

1 Rapid emergence of latent knowledge in the 2 sensory cortex drives learning

3 Céline Drieu^{1,2,*}, Ziyi Zhu³, Ziyun Wang¹, Kylie Fuller¹, Aaron Wang¹,
4 Sarah Elnozahy^{1,4}, Kishore Kuchibhotla^{1,2,3,*}

5 ¹Department of Psychological and Brain Sciences, Krieger School of Arts and Sciences, Johns Hopkins University, Baltimore, MD,
6 USA

7 ²Johns Hopkins Kavli Neuroscience Discovery Institute, Baltimore, MD, USA

8 ³Department of Neuroscience, School of Medicine, Johns Hopkins University, MD, USA

9 ⁴Present address: Sainsbury Wellcome Centre, London, UK

* e-mails: cdrieu1@jhu.edu; kkuchib1@jhu.edu

10 **Rapid learning confers significant advantages to animals in ecological environments.** Despite the need for speed, animals appear to only slowly learn to associate rewarded actions with predictive cues^{1–4}. This slow learning is thought to be supported by a gradual expansion of predictive cue representation in the sensory cortex^{2,5}. However, evidence is growing that animals learn more rapidly than classical performance measures suggest^{6–8}, challenging the prevailing model of sensory cortical plasticity. Here, we investigated the relationship between learning and sensory cortical representations. We trained mice on an auditory go/no-go task that dissociated the rapid acquisition of task contingencies (learning) from its slower expression (performance)⁷. Optogenetic silencing demonstrated that the auditory cortex (AC) drives both rapid learning and slower performance gains but becomes dispensable at expert. Rather than enhancement or expansion of cue representations⁹, two-photon calcium imaging of AC excitatory neurons throughout learning revealed two higher-order signals that were causal to learning and performance. First, a reward prediction (RP) signal emerged rapidly within tens of trials, was present after action-related errors only early in training, and faded at expert levels. Strikingly, silencing at the time of the RP signal impaired rapid learning, suggesting it serves an associative and teaching role. Second, a distinct cell ensemble encoded and controlled licking suppression that drove the slower performance improvements. These two ensembles were spatially clustered but uncoupled from underlying sensory representations, indicating a higher-order functional segregation within AC. Our results reveal that the sensory cortex manifests higher-order computations that separably drive rapid learning and slower performance improvements, reshaping

34 **our understanding of the fundamental role of the sensory cortex.**

35 Despite the value of rapid learning in ecological environments, most laboratory models of
36 rodent learning show that linking sensory cues with reinforced actions is a slow, gradual
37 process^{1-4,10}. An alternative view suggests that animals, including humans, rapidly infer
38 relationships between cues, actions, and reinforcement (i.e. learning)⁶ even if they continue
39 to make ongoing performance errors^{7,8,11}. Recent behavioral studies in rodents have begun to
40 reconcile these views, arguing that latent task knowledge (i.e. discriminative contingencies)
41 can emerge rapidly even though behavioral performance appears to improve only gradually⁷.
42 How are these two dissociable behavioral processes—rapid acquisition of contingencies versus
43 slower performance improvements—implemented in the brain?

44 An attractive brain region to consider is the sensory cortex as it is thought to subserve
45 instrumental learning by enhancing or attenuating the representation of sensory cues that
46 drive behavior. Plasticity of cue-related responses in the sensory cortex is thought to subserve
47 learning as it mirrors the slow and gradual improvements in behavioral performance^{1,2,5,10}.
48 This raises a fundamental challenge: if animals learn discriminative contingencies rapidly but
49 cue representations in the sensory cortex change slowly^{1,2,9}, the causal model linking cue-
50 related plasticity to learning becomes problematic. One possible solution is that the sensory
51 cortex plays a role beyond cue-related representational plasticity and directly represents high-
52 order signals that associate reinforced actions with predictive cues. Here we focus on the
53 auditory cortex (AC) and asked whether and how it plays a higher-order role in cue-guided
54 learning.

55 We trained head-fixed, water-restricted mice to lick to a target tone (S+) for water reward
56 and to withhold licking to a foil tone (S-) to avoid a timeout (auditory go/no-go task,
57 Fig.1a). We used simple pure tones to prevent the AC from being recruited for complex
58 sensory processing. To confirm this, two-photon imaging of AC excitatory neurons showed
59 that stimulus identity could accurately be decoded from AC activity from the first training
60 day with no subsequent improvement throughout training (Supplementary Figure 1), sug-
61 gesting that the AC was indeed not needed for perceptual sharpening in the task and thereby
62 allowing us to identify possible associative functions. Performance was evaluated in each ses-
63 sion in reinforced and non-reinforced ('probe') trials (Fig.1b). Performance in probe trials
64 revealed a rapid acquisition of task contingency knowledge which was only expressed much
65 later in reinforced trials (Fig.1c)⁷. Reinforcement feedback, although critical for learning,
66 paradoxically masked the underlying task knowledge. By combining this behavioral proce-
67 dure with optogenetics and longitudinal two-photon imaging, we aimed to determine how
68 quickly animals learn stimulus-action contingencies and to define the fundamental role of the
69 auditory cortex in sound-guided learning.

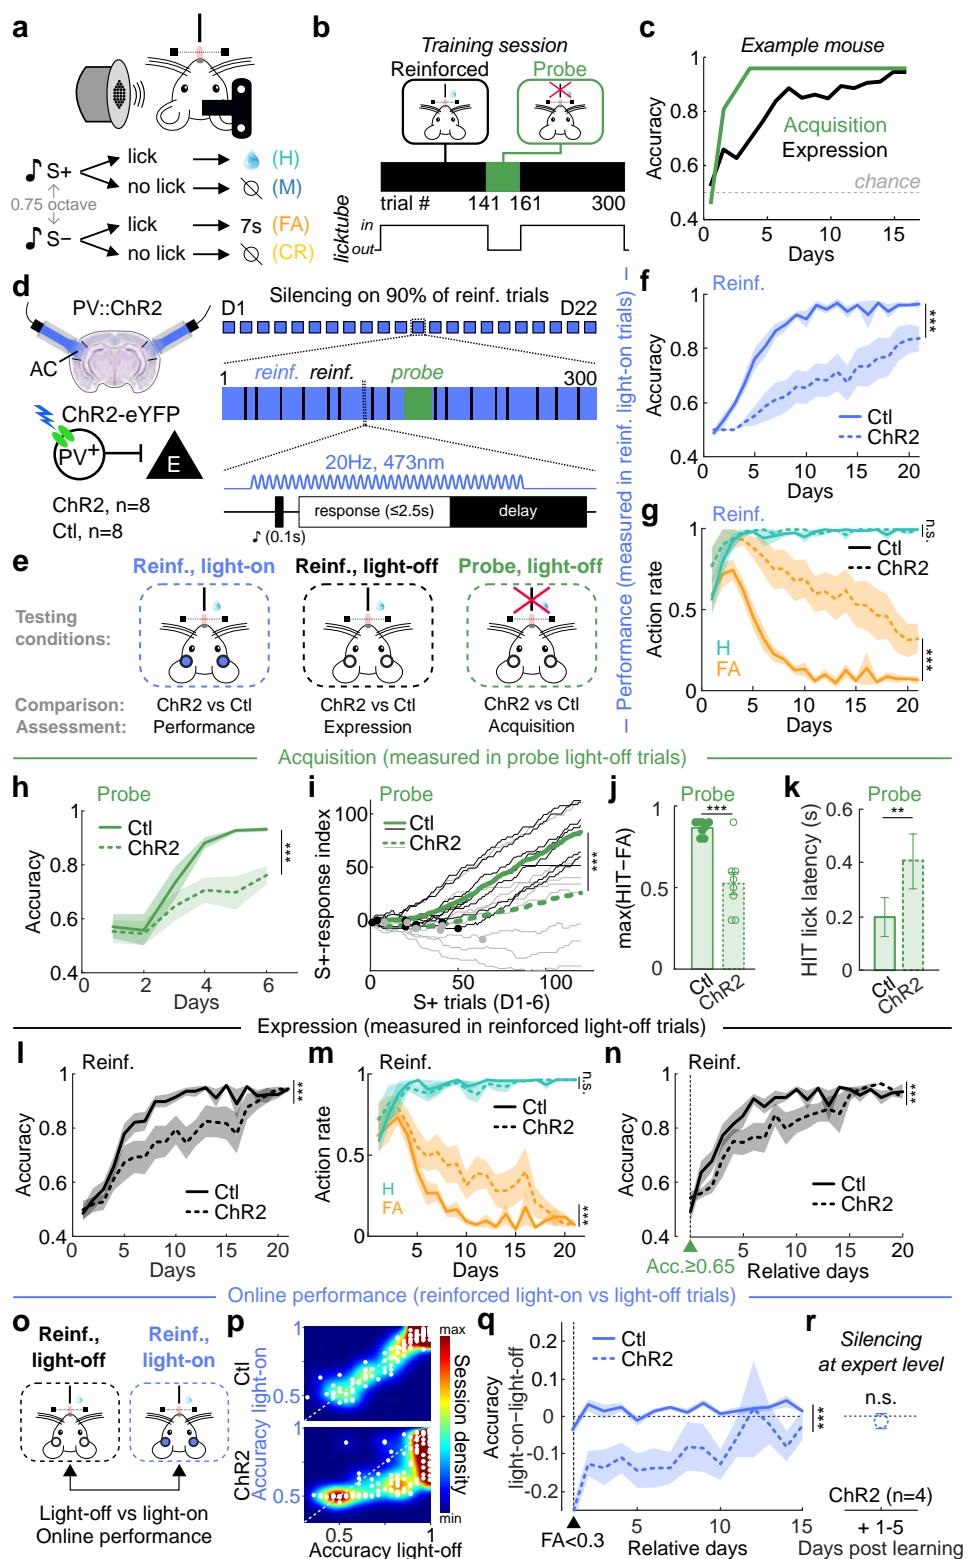


Fig.1. Auditory cortex silencing impairs sound-guided learning and performance during learning.

continued →

Fig.1 (continued).

