

Learning enhances behaviorally relevant representations in apical dendrites

Reviewed Preprint

v2 • December 3, 2024

Revised by authors

Reviewed Preprint

v1 • June 13, 2024

Sam E Benezra, Kripa B Patel, Citlali Pérez Campos, Elizabeth MC Hillman, Randy M Bruno 

Department of Neuroscience, Columbia University, New York, United States • Kavli Institute for Brain Science, Columbia University, New York, United States • Zuckerman Mind Brain Behavior Institute, Columbia University, New York, United States • Departments of Biomedical Engineering and Radiology, Columbia University, New York, United States • Department of Physiology, Anatomy & Genetics, University of Oxford, Oxford, United Kingdom

 https://en.wikipedia.org/wiki/Open_access
 Copyright information

eLife Assessment

This **important** study uses calcium imaging to show an increase in the selectivity of the sensory-evoked response in the apical dendritic tuft of layer 5 barrel cortex neurons as mice learn a whisker-dependent discrimination task. The evidence supporting the conclusions is **compelling**, and this work will be of great interest to neuroscientists working on reward-based learning and sensory processing.

<https://doi.org/10.7554/eLife.98349.2.sa3>

Abstract

Summary

Learning alters cortical representations and improves perception. Apical tuft dendrites in Layer 1, which are unique in their connectivity and biophysical properties, may be a key site of learning-induced plasticity. We used both two-photon and SCAPE microscopy to longitudinally track tuft-wide calcium spikes in apical dendrites of Layer 5 pyramidal neurons in barrel cortex as mice learned a tactile behavior. Mice were trained to discriminate two orthogonal directions of whisker stimulation. Reinforcement learning, but not repeated stimulus exposure, enhanced tuft selectivity for both directions equally, even though only one was associated with reward. Selective tufts emerged from initially unresponsive or low-selectivity populations. Animal movement and choice did not account for changes in stimulus selectivity. Enhanced selectivity persisted even after rewards were removed and animals ceased performing the task. We conclude that learning produces long-lasting realignment of apical dendrite tuft responses to behaviorally relevant dimensions of a task.

Introduction

Learning and memory depend on the ability of biological networks to alter their activity based on past experience. For example, as animals learn the behavioral relevance of stimuli in a sensory discrimination task, neural representations of those stimuli are enhanced^{1–7}, potentially improving the salience of information relayed to downstream areas. Studies in primary somatosensory (S1)⁸ and visual cortex² have revealed that top-down signals from distant cortical regions can modify sensory representations during learning, although the cellular and circuit mechanisms underlying this plasticity remain unclear.

Cortical layer 1, comprised mainly of apical tuft dendrites of layer 5 (L5) and layer 2/3 pyramidal neurons, may be a key site driving the enhancement of sensory representations during learning. Apical tufts are anatomically well positioned for learning, receiving top-down signals from numerous cortical and thalamic areas^{9–11}. While L5 distal tufts are electrically remote and far from the soma, they are in close proximity to the highly electrogenic calcium spike initiation zone at the main bifurcation of the apical dendrite, and form a separate biophysical and processing compartment from the proximal dendrites^{12–16}. Top-down signals arriving at the tuft can trigger tuft-wide dendritic calcium spikes in L5 neurons¹⁷, which can modulate synaptic plasticity across the entire dendritic tree¹⁸ and potentially drive somatic burst firing^{15,19–23}. Consistent with this observation, L5 apical dendrite activity is highly correlated with somatic activity^{24,25}. Therefore, by strongly influencing somatic activity, L5 apical dendritic calcium spikes can play an important role in modulating cortical output. Several neuromodulators can augment the excitability of the apical tuft and increase the likelihood of eliciting calcium spikes^{26,27}, which could be a substrate for control of plasticity by behavioral state. Consistent with these ideas, we recently demonstrated that during behavioral training with positive reinforcements, apical tufts in sensory cortex acquire associations that extend beyond their normal sensory modality²⁸. In mouse models of dementia and Alzheimer's disease^{29,30}, tuft dendrites exhibit degeneration which may contribute to the cognitive and memory deficits.

L5 pyramidal neurons are the major source of output from cortex, targeting numerous subcortical structures that affect behavior. The activity of apical dendrites is known to correlate with stimulus intensity, and manipulating L5 apical dendrites and their inputs impacts performance of sensory tasks^{17,31–33}. Apical dendritic calcium spikes of pyramidal cells could be a crucial cellular mechanism in learning-related plasticity and behavioral modification^{18,34,35}. However, sensory representations of apical tufts, as well as possible changes across learning, have received little attention.

To address this question, we used two-photon microscopy and a new high-speed volumetric imaging technique called Swept Confocally-Aligned Planar Excitation (SCAPE)^{36,37} to longitudinally track the activity of GCaMP6f-expressing L5 apical tufts in barrel cortex during a sensory discrimination task. We found that apical tufts underwent extensive dynamic changes in selectivity for task-relevant stimuli as performance improved, even though only one of the stimuli was unrewarded. These changes in responses persisted even after animals disengaged from the task, demonstrating that learning induced long-lasting changes in tuft sensory representations. Animals that were exposed to the same stimulation protocol without any reinforcement did not develop enhanced representations. Our results show for the first time that reinforcement learning expands apical tuft sensory representations along behaviorally relevant dimensions.

Results

Direction discrimination behavior

We devised an awake head-fixed mouse conditioning paradigm that enables controlled investigation of reinforcement effects across learning (**Fig.1A**, **B**). In addition to discriminating tactile objects, rodents are known to sense wind direction using their whiskers^{38,39} and can be trained to discriminate different directions of whisker deflections^{40,41}. With this in mind, we directed brief (100-ms) air puffs at the whiskers in either of two directions: rostrocaudal (backward) or ventrodorsal (upward). One of the directions was paired with a water reward delivered 500 ms after the air puff and thus constituted a conditioned stimulus (CS+). No reward was given for the other direction (CS-).

Licking and whisking were monitored throughout the session (**Fig.1C**, **D**). Stimuli elicited a brief passive whisker deflection followed by active whisking over the subsequent ~1.5 seconds (analyzed below, **Fig.6**). Any anticipatory licks prior to reward delivery were counted as a response. Typically, on the first session, mice exhibited few anticipatory licks to either stimulus (**Fig.1C**, top, grey shading). By session 2 or 3, mice had learned an association between whisker deflection and reward, but could not discriminate the CS+ and CS- (middle). Within a week (by sessions 7-9), every mouse we tested learned to reliably lick to the CS+ while withholding licks to the CS-, performing substantially above chance after a single week of training (**Fig.1C**, bottom; **Fig.1E**, **F**). Thus, mice rapidly learned to discriminate the direction of whisker stimuli in our behavioral task.

Overall stimulus-evoked activity is unbiased and stable across conditioning

To investigate the effects of reinforcement learning on apical tuft activity, we imaged apical tufts (433 x 433 μm field of view) across conditioning days as well as on an unrewarded pre-conditioning day to measure naïve stimulus responses and an unrewarded post-conditioning day to detect any long-lasting changes in responses (**Fig.1B**). Mice remained water-restricted on the post-conditioning day and continued licking for reward toward the beginning of the session (see below). We virally delivered the gene for Cre-dependent GCaMP6f⁴² in the barrel cortex of Rbp4-Cre mice, which labels a heterogeneous population of pyramidal neurons comprising approximately 50% of layer 5.^{28,43,44} By targeting our injections to layer 5B, we predominantly labeled thick-tufted pyramidal neurons (see Methods). Using intrinsic signal imaging, we mapped the location of the C2, D2, and gamma whisker barrel columns and identified an overlapping region in layer 1 with sufficient GCaMP6f expression (**Fig.2A**). The air puff nozzles were aimed toward the whiskers corresponding to this region. Dendritic activity was longitudinally recorded from the same field-of-view (horizontal location and depth) in layer 1 across all sessions (Supplementary Movie 1).

To extract calcium signals from individual cells, we segmented tufts using CalmAn, a sparse non-negative matrix factorization method that clusters pixels according to their temporal correlation⁴⁵ (see Methods), and analyzed regions of interest exhibiting apical tuft structure (**Fig.2B**; 65 ± 15 tufts per mouse; mean $\pm$ SD). Individual segmented tufts were substantial in their spatial extent ($>100 \mu\text{m}$), reflecting tuft-wide voltage-gated calcium spikes rather than branch-specific N-methyl-D-aspartate (NMDA) receptor-mediated spikes. All calcium analyses hereafter refer to tuft-wide calcium spikes. Average responses to an event include failures. In many tufts, the CS+ and CS- reliably evoked an influx of calcium that robustly activated the tuft (examples in **Fig.2C**). Successful calcium events across tufts averaged 28% $\Delta F/F$, consistent with previous studies of layer 5 apical dendrites.^{17,31} Interestingly, during intermediate but not




Fig. 1

Mice rapidly learn to discriminate stimulus direction in head-fixed paradigm.

a, A water droplet is paired with air puffs in one direction (CS+) but not the other (CS-). Licking in anticipation of water is assessed in the response window just after CS+ or CS- and prior to water delivery for the CS+ (grey bar). **b**, Experimental timeline. 2-3 weeks after virus injection, naive tuft responses to stimuli are recorded (pre). The CS+ is then paired with water for 8-9 days (blue). On the last day, stimuli are presented without reward (post). In a separate group of mice, the same stimuli are presented over 9 days in the absence of reward (unrewarded group). **c**, Lick rasters for three different sessions in one example mouse. On session 9, the CS+ but not the CS- reliably elicits licks. **d**, Mean baseline-subtracted whisking amplitude aligned to the CS+ (red) and CS- (navy) across sessions 1, 2, and 9 of an example mouse. **e**, Learning curve demonstrates rapid learning. Mean probability of at least one lick in the response window across sessions. **f**, Behavioral performance of each mouse in the rewarded group (M1 – M7).




Fig. 2

Overall tuft response to stimuli is unbiased and relatively stable across conditioning.

a, Dendritic activity was recorded in layer 1 (i) in the C1/C2 barrel columns (ii). (i) Two-photon image $\sim 60\ \mu\text{m}$ deep relative to pia. Dashed yellow lines denote C1 and C2 boundaries from intrinsic imaging. Single cell reconstruction in left panel from 500 nm. (ii) Tangential section through layer 4 showing barrels stained with streptavidin-Alexa 647 and GCaMP6f-expressing apical trunks. Red circles indicate location of 2-photon lesions to mark the imaging region for post-hoc analysis. **b**, Overlay of five segmented pseudo-colored tufts from imaging field in A(i). **c**, Time courses of calcium responses of example tufts in **b** to three air puffs (dashes). **d**, Amplitude for CS+ (red) and CS- responses (blue), computed for each segmented tuft in the first 1.5 s post-stimulus (grey points), do not differ within or across sessions. Colored lines indicate median. **e**, Same as in **d**, showing data for all conditioning sessions.

early learning, the average population response to the CS+ exhibited a two-peak structure (Supp Fig.1 [4](#), session 4) similar to tuft reward-related signals we observed previously in barrel cortex²⁸. By the last-rewarded and post sessions, the second CS+ peak was no longer visible, which could be an endpoint of mice learning that the conditioned stimulus predicts the upcoming reward.

Reward can alter somatic receptive fields in the auditory, visual, and somatosensory cortex of both rodents and non-human primates such that rewarded stimulus representations become more robust after learning^{4,5,28,46}, although cortical sensory responses can remain unchanged during learning⁴⁷. We investigated whether calcium responses to the CS+ increased in the tuft population as animals learned its association with reward (Fig.2 [4](#)). Average responses of tufts to the CS+ and CS- were similar during the pre-conditioning session (Fig.2D [4](#); $p = 0.20$, signed rank test, $n = 440$ pre tufts and 418 post tufts), indicating that there was no inherent bias in the population toward a particular stimulus in naïve animals. Surprisingly, even after learning, responses to the CS+ and CS- were similar on the last- and post-conditioning sessions ($p = 0.62$, 0.64 , respectively, signed rank test, Fig.2D [4](#), E [4](#)), revealing that no bias develops for the CS+ among dendritic tufts. Only a minority of tufts exhibited statistically significant (see Methods) average responses to air puff stimuli (CS+ responsive: $26 \pm 8\%$; CS-responsive: $25 \pm 8\%$; mean $\pm$ SD across all sessions). When we excluded responses that were not statistically significant, we again found no difference between the average response amplitudes to the CS+ and CS- on the pre, last-rewarded, and post sessions ($p = 0.65$, 0.31 , and 0.69 , respectively, rank sum test; data not shown). Similarly, the probability of transients in response to CS+ versus CS- (see Methods) did not differ during pre-conditioning or post-conditioning sessions ($p = 0.66$ and $p = 0.44$, respectively, data not shown). Therefore, reinforcement learning in our paradigm does not bias tuft representations toward the rewarded stimulus.

While a bias for the CS+ did not develop after learning, we wondered whether overall tuft responses to both conditioned stimuli increased as animals learned the task. Linear regression analysis revealed that conditioning session number was a poor predictor of both CS+ and CS- amplitudes (All tufts R^2 , CS+: 0.0064 , CS-: 0.0035 , Fig.2E [4](#); Significantly responding tufts R^2 , CS+: 0.014 , CS-: 0.014 , data not shown). We did find a small but significant decrease in amplitude from pre to last for CS+ ($p < 0.01$) and CS- ($p < 10^{-7}$), but this was not permanent: amplitudes did not significantly differ between the pre and post sessions (Fig.2D [4](#); $p = 0.53$, 0.33 , CS+ and CS- respectively, Wilcoxon rank sum test). Taken together, these findings demonstrate that reinforcement learning does not robustly bias the magnitudes of tuft calcium responses to either stimulus at the population level.

