimulations were delivered on pre-training Days -2 and -1, so what changes on Day 0 is the stimulus-reward contingency rather than the stimulus itself. This resembles contingency reassignment with reversal-like

properties (and possible habituation or latent inhibition) rather than de novo learning, and the licking response is already established during auditory training. The framing should be qualified accordingly.

(2) Barrel cortex dependence should be stated more narrowly: The data show that acute wS1 suppression prevents acquisition of this task, not that whisker detection in general requires barrel cortex; cortical dependence varies with task and manipulation (Hong et al., 2018; Ryan et al., 2022 vs Miyashita and Feldman, 2013). The near-threshold explanation would require psychometric or stimulus-intensity data.

(3) The passive block carries confounds: It is acquired after task disengagement, when satiety, arousal, and reward history differ across groups and days. The authors should report within-block response adaptation and the robustness of the main results to early versus late passive trials. Additionally, could passive presentation of the stimulus without reward delivery lead to devaluation of the stimulus, leading to additional behaviour and plasticity changes which are not addressed?

(4) Some statistics appear to be neuron-level rather than animal-level:

Very small p-values (e.g., the LMI-participation correlation $r = 0.24$, p on the order of 10^{-41}) suggest thousands of non-independent neurons treated as independent samples, risking pseudoreplication. Central claims should rest on hierarchical or animal-level statistics with effect sizes.

(5) The cosine-similarity decrease in R- animals needs clarification:

Because cosine similarity is scale-invariant, uniform suppression would leave it largely unchanged; the observed decrease therefore implies heterogeneous suppression, reduced signal-to-noise, or increased variability, and the favoured interpretation should be stated.

(6) The decoder requires cautious interpretation:

Training on passive trials and applying to active Day 0 trials could introduce a behavioural-state domain shift. The meaning of positive and negative values in Figure 4C should be defined, and near-zero early projections reflect the classifier boundary rather than a biological baseline.

(7) The reactivation analysis is the least conclusive and is susceptible to circularity:

The template and the LMI are both derived from the passive whisker response, predisposing responsive neurons to register as reactivation participants, and the Day 0 template is obtained after learning.

Leave-one-cell-out and pre-learning templates, cell-specific templates, and tests of whether reactivations predict subsequent trial responses would strengthen the claim; causal disruption would ultimately be required.

(8) Figure, sample-size, and specificity points:

The positive LMI shift in R+ animals is less visible than the R- shift in Figure 3F; the optogenetic cohort is small ($n = 6$ per group); and confirming that auditory detection was preserved during wS1 inactivation would establish whisker-specificity.

(9) The comparison to prior work is overly broad:

Banerjee et al. (2020) and Chéreau et al. (2020) are reversal learning and discrimination paradigms and may not be equated with simple whisker detection; the defensible novelty claim is the trial-resolved tracking within the first session and its concurrence with online reactivations.

<https://doi.org/10.7554/eLife.111818.1.sa0>