a, Head-fixed mice were trained on an auditory go/no-go task with $\frac{3}{4}$ -spaced pure tones. H: hit, M: miss, FA: false alarm, CR: correct reject. **b**, Every day during training, task knowledge is probed by omitting reinforcement for 20 trials. **c**, Two distinct learning trajectories are revealed: a fast acquisition of task contingencies (measured in probe trials; green) and a slower knowledge expression (measured in reinforced trials; black). **d**, Probabilistic optogenetic silencing of the auditory cortex over learning. **e**, Testing conditions. **f**, Accuracy in reinforced light-on trials (two-way ANOVA, $p < 10^{-8}$). **g**, Action rate in reinforced light-on trials (HIT, $p = 0.07$; FA, $p < 10^{-33}$). See also Supplementary Figure 4. **h**, Accuracy in probe light-off trials (two-way ANOVA, $p < 10^{-4}$). **i**, Tone response index in S+ trials (see Methods; two-way ANOVA, $p < 10^{-101}$). Black and gray lines are individual mice and dots indicate change points (see Methods). **j**, Maximal difference between hit and FA rates in probe light-off trials over the first 6 days (t-test, $p < 10^{-3}$). **k**, Hit lick latency in probe light-off trials (median \pm s.e.median; Wilcoxon test, $p = 0.007$). **l**, Accuracy in reinforced light-off trials (two-way ANOVA, $p < 10^{-8}$). **m**, Action rate in reinforced light-off trials (two-way ANOVA, HIT: $p = 0.57$, FA: $p < 10^{-8}$). **n**, Accuracy in reinforced light-off trials with inter-subject alignment to the day where probe accuracy ≥ 0.65 (green triangle) (two-way ANOVA, $p < 10^{-5}$). Supplementary Figure 3a-c. **o**, Comparison of light-off versus light-on trials to measure auditory cortex silencing effect on on-line performance. **p**, Session density plot of accuracy in reinforced light-on against light-off. Top, control; bottom, PV-ChR2. See also Supplementary Figure 3d-g. **q**, Within subject accuracy difference in reinforced light-on and light-off trials, aligned to the day where FA rate < 0.3 in reinforced light-off (two-way ANOVA, $p < 10^{-15}$). **r**, Within subject accuracy difference in reinforced light-on and light-off when silencing started at expert level ($n = 4$; t-test, $p = 0.58$). See also Supplementary Figure 6. mean \pm s.e.m.; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$, n.s.: not significant.

70 The auditory cortex is the default pathway for sound-guided learning

71 Lesion studies have suggested that the AC may not be *essential* to learn or execute cue-
72 guided tasks with simple sensory stimuli^{12–15}. However, permanent lesions cannot determine
73 whether the AC is normally *used* for, or causally produces¹⁶, learning in an intact brain.
74 To address this, we exploited a transient silencing approach to prevent the recruitment of
75 alternative pathways^{15,17–20} while also using a probabilistic design to allow assessment of
76 *learning* as distinct from *performance* by measuring behavior on non-silenced trials, thereby
77 avoiding direct effects of silencing on performance.

78 We examined the impact of bilateral cortical silencing of the AC throughout learning (Fig.1a).
79 We probabilistically silenced the AC on 90% of reinforced trials throughout learning ('light-on
80 reinforced', Fig.1d), leaving 10% of reinforced ('light-off reinforced') and 100% of probe trials
81 ('light-off probe') with intact AC activity. Silenced trials were pseudo-randomly sequenced
82 and equally split between S+ and S-. Silencing was achieved by shining blue light bilaterally
83 through cranial windows implanted above the AC of double transgenic mice (n=8) expressing
84 channel rhodopsin (ChR2) in parvalbumin (PV) interneurons^{14,21} (Fig.1d). We confirmed
85 that the excitatory network was effectively silenced using this approach by combining two-
86 photon calcium imaging of excitatory neurons and full-field optogenetic stimulation in PV-
87 ChR2 mice (Supplementary Figure 2). Control mice (n=8) received the same light stimulation
88 but did not express ChR2. This experimental design allowed us to assay the impact of cortical
89 silencing on performance (control vs PV-ChR2 performance on light-on reinforced trials)
90 versus acquisition learning (control vs PV-ChR2 performance on light-off probe trials) and
91 expression learning (control vs PV-ChR2 performance on light-off reinforced trials) (Fig.1e).

92 We first compared performance in light-on reinforced trials between PV-ChR2 and control
93 mice (Fig.1e) and observed a large performance impairment in PV-ChR2 mice (Fig.1f,g).
94 To address whether this performance reduction was accompanied by an impairment in rapid
95 learning, we compared performance in PV-ChR2 and control animals in light-off probe trials
96 (Fig.1e,h-k) when the AC was not silenced and knowledge acquisition can be accurately
97 measured⁷. Accuracy was lower during probe trials in PV-ChR2 mice (Fig.1h), with delayed
98 S+-response learning (Fig.1i), lower discrimination (Fig.1j), and longer lick latency on hit
99 trials (Fig.1k). Rapid acquisition of task knowledge was therefore impaired in PV-ChR2
100 mice.

101 Accuracy was also lower in reinforced light-off trials in PV-ChR2 mice (Fig.1l,m). This
102 remained true even after controlling for their slower task acquisition (Figs.1n, Supplementary
103 Figure 3a-c). These impairments were also apparent in response latency and response vigor
104 (Supplementary Figure 4). Together, these results suggest that the AC is the default pathway
105 for sound-guided reward learning, even when not needed for perceptual sharpening.

¹⁰⁶ **The auditory cortex is used during learning but becomes dispensable at expert**
¹⁰⁷ **levels**

¹⁰⁸ We next sought to understand the contribution of AC activity for the expression of the learned
¹⁰⁹ behavior as animals transitioned to expert performance. Transient inactivation of auditory
¹¹⁰ cortex in expert animals has led to conflicting results, with some reports showing degrada-
¹¹¹ tion of sound-guided behavior^{14,17,22,23} and others not^{14,24,25}. We exploited our probabilistic
¹¹² silencing strategy and compared performance in light-on (AC silenced) versus light-off (AC
¹¹³ functional) reinforced trials within subjects (Fig.1o). Performance on these two trial types
¹¹⁴ was similar at early periods of training, as performance was poor overall (Fig.1p). As train-
¹¹⁵ ing progressed, performance remained poor on light-on trials but improved on light-off trials
¹¹⁶ (Fig.1p), demonstrating that the AC is used for task performance at early and intermediate
¹¹⁷ time-point during learning. Surprisingly, this deficit in performance on light-on trials grad-
¹¹⁸ ually waned (Fig.1p,q), suggesting that while the AC was used during learning, it became
¹¹⁹ dispensable once the mice had mastered the task.

¹²⁰ These results could be explained by three alternative explanations. First, the optogenetic
¹²¹ manipulation *per se* may not be interfering with a task-relevant process but instead could be
¹²² ‘distracting’ the animal, necessitating more time to increase performance in light-on trials.
¹²³ We reasoned that bilateral silencing of another cortical region that is nominally unrelated
¹²⁴ to the task would serve as an important control. We bilaterally silenced the visual cortex
¹²⁵ throughout learning in PV-ChR2 mice and found no evidence of performance impairment in
¹²⁶ light-on trials (Supplementary Figure 5), demonstrating that the performance impairment
¹²⁷ was specific to AC silencing. Second, it is possible that AC silencing altered tone perception,
¹²⁸ increasing task difficulty at the perceptual level in light-on trials. Third, the reduction of
¹²⁹ impairment during light-on trials could be driven by a reduction of the silencing effect with
¹³⁰ time due, for example, to brain damage induced by repeated silencing. To address the
¹³¹ second and third possibilities, we trained a separate cohort of PV-ChR2 mice without daily
¹³² inactivation and, instead, inactivated the AC only after they reached expert performance
¹³³ (see Methods). We observed no impact from AC silencing (Figs.1r, Supplementary Figure
¹³⁴ 6)¹⁴.

¹³⁵ Altogether, these results show that the AC is engaged during learning but is dispensable at
¹³⁶ expert levels, potentially tutoring subcortical structures that take over once the associations
¹³⁷ are learned.

¹³⁸ **Unsupervised discovery of learning-related dynamics by low-rank tensor decom-
¹³⁹ position**

¹⁴⁰ We next sought to understand the nature and dynamics of auditory cortical activity under-
¹⁴¹ lying learning and performance. To do so, we performed longitudinal, two-photon calcium
¹⁴² imaging of thousands of excitatory neurons in mice learning the auditory go/no-go task
¹⁴³ ($n = 5$). A separate group of water-restricted mice was passively exposed to two pure tones

144 over the same duration but with no association with reinforcement ($n = 3$, see Methods;
145 Supplementary Figure 7). This design allowed us to use the passive network as a base-case
146 model to isolate learning-related neural dynamics.

147 We expressed the genetically encoded calcium indicator GCaMP6f under the CaMKII pro-
148 moter, targeting AC layer 2/3 pyramidal neurons. We imaged two planes $\sim 50\mu\text{m}$ apart
149 (Fig.2a), allowing us to record simultaneously hundreds of neurons per animal ($n=7,137$ neu-
150 rons in 8 mice). All mice were passively presented with a series of pure tones (4 to 64kHz,
151 quarter-octave spaced) to characterize auditory tuning properties within the local area of ex-
152 pression. We computed single-neuron tuning curves and then constructed a ‘best frequency’
153 map confirming the location in the AC (Fig.2b). For each mouse, we chose two stimuli that
154 were similarly represented in the recorded population and were 3/4 octaves apart (Fig.2c).
155 We used a custom head-fixation system that allowed for kinematic registration and tracked
156 the activity of the same neurons across weeks, including pre- and post-learning tuning curve
157 sessions ($n = 4,643$ neurons in 8 mice, see Methods; Fig.2d-g).

158 From this high-dimensional dataset, we sought to identify single neurons and neuronal ensem-
159 bles carrying learning-related information, resolve stimulus and non-stimulus related activity
160 within a given trial, identify changes in representation across trials, and determine outcome-
161 specific dynamics. To do so, we organized our data into a 4-dimensional array containing
162 neurons \times time in trial \times trials across learning \times trial outcome (Fig.2h). To identify shared
163 and distinct variability in neuronal populations recorded in passive mice ($n = 2,339$, ‘passive
164 network’) and in learning mice ($n = 2,304$, ‘learning network’), we created a ‘megamouse’ by
165 combining data from all mice and aligning neural activity to learning phase ($n=4,643$ neurons,
166 see Methods; Fig.2i; Supplementary Figure 8). We then used low-rank tensor decomposi-
167 tion to allow unsupervised identification of demixed, low-dimensional neural dynamics across
168 multiple (> 2) dimensions^{26,27} (Supplementary Figure 9 and Supplementary Figure 10a,b;
169 see Methods). The tensor decomposition revealed six neuronal dynamics, each characterized
170 by the four factors of the original tensor (see Methods; Figs.2j, Supplementary Figure 10c,d,
171 Supplementary Figure 11d). These six dynamics represented independent computations per-
172 formed by the auditory cortical networks.