Development of tuft selectivity with task learning

While learning produced no bias in overall tuft activity, learning might enhance selectivity for conditioned stimuli. Barrel cortex neurons are tuned to the angle of whisker deflection^{48–50}, indicating that the sets of synaptic connections activated by the CS+ and CS- may be overlapping but should not be identical. Therefore, the possibility exists that responses to the CS+ and CS- can change independently of each other. To examine this, we compared the amplitude of the average response to CS+ and CS- trials for all segmented tufts on the pre, last-rewarded, and post sessions (Fig.3A [4](#); $n = 7$ mice; 465 pre, 442 last-rewarded, and 430 post tufts). In agreement with our previous analysis, we found no significant bias in response amplitude toward CS+ or CS- during any of the three sessions (Fig.3A [4](#); Pre: $p = 0.20$; last-rewarded: $p = 0.43$; Post: $p = 0.64$, sign-rank test). Under naïve conditions during the pre session, most tufts that responded to air puff stimuli did not strongly prefer the CS+ or CS- (Fig.3A [4](#), left). Surprisingly, on the last-rewarded session and the unrewarded post-conditioning session, we observed a prominent shift in the response distribution, where many tufts exhibited more selective responses to one stimulus or the other (Fig.3A [4](#), middle and right).




Fig. 3

Reinforcement learning, but not stimulus exposure, enhances tuft selectivity for CS+ and CS- stimuli.

a, Across the indicated sessions, individual tufts (circles) exhibit larger biases to CS+ or CS- (pooled across all conditioned mice). **b**, Repeated exposure to stimuli does not bias individual tufts to CS+ or CS-. **c**, Conditioning reshapes distribution of selectivity indices for tufts from Normal on pre-conditioning session to uniform on post-conditioning session. **d**, Distribution of tuft selectivity indices remains Normal throughout all repeated exposure sessions. **e**, Selectivity (median SI magnitude of tufts for each session) increases with behavioral performance of 6 animals. **f**, Same as e, but with neural discriminability plotted on the y axis. **g**, Neural discriminability (mean $\pm$ sem) of tufts, pooled across all animals on each session, increases with conditioning and decreases with repeated exposure.

Plasticity can occur after repeated exposure to stimuli even in the absence of reinforcements^{51–55}. To test whether enhanced selectivity depended on reinforcement, we imaged a separate group of similarly water-restricted mice that were repeatedly exposed to the same stimuli for the same number of days but without any reward. These mice only received water in their home cage following each imaging session, but never during stimulus presentation. Repeated exposure mice exhibited a stable distribution of response selectivity over time (**Fig.3B**; a separate cohort of 7 mice; 317, 313, and 321 tufts on Day 1, Day 8, and Day 9, respectively). These results suggest that reinforcement learning, and not simply repeated stimulus exposure, drives apical tufts to become more selective for either the CS+ or CS-.

To directly quantify the response selectivity of tufts, we computed a selectivity index (SI; see Methods) ranging from -1 (exclusively CS-responsive) to 1 (exclusively CS+ responsive) for each tuft. Initially in both the conditioned and repeated exposure mice, the SI distribution was centered around zero, indicating that most tufts in naïve animals did not strongly prefer either stimulus (**Fig.3C,D**, left panels). Consistent with our other analyses (**Fig.2D**), the mean SI remained close to zero for each of the three sessions (**Fig.3C** and Supp.Fig.2D; -0.049, -0.001, and 0.003 for pre-conditioning, last rewarded, and post-conditioning days, respectively; one-way ANOVA $p = 0.37$), confirming that learning produced no overall bias toward one particular stimulus among the population. During learning, the SI distribution of conditioned but not repeated exposure mice shifted markedly, whereby a much greater proportion of neurons were highly selective for either the CS+ or CS- (**Fig.3C,D**, middle and right panels, $|SI|$ pre versus last-rewarded: $p < 10^{-6}$, $|SI|$ pre versus post: $p < 10^{-5}$; Wilcoxon rank sum test). These effects can even be observed within individual mice (Supp.Fig.2). Notably, different tufts within the same animal exhibited opposite changes in selectivity (Supp.Fig.2A,B). Learning significantly increased tuft selectivity in individual conditioned mice, but not repeated exposure mice (Supp.Fig.2C). The degree of enhancement in tuft selectivity was closely correlated with conditioned animals' ability to discriminate stimuli across sessions (**Fig.3E**; Pearson's $R = 0.60$, $p < 10^{-5}$).

Whereas selectivity magnitude ($|SI|$) only considers the amplitude of tuft responses to CS+ and CS-, their discriminability also depends on their variability. For example, a large difference in CS+ and CS-responses would not be discriminable if the variability of those responses were very high; a small difference might be discriminable if the variability were low. We therefore additionally calculated a d-prime metric of neural discriminability that normalizes differences in response magnitudes to each stimulus by their variability (see Methods). Similar to selectivity magnitude, we found that neural discriminability was correlated with behavioral performance (**Fig.3F**). In conditioned animals, neural discriminability of CS+ and CS-responses of tufts increased significantly across learning (**Fig.3G**, blue; first-rewarded versus last-rewarded: $p < 10^{-3}$, pre versus post: $p < 10^{-4}$; Wilcoxon rank sum test). By contrast, neural discriminability of tuft responses in the repeated exposure mice decreased slightly with progressive exposure to the stimuli (**Fig.3G**, gray; Day 1 versus Final: $p < 0.01$). Finally, we asked whether the ability to decode stimulus identity on a trial-by-trial basis increased after learning. To test this, we trained a support vector machine (SVM) to decode stimulus identity from tuft population activity (see Methods). We found that decoder performance increased significantly when comparing Pre and First sessions to Post and Last sessions (Supp.Fig.3A; sign-rank test, $p = 0.002$), whereas decoder performance did not improve over time in the repeated exposure mice (Supp.Fig.3B; sign-rank test, $p = 0.22$). Taken together, these results show that enhanced stimulus representations can emerge in apical tufts, but require reinforcement.

The above analyses rely on the accurate measurement of calcium spikes from individual tufts. While two-photon microscopy acquires images with high resolution and speed, the imaging field is restricted to a single focal plane. This method can only measure calcium signals from a thin cross-section of the three-dimensionally complex apical structures. Indeed, many of the spatial components extracted from our two-photon data were comprised of dendritic branches that cross the imaging plane at different locations (Supp.Fig.4A), which makes it difficult to determine

whether the segmentation software accurately extracted signals from one tuft or erroneously merged multiple tufts. For the same reasons, a single apical tuft could be falsely classified as two different tufts. Such errors could mislead our interpretation of selectivity in the population, especially given that a single apical tuft can exhibit non-homogenous branch-specific events^{15,56,57}.

To confirm that our interpretation was not due to segmentation errors, we repeated the conditioning experiment using a new, high-speed volumetric imaging approach called SCAPE^{36,37}, which allowed us to monitor calcium across entire apical tufts (Supplementary Movie 2). These three-dimensional datasets (300 × 1050 × 234 μm field of view) encompassed large portions of the apical tree which included branches converging on their bifurcation points in layer 2, enabling us to identify whole apical trees unambiguously (Fig.4A,B; Supp.Fig.4B).

CaImAn effectively demixed overlapping trees in these three-dimensional volumes. Using SCAPE microscopy, we imaged tuft activity of two additional mice conditioned with the same behavioral paradigm (Fig.4C). Comparison of tuft responses to the CS+ and CS- on the pre, last-rewarded, and post sessions (Fig.4D; 241 pre, 215 last-rewarded, 150 post tufts in 2 mice) revealed again that task learning induced significant increases in tuft selectivity (Fig.4E; pre versus last-rewarded: $p < 10^{-5}$, pre versus post: $p < 10^{-4}$, Wilcoxon rank sum test of $|SI|$). On average, the SI magnitudes were similar between tufts imaged using 2-photon microscopy and SCAPE (mean ± s.e.m. $|SI|$ for 2-photon versus SCAPE; pre: 0.41±0.01 versus 0.40±0.02; last-rewarded: 0.54±0.02 versus 0.54±0.02; post: 0.51±0.02 versus 0.53±0.03). These data demonstrate that the effects in our two-photon dataset are not caused by errors in segmentation, but rather reflect changes at the level of individual dendritic tufts. Our results, based on two different imaging approaches, clearly demonstrate that reinforcement increases stimulus selectivity at the level of the entire apical tuft.

Selective tufts emerge from both initially unresponsive and responsive populations

The striking effect of reinforcement learning on tuft response selectivity could develop in several ways. For example, initially unresponsive tufts could develop a robust response to either stimulus after learning (e.g., Fig.5A, top). Conceivably, tufts that were initially unselective in naïve animals could also maintain their response to one stimulus while losing their response to the other (e.g., Fig.5A, middle). Either or both scenarios could lead to the increase in neurons that are selective for stimulus direction. To investigate which changes in individual tufts underlie population-wide improvements in stimulus selectivity, we longitudinally tracked the same set of tufts across all sessions and compared their selectivity in pre- and post-conditioning sessions for both conditioned and repeated exposure mice.

First, we categorized tufts that were unresponsive to either stimulus on the first imaging session, which accounted for the large majority of tufts (Fig.5E; conditioned: 458/603; repeated exposure: 334/457), and compared their response to the CS+ and CS- on the last session to determine if they became selective (Fig.5B, see Methods). Stimulus-unresponsive tufts, while on average less active than responsive ones (median calcium events per minute: 2.65 versus 3.66 for stimulus-unresponsive and responsive tufts, respectively; $p < 10^{-40}$, Wilcoxon rank sum test; Supplementary Fig.4), were not silent, with many undergoing tuft-wide calcium influx several times per minute. Silent tufts that are never active during the session may not have been detected in our imaging, but we were able to detect tufts that discharged as few as 3 voltage-gated calcium spikes over a 30-minute behavioral session. Interestingly, in both the conditioned and repeated exposure mice, approximately 40% of initially unresponsive tufts developed a response to at least one stimulus by the last session, becoming either selective or unselective (Fig.5B). However, in conditioned animals, the proportion of initially unresponsive tufts that became selective was significantly larger than in repeated exposure mice (Fig.5B; $p = 0.04$, 2-sample t-test comparing




Fig. 4

High-speed volumetric imaging of apical tufts confirms the emergence of enhanced selectivity after learning.

a, Top and side view of four example tufts segmented from volumetric SCAPE imaging. **b**, Time courses of calcium activity from example tufts in **a** during five presentations of air puff stimuli (dashes). **c**, Performance across all conditioning sessions of two mice that were imaged with SCAPE. **d**, Across the indicated sessions, individual SCAPE-imaged tufts (circles) exhibit larger biases to CS+ or CS-. **e**, Conditioning reshapes selectivity distribution from Normal to uniform.




Fig. 5

Longitudinal tracking reveals that reward enhances the selectivity of both initially unresponsive and responsive tufts.

a, Three example tufts that were longitudinally tracked across learning. Top row: An initially unresponsive tuft develops a robust response to the CS+ but not the CS- after learning. Middle row: A responsive but unselective tuft loses its robust CS+ response and becomes selective for the CS-. Bottom row: A CS-selective neuron becomes unresponsive to both stimuli. **b**, Tufts that were unresponsive during the first session were longitudinally tracked to the last session. Plotted is the mean proportion of selective and unselective neurons across all animals in the conditioned (black bars) and repeated exposure (grey bars) groups. **c,d**, Same analysis as **b** for initially selective (**c**) and unselective (**d**) tufts. Two-sample t-test was used for comparisons between conditioned and repeated exposure groups. Paired t-test was used for comparisons within a group. * $p < 0.05$. **e**, Total tuft counts from first to last session within the 3 response categories for either conditioned (left) or repeated exposure (right) groups. **f**, SI of responsive tufts on the last session that were initially unresponsive during the first session. Conditioned tufts have enhanced selectivity compared to repeated exposure. **g**, Tufts that were selective on the last session are more selective if conditioned (black) rather than undergoing repeated exposure (grey). **h**, Tufts that responded on both pre and post sessions tend to have higher selectivity if conditioned rather than undergoing repeated exposure. **i**, SI of responsive tufts on the first session that later became unresponsive during the last session.

mice). Furthermore, while the proportion of selective and unselective tufts in this category was similar for conditioned animals, unselective tufts were more common in repeated exposure mice (**Fig.5B** [↗](#); $p = 0.03$, paired t-test).

Next, we analyzed tufts that were initially responsive and either selective (**Fig.5C** [↗](#); conditioned: 56/603, RE: 43/457) or unselective (**Fig.5D** [↗](#); conditioned: 89/603, repeated exposure: 80/457). In these smaller categories, we found no significant differences in the outcome of selectivity between the two groups of animals. Together, these results indicate that, while both stimulus exposure and reinforcement can alter tuft tuning, the presence of reward increases the likelihood that initially unresponsive tufts develop selectivity for either the CS+ or CS- (summarized in **Fig.5E** [↗](#)).

While a greater proportion of tufts from the conditioned animals were selective during the final session (20.2% versus 10.3% of tufts from conditioned and repeated exposure mice, respectively), we wondered whether conditioning also impacted the degree of selectivity. Note that some tufts had very small yet statistically different CS+ and CS-response amplitudes and were thus classified as selective despite a small SI. First, we compared the SI of initially unresponsive tufts on the final imaging session (**Fig. 5F** [↗](#)). Supporting our results in **Fig. 5B** [↗](#), the SI distribution was shifted toward the tails in conditioned, but not repeated exposure mice, indicating that reward enhances selectivity for either the CS+ or CS- in this subset ($|SI|$ conditioned versus repeated exposure: $p < 10^{-5}$ [↗](#), Wilcoxon rank sum test, $n = 199$ and 110 tufts, respectively).

Next, we compared the $|SI|$ of all tufts that were categorized as selective during the last imaging session in conditioned and repeated exposure mice (**Fig. 5G** [↗](#)). Interestingly, we found that even among selective tufts, the $|SI|$ distribution in conditioned mice was significantly greater than in repeated exposure mice ($p = 0.006$, Wilcoxon rank sum test, $n = 122$ and 47 tufts, respectively), indicating that while selective tufts are present after both conditioning and repeated stimulus exposure, the magnitude of selectivity is stronger after conditioning.

We then quantified the change in $|SI|$ of all tufts that were responsive in both the first and last sessions by computing the difference between the two sessions (**Fig. 5H** [↗](#)). Tufts in conditioned mice exhibited a greater increase in $|SI|$ across sessions compared to repeated exposure mice ($p = 0.01$, Wilcoxon rank sum test, $n = 48$ and 42 tufts, respectively), demonstrating that the magnitude of selectivity in initially responsive tufts increases after reinforcement learning.