173 Projecting the product of the decomposition into principal component subspace showed that
174 learning and passive networks exhibit almost orthogonal dynamics (Fig.2k; Supplementary
175 Figure 10f,g) and that the neural dynamics of different trial types evolved further apart in
176 the learning network than in the passive network (Supplementary Figure 10h,i). Importantly,
177 we ensured that the identified dynamics were not driven by isolated mice (Supplementary
178 Figure 10e). Therefore, decomposition of the megamouse tensor discovered distinct dynamics
179 exhibited by passive versus learning networks.

180 For further analyses, we attributed each dynamic to individual neurons based on the neuron’s
181 maximum weight (‘unique participation’; Fig.2l; see Methods and Supplementary Figure 11).
182 This allowed us to map the six dynamics onto six distinct cell ensembles, i.e. groups of neu-

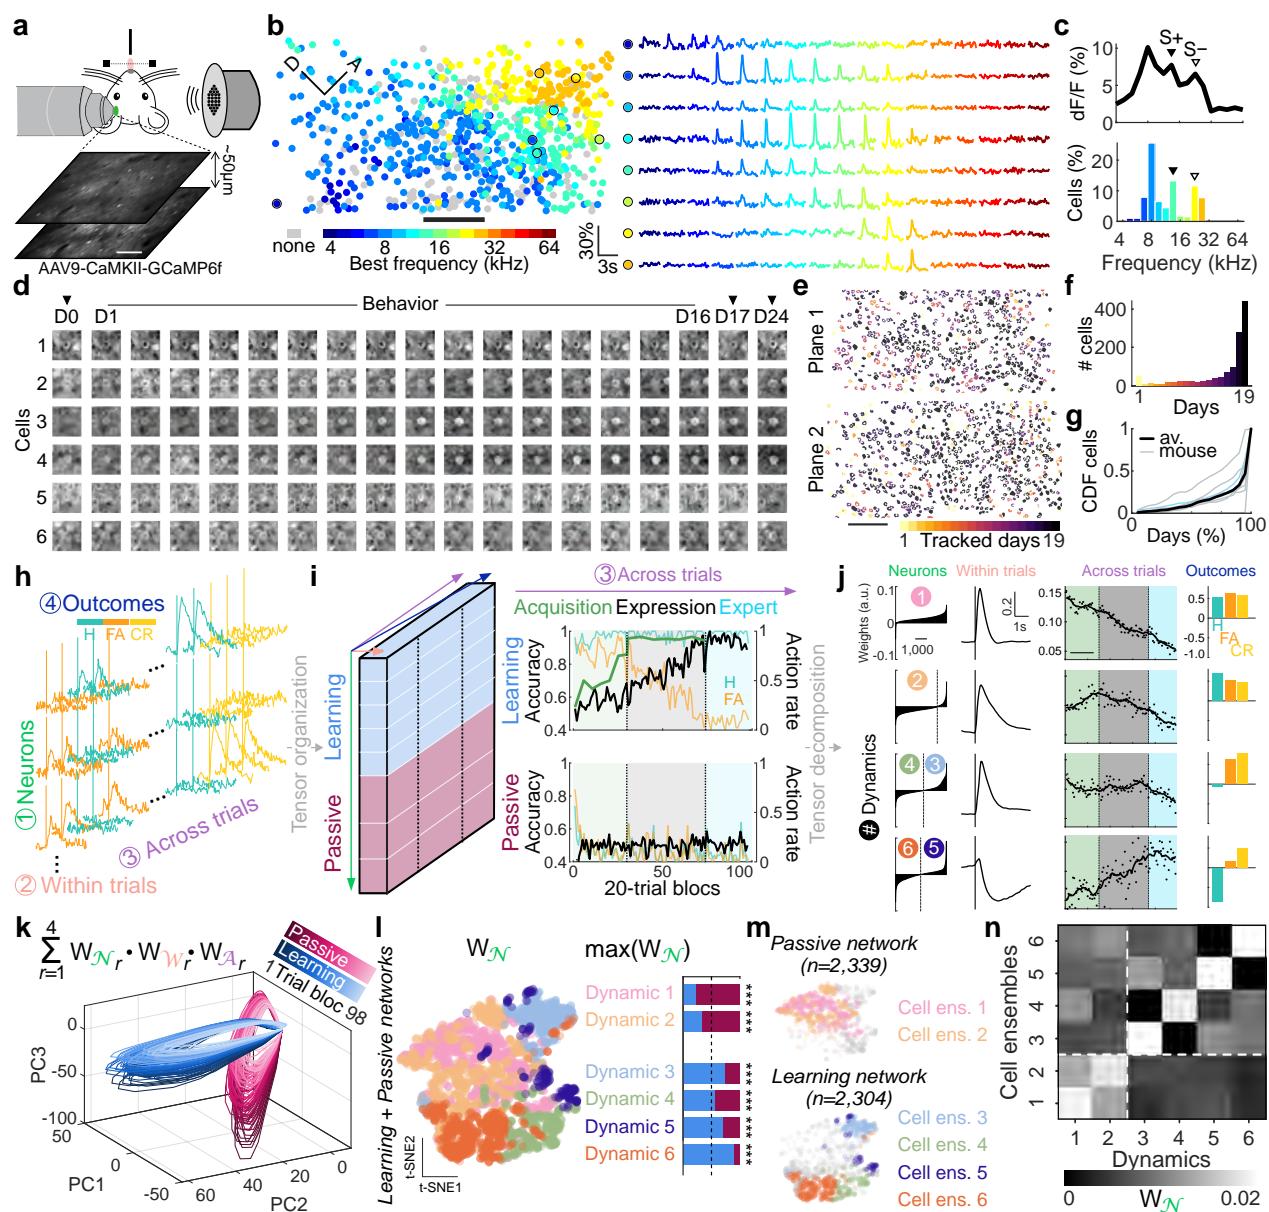


Fig.2. Low-rank tensor decomposition reveals learning-related network dynamics.

a, Multi-plane, longitudinal two-photon calcium imaging of layer 2/3 excitatory network in the auditory cortex during learning ($n = 5$ mice) or passive exposure ($n = 3$ mice; see Methods). **b**, Tonotopic organization of the field of view of one example mouse before learning (left). Cells are colored according to their best frequency and tone-evoked responses of example cells circled in black to 17 pure tones ranging from 4 to 64 kHz are displayed on the right. **c**, Tone-evoked activity (top) and proportion of responsive cells (bottom) to pure tones. S+ and S- (filled and unfilled triangles, respectively) are chosen for training in the task based on their equal representation in the field of view in **b**. **d**, Six example cells tracked everyday across weeks. **e**, Two planes recorded in one example mouse. Cells are colored according to the number of days tracked among the 19 recording sessions in this mouse. **f**, Distribution of number of tracked days per cells in **e**. **g**, Cumulative distribution of tracked cells according to the percentage of recording sessions. Data for mouse in **e** is the light blue line. **h**, Calcium data is arranged by neurons \times time within trial (-1 to $+4s$ relative to tone onset, vertical line) \times trials over time \times trial outcomes.

continued →

Fig.2 (continued).

i, Activity from all Learning and Passive cells are concatenated together to create a fourth-order tensor (megamouse; left). In the 3rd, ‘across trials’ dimension, data is aligned across mice according to learning phases: Acquisition (performance increases in probe trials), Expression (performance increases in reinforced trials), and Expert (high, stable performance in reinforced trials; see Methods and Supplementary Figure 8). **j**, Megamouse tensor decomposition identifies six neuronal dynamics (numbered; see Methods) that are characterized by a set of four factors: Neuron, Within trial, Across trial, and Outcome (see also Supplementary Figure 10). **k**, Projection of the tensor decomposition output onto principal subspace. W_{N_r} , W_{W_r} and W_{A_r} indicate neuronal, within trial and across trial weights for a component r , respectively. **l**, t-distributed stochastic neighbor embedding (t-SNE) projections of neuronal weights. Each dot represents a cell, colored according to the neuronal dynamic it contributed in the most. Bars (right) display the proportion of learning and passive cells among the highest contributors for each dynamic. Dynamics 1 and 2 are driven by the passive network (burgundy), while Dynamics 3 to 6 are driven by the learning network (blue). **m**, In the passive network, the highest contributing cells in Dynamic 1 define cell ensemble 1, and highest contributing cells in Dynamic 2 define cell ensemble 2. Similarly, in the learning network, cell ensembles 3 to 6 are constituted of the highest contributing cells to Dynamics 3 to 6, respectively. **n**, Absolute weights of cell ensembles across the six identified dynamics. Neurons can participate in more than one dynamic.

183 rons maximally encoding a particular network-specific dynamic (Fig.2m and Supplementary
184 Figure 11d). It is important to note that individual neurons (and corresponding ensembles)
185 could exhibit mixed selectivity for the six dynamics, which allows an individual neurons to
186 contribute to multiple, independent computations (Fig.2n).

187 **Learning counteracts tone-evoked habituation by maintaining stimulus selectivity
188 in distinct cell populations**

189 A prevailing view in sensory systems holds that sensory cortices subserve associative learning
190 through plasticity of the cue representation^{5,28–36}. This model posits that individual neurons
191 (via changes in sensory tuning) and neural populations (via cortical map expansion) enhance
192 the representation of behaviorally relevant cues for use by downstream regions^{37–39}. These
193 studies, however, measure neural tuning and map expansion outside of the task context in
194 a ‘pre’ and ‘post’ learning design and infer that plasticity of cue representations reflects the
195 mechanistic role of the sensory cortex. To assess this model, we initially focused on the cell
196 ensembles that exhibited classical stimulus-evoked activity (Fig.2j), namely cell ensembles
197 1–4.