Finally, we found that the degree of selectivity of tufts that eventually became unresponsive on the last session was overall similar between the two groups (**Fig.5I** [↗](#), $|SI|$ conditioned versus repeated exposure: $p = 0.06$, Wilcoxon rank sum test, $n = 97$ and 81 tufts, respectively). However, tufts that became unresponsive were more likely to be initially highly selective in the conditioned group than in the repeated exposure group (19 tufts with initial $|SI| > 0.75$ / 97 tufts ending as unresponsive in the conditioned group versus 3/81 in the repeated exposure group; $p = 0.0013$, Z approximation to binomial). Therefore, learning can involve a loss of responsivity in a small subset of well-tuned tufts.

In summary, our longitudinal analyses revealed that reinforcement learning biases initially unresponsive tufts toward becoming selective and enhances the selectivity of tufts that are initially responsive.

Neither movement nor behavioral choice account for enhanced selectivity

Several plausible factors could underlie the changes in selectivity we observed across learning. For instance, movements like whisking are correlated with layer 5 somatic action potentials⁵⁸ [↗](#)–⁶⁰ [↗](#) and might have impacted calcium activity in the apical tuft. To investigate whether whisking could account for the changes in tuft selectivity, we imaged the whiskers with a high-speed camera and

computed whisking amplitude (see Methods) while mice underwent conditioning and two-photon imaging (**Fig.6A**). First, we considered whether animals changed their whisker movements in response to conditioned stimuli over the course of learning. We computed the peak of the mean stimulus-aligned whisking amplitude for the CS+ and CS- (**Fig.1C**, left; **Fig.6B**) for each session in five mice. Although conditioning alters licking behavior (**Fig.1C**, E), the magnitudes of whisker movements following both stimuli were stable across sessions (**Fig.6B**; CS+: $p = 0.44$; CS-: $p = 0.45$; linear regression). We also computed the standard deviation (SD) of stimulus-evoked whisker amplitude across trials for all sessions (**Fig. 6C**). While the whisking amplitude became slightly more reliable (decreased SD) across sessions ($p < 10^{-4}$), the change in reliability across sessions was similar for CS+ and CS- ($p = 0.53$). Therefore, whisking is similar on both trial types throughout learning.

We next examined whether whisking was correlated with tuft calcium activity by comparing stimulus-triggered averages and intertrial interval (ITI) whisk-triggered averages of all tufts during post-conditioning. Whisking amplitude was similar between spontaneous ITI whisking bouts and evoked whisking responses to stimuli ($n = 115$ and 617 events, respectively; $p = 0.53$, Wilcoxon rank sum test). In contrast to air puff stimuli, ITI whisking bouts were not associated with a robust calcium response (**Fig.6D**).

To quantify the relationship of whisking and sensory stimuli to tuft calcium spikes, we performed a linear regression analysis (see Methods) on 322 tufts using calcium influx as the response variable and either stimulus or whisking amplitude as a single predictor variable (**Fig.6E**). Air puff stimuli more reliably predicted calcium influx than whisking amplitude for each of virtually all tufts ($p < 10^{-12}$, sign rank test). These results are consistent with other studies that found either only weak or no correlation between whisking and L5 tuft calcium spikes in S1^{28,31,32}. Furthermore, we found no relationship between the whisking response and the median SI magnitude on a given session (**Fig.6F**, whisking to CS+ $p = 0.22$, CS- $p = 0.78$). Therefore, changes in whisker movement cannot account for the changes in selectivity during learning that we observed.

Finally, the possibility remains that other task-related signals relaying information about reward expectation and behavioral choice could impact apical tuft activity and drive increases in selectivity. To test this, we compared tuft responses to the CS-in false alarm trials (FA; mouse incorrectly licked for reward) and correct rejection trials (CR; mouse correctly withheld licks) to determine if their activity was modulated by behavioral choice. Notice that these two trial types have the same sensory input but involve different choices. (The corresponding analysis for CS+ trials is not technically possible for lack of sufficient Miss trials after the first conditioning day, an issue also observed in¹. A future experiment in which the stimulus strengths are substantially reduced would drastically increase the error rates, enabling a comparison between Hit and Miss trials.) Tufts were classified as behaviorally modulated if the FA response was significantly different from the CR response, and were not behaviorally modulated if CR and FA responses were statistically indistinguishable (e.g. **Fig.7A**). Behaviorally modulated tufts accounted for only ~10% of the total tuft population in both early and late learning (50/395 in early; 35/406 in late learning).

To test whether these behaviorally modulated tufts contributed to increased selectivity during learning, we excluded them and compared selectivity of the remaining behaviorally-insensitive tufts. We found that selectivity increased significantly from early to late learning (**Fig.7B**, C; median $|SI|$ of 345 tufts early versus 371 tufts late learning: 0.38 versus 0.47, $p = 0.02$, Wilcoxon rank sum test), similar to our previous analysis of the entire population. Licking, like whisking, was a relatively poor predictor of tuft calcium influx (Supp.Fig.6A,B). Because some behaviorally modulated tufts may not have been statistically detectable, we used multivariate linear regression to disentangle stimulus responses from licking and whisking, which may have been confounded with choice. Median coefficients for licking and whisking were on average 3.3 times smaller than




Fig. 6

Whisking is only weakly correlated with tuft activity and cannot account for changes in selectivity during learning.

a, Whisking amplitude aligned to calcium activity of three example tufts in one session. Green shading indicates periods of whisking. Red and navy ticks indicate CS+ or CS-delivery, respectively. **b**, Mean whisking response of five mice to CS+ (red) and CS- (navy) does not change across sessions during learning (mean $\pm$ s.e.m.). **c**, Mean standard deviation of whisking decreases for both CS+ and CS- across learning, but CS+ and CS- do not differ. **d**, Event-triggered averages of 322 tufts on the post-conditioning day (grey traces - individual tufts, black inset - population average) are responsive to stimuli but relatively unmodulated by whisking. **e**, R^2 values for linear models predicting calcium from stimuli (y axis) are consistently greater than those predicting calcium from whisking (x axis). Each circle represents a tuft. ($n = 322$ tufts) **f**, Magnitude of tuft selectivity does not correlate with mean whisking amplitude during CS+ (left) and CS- trials (right) on that session.




Fig. 7

Behavioral responses do not account for enhancement of stimulus selectivity during learning.

a, Mean stimulus responses of four tufts during hit (red), CR (cyan), and FA (black) trials. Top row: Example tufts whose responses are not behaviorally modulated (CR is similar to FA). Bottom row: Example tufts with behaviorally modulated responses (CR and FA differ). **b**, Selectivity index (SI) distribution changes from early (left) and late learning sessions (right) even when tufts with behaviorally modulated responses ($CR \neq FA$) are excluded. **c**, Median SI magnitude of tufts in each of six animals (from panel **b**) increases from early to late learning sessions.

median stimulus coefficients for the first rewarded, last rewarded, and post sessions (all $p < 10^{-6}$, Wilcoxon rank sum test). Even after we factored out possible effects of movements, CS+ and CS- coefficients were enhanced by learning but not repeated exposure (Supp.Fig.6C,D), consistent with our other analyses. Together, these results demonstrate that enhanced selectivity during learning cannot be explained by non-sensory signals related to the animals' behavior.

Enhanced selectivity in barrel cortex is long-lasting when mice exclusively use whiskers

Mice could conceivably exploit other sensory cues to learn and perform the task, such as auditory cues from the air nozzles or non-whisker tactile cues from air current eddies contacting the fur or skin. To determine which mice exclusively used their whiskers to distinguish the CS+ and CS-, we trimmed all whiskers after the post-conditioning session and assessed performance in five mice (Figure 8). Performance in each of the five mice decreased after whisker trimming, indicating that each used some whisker information. Three mice performed the task exclusively with their whiskers, falling to chance levels after the whisker trim ("whiskers only"). Two other mice still performed the task above chance after the whisker trim, indicating that they were not exclusively using their whiskers and exploited information from multiple sensory streams ("whiskers + other senses").

We examined whether these two different behavioral strategies impacted tuft selectivity. Both the "whiskers only" and "whiskers + other senses" groups exhibited enhanced tuft selectivity in the last-rewarded session relative to pre-conditioning. This effect was more pronounced in the "whiskers only" mice (Fig.8A,B, left and middle; whiskers only: median $|SI|$ of 180 pre tufts versus 169 last-rewarded tufts: 0.36 versus 0.59, $p < 10^{-3}$; "whiskers + other senses": median $|SI|$ of 144 pre tufts versus 155 last-rewarded tufts: 0.39 versus 0.50, $p = 0.01$). Surprisingly, enhanced selectivity persisted during the post-conditioning session for the "whiskers only" group but not the "whiskers + other senses" group (Fig.8A,B right panels; whiskers only: median $|SI|$ of pre versus 167 tufts post: 0.36 versus 0.58; $p < 10^{-3}$; whiskers + other senses: median $|SI|$ of 155 pre versus post tufts: 0.39 versus 0.42; $p = 0.45$). Therefore, tuft selectivity in barrel cortex is enhanced regardless of behavioral strategy, but outlasts conditioning only when mice rely solely on their whiskers to perform the task.

We further examined this persistence of enhanced tuft selectivity as experienced mice stopped performing the task. While the entire post-conditioning session was unrewarded, mice initially expected rewards and licked for many CS+ trials in the first half of the session. By the second half of the session, the probability of a lick occurring during the CS+ extinguished, approaching zero (Fig.8C). We compared the selectivity of tufts during the first and second halves of the post-conditioning sessions of mice that exclusively used their whiskers and found no difference in the two distributions (Fig.8D, $p = 0.94$, Wilcoxon rank sum test of $|SI|$), demonstrating that selectivity of the population remained stable throughout the session. Taken together, these results demonstrate that enhanced stimulus selectivity of apical tuft dendrites after reinforcement learning is long lasting, persisting even after mice cease performing the task and expecting reward.

Discussion

Our study is the first to investigate how learning a discrimination task alters apical tuft activity. Using both novel volumetric whole-tuft imaging and conventional planar microscopy, we discovered that L5 apical tufts acquire enhanced representations of multiple stimuli during learning. Rather than simply retuning tufts toward the rewarded stimulus, learning enhanced selectivity for both stimuli, suggesting that tufts are aligning themselves to the behaviorally relevant stimulus dimensions. These enhanced sensory representations persist even after mice




Fig. 8

Apical tufts in barrel cortex of mice performing the task exclusively with their whiskers undergo long-lasting changes in selectivity.

a, SI histograms of mice performing the task exclusively with their whiskers exhibit increased selectivity across pre-conditioning, last-rewarded, and post-conditioning sessions. **b**, Relative to pre-conditioning, mice using their whiskers and other sensory cues to perform the task have increased selectivity during the last rewarded session, but not the post-conditioning session. **c**, The probability of anticipatory licks in response to the CS+ extinguishes across post-conditioning blocks (of 20 trials each). **d**, Tuft selectivity remains uniformly distributed during post-conditioning trial blocks 1-2 (top) while licking is extinguishing, and blocks 3-4 (bottom) in which licking is extinguished.

cease performing the task. In contrast, representations are slightly degraded by mere repeated exposure to stimuli outside of a task. Consistent with previous studies^{28,31}, we found that movement in and of itself has little direct impact on tuft spikes, indicating that increased selectivity of apicals reflects alterations in sensory coding as animals learn. This sensitization of tufts to behaviorally relevant sensory dimensions may be a general feature of all sensory cortical areas.

Tuft spikes enhance plasticity of synaptic inputs that occur over behavioral (seconds-long) timescales^{18,34}. These new behaviorally relevant tuft representations may therefore prime subsequent plasticity of synapses across the entire pyramidal neuron. Additionally, tuft events potentially modulate somatic burst firing and enhance how somata respond to their basal inputs^{15,61}. As learning and plasticity increase apical selectivity for a behaviorally relevant axis, tuft events will unavoidably amplify somatic burst output along the same axis. This could enable action potential output of L5 cells in primary sensory cortex to directly drive behavioral responses via projections to movement related areas, such as the corticostriatal, corticopontine, and corticotrigeminal pathways. Thus, tuft spikes have the potential to modify somatic output, both in the present and in the future.

An open question is whether enhanced stimulus representations in apical tufts are required for learning this task. One way to address this question would be to silence tuft activity during and after learning by optogenetically activating NDNF-positive interneurons in layer 1⁶². This approach is not ideal as NDNF interneurons also inhibit other cells such as Layer 2/3 pyramidal cells, PV interneurons⁶³ and possibly the axons of Layer 5 pyramidal cells, which are known to densely innervate layer 1. Because this manipulation is not specific to layer 5 apicals, the results would be difficult to interpret. Focal illumination of inhibitory opsins in tufts has also been used to assess tuft function⁶⁴, but balancing tuft against soma silencing remains challenging and complicates interpretation. Better tools for selective targeting of apicals would be extremely useful for addressing such issues.

Enhanced Representation of Behaviorally Relevant Stimuli

Enhancing the representation of relevant stimulus dimensions rather than a singularly important stimulus, such as a rewarded event, has multiple benefits for behavior. In our paradigm, both the CS+ and CS- are predictive of whether or not a reward will occur in the future. Explicitly encoding both stimuli could allow sensory cortical areas to directly elicit actions. In the context of this task, CS+ preferring tufts in barrel cortex may trigger anticipatory licking while CS- preferring tufts could suppress licking. L5 cells in sensory cortex via their output to striatum, pons, brain stem, and spinal cord would thereby be able to directly and rapidly drive action without further cortical processing, such as by frontal areas including motor cortex^{32,65}. Such rapid sensory-motor transformations by primary sensory areas may be critical for natural time-constrained behavior.

Furthermore, learning produced a representation in which the degree of selectivity for the two stimuli was continuous and uniformly distributed. Exclusively CS+ or CS- selective apicals never dominated the population. Continuous degrees of selectivity across the population, rather than discrete representations, may allow the system to be more robust to the variability caused by active movements that alter sensory input. A continuous distribution may also facilitate future adjustments of neural representations as subjects continue to learn a task or encounter new tasks. The uniformity we observed may reflect that neurons are high-dimensional, being sensitive to mixtures of variables^{60,66–68}, only one of which might be altered here by learning. The uniform distribution of selectivity corresponds to a full range of pessimism to optimism concerning stimulus predictions of upcoming rewards. Recent work shows that behavioral performance benefits from reinforcement learning that incorporates the distribution of reward probabilities rather than just the average expected reward value⁶⁹. L5 corticostriatal synapses could theoretically afford a plastic substrate for acquiring the necessary distribution of reward probabilities.