198 We observed a prominent signature of stimulus-evoked habituation over hundreds to thou-
199 sands of trials. This habituation dominated activity in passive networks, as seen in cell en-
200 sembles 1 and 2 which represented ~77% (1,803/2,339) of all passive cells (Fig.3a,d). These
201 neurons exhibited stimulus-evoked activation (cell ensemble 1) or suppression (cell ensemble
2), both of which decreased in amplitude over time (Fig.3b-c,e-f). These cell ensembles were

203 not stimulus selective and displayed the same dynamic in both stimulus 1 (S1) and stimulus
204 2 (S2) trials (Fig.3b,e). These ensembles thus reflected the broad-based suppression of
205 non-selective neurons after long-term repeated presentation of the same sounds.

206 Stimulus-evoked responses in learning networks were observed in cell ensembles 3 and 4
207 (Fig.3g-j). This includes a high selectivity for the S- (cell ensemble 3) or S+ (cell ensemble 4)
208 cues (Fig.3g-j). Cell ensemble 3 consisted of 19% of the Learning cell population (Fig.3g), and
209 displayed a slight habituation but mainly a strong preference for the S- throughout learning
210 (Fig.3h), while cell ensemble 4 (12% of total learning cells; Fig.3j) exhibited S+ selectivity
211 throughout learning (Fig.3j). Cell ensembles 3 and 4 were more tone responsive and tone
212 selective than cell ensembles 1 and 2 (Fig.3k,l). Stimulus-evoked activity analyses across
213 days of all recorded neurons ($n = 7,137$) also support these results (Supplementary Figure
214 12, Supplementary Figure 13). Therefore, learning counteracted tone-evoked habituation by
215 maintaining distinct ensembles that encoded either the S+ or S- selectively.

216 **Learning was not associated with cortical map expansion**

217 To directly test representational expansion and tuning shifts, we conducted a series of anal-
218 yses focusing on stimulus-evoked responses before (pre-task) and after (post-task) learning,
219 akin to classical measures of tuning and tonotopy. We computed the change in surface area
220 occupied by S+ and S- preferring cells in tuning curve sessions, outside the task (Fig.3m).
221 Surprisingly, we observed no increase in the map-level representation of the S+ or S- af-
222 ter learning, and instead, observed a modest decrease (Fig.3m-n). In addition to the best
223 frequency representation, the fraction of neurons responding to the S+ and S- decreased
224 (Fig.3o) and the response amplitude of neurons that were initially tuned to the S+ and S-
225 was lower after learning (Fig.3p). Interestingly, while we observed no increase in representa-
226 tion to the S+ and S-, learning networks favored the representation of frequencies in between
227 S+ and S-, but not higher or lower as seen in passive networks (Fig.3n). Finally, using our
228 passive networks as a base-case comparison, we calculated the local changes in the tonotopic
229 map structure (Fig.3q). Learning networks were surprisingly stable and exhibited less local
230 changes than passive networks (Fig.3r). These pre- vs post-learning changes in responsiveness
231 and tonotopy thus mirrored the responsiveness observed online during learning (in dynamics
232 1 and 2) in a stable, tracked network ($n=4,643$ neurons, Fig.3a-l), as well as when we include
233 all neurons from each session ($n=7,137$ neurons) (Supplementary Figure 13). Altogether, our
234 results suggest that cortical map expansion and changes in single-neuron tuning are unlikely
235 to be the substrate for associative learning^{40,41}.

236 **Tone-restricted silencing only partially impairs learning and performance**

237 We next sought to understand the extent to which the maintenance of stimulus-selectivity
238 by learning networks was important to learning and performing the task. We performed
239 daily bilateral silencing of AC during stimulus presentation throughout learning (Supple-

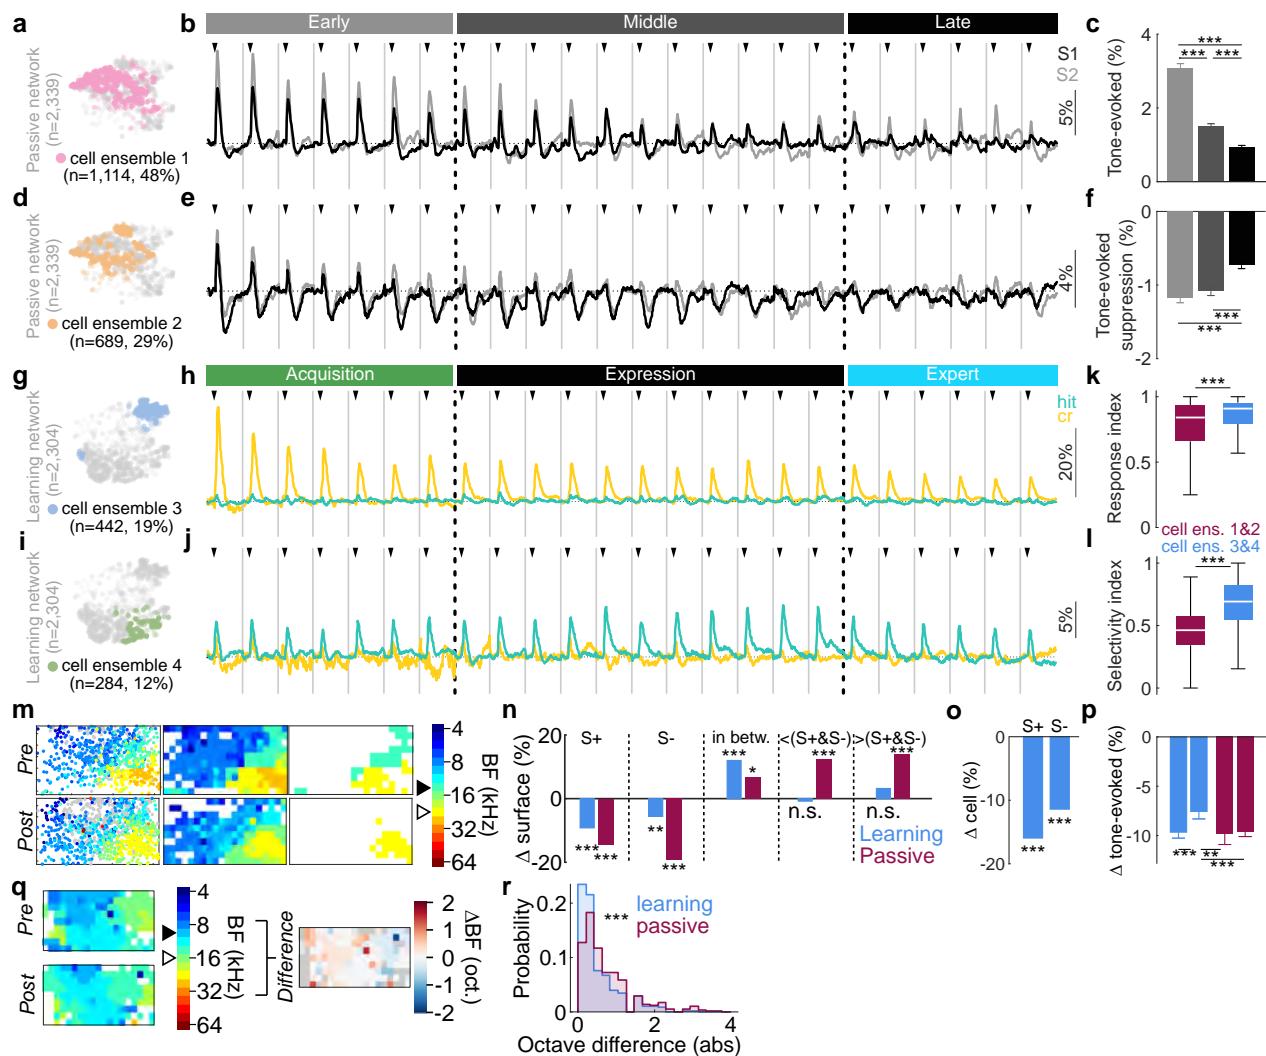


Fig.3. Learning counteracts tone-evoked habituation by maintaining stimulus selectivity in distinct populations.

a, Representation of cell ensemble 1 in the Passive network. **b**, Average activity of cell ensemble 1 in S1 (black) and S2 (gray) trials across time in 80-trial blocks. Black triangles indicate tone onset, gray lines delimit averaged trial blocks. Black dashed lines separate time phases indicated by light to dark gray rectangles at the top: early, middle and late (see Methods). **c**, Cell ensemble 1 tone-evoked calcium responses across time phases for S1 and S2 trials combined (Friedman test, $p = 1.26 \cdot 10^{-291}$). **d**, Representation of cell ensemble 2 in the Passive network. **e**, Average activity of cell ensemble 2 in S1 and S2 trials across time. **f**, Cell ensemble 2 tone-evoked calcium responses across time phases for S1 and S2 trials combined (Friedman test, $p = 7.32 \cdot 10^{-121}$). **g**, Representation of cell ensemble 3 in the Learning network. **h**, Average activity of cell ensemble 3 in hit (green) and CR (yellow) trials across learning in 80-trial blocks. Black triangles indicate tone onset, gray lines delimit averaged trial blocks. Black dashed lines separate learning phases indicated by colored rectangles at the top: Acquisition, Expression and Expert (see Methods). **i**, Representation of cell ensemble 4 in the Learning network. **j**, Average activity of cell ensemble 4 in hit and CR trials across learning. **k**, Response index (response probability over learning; see Methods) of cell ensembles 1 and 2 (red) vs cell ensembles 3 and 4 (blue) (Wilcoxon test, $p = 1.23 \cdot 10^{-30}$). **l**, Selectivity index. **m**, **Post** vs **Pre** BF heatmaps. **n**, Δ surface (%) for S+, S-, in betw., <(S+&S-)>(S+&S-), and Learning vs Passive. **o**, Δ cell (%) for S+ and S-. **p**, Δ tone-evoked (%). **q**, **Post** vs **Pre** BF heatmaps. **r**, Probability vs Octave difference (abs).

continued →

Fig.3 (continued).

l, Selectivity index (see Methods) of cell ensembles 1 and 2 (red) vs cell ensembles 3 and 4 (blue) (Wilcoxon test, $p = 1.37 \cdot 10^{-94}$). **m**, Pre (top raw) and post (bottom raw) learning tonotopic maps (left), after spatial binning (middle) and restricted to surface with S+ (filled triangle) and S- (open triangle) best frequency (right) of one example mouse. **n**, Change in surface representation of S+ and S- pre- vs post-task learning (Learning) or pre- vs post-passive exposure (Passive) (binomial proportion tests). **o**, Pre vs post-learning change in percentage of neurons responsive to S+ and S- (binomial proportion tests). **p**, Pre vs post-learning change in tone-evoked responses of pre-task S+ and S- responsive neurons (KW test, $p = 2.77 \cdot 10^{-5}$). **q**, Pre- vs post-learning comparison of local best frequency differences in tonotopic maps. **r**, Distribution of local differences (from difference maps in **q**) in Learning versus Passive. median \pm s.e.median; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$, n.s.: not significant.