Surprisingly, past studies in which mice were trained to associate one or more stimuli with a reward typically show that cortical representations are stronger for the rewarded stimulus [1](#), [3](#), [5](#). In contrast to these studies of layer 2/3 somatic activity, our experiments revealed that the overall tuft calcium response to the CS+ and CS- at the population level did not change significantly after animals learned the task (**Fig.2**). Instead, representations for both stimuli were enhanced by individual tufts developing selectivity for either the CS+ or the CS- (**Fig.3**). This divergence in phenomena may result from several important differences between our work and the aforementioned studies.

First, enhanced selectivity for both rewarded and unrewarded stimuli could be a phenomenon that is unique to the apical dendritic tufts. In addition to local inputs, the apical tufts of pyramidal cells in S1 receive long-range top-down input from several sources, including motor cortex [31](#), [70](#), secondary somatosensory cortex [11](#), and secondary thalamus [9](#), [10](#), [71](#). Frontal areas, such as prefrontal cortex, indeed have enhanced representations of the CS+ and CS- after learning [47](#). In contrast, input to the somata is dominated by the local cortical area and primary thalamus [72](#), [73](#). While somato-dendritic coupling can be strong in L5 neurons [25](#), it is asymmetric; at least 40% of somatic transients attenuate in a distance-dependent manner along the apical trunk and distal tufts [24](#). The non-overlapping anatomical inputs and asymmetric coupling together could produce different learning-related effects on apical tuft and somatic stimulus representations.

Second, learning-related changes may manifest differently in layer 2/3, the usual focus of previous studies [1](#), [3](#), and layer 5 pyramidal cells, the tufts of which we studied. With the exception of a small population of corticostriatal cells, most excitatory cells in layer 2/3 project to other cortical areas to affect further cortical processing [74](#), [75](#). In contrast, many L5 cells project to subcortical structures including the thalamus, superior colliculus, and brainstem, which may directly trigger behavioral responses [76](#)–[78](#). In discrimination paradigms, both stimuli are relevant to behavior. In our task, the CS+ prompted licking to obtain a reward, and the CS- suppressed licking that would have no benefit. Thus, an enhanced representation of both stimuli in layer 5 would be advantageous for animals to perform the task efficiently. Recently, it was shown that apical dendrite activation of subcortical-targeting pyramidal tract L5 cells, but not intratelencephalic L5 cells that are more like L2/3 cells in their connectivity, determines the detection of tactile stimuli [32](#). The Rbp4-Cre mice we used in this study labels a heterogeneous population of layer 5 pyramidal cells, comprising both pyramidal tract and intratelencephalic neurons. In the future, it would be interesting to examine whether learning has different effects on the sensory representations of these two populations. Moreover, direct comparisons of the layers would be particularly informative.

Finally, it is possible that learning-related changes in sensory representations manifest differently between a somatosensory modality and a visual modality, the latter being the focus of previous studies. To our knowledge, we are the first to show changes of sensory representations in somatosensory cortex within a discrimination paradigm. Mice are known to rely more heavily on their tactile senses than vision [79](#). Their heavy reliance on whisker-mediated touch may make it advantageous to develop sensory representations of a larger variety of relevant tactile stimuli, in this case, both the CS+ and CS-.

Candidate Plasticity Mechanisms

Enhanced selectivity could be due to changes in local synaptic connectivity, long-range inputs, or both. Learning may strengthen and weaken synapses onto barrel cortex neurons from ascending thalamocortical input or from neighboring cells. Such local plasticity could enhance CS+ or CS- responsiveness. Alternatively or additionally, other cortical regions encoding task context could via long-range inputs reconfigure barrel cortex to respond more strongly to these stimuli. The present results do not completely distinguish between these two scenarios because long-range inputs may still encode the context while the mouse is in the behavioral apparatus. However, we

found that enhanced representations persist after mice are no longer engaged in the task and receiving rewards. This result suggests that enhanced representations may be a product of local plasticity in sensory cortex that alters receptive fields.

Even in the absence of reward, repeated exposure to stimuli can drive plasticity in sensory cortex and alter response tuning. For instance, repeated exposure to oriented gratings can alter the orientation tuning of cells in primary visual cortex^{51–53}, and overstimulation of whiskers induces plasticity at dendritic spines and alters whisker representations in somatosensory cortex^{54,80}. Our results demonstrate that at the population level enhanced representations developed only when stimuli were behaviorally relevant. Our longitudinal analysis revealed that while the response dynamics of some tufts changed after repeated stimuli presentations, overall selectivity of the population did not increase when rewards were omitted (Figs.3&5). This raises the question: What are the mechanisms that drive enhanced selectivity under rewarded conditions? In one possible scenario, reward delivery causes the release of neuromodulators that augment the activity of apical tufts. Cortical layer 1 is innervated by cholinergic afferents from the nucleus basalis⁸¹ and adrenergic afferents from the locus coeruleus⁸², the main source of acetylcholine and norepinephrine, respectively. Salient events such as reward and arousal lead to the release of these neuromodulators in cortex^{83,84}, which could increase the excitability of apical dendrites by recruiting disinhibitory circuits or directly influencing dendritic currents^{26,27,83,85}. In this model, the release of reward-driven neuromodulators promotes plasticity and an enhanced representation of temporally aligned sensory inputs. This phenomenon was demonstrated in auditory cortex, where tones paired with stimulation of the nucleus basalis shifted the tuning of neurons toward the frequency of the paired stimulus⁸⁶.

Why are representations of the CS-equally enhanced when there is no associated reward? One explanation is that, as mice learn that the CS-indicates absence of reward, the CS-effectively signals punishment and acquires negative value. Acetylcholine is released in response to aversive stimuli, and can activate disinhibitory microcircuits that reduce inhibition onto pyramidal cells and may be essential for learning^{87,88}. Thus, it is possible that both the CS+ and CS-representations are enhanced by neuromodulatory mechanisms tied to reward and punishment, respectively. An open question is whether the outcome is due to reinforcement learning or the behavioral state brought on by the reinforcers rather than their valence. Sensory cortical plasticity may not be tied to reinforcer valence. Our paradigm creates an environment where mice benefit from being attentive and engaged in order to maximize reward while minimizing effort. Previous work has shown that active engagement in a visual discrimination task was associated with significantly higher selectivity in layer 2/3 cells in visual cortex¹. Task engagement may lead to a sustained increase in neuromodulator release throughout the conditioning session, priming the apical dendrites for plasticity and the development of selective responses for task-relevant stimuli as they learn.

What determines whether a particular tuft eventually becomes selective for the CS+ or CS-? Our longitudinal analysis revealed that many tufts that were initially unresponsive to either stimulus developed a highly selective response to either the CS+ or the CS- (**Fig.5**). In these tufts, stimulus preference after learning might be seeded by initially weak, directionally selective inputs on to the neuron that already exist prior to conditioning and that are potentiated by the learning process. We also found tufts that initially exhibited robust responses to both stimuli and either lost or significantly reduced their response to one stimulus after learning. The reduction of an apical response to a particular stimulus could be driven by local disinaptic inhibition between L5 pyramidal cells mediated by the apical-targeting Martinotti cells^{89–91}. Through this mechanism, L5 neurons that are selective for a particular stimulus could inhibit responses to that stimulus in neighboring L5 apical tufts. Experiments that assess the tuning of excitatory and inhibitory inputs onto apical dendrites as a function of learning could test such mechanisms.

In addition to demonstrating increased tuft selectivity with learning, we replicated a surprising phenomenon in a previous instrumental behavior in which a population of apical tufts exhibit activity around the time of reward²⁸. This reward-related activity was observed in four out of the seven conditioned animals only during CS+ trials and was most prominent during intermediate conditioning sessions, when most animals were still performing at chance levels, and disappeared completely by the final conditioning session (Supp.Fig.1). Other than this transient effect, unconditioned stimuli did not appear to elicit calcium responses, consistent with our previous findings²⁸. The disappearance of this reward-related peak might be attributable to the reward becoming predictable in later stages of learning. In previous classical conditioning experiments, dopaminergic cells exhibit responses to rewards early in learning due to the novelty of an unexpected stimulus. These responses are lost after extended training, as animals learn the association between the CS and reward^{92,93}. While dopaminergic terminals are sparse in primary sensory areas, they are not entirely absent, nor are dopaminergic receptors. Furthermore, the excitability of the apical tuft is sensitive to noradrenaline²⁶. Interestingly, noradrenergic neurons in the locus coeruleus exhibit a similar phenomenon to dopaminergic neurons, where responses shift from temporal alignment with the reward to a predictive conditioned stimulus after learning⁹⁴. Such mechanisms could explain why reward-related activity is restricted to early-to-intermediate learning in our paradigm.

Global versus local dendritic spikes

Apical dendrites exhibit not only global spikes that elicit calcium influx across the entire tuft, which we exclusively analyzed here, but also local events known as NMDA spikes, which typically engage short (<30-μm) segments of individual dendritic branches^{15,31,57}. These local, NMDA receptor-dependent events can promote prolonged plasticity within individual dendritic branches in the absence of backpropagating actions potentials, a feature that is unique to the apical dendrites¹⁶. In motor cortex, branch-specific NMDA spikes are crucial for establishing the long-lasting plasticity necessary for learning⁵⁶, and depolarization provided by multiple local NMDA spikes is thought to be essential for the generation of a global calcium spike triggered by distal synaptic inputs¹⁵. We focused this study on global tuft-wide calcium events, rather than local events. Local events are more difficult to unambiguously identify in planar imaging⁹⁵, and their existence *in vivo* is still an open question for L5 apicals in barrel cortex^{31,57}. Nonetheless, they may play important roles in plasticity processes that eventually lead to the emergence of global tuft spike selectivity for stimuli. Volumetric microscopy studies, the feasibility of which we showed here, are needed to further investigate the existence of local events in such behaviors as well as examine possible relationships between local and global tuft events during reinforcement learning. However, it would be essential to verify that seemingly spatially overlapping local and global events derive from the same dendritic tree, which requires greater resolution than was practical for the present study.

To analyze activity of individual tufts, we segmented these structures based on spatiotemporal covariance⁴⁵. This method does not discount the possibility of errors where one tuft is split erroneously into two trees, or where two highly correlated tufts are merged. With this in mind, we used volumetric imaging SCAPE microscopy, which allowed us to visualize the apicals in three dimensions and unambiguously screen for such artifacts. The results from SCAPE are quantitatively similar to those from two-photon microscopy, and confirm that our observation of enhanced selectivity with learning is not an artifact of planar imaging.

Stability of learned tuft representations

In contrast to previous studies of discrimination learning^{1,3}, we included an unrewarded post-conditioning session to examine whether learning-related effects persisted through extinction. Our results show that post-conditioning selectivity of the apical population remains significantly higher than pre-conditioning, even after animals stop licking in response to the CS+ (Fig.8).

Interestingly, the effects of learning are much more pronounced in animals that relied exclusively on their whiskers to perform the task. In animals that apparently used other sensory modalities, we observed a modest increase from the pre to last-rewarded session, which seemed to be largely absent by the post-conditioning session. Considering that these animals were additionally exploiting other sensory areas to perform, selectivity may have been more widely distributed and thus diluted in barrel cortex, diminishing the effect and its stability. How long selectivity persists in the neuronal population after conditioning and which factors influence stability are interesting open questions for future study.

Conclusion

In summary, we have shown for the first time that reinforcement learning enhances representations along behaviorally relevant dimensions in apical tufts. Our results suggest that dendritic calcium spikes are an important cellular mechanism underlying the changes in sensory encoding that occur with learning, and provide an avenue for further investigation of cellular and circuit mechanisms underlying plasticity induced by perceptual experience and reinforcement. This cellular compartment may be key to understanding pathology in some cognitive, memory, and learning disorders.

Additional information

Acknowledgements

We thank Venkatakaushik Voleti for help with the design, construction, and alignment of the SCAPE microscope; Dan Kato, Georgia Pierce, and Jung Park for help with pilot experiments; Eftychios Pnevmatikakis and Johannes Friedrich for advice on dendrite segmentation; and Larry Abbott, Stefano Fusi, Ashok Litwin-Kumar, Chris Rodgers, Georgia Pierce, Gordon Petty, and Dan Kato for comments on the manuscript. Funding was provided by a Wellcome Trust Discovery Award, an Academy of Medical Sciences Professorship, NIH/NINDS R01 NS069679, and NIH/NINDS R01 NS094659 (RMB); a Kavli Institute for Brain Science Postdoctoral Fellowship (SEB); NIH/NINDS/NIMH/BRAIN U01 NS094296, UF1 NS108213, U19 NS104649, and RF1 MH114276 (EMCH).

Author contributions

SEB and RMB conceived of the behavioral and two-photon imaging experiments. EMCH and RMB conceived of the SCAPE imaging experiments. SEB built the behavioral apparatus, EMCH, KBP, and CC designed, built, and maintained the SCAPE microscope, and RMB built and maintained the two-photon and intrinsic signal microscopes. SEB performed the experiments and analyzed the data with input from RMB and EMCH. SEB and RMB wrote the manuscript.

Data availability

Due to the large volume of data (~80TB), data are maintained by the authors and available upon request.

Methods

All experiments complied with the NIH Guide for the Care and Use of Laboratory Animals and were approved by the Institutional Animal Care and Use Committee of Columbia University.

Sixteen C57BL/6 mice ranging in age from 77 to 316 days old (mean of 123 days at the time of imaging) were used in these experiments. Six were male, and 10 female. Our results were observed in both male and female individuals, and no sex difference was detected.

Surgery

Animals were administered dexamethasone (1 mg/kg) via intramuscular injection 1-4 hours prior to surgery to reduce edema. Anesthesia was induced with 3% isoflurane in oxygen and maintained at 1%. Mice were head-fixed in a stereotax, and a subcutaneous injection of bupivacaine (0.5%, 0.1 mL) was administered under the scalp. Buprenorphine (0.05 mg/kg) was injected subcutaneously on the back. The scalp was cut, and the skull was covered with a thin layer of Vetbond. A circular craniotomy (3-mm diameter) centered at 1.5 mm posterior and 3.5 mm lateral to bregma was made using a dental drill. The dura was kept moist using artificial cerebrospinal fluid.