240 mentary Figure 14a). Tone-restricted AC silencing impaired task performance throughout
241 learning (Supplementary Figure 14b-e), task acquisition (Supplementary Figure 14f-i), and
242 online performance during learning, with gradual fading of the effect at expert performance
243 (Supplementary Figure 14n-q). Accuracy and action rate were not affected in reinforced
244 light-off trials (Supplementary Figure 14j-k), but PV-ChR2 mice lick more and faster to the
245 S- (Supplementary Figure 14l-m), suggesting that tone-restricted AC silencing also impaired
246 expression, but to a lesser extent than full-trial silencing. Altogether, these results showed
247 that information carried by the AC network in the tone-evoked window is used during learn-
248 ing. Interestingly, tone-restricted silencing impacted learning less than full trial silencing
249 across nearly all measures (Fig.1, Supplementary Figure 14), suggesting that activity *af-*
250 *ter* the tone-evoked window was critical for rapid contingency acquisition and performance
251 during learning.

252 **Rapid emergence of reward prediction activity in the auditory cortex**

253 The sensory cortex is widely considered to be specialized for perception by interpreting com-
254 plex sensory objects^{42,43} or adjusting representations of behaviorally-relevant stimuli^{2,33,37,44,45}.
255 Recent evidence, however, suggests that sensory cortical neurons directly encode non-sensory
256 variables such as movement⁴⁶⁻⁴⁹, reward timing⁵⁰⁻⁵³, expectation^{54,55}, and context^{23,45,56-63}.
257 Conjoint representations of sensory and non-sensory variables in the same network could
258 further hone perception or, alternatively, subserve more integrative associative processes.

259 Inspection of the within-trial dynamics of learning-driven cell ensembles 5 and 6 suggested
260 that these neurons exhibited non-canonical activity in the form of a signal that occurred
261 late in the trial, delayed from the tone-evoked response (Fig.2j). This late-in-trial signal
262 increased over learning and was trial type selective (Fig.2j). We next sought to further
263 explore the encoding properties of these two cell ensembles. Cell ensemble 5 ($n = 155$ cells
264 from the learning network), exhibited late-in-trial activity on hit trials (licking to the S+)
265 that increased with learning (Fig.4a). This delayed activity was not apparent on correct S-

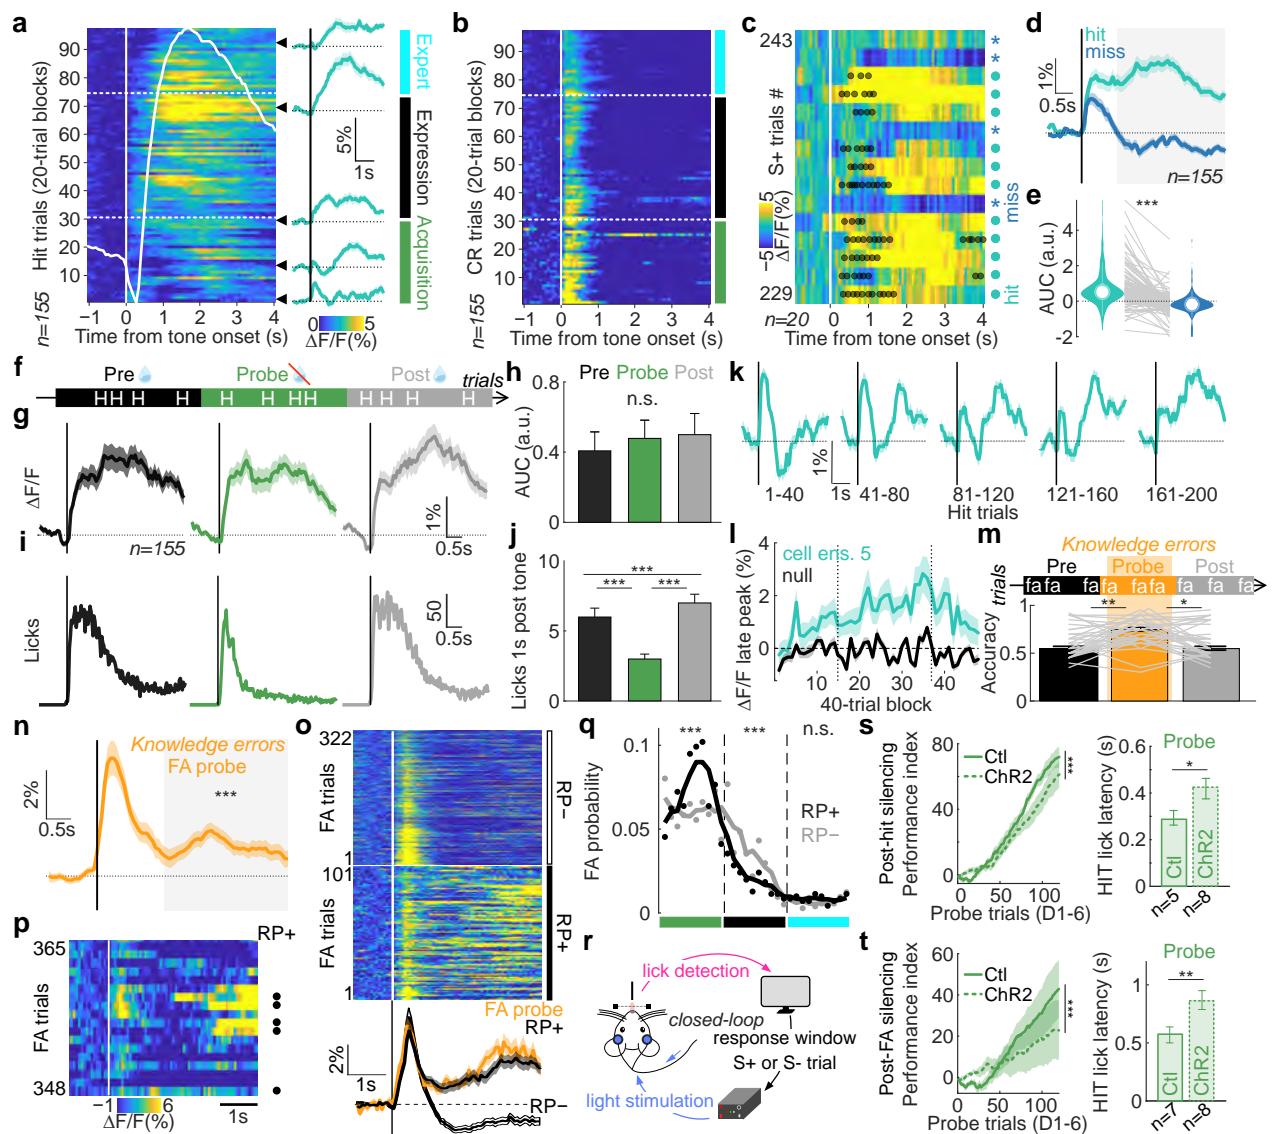


Fig.4. Rapid emergence of reward prediction encoding drives learning.

a, Heat map of cell ensemble 5 activity ($n = 155$ cells) across learning phases (delimited by horizontal white dashed lines) in hit trials (20-trial blocks). White trace represents the average trial trace. Inserts (right) show average activity at time indicated by black triangles. Colored rectangles indicate learning phases: Acquisition (green), Expression (black) and Expert (blue). **b**, Heat map of cell ensemble 5 activity across learning phases (delimited by horizontal white dashed lines) in CR trials (20-trial blocks). **c**, Heat map of the activity of a fraction of cells from cell ensemble 5 ($n = 20$ cells) from one example mouse across consecutive S+ trials. Black dots indicate licks. Trial outcome is represented on the right (green circle: hit; blue stars: miss). **d**, Cell ensemble 5 activity in hit vs miss trials (time and number matched, see Methods and Supplementary Figure 15a). **e**, Area under the curve (AUC) quantification of data in gray rectangle in **d** (Wilcoxon signed rank test, $p = 6.78 \cdot 10^{-21}$). **f**, Procedure of reinforced and probe hit trial (H) matching. **g**, Average cell ensemble 5 activity in reinforced hit trials immediately before (black) or after (gray) probe hit trials (green). **h**, AUC quantification of data in **h** (Friedman test, $p = 0.3071$). **i**, Lick PSTHs in reinforced hit trials immediately before (black) or after (gray) probe hit trials (green). **j**, Quantification of number of licks in 1-s window post-tone (KW test, $p = 3.18 \cdot 10^{-56}$).

continued →

Fig.4 (continued).

k, Average activity of cell ensemble 5 over the first five blocks of 40-reinforced hit trials in learning. **l**, Late peak activity in HIT trials across learning phases of cell ensemble 5 (green) and low weighted cells (null, black). **m**, Procedure of reinforced and probe FA trial (fa) matching (top) and corresponding local accuracy quantification (bottom; see Methods; repeated measures ANOVA, $p = 3.16 \cdot 10^{-4}$). **n**, Average cell ensemble 5 activity in FA trials in the probe, non-reinforced context (orange). AUC late-in-trial (gray rectangle) compared to zero (Wilcoxon signed rank test, $p = 1.46 \cdot 10^{-8}$). **o**, Average activity of cell ensemble 5 ($n = 51$ cells) from one example mouse in FA trials in the reinforced context ($n = 423$) after classification based on the detection of a reward prediction signal. Bottom, average activity of FA trials with (RP+, $n = 101$) or without (RP-, $n = 322$) reward prediction signal, and activity during FA trials in the probe context ($n = 19$ trials, orange) reflecting ‘knowledge’ errors (see also Supplementary Figure 16). **p**, Heat map of the activity of a fraction of cells from cell ensemble 5 ($n = 51$ cells) from the same example mouse in **o** across consecutive FA trials in the reinforced context. Identification of a RP signal is represented by a black dot (right). **q**, Distribution of RP+ and RP- FA trials over learning in learning mice (binomial proportion tests, Acquisition, $p = 1.65 \cdot 10^{-7}$, Expression, $p = 3.32 \cdot 10^{-10}$, Expert, $p = 0.22$). **r**, Trial-specific closed-loop optogenetic AC inactivation over learning. **s**, Performance index (left, see Methods; two-way ANOVA, $p = 2.11 \cdot 10^{-21}$) and hit lick latency (right; two-way ANOVA, $p = 0.013$) in probe context in post-hit silencing experiments. **t**, Performance index (left, see Methods; two-way ANOVA, $p = 6.36 \cdot 10^{-5}$) and hit lick latency (right; two-way ANOVA, $p = 0.008$) in probe context in post-FA silencing experiments.

²⁶⁶ trials (correct reject, CR), where neurons exhibited classical stimulus-evoked response that
²⁶⁷ habituated over learning (Fig.4b).

²⁶⁸ To understand the nature of the late-in-trial activity, we exploited our multiple trial types to
²⁶⁹ disambiguate the contribution of sensory, motor, and reward signals. To assess whether the
²⁷⁰ late-in-trial signal was a delayed form of sensory activity, we compared activity in hit trials to
²⁷¹ activity in trials where the same stimulus was presented but the mice did not lick and did not
²⁷² get rewarded (miss trials, Figs.1a and 4c-e). To ensure an appropriate comparison between
²⁷³ hit and miss trials, we generated a balanced set of trials that were matched in number (given
²⁷⁴ that miss trials were less frequent) and occurred within the same time period (given that
²⁷⁵ the signal amplitude evolved with learning) (Supplementary Figure 15a). Cell ensemble 5
²⁷⁶ did not exhibit late-in-trial activity on miss trials (Fig.4c-e), discarding the possibility that
²⁷⁷ it reflected a delayed sensory response. We then asked whether this activity reflected reward
²⁷⁸ consumption. We compared cell ensemble activity during hit trials in the reinforced context
²⁷⁹ to the activity during hit trials in the probe context (Fig.4f), where the mice expected reward
²⁸⁰ and thus correctly licked to the S+ but the reward was omitted (Fig.1b). We matched the
²⁸¹ number of trials between reinforced and probe contexts and controlled for within-session
²⁸² and across-session changes by comparing probe hit trials to reinforced hit trials immediately
²⁸³ before and after the probe block (Fig.4f). Strikingly, late-in-trial activity was preserved
²⁸⁴ in probe trials (Fig.4g,h), indicating that it did not reflect reward consumption. Finally,
²⁸⁵ although movement has been reported to decrease auditory cortical activity^{46,64–66}, we sought

286 to understand the degree to which this late-in-trial signal could be driven by licking itself. To
287 do this, we first exploited probe hit trials where the lick rate was strongly reduced compared
288 to reinforced hit trials (Fig.4i,j). We observed no difference in the late-in-trial neural signal
289 and could thus conclude that the signal was not due to ongoing licking (Fig.4i,j). Second, we
290 tested the possibility that this late-in-trial signal was driven by the initiation of a lick bout
291 as compared to the ongoing licking activity. We isolated spontaneous lick bouts in between
292 training blocks and observed that the cell ensemble was not lick-responsive (Supplementary
293 Figure 15b,c). In addition, if lick initiation drove this activity, we would also expect to
294 see it on false alarm trials (incorrect licking to the S-). For this analysis, we focused on
295 false alarms that occurred after task acquisition, as these errors are unlikely to be errors
296 due to imperfect task knowledge. We observed no systematic late-in-trial activity on these
297 trials (Supplementary Figure 15d) even though the licking pattern in false alarm trials was
298 similar to that during probe hit trials (Supplementary Figure 15e). Taken together, the
299 late-in-trial activity did not reflect stimulus, reward consumption, licking, nor lick initiation.
300 Instead, these results showed that cell ensemble 5 encoded the higher-order process of reward
301 prediction (RP).

302 We next sought to identify the precise moment when a contingency is formed by identify-
303 ing the trials when this reward prediction signal emerged. Initially, these neurons exhibited
304 classical tone-evoked responses but then abruptly and within only 40 hit trials, developed a
305 robust reward prediction activity (Fig.4k, Supplementary Figure 15f). This reward predic-
306 tion signal continued to develop over Acquisition, strengthened during Expression, and then
307 surprisingly receded at Expert level when learning is nominally complete (Fig.4a,l, Supple-
308 mentary Figure 15g). This longitudinal temporal dynamic mirrored our optogenetic results
309 which demonstrates that the AC is the default pathway for learning but then becomes dis-
310 pensable at expert levels. Altogether, these results show that a reward prediction signal
311 rapidly emerges at the timescale of Acquisition in auditory cortical networks.

312 **Revealing the underlying cognitive drivers of errors**

313 Identifying the cognitive drivers of errors is particularly challenging during learning⁴. Errors
314 during learning are typically considered ‘mistakes’ while discriminative contingencies (task
315 knowledge) are still forming. However, errors arise not only from knowledge-related mistakes
316 (for which animals incorrectly expect reward), but also from factors such as impulsivity,
317 disengagement, and exploration (for which animals do not expect reward). While detailed
318 behavioral inspection has been a promising route to uncover the nature of errors¹¹, an al-
319 ternative approach is to use neural activity itself. Given our findings of reward prediction
320 encoding on correct trials, we hypothesized that the same signal would be present when an-
321 imals make ‘knowledge-related’ errors, when animals incorrectly ‘expected’ rewards on S-
322 trials. To address this, we first focused on the occasional false alarms (FA) that occurred
323 during probe trials, as they reflected errors of task knowledge (Fig.4m)⁷. Strikingly, we ob-
324 served a robust reward prediction activity in these trials (Fig.4n), strongly suggesting that

325 animals were indeed expecting reward. We next reasoned that such knowledge errors should
326 be present not only on probe trials, but also in a subset of reinforced trials, interspersed with
327 non-knowledge errors. We classified individual FA trials in the reinforced context based on
328 the presence of a reward prediction signal (see Methods; Supplementary Figure 16a). We
329 identified a significant proportion of trials that exhibited robust reward prediction activity,
330 but also many that did not (Fig.4o, Supplementary Figure 16b). The reward prediction
331 signal was identical to that observed in probe trials (Fig.4o, Supplementary Figure 16d),
332 providing further confidence that these were indeed knowledge errors. These data suggest
333 that we could isolate knowledge errors using neural data, which was not possible from behav-
334 ioral inspection alone (Supplementary Figure 16c). Interestingly, we found that knowledge
335 errors were interspersed with errors that did not elicit reward prediction activity (Fig.4p).
336 Finally, we hypothesized that knowledge errors should predominantly occur during the Ac-
337 quisition phase of behavior, when animals are still learning the discriminative contingencies.
338 We computed the fraction of RP+ (knowledge-related errors) and RP- (non-knowledge er-
339 rors) over time and found that RP+ errors peaked during the Acquisition phase of learning,
340 and rarely occurred during Expression or Expert phases of behavior (Fig.4q, Supplementary
341 Figure 15d). These results demonstrate that the internal cognitive drivers of errors may be
342 accessible from neural data, which is particularly striking when behavior alone is insufficient.

343 **Reward prediction activity provides the core teaching signal**

344 Learning theory proposes that animals learn from correct actions that are rewarded but also
345 from incorrect actions that are not rewarded⁶⁷. This allows animals to select the appropriate
346 action after reward-predictive (S+) versus non-predictive (S-) cues. Given the presence of
347 the reward prediction activity on correct S+ trials (throughout learning) and incorrect S-
348 trials (early in learning), we reasoned that silencing auditory cortical activity during the
349 post-response period could impact learning and/or performance. To test this, we performed
350 closed-loop probabilistic optogenetic silencing of the AC whereby light was delivered upon
351 lick detection in 90% of either S+ reinforced trials ($n = 5$ control, $n = 8$ PV-ChR2 mice) or,
352 in a separate cohort, S- reinforced trials ($n = 7$ control, $n = 8$ PV-ChR2 mice; see Methods;
353 Fig.4r, Supplementary Figure 17a, Supplementary Figure 18a). No light was delivered in
354 10% of S+ reinforced trials and 100% of probe trials. Given that the light was delivered after
355 the instrumental lick response, the effect of the manipulation could not affect the instrumen-
356 tal behavior on the current trial, only on subsequent ones. To confirm this, we calculated
357 the difference in performance between light-on and light-off trials and observed no differ-
358 ence (Supplementary Figure 17b-d and Supplementary Figure 18b-d). In the S+ cohort,
359 post-hit silencing weakened the stimulus-action association (Fig.4s), delayed cue-response
360 discrimination (Figs.4s), but did not impact probe accuracy over the first 6 days (Supple-
361 mentary Figure 17e-g). Importantly, the same silencing protocol above the visual cortex
362 ($n = 6$ PV-ChR2 mice) had no effect on behavior, confirming that these effects were specific
363 to the auditory cortex (Supplementary Figure 17k,l). In the S- cohort, post-FA silencing

364 weakened the stimulus-action association as measured on hit trials (Fig.4t), robustly de-
 365 layed cue-response discrimination (Fig.4t, Supplementary Figure 17g), and impaired probe
 366 accuracy over the first 6 days (Supplementary Figure 17e,f). Accuracy of PV-ChR2 mice
 367 was lower than control in the reinforced context in both experiments (Supplementary Figure
 368 17h and Supplementary Figure 18h), with lower hit rate and higher FA rate (Supplemen-
 369 tary Figure 17i and Supplementary Figure 18i), and longer response latencies on hit trials
 370 (Supplementary Figure 17j and Supplementary Figure 18j), suggesting an impairment of ex-
 371 pression. Overall, these closed-loop manipulations showed that AC activity at the time of
 372 the reward prediction signal in both hit and FA trials was used by the animal for the task
 373 acquisition and expression. These data also demonstrate that learning is sensitive to cortical
 374 silencing on mistakes (FA trials) suggesting that in a go/no-go paradigm, reward feedback on
 375 error trials is crucial to the learning process. Altogether, these results suggest that reward
 376 prediction activity in auditory cortical networks is used as a teaching signal during learning.