For both two-photon and SCAPE microscopy, Rbp4-Cre_KL100 mice were injected with 100 nL of virus (initial titer $\sim 2 \times 10^{13}$ cfu/mL, diluted 1:4 in artificial cerebrospinal fluid) encoding GCaMP6f in a Cre recombinase-specific manner (AAV1-CAG-flex-GCaMP6f, UPenn Vector Core). The virus was injected in layer 5B of the barrel cortex (1.0 mm deep to the pia) using a pulled pipette (20-30 μ m ID) fastened on a Nanoject III, which was mounted on a manipulator angled at $\sim 30^\circ$ from vertical. The virus was delivered via four injections of 100 nL each, spaced at least 400 μ m apart. The depth was chosen to maximize labeling of thick-tufted pyramidal neurons. In pilot experiments, we found that placing injections 1.0 mm deep resulted primarily in thick-tufted labeling whereas at more superficial depths (e.g., 0.8 mm deep) we obtained mainly thin-tufted tufts, consistent with ref ⁹⁶. The dura was then removed, and a thin cover glass was implanted and sealed using superglue. A custom metal head plate was implanted on the skull using dental cement. Twenty-four hours after surgery, carprofen (5 mg/kg) was administered subcutaneously. Imaging and behavioral training commenced 3 weeks after surgery.

Behavior

Animals in both rewarded ‘conditioning’ and unrewarded ‘repeated exposure’ groups were water restricted for 2 days prior to starting imaging and habituated to head fixation for ~ 10 minutes on each of these 2 days. They were subsequently given ~ 1 mL of water per day for 9 days either by pairing water rewards with a specific stimulus (conditioning group), or in their cage following the imaging session (repeated exposure group). Mice were head restrained in a custom-made behavioral apparatus by positioning the body in a 3D-printed chamber and fastening the head plate to metal posts flanking the chamber. Air puff stimuli (10 psi measured before a control solenoid, 100 ms) were delivered from two nozzles (cut P200 pipette tips) positioned toward the distal tips of the whiskers, in either the rostrocaudal or ventrodorsal direction. Nozzles were oriented to prevent air jets from stimulating other parts of the face. One of these directions (CS+) was paired with a water reward (10 μ L), delivered through a lick port 0.5 seconds after the stimulus onset. The particular direction (rostrocaudal vs ventrodorsal) used as the CS+ was randomized and counterbalanced across mice. Approximately 180 stimuli were presented over the course of a 30-minute imaging session (8-12-s intertrial interval). The probability of CS+ or CS- delivery was 50%. In preliminary experiments, we found that an auditory mask helped prevent mice from exploiting auditory cues to discriminate the two stimuli: a third air nozzle was positioned close to the mouse and was active throughout the session.

During the first session (pre-conditioning), stimuli were delivered in the absence of reward to assess neural and behavioral responses in naïve animals. In the following 7-9 days, the CS+ was paired with reward. Licks for rewards were detected with a capacitance-based touch sensor (Sparkfun). A trial response was registered when one or more licks were elicited within a 0.5-second response window following the stimulus and before reward delivery. To determine whether behavioral performance was above chance, we computed 95% confidence intervals using the ‘binofit’ function in MATLAB. During the final session (post-conditioning), stimuli were

delivered in the absence of reward. Animals in the unrewarded group received the same two stimuli across 9 days without reward pairing. Behavioral experiments were performed with the Arduino-based OpenMaze open-source behavioral system, whose designs are fully described at www.openmaze.org ⁹⁷. Whisking was monitored at 125 fps with a camera (Sony PS3eye) and automatically tracked using published software ⁹⁷.

Intrinsic signal optical imaging and two-photon imaging

Intrinsic signal optical imaging and two-photon imaging were performed on a Sutter movable objective microscope. The locations of whisker barrels in S1 were identified using intrinsic signal optical imaging. Single whiskers in isoflurane-anesthetized mice were stimulated at 5 Hz using a piezoelectric bimorph while recording the reflectance of 700-nm long-pass incandescent light with a Rolera CCD camera (QImaging) through a low-magnification objective (Zeiss 5X/0.16NA). Movies were collected using software custom-written in Labview (National Instruments). Regions of reflectance change were referenced to an image acquired under green illumination.

Two-photon imaging was conducted on the same microscope under the control of the ScanImage software package (V. Iyer, Janelia Farms). All calcium imaging data was collected by two-photon microscopy except for those in **figure 4**. Scanning during awake conditions was performed at 30 fps using a Chameleon Ultra II laser (Coherent) tuned to 920 nm, precompensated for group velocity dispersion and focused through a 20x/1.0NA water immersion lens (Zeiss). Aquasonic clear ultrasound gel was used for the immersion medium. Emitted light was collected with an HQ535/50 filter (Chroma) and GaAsP photomultiplier tubes (Hamamatsu Photonics). Apical tuft tufts in Layer 1 were imaged at depths of 40-80 μm from the pial surface (1.5x digital zoom in ScanImage which yielded a 433 x 433 μm field of view, 512 x 512 pixels).

SCAPE imaging

High-speed volumetric imaging was performed using a custom SCAPE microscope as previously described, including for dendritic tufts ^{36,37,98}. Briefly, the cortex was illuminated with an oblique light sheet through a Olympus XLUMPLFLN 20XW 1.0 NA water immersion objective with a 2-mm working distance. Fluorescence excited by this sheet (extending in the y-z' direction) was collected by the same objective lens. A galvanometer mirror in the system was positioned to both cause the oblique light sheet to scan from side to side across the sample (in the x direction) but also to de-scan returning fluorescence light. This optical path results in an intermediate, de-scanned oblique image plane that is stationary yet always co-aligned with the plane in the sample that is being illuminated by the scanning light sheet. Image rotation optics and a fast sCMOS camera (Andor Zyla 4.2+) were then focused to capture these y-z' images (750 x 200 pixels) at >1000 frames per second as the sheet was repeatedly scanned across the cortex in the x direction. All other system parts, including the objective and sample stage, were stationary during high-speed 3D image acquisition. Data were reformed into a 3D volume by stacking successive y-z' planes according to the scanning mirror's x position and de-skewing to correct for the oblique sheet angle. This rotation of the image volume is responsible for its rectangular appearance despite the camera's square frames. The resulting volumes were large enough to encompass many GCaMP6f-labeled tufts in barrel cortex,

In this study, the stationary objective lens in SCAPE was configured on a manual rotation mount and set to 20°-30° away from the standard upright configuration, so the optical axis was perpendicular to the cranial window to achieve optimal performance without tilting the head of the animal. A 488-nm laser (Coherent OBIS) was used for excitation (<10 mW at the sample) with a 500-nm long-pass filter in the emission path. To achieve optimal spatiotemporal resolution and volume rate, the sample was imaged with an x-direction scanning step of 3 μm over a 300 × 1050 × 234 μm field of view (x-y-z, 3.0 × 1.40 × 1.17 μm per voxel, 100 x 750 x 200 voxels) at 10 volumes per

second (VPS). Our imaging involves no special practical considerations or limitations of field of view or resolution, beyond the usual imaging goal of maximizing FOV while maintaining sufficient resolution to discern structures of interest (dendrites).

Analysis

Two-photon movies were motion corrected using the NormCorre package⁹⁹ in MATLAB. Spatial and temporal components for individual tufts imaged by two-photon and SCAPE were segmented using CalmAn v1.8.3, which employs large-scale sparse non-negative matrix factorization^{45,100}. CalmAn inherently corrects for background signal. All further analyses used custom-written routines implemented in MATLAB. Spatial components with tuft structural characteristics were identified and analyzed, while neuropil components were discarded.

To quantify a tuft's response to stimuli, the mean stimulus-aligned $\Delta F/F$ was computed across all CS+ or CS- trials and corrected by the mean $\Delta F/F$ of the second before the trial. Probability of transients was obtained by taking each trial's $\Delta F/F$ in the first 1.5 seconds following either the CS+ or CS- and fitting these data with a univariate mixture of two Normal distributions: $(1-p)N(\mu_1, \sigma_1) + pN(\mu_2, \sigma_2)$. The smaller Normal reflects the distribution of failures, and the larger Normal the distribution of transient amplitudes following the stimulus. The parameter p captures the probability of transients.

From these data, a selectivity index (SI) was defined as $(F_{CS+} - F_{CS-}) / (F_{CS+} + F_{CS-})$, in which F_{CS+} and F_{CS-} are the mean stimulus-aligned amplitudes ($\Delta F/F$) to the CS+ and CS- within the first 1.5 seconds, respectively. This yielded values that range from -1 (exclusively CS-responsive) to 1 (exclusively CS+ responsive). Neural discriminability was defined as $d' = |F_{CS+} - F_{CS-}| / \sqrt{(\sigma_{CS+}^2 + \sigma_{CS-}^2)/2}$ where σ_{CS+}^2 is the variance of the response amplitudes in F_{CS+} and σ_{CS-}^2 is the variance of the response amplitudes in F_{CS-} .

For longitudinal analysis, tufts were categorized as stimulus responsive if they met two criteria: 1) Across all trials, the mean $\Delta F/F$ 1.5 seconds before and 1.5 seconds after the stimulus were significantly different according to the Wilcoxon rank sum test, for either the CS+ or CS-, and 2) the average response amplitude for that stimulus was greater than 0.04 $\Delta F/F$. Tufts with a significant response to only one stimulus were categorized as highly selective and their $|SI|$ was set to 1. To classify tufts as behaviorally modulated, the mean $\Delta F/F$ of the first 1.5 seconds after the stimulus was computed for false alarm and correct rejection trials and compared with a rank sum test. Only sessions with at least 12 false alarm trials were used for this analysis. If the two distributions were significantly different, the tuft was classified as behaviorally modulated.

Custom MATLAB software was used to compute the median whisker angle, and whisking amplitude was computed as described previously¹⁰¹. The median angle was bandpass filtered from 4 to 30 Hz and passed through a Hilbert transform to calculate phase. We defined the upper and lower envelopes of the unfiltered median whisking angle as the points in the whisk cycle where phase equaled 0 (most protracted) or π (most retracted), respectively. Whisking amplitude was defined as the difference between these two envelopes. Periods of whisking were defined as times where whisking amplitude exceeded 20% of maximum for at least 250 ms. Periods of time where amplitude exceeded this threshold for less than 250 ms were considered ambiguous and excluded from analysis of whisking versus quiescence. The whisking-triggered average for each tuft was computed by aligning the calcium signal to the start times of whisking periods during inter-trial intervals (2-8 seconds after stimulus delivery).

For the linear regression analysis, we excerpted the calcium timeseries 2 seconds before and 6 seconds after each stimulus onset. The whisking amplitude signal was frame aligned to the calcium signal according to the lag of the calcium-whisking cross-correlation peak for each tuft. Whisking amplitude was then normalized to the max, yielding values that ranged from 0 to 1. The stimulus predictor variable was a binary vector with an 800-msec 'on' period (24 frames) centered

at the stimulus time. The timing of the stimulus variable was then aligned to the calcium signal according to the latency of peak of the mean $\Delta F/F$ of the first 1.5 seconds relative to the stimulus. The lick predictor variable was a binary vector with 'on' periods denoting lick bouts. Lick bouts were defined as periods of time where the mouse elicited at least 2 licks, with a maximum gap of 200 ms, and therefore had variable lengths.

For support vector machine (SVM) analysis, the mean $\Delta F/F$ was computed for a pre-stimulus epoch (1 second immediately preceding the stimulus, used as a negative control) and a post-stimulus epoch (0.1 – 1.1 seconds after the stimulus) for each trial. Binary SVMs were trained separately for each epoch using the MATLAB function `fitsvm`. For each iteration, 75% of trials were randomly chosen to train the SVM, and decoder performance was tested on the remaining 25% of trials. Decoder performance for each session was averaged across 10 iterations. All statistical tests were two-sided. T-tests were used for Normally distributed data. Otherwise non-parametric tests were applied.

References

- 1 Poort J., 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>
- 2 Liu D., et al. (2020) **Orbitofrontal control of visual cortex gain promotes visual associative learning** *Nat Commun* **11** <https://doi.org/10.1038/s41467-020-16609-7>
- 3 Henschke J. U., et al. (2020) **Reward Association Enhances Stimulus-Specific Representations in Primary Visual Cortex** *Curr Biol* **30**:1866–1880 <https://doi.org/10.1016/j.cub.2020.03.018>
- 4 Beitel R. E., Schreiner C. E., Cheung S. W., Wang X., Merzenich M. M (2003) **Reward-dependent plasticity in the primary auditory cortex of adult monkeys trained to discriminate temporally modulated signals** *Proc Natl Acad Sci U S A* **100**:11070–11075 <https://doi.org/10.1073/pnas.1334187100>
- 5 Goltstein P. M., Coffey E. B., Roelfsema P. R., Pennartz C. M (2013) **In vivo two-photon Ca²⁺ imaging reveals selective reward effects on stimulus-specific assemblies in mouse visual cortex** *J Neurosci* **33**:11540–11555 <https://doi.org/10.1523/JNEUROSCI.1341-12.2013>
- 6 David S. V., Fritz J. B., Shamma S. A (2012) **Task reward structure shapes rapid receptive field plasticity in auditory cortex** *Proc Natl Acad Sci U S A* **109**:2144–2149 <https://doi.org/10.1073/pnas.1117717109>
- 7 Fritz J., Shamma S., Elhilali M., Klein D (2003) **Rapid task-related plasticity of spectrotemporal receptive fields in primary auditory cortex** *Nat Neurosci* **6**:1216–1223 <https://doi.org/10.1038/nn1141>
- 8 Banerjee A., et al. (2020) **Value-guided remapping of sensory cortex by lateral orbitofrontal cortex** *Nature* **585**:245–250 <https://doi.org/10.1038/s41586-020-2704-z>
- 9 Zhang W., Bruno R. M (2019) **High-order thalamic inputs to primary somatosensory cortex are stronger and longer lasting than cortical inputs** *eLife* **8** <https://doi.org/10.7554/eLife.44158>
- 10 Rubio-Garrido P., Perez-de-Manzo F., Porrero C., Galazo M. J., Clasca F (2009) **Thalamic input to distal apical dendrites in neocortical layer 1 is massive and highly convergent** *Cereb Cortex* **19**:2380–2395 <https://doi.org/10.1093/cercor/bhn259>
- 11 Cauller L. J., Clancy B., Connors B. W (1998) **Backward cortical projections to primary somatosensory cortex in rats extend long horizontal axons in layer I** *J Comp Neurol* **390**:297–310
- 12 Amitai Y., Friedman A., Connors B. W., Gutnick M. J (1993) **Regenerative activity in apical dendrites of pyramidal cells in neocortex** *Cereb Cortex* **3**:26–38