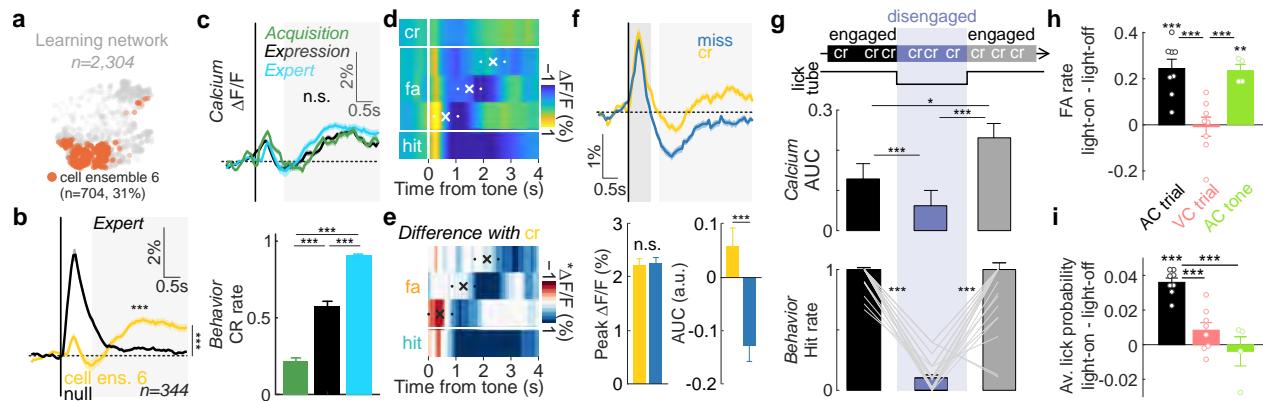


Fig.5. Action suppression signals in the AC induce suppression of licking.

a, Representation of cell ensemble 6 ($n = 704$ cells) in the Learning network. **b**, Average activity of cell ensemble 6 (yellow) versus cells that do not contribute to this dynamic (null, black) in CR trials in Expert phase (Wilcoxon test, $p = 7.44 \cdot 10^{-17}$). **c**, Average activity of cell ensemble 6 in CR trials (top) and CR rate (bottom) during Acquisition (green), Expression (black) and Expert (blue) phases (KW test, $p = 0.09$). **d**, Heat map of cell ensemble 6 activity in hit, FA and CR trials. FA trials are binned according to lick latencies (white dots, latency range extrema; white cross, mean latency). **e**, Heat map of cell ensemble 6 activity in hit and FA trials significantly different from CR trials (Wilcoxon tests, red, higher; blue, lower; white, n.s.). **f**, Average cell ensemble 6 activity in miss and CR trials (time and number matched, see Methods; middle). Quantifications of tone-evoked activity (bottom left; Wilcoxon signed rank test, $p = 0.84$) and late-in-trial AUC (bottom right; Wilcoxon signed rank test, $p = 5.24 \cdot 10^{-26}$). **g**, Procedure of reinforced and probe CR trial matching (top) and corresponding calcium activity (middle; Friedman test, $p = 1.36 \cdot 10^{-11}$) and local hit rate (bottom; Friedman test, $p = 3.45 \cdot 10^{-11}$). **h**, FA rate difference between light-on and light-off trials in PV-ChR2 mice (two-way ANOVA, $p = 7.20 \cdot 10^{-16}$; t-tests compared to 0, $p = 4.96 \cdot 10^{-4}$, $p = 0.96$, $p = 0.002$). Auditory or visual cortex were inhibited during the full trial (AC trial, $n = 8$; VC trial, $n = 8$) or AC was silenced during tone presentation only (AC tone, $n = 4$). **i**, Average lick probability in FA light-on versus FA light-off trials (two-way ANOVA, $p = 1.18 \cdot 10^{-5}$; t-tests compared to 0, $p = 1.94 \cdot 10^{-6}$, $p = 0.10$, $p = 0.68$).

377 **Encoding of action suppression enables task performance**

378 A critical requirement in a go/no-go task is the ability to suppress responding to the non-
379 rewarded, S- cue. In our task, we demonstrate that mice have the capacity to withhold
380 licking to the S- very early in learning (as shown in probe trials during the acquisition
381 phase) but continue to lick for hundreds to thousands of trials when being reinforced and
382 throughout Expression. Here, we ask the extent to which the AC mediates this form of action
383 suppression. Neurons in cell ensemble 6 ($n = 704$, 31% of learning networks; Fig.5a), but
384 not non-member cells, exhibited late-in-trial activity when animals correctly withheld from
385 licking on S- trials (correct rejects, CR; Fig.5b, Supplementary Figure 19a-b). This signal
386 was stable throughout training despite the strong increase of CR rate over learning (Fig.5c,
387 Supplementary Figure 19c-d). This all-or-none attribute suggested that this late-in-trial
388 activation was tied to performance rather than being a signal used for learning. Once mice
389 acquired the task contingencies, they essentially learned to inhibit a licking response to the
390 S- tone. We therefore thought to test the hypothesis that late-in-trial activation in CR trials
391 reflected action suppression. First, we reasoned that activity in FA and CR trials should be
392 similar until the moment of suppression failure (i.e. first lick). We compared the activity
393 of cell ensemble 6 in CR vs FA trials, i.e. when mice fail to withhold licking (see Methods)
394 exploiting the different first lick latencies in FA trials (Fig.5d). We observed that calcium
395 activity dropped abruptly in FA trials at the time of the first lick compared to CR trials
396 (Fig.5d,e, Supplementary Figure 19e). Second, if lick suppression is an active contingency-
397 specific process, the late-in-trial activation should be specific for correct rejections for the
398 S- tone, and not observed when the animal did not lick in response to the S+ tone (miss
399 trials). Given that miss trials were rare and sporadic, we controlled for the effect of time
400 over learning and difference in the number of trials for each outcome type (see Methods) and
401 did not observe late-in-trial activation on miss trials despite similar peak activity after tone
402 onset in miss and CR trials (Fig.5f). Third, we reasoned that if this activity reflects the
403 active process of action suppression, the signal should decrease when the animal disengaged
404 from the task. We therefore compared late-in-trial activity in CR trials immediately before,
405 during and after short blocks of disengagement (see Methods) and observed that the activity
406 dropped significantly when mice transiently disengaged from the task (Fig.5g). These data
407 suggest that the auditory cortex integrates a higher-order action suppression signal.

408 Finally, we wondered whether the action suppression activity in the AC was causal to perfor-
409 mance during learning. To test this, we reasoned that silencing the AC network throughout
410 S- trials should increase the FA rate but also the lick probability (since the action suppres-
411 sion neurons are silenced during this period). In contrast, silencing the AC network only
412 during the stimulus period should increase the FA rate but not impact the lick probability
413 when the light is off (Supplementary Figure 19f). We first compared the FA rate between
414 light-on and light-off trials in PV-ChR2 mice during full trial silencing (Fig.5h) and observed
415 a marked increase in FA rate and lick probability (Fig.5h,i). Importantly, this effect was not
416 the result of the perception of optogenetic manipulation *per se* as suppression of the visual

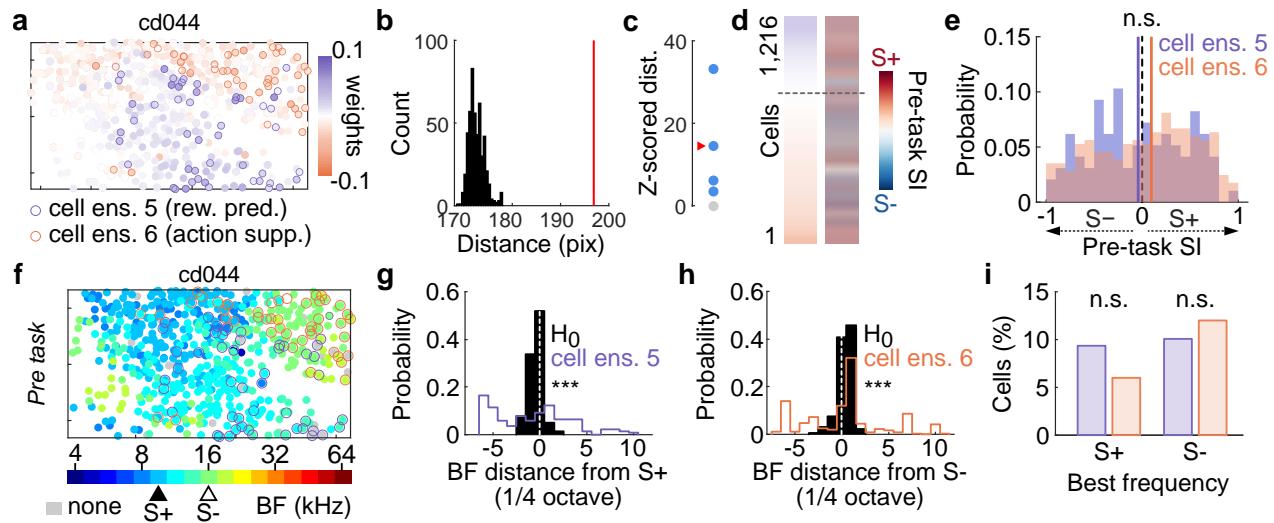


Fig.6. Reward prediction and action suppression signals emerged in segregated neuronal populations and do not rely on underlying stimulus selectivity

a, Spatial distribution of reward prediction (purple circles) and behavioral inhibition (orange circles) cell ensembles in an example mouse. Color scale indicates neuronal weights in Dynamics 5 (purple) and 6 (orange). **b**, Median of cell distance between cell ensembles compared to shuffle distribution ($n = 500$) for example mouse in **a**. The null hypothesis is that the distance between the two ensembles is no different than chance (i.e. no spatial organization). **c**, Z-scored distances between clusters per mouse (blue: significant; gray: non-significant). Red arrow points to example mouse in **a**. **d**, Neuronal weights in Dynamics 5 and 6 of cells from learning mice ($n = 1,216$, left) and their pre-task stimulus selectivity index (right). **e**, Distributions of pre-task stimulus selectivity of cell ensembles 5 and 6 (KS test, $p = 0.25$, Wilcoxon test, $p = 0.18$). **f**, Pre-task tonotopic map of the example mice in **a**. Cells are colored according to their best frequency (BF). Frequencies used as S+ and S- for training are indicated by full and empty triangles, respectively. **g**, Distribution of BF distance from S+ for reward prediction cell ensemble (purple). Null hypothesis is that reward prediction cells have a BF as close to S+ as possible (black; see Methods; KS test, $p = 3.81 \cdot 10^{-9}$). **h**, Distribution of BF distance from S- for action suppression cell ensemble (orange). Null hypothesis is that action suppression cells had a BF as close to S- as possible (black; see Methods; KS test, $p = 9.21 \cdot 10^{-16}$). **i**, Proportions of S+ and S--preferring cells in reward prediction and action suppression cell ensembles (binomial proportion tests, S+, $p = 0.17$, S-, $p = 0.53$)

417 cortex in PV-ChR2 mice did not have this effect (Fig.5h,i). In contrast, restricting silenc-
418 ing to the stimulus period increased FA rate while not affecting lick probability (Fig.5h,i),
419 suggesting that the late-in-trial activity in CR trials was critical for the maintenance of ac-
420 tion suppression. Altogether, these results showed that action suppression is encoded in the
421 auditory cortex and is instrumental for performance during learning.