- 13 Yuste R., Gutnick M. J., Saar D., Delaney K. R., Tank D. W (1994) **Ca²⁺ accumulations in dendrites of neocortical pyramidal neurons: an apical band and evidence for two functional compartments** *Neuron* **13**:23–43
- 14 Schiller J., Schiller Y., Stuart G., Sakmann B (1997) **Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons** *J Physiol* **505**:605–616
- 15 Larkum M. E., Nevian T., Sandler M., Polsky A., Schiller J (2009) **Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle** *Science* **325**:756–760 <https://doi.org/10.1126/science.1171958>
- 16 Sandler M., Shulman Y., Schiller J (2016) **A Novel Form of Local Plasticity in Tuft Dendrites of Neocortical Somatosensory Layer 5 Pyramidal Neurons** *Neuron* **90**:1028–1042 <https://doi.org/10.1016/j.neuron.2016.04.032>
- 17 Manita S., et al. (2015) **A Top-Down Cortical Circuit for Accurate Sensory Perception** *Neuron* **86**:1304–1316 <https://doi.org/10.1016/j.neuron.2015.05.006>
- 18 Roelfsema P. R., Holtmaat A (2018) **Control of synaptic plasticity in deep cortical networks** *Nat Rev Neurosci* **19**:166–180 <https://doi.org/10.1038/nrn.2018.6>
- 19 Larkum M. E., Zhu J. J (2002) **Signaling of layer 1 and whisker-evoked Ca²⁺ and Na⁺ action potentials in distal and terminal dendrites of rat neocortical pyramidal neurons in vitro and in vivo** *J Neurosci* **22**:6991–7005
- 20 Larkum M. E., Senn W., Luscher H. R (2004) **Top-down dendritic input increases the gain of layer 5 pyramidal neurons** *Cereb Cortex* **14**:1059–1070 <https://doi.org/10.1093/cercor/bhh065>
- 21 Schwindt P., Crill W (1999) **Mechanisms underlying burst and regular spiking evoked by dendritic depolarization in layer 5 cortical pyramidal neurons** *J Neurophysiol* **81**:1341–1354 <https://doi.org/10.1152/jn.1999.81.3.1341>
- 22 Larkum M. E., Zhu J. J., Sakmann B (2001) **Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons** *J Physiol* **533**:447–466
- 23 Manita S., Miyakawa H., Kitamura K., Murayama M (2017) **Dendritic Spikes in Sensory Perception** *Front Cell Neurosci* **11** <https://doi.org/10.3389/fncel.2017.00029>
- 24 Francioni V., Padamsey Z., Rochefort N. L (2019) **High and asymmetric somato-dendritic coupling of V1 layer 5 neurons independent of visual stimulation and locomotion** *Elife* **8** <https://doi.org/10.7554/eLife.49145>
- 25 Beaulieu-Laroche L., Toloza E. H. S., Brown N. J., Harnett M. T (2019) **Widespread and Highly Correlated Somato-dendritic Activity in Cortical Layer 5 Neurons** *Neuron* **103**:235–241 <https://doi.org/10.1016/j.neuron.2019.05.014>
- 26 Labarrera C., et al. (2018) **Adrenergic Modulation Regulates the Dendritic Excitability of Layer 5 Pyramidal Neurons In Vivo** *Cell Rep* **23**:1034–1044 <https://doi.org/10.1016/j.celrep.2018.03.103>
- 27 Brombas A., Fletcher L. N., Williams S. R (2014) **Activity-dependent modulation of layer 1 inhibitory neocortical circuits by acetylcholine** *J Neurosci* **34**:1932–1941 <https://doi.org/10.1523/JNEUROSCI.4470-13.2014>

- 28 Lacefield C. O., Pnevmatikakis E. A., Paninski L., Bruno R. M (2019) **Reinforcement Learning Recruits Somata and Apical Dendrites across Layers of Primary Sensory Cortex** *Cell Rep* **26**:2000–2008 <https://doi.org/10.1016/j.celrep.2019.01.093>
- 29 Luebke J. I., et al. (2010) **Dendritic vulnerability in neurodegenerative disease: insights from analyses of cortical pyramidal neurons in transgenic mouse models** *Brain Struct Funct* **214**:181–199 <https://doi.org/10.1007/s00429-010-0244-2>
- 30 Tsai J., Grutzendler J., Duff K., Gan W. B (2004) **Fibrillar amyloid deposition leads to local synaptic abnormalities and breakage of neuronal branches** *Nat Neurosci* **7**:1181–1183 <https://doi.org/10.1038/nn1335>
- 31 Xu N. L., et al. (2012) **Nonlinear dendritic integration of sensory and motor input during an active sensing task** *Nature* **492**:247–251 <https://doi.org/10.1038/nature11601>
- 32 Takahashi N., et al. (2020) **Active dendritic currents gate descending cortical outputs in perception** *Nat Neurosci* <https://doi.org/10.1038/s41593-020-0677-8>
- 33 Takahashi N., Oertner T. G., Hegemann P., Larkum M. E (2016) **Active cortical dendrites modulate perception** *Science* **354**:1587–1590 <https://doi.org/10.1126/science.aah6066>
- 34 Bittner K. C., Milstein A. D., Grienberger C., Romani S., Magee J. C (2017) **Behavioral time scale synaptic plasticity underlies CA1 place fields** *Science* **357**:1033–1036 <https://doi.org/10.1126/science.aan3846>
- 35 Doron G., et al. (2020) **Perirhinal input to neocortical layer 1 controls learning** *Science* **370** <https://doi.org/10.1126/science.aaz3136>
- 36 Bouchard M. B., et al. (2015) **Swept confocally-aligned planar excitation (SCAPE) microscopy for high speed volumetric imaging of behaving organisms** *Nat Photonics* **9**:113–119 <https://doi.org/10.1038/nphoton.2014.323>
- 37 Hillman E. M., et al. (2018) **High-speed 3D imaging of cellular activity in the brain using axially-extended beams and light sheets** *Curr Opin Neurobiol* **50**:190–200
- 38 Yu Y. S., Graff M. M., Bresee C. S., Man Y. B., Hartmann M. J (2016) **Whiskers aid anemotaxis in rats** *Sci Adv* **2** <https://doi.org/10.1126/sciadv.1600716>
- 39 Yu Y. S., Graff M. M., Hartmann M. J (2016) **Mechanical responses of rat vibrissae to airflow** *J Exp Biol* **219**:937–948 <https://doi.org/10.1242/jeb.126896>
- 40 Nakamura S., Narumi T., Tsutsui K., Iijima T (2009) **Difference in the functional significance between the lemniscal and paralemniscal pathways in the perception of direction of single-whisker stimulation examined by muscimol microinjection** *Neurosci Res* **64**:323–329 <https://doi.org/10.1016/j.neures.2009.04.005>
- 41 Bernhard S. M., et al. (2020) **An automated homecage system for multiwhisker detection and discrimination learning in mice** *PLoS One* **15** <https://doi.org/10.1371/journal.pone.0232916>
- 42 Chen T. W., et al. (2013) **Ultrasensitive fluorescent proteins for imaging neuronal activity** *Nature* **499**:295–300 <https://doi.org/10.1038/nature12354>

- 43 Kozorovitskiy Y., Saunders A., Johnson C. A., Lowell B. B., Sabatini B. L (2012) **Recurrent network activity drives striatal synaptogenesis** *Nature* **485**:646–650 <https://doi.org/10.1038/nature11052>
- 44 Glickfeld L. L., Andermann M. L., Bonin V., Reid R. C (2013) **Cortico-cortical projections in mouse visual cortex are functionally target specific** *Nat Neurosci* **16**:219–226 <https://doi.org/10.1038/nn.3300>
- 45 Giovannucci A., et al. (2019) **CaImAn an open source tool for scalable calcium imaging data analysis** *Elife* **8** <https://doi.org/10.7554/eLife.38173>
- 46 Pakan J. M. P., Currie S. P., Fischer L., Rochefort N. L (2018) **The Impact of Visual Cues, Reward, and Motor Feedback on the Representation of Behaviorally Relevant Spatial Locations in Primary Visual Cortex** *Cell Rep* **24**:2521–2528 <https://doi.org/10.1016/j.celrep.2018.08.010>
- 47 Wang P. Y., et al. (2020) **Transient and Persistent Representations of Odor Value in Prefrontal Cortex** *Neuron* **108**:209–224 <https://doi.org/10.1016/j.neuron.2020.07.033>
- 48 Bruno R. M., Khatri V., Land P. W., Simons D. J (2003) **Thalamocortical angular tuning domains within individual barrels of rat somatosensory cortex** *J Neurosci* **23**:9565–9574
- 49 Bruno R. M., Sakmann B (2006) **Cortex is driven by weak but synchronously active thalamocortical synapses** *Science* **312**:1622–1627 <https://doi.org/10.1126/science.1124593>
- 50 Ramirez A., et al. (2014) **Spatiotemporal receptive fields of barrel cortex revealed by reverse correlation of synaptic input** *Nat Neurosci* **17**:866–875 <https://doi.org/10.1038/nn.3720>
- 51 Yao H., Dan Y (2001) **Stimulus timing-dependent plasticity in cortical processing of orientation** *Neuron* **32**:315–323 [https://doi.org/10.1016/s0896-6273\(01\)00460-3](https://doi.org/10.1016/s0896-6273(01)00460-3)
- 52 Dragoi V., Sharma J., Miller E. K., Sur M (2002) **Dynamics of neuronal sensitivity in visual cortex and local feature discrimination** *Nat Neurosci* **5**:883–891 <https://doi.org/10.1038/nn900>
- 53 Dragoi V., Sharma J., Sur M (2000) **Adaptation-induced plasticity of orientation tuning in adult visual cortex** *Neuron* **28**:287–298 [https://doi.org/10.1016/s0896-6273\(00\)00103-3](https://doi.org/10.1016/s0896-6273(00)00103-3)
- 54 Zhang Y., Cudmore R. H., Lin D. T., Linden D. J., Huganir R. L (2015) **Visualization of NMDA receptor-dependent AMPA receptor synaptic plasticity in vivo** *Nat Neurosci* **18**:402–407 <https://doi.org/10.1038/nn.3936>
- 55 Chu M. W., Li W. L., Komiyama T (2016) **Balancing the Robustness and Efficiency of Odor Representations during Learning** *Neuron* **92**:174–186 <https://doi.org/10.1016/j.neuron.2016.09.004>
- 56 Cichon J., Gan W. B (2015) **Branch-specific dendritic Ca(2+) spikes cause persistent synaptic plasticity** *Nature* **520**:180–185 <https://doi.org/10.1038/nature14251>
- 57 Palmer L. M., et al. (2014) **NMDA spikes enhance action potential generation during sensory input** *Nat Neurosci* **17**:383–390 <https://doi.org/10.1038/nn.3646>

- 58 Derdikman D., et al. (2006) **Layer-specific touch-dependent facilitation and depression in the somatosensory cortex during active whisking** *J Neurosci* **26**:9538–9547 <https://doi.org/10.1523/JNEUROSCI.0918-06.2006>
- 59 de Kock C. P., Sakmann B (2009) **Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific** *Proc Natl Acad Sci U S A* **106**:16446–16450 <https://doi.org/10.1073/pnas.0904143106>
- 60 Rodgers C. C., et al. (2021) **Sensorimotor strategies and neuronal representations for shape discrimination** *Neuron* **109**:2308–2325 <https://doi.org/10.1016/j.neuron.2021.05.019>
- 61 Larkum M (2013) **A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex** *Trends Neurosci* **36**:141–151 <https://doi.org/10.1016/j.tins.2012.11.006>
- 62 Abs E., et al. (2018) **Learning-Related Plasticity in Dendrite-Targeting Layer 1 Interneurons** *Neuron* **100**:684–699 <https://doi.org/10.1016/j.neuron.2018.09.001>
- 63 Cohen-Kashi Malina K., et al. (2021) **NDNF interneurons in layer 1 gain-modulate whole cortical columns according to an animal's behavioral state** *Neuron* **109**:2150–2164 <https://doi.org/10.1016/j.neuron.2021.05.001>
- 64 Ranganathan G. N., et al. (2018) **Active dendritic integration and mixed neocortical network representations during an adaptive sensing behavior** *Nat Neurosci* **21**:1583–1590 <https://doi.org/10.1038/s41593-018-0254-6>
- 65 Park J. M., et al. (2020) **Deep and superficial layers of the primary somatosensory cortex are critical for whisker-based texture discrimination in mice** *bioRxiv* <https://doi.org/10.1101/2020.08.12.245381>
- 66 Rigotti M., et al. (2013) **The importance of mixed selectivity in complex cognitive tasks** *Nature* **497**:585–590 <https://doi.org/10.1038/nature12160>
- 67 Stringer C., Pachitariu M., Steinmetz N., Carandini M., Harris K. D (2019) **High-dimensional geometry of population responses in visual cortex** *Nature* **571**:361–365 <https://doi.org/10.1038/s41586-019-1346-5>
- 68 Kim J., Erskine A., Cheung J. A., Hires S. A (2020) **Behavioral and Neural Bases of Tactile Shape Discrimination Learning in Head-Fixed Mice** *Neuron* **108**:953–967 <https://doi.org/10.1016/j.neuron.2020.09.012>
- 69 Dabney W., et al. (2020) **A distributional code for value in dopamine-based reinforcement learning** *Nature* **577**:671–675 <https://doi.org/10.1038/s41586-019-1924-6>
- 70 Petreanu L., et al. (2012) **Activity in motor-sensory projections reveals distributed coding in somatosensation** *Nature* **489**:299–303 <https://doi.org/10.1038/nature11321>
- 71 Wimmer V. C., Bruno R. M., de Kock C. P., Kuner T., Sakmann B (2010) **Dimensions of a Projection Column and Architecture of VPM and POm Axons in Rat Vibrissal Cortex** *Cereb Cortex* **20**:2265–2276 <https://doi.org/10.1093/cercor/bhq068>
- 72 Feldmeyer D (2012) **Excitatory neuronal connectivity in the barrel cortex** *Front Neuroanat* **6** <https://doi.org/10.3389/fnana.2012.00024>

- 73 Constantinople C. M., Bruno R. M (2013) **Deep cortical layers are activated directly by thalamus** *Science* **340**:1591–1594 <https://doi.org/10.1126/science.1236425>
- 74 Petersen C. C., Crochet S (2013) **Synaptic computation and sensory processing in neocortical layer 2/3** *Neuron* **78**:28–48 <https://doi.org/10.1016/j.neuron.2013.03.020>
- 75 Yamashita T., et al. (2018) **Diverse Long-Range Axonal Projections of Excitatory Layer 2/3 Neurons in Mouse Barrel Cortex** *Front Neuroanat* **12** <https://doi.org/10.3389/fnana.2018.00033>
- 76 Llinas R., Ribary U., Contreras D., Pedroarena C (1998) **The neuronal basis for consciousness** *Philos Trans R Soc Lond B Biol Sci* **353**:1841–1849 <https://doi.org/10.1098/rstb.1998.0336>
- 77 Krauzlis R. J., Lovejoy L. P., Zenon A (2013) **Superior colliculus and visual spatial attention** *Annu Rev Neurosci* **36**:165–182 <https://doi.org/10.1146/annurev-neuro-062012-170249>
- 78 Parvizi J., Damasio A. (2001) **Consciousness and the brainstem** *Cognition* **79**:135–160 [https://doi.org/10.1016/S0010-0277\(00\)00127-X](https://doi.org/10.1016/S0010-0277(00)00127-X)
- 79 Petty G. H., Bruno R. M (2024) **Attentional modulation of secondary somatosensory and visual thalamus of mice** *eLife* **13**
- 80 Feldman D. E., Brecht M (2005) **Map plasticity in somatosensory cortex** *Science* **310**:810–815 <https://doi.org/10.1126/science.1115807>
- 81 Mechawar N., Cozzari C., Descarries L (2000) **Cholinergic innervation in adult rat cerebral cortex: a quantitative immunocytochemical description** *J Comp Neurol* **428**:305–318
- 82 Freedman R., Foote S. L., Bloom F. E (1975) **Histochemical characterization of a neocortical projection of the nucleus locus coeruleus in the squirrel monkey** *J Comp Neurol* **164**:209–231 <https://doi.org/10.1002/cne.901640205>
- 83 Chubykin A. A., Roach E. B., Bear M. F., Shuler M. G (2013) **A cholinergic mechanism for reward timing within primary visual cortex** *Neuron* **77**:723–735 <https://doi.org/10.1016/j.neuron.2012.12.039>
- 84 Thiele A., Bellgrove M. A. (2018) **Neuromodulation of Attention** *Neuron* **97**:769–785 <https://doi.org/10.1016/j.neuron.2018.01.008>
- 85 Hangya B., Ranade S. P., Lorenc M., Kepecs A (2015) **Central Cholinergic Neurons Are Rapidly Recruited by Reinforcement Feedback** *Cell* **162**:1155–1168 <https://doi.org/10.1016/j.cell.2015.07.057>
- 86 Froemke R. C., Merzenich M. M., Schreiner C. E (2007) **A synaptic memory trace for cortical receptive field plasticity** *Nature* **450**:425–429 <https://doi.org/10.1038/nature06289>
- 87 Letzkus J. J., et al. (2011) **A disinhibitory microcircuit for associative fear learning in the auditory cortex** *Nature* **480**:331–335 <https://doi.org/10.1038/nature10674>
- 88 Gasselin C., Hohl B., Vernet A., Crochet S., Petersen C. C. H (2021) **Cell-type-specific nicotinic input disinhibits mouse barrel cortex during active sensing** *Neuron* <https://doi.org/10.1016/j.neuron.2020.12.018>

- 89 Berger T. K., Silberberg G., Perin R., Markram H (2010) **Brief bursts self-inhibit and correlate the pyramidal network** *PLoS Biol* **8** <https://doi.org/10.1371/journal.pbio.1000473>
- 90 Kapfer C., Glickfeld L. L., Atallah B. V., Scanziani M (2007) **Supralinear increase of recurrent inhibition during sparse activity in the somatosensory cortex** *Nat Neurosci* **10**:743–753 <https://doi.org/10.1038/nn1909>
- 91 Naka A., Adesnik H (2016) **Inhibitory Circuits in Cortical Layer 5** *Front Neural Circuits* **10** <https://doi.org/10.3389/fncir.2016.00035>
- 92 Ljungberg T., Apicella P., Schultz W (1992) **Responses of monkey dopamine neurons during learning of behavioral reactions** *J Neurophysiol* **67**:145–163 <https://doi.org/10.1152/jn.1992.67.1.145>
- 93 Pan W. X., Schmidt R., Wickens J. R., Hyland B. I (2005) **Dopamine cells respond to predicted events during classical conditioning: evidence for eligibility traces in the reward-learning network** *J Neurosci* **25**:6235–6242 <https://doi.org/10.1523/JNEUROSCI.1478-05.2005>
- 94 Bouret S., Sara S. J (2004) **Reward expectation, orientation of attention and locus coeruleus-medial frontal cortex interplay during learning** *Eur J Neurosci* **20**:791–802 <https://doi.org/10.1111/j.1460-9568.2004.03526.x>
- 95 Sheffield M. E., Dombeck D. A (2015) **Calcium transient prevalence across the dendritic arbour predicts place field properties** *Nature* **517**:200–204 <https://doi.org/10.1038/nature13871>
- 96 Oberlaender M., et al. (2012) **Cell type-specific three-dimensional structure of thalamocortical circuits in a column of rat vibrissal cortex** *Cereb Cortex* **22**:2375–2391 <https://doi.org/10.1093/cercor/bhr317>
- 97 Clack N. G., et al. (2012) **Automated tracking of whiskers in videos of head fixed rodents** *PLoS Comput Biol* **8** <https://doi.org/10.1371/journal.pcbi.1002591>
- 98 Voleti V., et al. (2019) **Real-time volumetric microscopy of in vivo dynamics and large-scale samples with SCAPE 2.0** *Nat Methods* **16**:1054–1062 <https://doi.org/10.1038/s41592-019-0579-4>
- 99 Pnevmatikakis E. A., Giovannucci A (2017) **NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data** *J Neurosci Methods* **291**:83–94 <https://doi.org/10.1016/j.jneumeth.2017.07.031>
- 100 Pnevmatikakis E. A., et al. (2016) **Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data** *Neuron* **89**:285–299 <https://doi.org/10.1016/j.neuron.2015.11.037>
- 101 Petty G. H., Kinnischtzke A. K., Hong Y. K., Bruno R. M. (2020) **Effects of arousal and movement on secondary somatosensory and visual thalamus** *bioRxiv* <https://doi.org/10.1101/2020.03.04.977348>

Author information

Sam E Benezra

Department of Neuroscience, Columbia University, New York, United States, Kavli Institute for Brain Science, Columbia University, New York, United States, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, United States

Kripa B Patel

Departments of Biomedical Engineering and Radiology, Columbia University, New York, United States, Kavli Institute for Brain Science, Columbia University, New York, United States, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, United States

Citlali Pérez Campos

Departments of Biomedical Engineering and Radiology, Columbia University, New York, United States, Kavli Institute for Brain Science, Columbia University, New York, United States, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, United States

Elizabeth MC Hillman

Department of Neuroscience, Columbia University, New York, United States, Departments of Biomedical Engineering and Radiology, Columbia University, New York, United States, Kavli Institute for Brain Science, Columbia University, New York, United States, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, United States
ORCID iD: [0000-0001-5511-1451](https://orcid.org/0000-0001-5511-1451)

Randy M Bruno

Department of Neuroscience, Columbia University, New York, United States, Kavli Institute for Brain Science, Columbia University, New York, United States, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, United States, Department of Physiology, Anatomy & Genetics, University of Oxford, Oxford, United Kingdom
ORCID iD: [0000-0002-5122-4632](https://orcid.org/0000-0002-5122-4632)

For correspondence: randy.bruno@dpag.ox.ac.uk

Editors

Reviewing Editor

Brice Bathellier

Centre National de la Recherche Scientifique, Paris, France

Senior Editor

John Huguenard

Stanford University School of Medicine, Stanford, United States of America

Reviewer #1 (Public review):

What neurophysiological changes support the learning of new sensorimotor transformations is a key question in neuroscience. Many studies have attempted to answer this question at the neuronal population level - with varying degrees of success - but few, if any, have studied the change in activity of the apical dendrites of layer 5 cortical neurons. Neurons in the layer 5 of the sensory cortex appear to play a key role in sensorimotor transformations, showing important decision and reward-related signals, and being the main source of cortical and subcortical projections from the cortex. In particular, pyramidal track (PT) neurons project directly to subcortical regions related to motor activity, such as the striatum and brainstem, and could initiate rapid motor action in response to given sensory inputs. Additionally, layer

5 cortical neurons have large apical dendrites that extend to layer 1 where different neuromodulatory and long-range inputs converge, providing motor and contextual information that could be used to modulate layer 5 neurons output and/or to establish the synaptic plasticity required for learning a new association.

In this study, the authors aimed to test whether the learning of a new sensorimotor transformation could be supported by a change in the evoked response of the apical dendrites of layer 5 neurons in the mouse whisker primary somatosensory cortex. To do this, they performed longitudinal functional calcium imaging of the apical dendrites of layer 5 neurons while mice learned to discriminate between two multiwhiskers stimuli. The authors used a simple conditioning task in which one whisker stimulus (upward or backward air puff, CS+) is associated with reward after a short delay, while the other whisker stimulus (CS-) is not. They found that task learning (measured by the probability of anticipatory licking just after the CS+) was not associated with a significant change of the average population response evoked by the CS+ or the CS-, nor change in the average population selectivity. However, when considering individual dendritic tufts, they found interesting changes in selectivity, with approximately equal numbers of dendrites becoming more selective for CS+ and dendrites becoming more selective for CS-.

One of the major challenges when assessing changes in neural representation during the learning of such Go/NoGo tasks is that the movements and rewards themselves may elicit strong neural responses that may be a confounding factor, that is, inexperienced mice do not lick in response to the CS+, while trained mice do. In this study, the authors addressed this issue in three ways: first, they carefully monitor the orofacial movements of mice and show that task learning is not associated with changes in evoked whisker movements. Second, they show that whisking or licking evokes very little activity in the dendritic tufts compared to whisker stimuli (CS+ and CS-). Finally, the authors introduced into the design of their task a post-conditioning session after the last conditioning session during which the CS+ and the CS- are presented but no reward is delivered. During this post-session, the mice gradually stopped licking in response to the CS+. A better design might have been to perform the pre-conditioning and post-conditioning sessions in non-water-restricted, unmotivated mice to completely exclude any lick response, but the fact that the change in selectivity persists after the mice stopped licking in the last blocks of the post-conditioning session (in mice relying only on their whiskers to perform the task) is convincing.

The clever task design and careful data analysis provide compelling evidence that learning this whisker discrimination task does not result in a massive change in sensory representation in the apical dendritic tufts of layer 5 neurons in the primary somatosensory cortex on average. Nevertheless, individual dendritic tufts do increase their selectivity for one or the other sensory stimulus, likely enhancing the ability of S1 neurons to accurately discriminate the two stimuli and trigger the appropriate motor response (to lick or not to lick).

One limitation of the present study is the lack of evidence for the necessity of the primary somatosensory cortex in the learning and execution of the task. As the authors have strongly emphasized in their previous publications, the primary somatosensory cortex may not be necessary for the learning and execution of simple whisker detection tasks, especially when the stimulus is very salient. Although this new task requires the discrimination between two whisker stimuli, the simplicity and salience of the whisker stimuli used could make this task cortex independent. Especially when considering that some mice seem to not rely entirely on their whiskers to execute the task.

Nevertheless, this is an important result that shows for the first-time changes in the selectivity to sensory stimuli at the level of individual apical dendritic tufts in correlation with the learning of a discrimination task. This study sheds new light on the cortical cellular substrates of reward-based learning, and opens interesting perspectives for future research

in this area. In future studies, it will be important to determine whether the change in selectivity of dendritic calcium spikes is causally involved in the learning the task or whether it simply correlates with learning, as a consequence of changes in synaptic inputs caused by reward. The dendritic calcium spikes may be involved in the establishment of synaptic plasticity required for learning and impact the output of layer 5 pyramidal neurons to trigger the appropriate motor response. It would be important also to study the changes in selectivity in the apical dendrite of the identified projection neurons.

Comments on revisions:

The authors have addressed all my questions. I have no further recommendations.

<https://doi.org/10.7554/eLife.98349.2.sa2>

Reviewer #2 (Public review):

Summary:

The authors did not find an increased representation of CS+ throughout reinforcement learning in the tuft dendrites of Rbp4-positive neurons from layer 5B of the barrel cortex, as previously reported for soma from layer 2/3 of the visual cortex.

Alternatively, the authors observed an increased selectivity to both stimuli (CS+ and CS-) during reinforcement learning. This feature 1) was not present in repeated exposures (without reinforcement), 2) was not explained by animal's behaviour (choice, licking and whisking) and 3) was long-lasting, being present even when the mice disengaged from the task.

Importantly, increased selectivity was correlated with learning (% correct choices), and neural discriminability between stimuli increased with learning.

In conclusion, the authors show that tuft dendrites from layer 5B of the barrel cortex increase the representation of conditioned (CS+) and unconditioned stimuli (CS-) applied to the whiskers, during reinforcement learning.

Strengths:

The results presented are very consistent throughout the entire study, and therefore very convincing:

- (1) The results observed are very similar using two different imaging techniques (using 2-photon -planar imaging- and SCAPE - volumetric imaging). Fig. 3 and Fig.4 respectively.
- (2) The results are similar using "different groups" of tuft dendrites for the analysis (e.g. initially unresponsive and responsive pre- and post-learning). Fig. 5.
- (3) The results are similar from a specific set of trials (with the same sensory input, but different choices). Fig.7.
- (4) Additionally, the selectivity of tuft dendrites from layer 5B of the barrel cortex was higher in the mice that exclusively used the whisker to respond to the stimuli (CS+ and CS-).

The results presented are controlled against a group of mice that received the same stimuli presentation, except the reinforcement (reward).

Additionally, the behaviour outputs, such as choice, whisking and licking could not account for the results observed.

Although there are no causal experiments, the correlation between selectivity and learning (% of correct choices), as well as the increased neural discriminability with learning, but not in repeated exposure, are very convincing.

Weaknesses:

The biggest weakness is the absence of causality experiments. Although inhibiting specifically tuft dendritic activity in layer 1 from layer 5 pyramidal neurons is very challenging, tuft dendritic activity in layer 1 could be silenced through optogenetic experiments as in Abs et al. 2018. By manipulating NDNF-positive neurons the authors could specifically modify tuft dendritic activity in the barrel cortex during CS presentations, and test if silencing tuft dendritic activity in layer 1 would lead to the lack of selectivity and an impairment of reinforcement learning. Additionally, this experiment will test if the selectivity observed during reinforcement learning is due to changes in the local network, namely changes in local synaptic connectivity, or solely due to changes in the long-range inputs.

<https://doi.org/10.7554/eLife.98349.2.sa1>

Author response:

The following is the authors' response to the original reviews.

Reviewer #1 (Public Review):

What neurophysiological changes support the learning of new sensorimotor transformations is a key question in neuroscience. Many studies have attempted to answer this question at the neuronal population level - with varying degrees of success - but few, if any, have studied the change in activity of the apical dendrites of layer 5 cortical neurons. Neurons in layer 5 of the sensory cortex appear to play a key role in sensorimotor transformations, showing important decision and reward-related signals, and being the main source of cortical and subcortical projections from the cortex. In particular, pyramidal track (PT) neurons project directly to subcortical regions related to motor activity, such as the striatum and brainstem, and could initiate rapid motor action in response to given sensory inputs. Additionally, layer 5 cortical neurons have large apical dendrites that extend to layer 1 where different neuromodulatory and long-range inputs converge, providing motor and contextual information that could be used to modulate layer 5 neurons output and/or to establish the synaptic plasticity required for learning a new association.

In this study, the authors aimed to test whether the learning of a new sensorimotor transformation could be supported by a change in the evoked response of the apical dendrites of layer 5 neurons in the mouse whisker primary somatosensory cortex. To do this, they performed longitudinal functional calcium imaging of the apical dendrites of layer 5 neurons while mice learned to discriminate between two multi-whisker stimuli. The authors used a simple conditioning task in which one whisker stimulus (upward or backward air puff, CS+) is associated with a reward after a short delay, while the other whisker stimulus (CS-) is not. They found that task learning (measured by the probability of anticipatory licking just after the CS+) was not associated with a significant change in the average population response evoked by the CS+ or the CS-, nor a change in the average population selectivity. However, when considering individual dendritic tufts, they found interesting changes in selectivity, with approximately equal numbers of dendrites becoming more selective for CS+ and dendrites becoming more selective for CS-.

One of the major challenges when assessing changes in neural representation during the learning of such Go/NoGo tasks is that the movements and rewards themselves may elicit strong neural responses that may be a confounding factor, that is, inexperienced mice do not lick in response to the CS+, while trained mice do. In this study, the authors addressed this issue in three ways: first, they carefully monitored the orofacial movements of mice and showed that task learning is not associated with changes in evoked whisker movements. Second, they show that whisking or licking evokes very little

activity in the dendritic tufts compared to whisker stimuli (CS+ and CS-). Finally, the authors introduced into the design of their task a post-conditioning session after the last conditioning session during which the CS+ and the CS- are presented but no reward is delivered. During this post-session, the mice gradually stopped licking in response to the CS+. A better design might have been to perform the pre-conditioning and post-conditioning sessions in nonwater-restricted, unmotivated mice to completely exclude any lick response, but the fact that the change in selectivity persists after the mice stopped licking in the last blocks of the post-conditioning session (in mice relying only on their whiskers to perform the task) is convincing.

The clever task design and careful data analysis provide compelling evidence that learning this whisker discrimination task does not result in a massive change in sensory representation in the apical dendritic tufts of layer 5 neurons in the primary somatosensory cortex on average. Nevertheless, individual dendritic tufts do increase their selectivity for one or the other sensory stimulus, likely enhancing the ability of S1 neurons to accurately discriminate the two stimuli and trigger the appropriate motor response (to lick or not to lick).

One limitation of the present study is the lack of evidence for the necessity of the primary somatosensory cortex in the learning and execution of the task. As the authors have strongly emphasized in their previous publications, the primary somatosensory cortex may not be necessary for the learning and execution of simple whisker detection tasks, especially when the stimulus is very salient. Although this new task requires the discrimination between two whisker stimuli, the simplicity and salience of the whisker stimuli used could make this task cortex-independent. Especially when considering that some mice seem to not rely entirely on their whiskers to execute the task.

Nevertheless, this is an important result that shows for the first time changes in the selectivity to sensory stimuli at the level of individual apical dendritic tufts in correlation with the learning of a discrimination task. This study sheds new light on the cortical cellular substrates of reward-based learning and opens interesting perspectives for future research in this area. In future studies, it will be important to determine whether the change in selectivity of dendritic calcium spikes is causally involved in the learning of the task or whether it simply correlates with learning, as a consequence of changes in synaptic inputs caused by reward. The dendritic calcium spikes may be involved in the establishment of synaptic plasticity required for learning and impact the output of layer 5 pyramidal neurons to trigger the appropriate motor response. It would be important also to study the changes in selectivity in the apical dendrite of the identified projection neurons.

Reviewer #2 (Public Review):

Summary:

The authors did not find an increased representation of CS+ throughout reinforcement learning in the tuft dendrites of Rbp4-positive neurons from layer 5B of the barrel cortex, as previously reported for soma from layer 2/3 of the visual cortex.

Alternatively, the authors observed an increased selectivity to both stimuli (CS+ and CS-) during reinforcement learning. This feature:

- (1) was not present in repeated exposures (without reinforcement),
- (2) was not explained by the animal's behaviour (choice, licking, and whisking), and
- (3) was long-lasting, being present even when the mice disengaged from the task.

Importantly, increased selectivity was correlated with learning (% correct choices), and neural discriminability between stimuli increased with learning.

In conclusion, the authors show that tuft dendrites from layer 5B of the barrel cortex increase the representation of conditioned (CS+) and unconditioned stimuli (CS-) applied to the whiskers, during reinforcement learning.

Strengths:

The results presented are very consistent throughout the entire study, and therefore very convincing:

(1) The results observed are very similar using two different imaging techniques (2-photon planar imaging- and SCAPE-volumetric imaging). Figure 3 and Figure 4 respectively.

(2) The results are similar using "different groups" of tuft dendrites for the analysis (e.g. initially unresponsive and responsive pre- and post-learning). Figure 5.

(3) The results are similar from a specific set of trials (with the same sensory input, but different choices). Figure 7.

(4) Additionally, the selectivity of tuft dendrites from layer 5B of the barrel cortex was higher in the mice that exclusively used the whisker to respond to the stimuli (CS+ and CS-). The results presented are controlled against a group of mice that received the same stimuli presentation, except for the reinforcement (reward).

Additionally, the behaviour outputs, such as choice, whisking, and licking could not account for the results observed.

Although there are no causal experiments, the correlation between selectivity and learning (percentage of correct choices), as well as the increased neural discriminability with learning, but not in repeated exposure, are very convincing.

Weaknesses:

The biggest weakness is the absence of causality experiments. Although inhibiting specifically tuft dendritic activity in layer 1 from layer 5 pyramidal neurons is very challenging, tuft dendritic activity in layer 1 could be silenced through optogenetic experiments as in Abs et al. 2018. By manipulating NDNF-positive neurons the authors could specifically modify tuft dendritic activity in the barrel cortex during CS presentations, and test if silencing tuft dendritic activity in layer 1 would lead to the lack of selectivity and an impairment of reinforcement learning. Additionally, this experiment will test if the selectivity observed during reinforcement learning is due to changes in the local network, namely changes in local synaptic connectivity, or solely due to changes in the long-range inputs.

We agree that such causal manipulations are a logical next step. Such manipulations are unfortunately not specific to layer 5 apicals, so the results would be difficult to interpret. We now discuss the challenge of such manipulations in the Discussion section.

Recommendations for the authors:

Reviewer #1 (Recommendations For The Authors):

Overall, the study is solid and the article is well and clearly written. I have no suggestion for other experiments that would fall within the scope of this article. I would like only to

suggest some additional analyses and clarifications in the writing.

Additional analyses:

Obviously, the main confounding factor in this type of data comes from the acquired motor response which follows - with a short latency - the sensory stimulus. This is particularly problematic for functional calcium imaging which has very low temporal resolution. The authors have addressed this question to some extent by showing that motor-evoked activity does not account for the change in selectivity acquired with learning and through the use of a post-conditioning session during which no reward was delivered. Figures 8C-D show that mice gradually stop licking in response to CS+ in this session and that the distribution of the selectivity index remains similar in these last blocks. Perhaps a more convincing analysis would be to simply select Miss and Correct rejection trials in which mice did not lick in response to the CS+ and CS-, respectively. Ideally, if the number of trials is sufficient, one could even select trials devoid of any evoked movement (no licking and no whisking).

We agree it would be interesting to compare Miss and Correct rejection trials to further rule out effects of a motor response, but there were never enough Miss trials to conduct such an analysis. Even in very early learning, there are few Miss trials (see Figure 1, session 2). We found that in early learning, animals would lick in most trials. Then, over the course of conditioning, they would learn to withhold licks during CS- presentation. Thus, we were able to examine Hits, Correct rejections, and False alarms (Figure 7), but not Miss trials. We have added text suggesting a future experiment in which the stimulus strengths are substantially reduced to drastically increase the error rates.

The fact that changes in selectivity occur in both directions overall is really interesting. However, in the way the data are presented currently, one may wonder about mice/field of view vs single cell effect. i.e., do different dendritic tufts in the same field of view show opposite changes in selectivity? If we were to replot Figure 3A for a single mouse, would we obtain the same picture?

We appreciate this very good suggestion and have added scatter plots and selectivity index histograms for individual conditioned animals in Supplementary figure 2. These data demonstrate that different dendritic tufts in the same field of view exhibit opposite changes in selectivity.

The authors point out that they observed no change in the mean response or selectivity during learning, but did find changes in selectivity at the level of individual dendritic tufts. This suggests that, at the population level, the ability to discriminate between the two stimuli should improve. A possible complementary analysis would be to show that the ability to decode stimulus identity from dendritic tuft population activity increases with learning.

Given the substantial change in individual tuft selectivity and that the tuft events occur are not rare, the population result is guaranteed. If individual tufts increase selectivity, the population will also increase its selectivity on a trial-by-trial basis. We have nevertheless included a new supplementary figure with a population analysis using SVMs to demonstrate this.

Clarification:

The authors should make it clear from the beginning that mice are still water-restricted during the post-conditioning session and actually do keep licking for many CS+ trials. Therefore, this session is not devoid of motor response.

We have clarified this in the text.

Did mice in the repeated exposure condition receive any reward during the recording sessions? If so when were rewards delivered?

We previously described in the Methods that these mice received water in their home cage, but we now additionally clarify this in the Results section.

Minor:

Figure 2Aii, the labels of the Alpha and Beta barrels should be swapped.

Fixed

Line 218: I believe this sentence should read "Using SCAPE microscopy, ...".

Corrected.

Line 665: 'Reconstruction from 50' does that refer to the single cell reconstruction on the left panel?

Yes – Clarified in legend

Reviewer #2 (Recommendations For The Authors):

Minor suggestions:

The 'summary' should mention from which brain area the results were acquired. Otherwise, it is misleading, giving the idea that the results described a generic feature, which is still unknown.

Added to the text.

Please correct sentence 219: "SCAPE microscopy, we image tuft activity of additional mice..."

Added to the text.

In the same sentence (219) it would be good to provide the number of additional mice imaged (2).

Added to the text.

Regarding Supplementary Figure 1, it would be interesting to correlate the second peak after reward and learning rate, to provide further support to the sentences 109 to 113.

We agree this would be interesting to examine, but only four animals exhibited this second peak, which is too small of a sample to observe a meaningful correlation. We now clarify this in the text.

In Figure 3, why not present the correlation between 'neural discriminability' and % of correct choices?

We appreciate the suggestion and have added this plot to Figure 3.

The 'results' section will benefit tremendously if the authors consistently indicate the figures to which the results are being described, or 'data not shown' if it is the case. To give a few examples:

Sentence 108 - "averaged 28% $\Delta F/F$ " - From which figure is this result coming from?

Sentence 123 - " $p = 0.62, 0.64$, respectively)" - comparison not shown, but see Figures 2E and D respectively?

Sentence 125 - "(CS+ responsive (...) across all sessions)" - From which figure is this result coming from?

Sentence 130 - "during pre-conditioning ($p=0.66$) or post-conditioning sessions ($p=0.44$) - From which figure?

Sentence 154 - "(Pre: $p=0.20$; last rewarded: $p=0.43$; Post: $p=0.64$, sign-rank test)" - From which figure?

Sentence 175 - "(-0.049, -0.001, and 0.003" - From which figure? Please show the graph that shows that the mean SI is not different. It can be supplementary. The distribution of SI will be strengthened by it.

We added this plot to supplementary figure 2.

Sentence 244 - "(conditioned: 458/603; repeated exposure: 334/457) - From Figure 5E.

Sentence 256 - " $p=0.04$, 2-sample t-test comparison mice) - From Figure 5B. Sentence

258 - " $p=0.03$, paired t-test) - from Figure 5B Sentences 370 to 378 - No reference to the figure.

The 'discussion' section (sentences 459 to 494) refers to the differences between the current and previous studies (references 1,3,5), namely soma vs. dendrites and layer 2/3 vs. layer 5. However, it should also mention the difference between the nature of the stimuli and the brain area recorded (visual cortex vs. barrel cortex).

We have addressed these issues in the text.

<https://doi.org/10.7554/eLife.98349.2.sa0>