422 **Higher-order contingency ensembles are spatially clustered and uncoupled from**
423 **sensory representations**

424 We next asked the extent to which reward prediction and action suppression ensembles
425 mapped onto the underlying stimulus properties of the AC. We exploited the spatial res-
426 olution of two-photon imaging to characterize the spatial distribution of reward prediction
427 and action suppression neurons in the AC network. Strikingly, we observed that the two cell
428 ensembles were spatially clustered (Fig.6a-c). To determine whether this organization was
429 driven by the neuron's pre-learning stimulus selectivity, we calculated the selectivity index
430 (SI) of each neuron before training to test whether neurons selective for the S+ preferentially
431 became reward prediction neurons and S- selective neurons preferentially became action
432 suppression neurons. We observed no difference in SI distribution between reward predictive
433 and action suppression neurons (Fig.6d,e), suggesting that pre-task stimulus selectivity was
434 not predictive of either reward prediction or action suppression. We then asked whether the
435 spatial location of reward prediction and action suppression neurons aligned with the under-
436 lying tonotopic map. In other words, did action suppression neurons have S- tone for best
437 frequency, and were reward prediction neurons preferentially responsive to S+ tone (Fig.6f;
438 see Methods)? We found that this was not the case (Fig.6g,h), with similar proportion
439 of S+- and S--preferring neurons in reward prediction and action suppression cell ensem-
440 bles (Fig.6i). Therefore, contingency-related ensembles clustered into spatial domains that
441 were uncoupled from underlying stimulus selectivity and tonotopy, indicating a higher-order
442 functional segregation within the AC.

443 **Discussion**

444 Learning-related neural dynamics are traditionally defined as task-specific neural activity
445 changes that occur at the timescale of an animal's performance improvements in the learn-
446 ing, i.e. a reinforced context⁸. Using this conceptual and experimental framework, perceptual
447 and instrumental (reward-based) learning and their underlying neural dynamics have been
448 described as slow and gradual e.g^{2,44}, with animals requiring thousands of trials to learn
449 low information-content tasks^{3,4}. We took advantage of a recent behavioral paradigm⁷ that
450 uses non-reinforced probe trials to show that task knowledge emerges more rapidly and ear-
451 lier than behavioral performance improvement in the learning, reinforced context. Using this
452 powerful behavioral manipulation to quantitatively assess when the animals acquired the task
453 contingencies, we aligned our neuronal recordings to learning stages between animals while

454 preserving trial-based resolution, and took advantage of an unsupervised, dimensionality re-
455 duction method across multiple timescales²⁶ to identify learning-specific neural dynamics.
456 We observed that reward prediction activity emerged remarkably fast - within tens of trials
457 and on the first day of training - in the AC, hundreds to thousands of trials before noticeable
458 performance improvements. The AC thus exhibits latent knowledge of the task (encoded in
459 the network but not behaviorally apparent) with animals experiencing periods when knowl-
460 edge of environmental contingencies (between cues, actions, and rewards) becomes rapidly
461 encoded in the brain, perhaps reflecting an insight-like moment. The latent task knowledge
462 was manifested not as changes in sensory representations, but as the emergence of discrete
463 ensembles encoding reward prediction (needed for identifying that a particular cue signals
464 reward availability) and action suppression (needed for suppressing licking on S- trials).
465 These computations were spatially clustered and developed in a manner that was uncoupled
466 from the underlying stimulus-related processing that takes place in the AC, suggesting a
467 higher-order functional organization. Overall, we find that AC contains separable and causal
468 neural dynamics for both learning and performance.

469 Our results call for a revision of the classical view of the sensory cortex, according to which
470 its primary role is to process and interpret sensory stimuli. We propose instead that the sen-
471 sory cortex is better described as a sensory-enriched associative cortex, driving rapid forms
472 of associative learning and where sensory and associative functions are intrinsically inter-
473 mingled (i.e. co-exist within the same network) but computationally separable (Fig.6). This
474 function of the sensory cortex may have thus far been obscured by the use of complex sensory
475 objects that recruit the sensory cortex for object-level processing, making it difficult to iso-
476 late non-perceptual learning computations. Finally, it is important to note that our results
477 do not contradict studies that demonstrate single-neuron tuning curve shifts and tonotopic
478 map plasticity when animals learn perceptually challenging tasks. Our revised model of the
479 sensory cortex would suggest that perceptual sharpening and complex object processing can
480 be subserved by stimulus-related plasticity while the higher-order computations related to
481 associative learning and performance occur in parallel. We expect this view will apply beyond
482 rodents, as rich encoding of non-sensory and task-relevant variables has also been described
483 in human and non-human primate sensory cortical areas^{68–70}.

484 The detailed input-output circuit that enables reward prediction and action suppression
485 computations remains an important area for future exploration. One possibility is that
486 ascending neuromodulatory inputs^{23,29,71–75} and top-down projections from motor and frontal
487 regions^{44,76–78} serve as critical non-sensory inputs to the sensory cortex. The sensory cortex
488 may then integrate and generate higher-order computations that are incorporated by broader
489 decision-related circuits (e.g. frontal cortex, striatum and amygdala) to enable rapid learning
490 and ongoing performance.

491 References

492 [1] Blake, D. T., Strata, F., Churchland, A. K. & Merzenich, M. M. Neural correlates of
493 instrumental learning in primary auditory cortex. *Proceedings of the National Academy of
494 Sciences of the United States of America* **99**, 10114–10119 (2002).

495 [2] Poort, J. *et al.* Learning Enhances Sensory and Multiple Non-sensory Representations
496 in Primary Visual Cortex. *Neuron* **86**, 1478–1490 (2015). URL <http://dx.doi.org/10.1016/j.neuron.2015.05.037>.

498 [3] Meister, M. Learning, fast and slow. *Current Opinion in Neurobiology* **75**, 102555
499 (2022). URL <https://doi.org/10.1016/j.conb.2022.102555>.

500 [4] Carandini, M. & Churchland, A. K. Probing perceptual decisions in rodents. *Nature
501 Neuroscience* **16**, 824–831 (2013).

502 [5] Polley, D. B., Heiser, M. A., Blake, D. T., Schreiner, C. E. & Merzenich, M. M. Associative
503 learning shapes the neural code for stimulus magnitude in primary auditory cortex.
504 *Proceedings of the National Academy of Sciences of the United States of America* **101**,
505 16351–16356 (2004).

506 [6] Gallistel, C. R., Fairhurst, S. & Balsam, P. The learning curve: Implications of a
507 quantitative analysis. *Proceedings of the National Academy of Sciences of the United
508 States of America* **101**, 13124–13131 (2004).

509 [7] Kuchibhotla, K. V. *et al.* Dissociating task acquisition from expression during learning
510 reveals latent knowledge. *Nature Communications* **10** (2019).

511 [8] Moore, S. & Kuchibhotla, K. V. Slow or sudden: Re-interpreting the learning curve
512 for modern systems neuroscience. *IBRO Neuroscience Reports* **13**, 9–14 (2022). URL
513 <https://doi.org/10.1016/j.ibneur.2022.05.006>.

514 [9] Schreiner, C. E. & Polley, D. B. Auditory map plasticity: diversity in causes and
515 consequences. *Current Opinion in Neurobiology* **24**, 143–156 (2014). URL <http://dx.doi.org/10.1016/j.conb.2013.11.009>. NIHMS150003.

517 [10] Sachidhanandam, S., Sreenivasan, V., Kyriakatos, A., Kremer, Y. & Petersen, C. C.
518 Membrane potential correlates of sensory perception in mouse barrel cortex. *Nature
519 Neuroscience* **16**, 1671–1677 (2013).

520 [11] Zhu, Z. & Kuchibhotla, K. V. Performance errors during learning reflect a dynamic
521 choice strategy. *Current Biology* **34**, 1–11 (2024).

522 [12] Ohl, F. W., Wetzel, W., Wagner, T., Rech, A. & Scheich, H. Bilateral Ablation of
523 Auditory Cortex in Mongolian Gerbil Affects Discrimination of Frequency Modulated
524 Tones but not of Pure Tones. *Learning* **6**, 347–362 (1999).

525 [13] Romanski, L. M. & LeDoux, J. E. Bilateral destruction of neocortical and perirhinal
526 projection targets of the acoustic thalamus does not disrupt auditory fear conditioning.
527 *Neuroscience Letters* **142**, 228–232 (1992).

528 [14] Ceballo, S., Piwkowska, Z., Bourg, J. & Bathellier, B. Targeted Cortical Manipulation.
529 *Neuron* **104**, 1–12 (2019).

530 [15] O’sullivan, C., Weible, A. P. & Wehr, M. Auditory cortex contributes to discrimination
531 of pure tones. *eNeuro* **6**, 1–14 (2019).

532 [16] Barack, D. L. *et al.* A call for more clarity around causality in neuroscience. *Trends*
533 in *Neurosciences* **45**, 654–655 (2022). URL <https://doi.org/10.1016/j.tins.2022.06.003>.

535 [17] Talwar, S. K., Musial, P. G. & Gerstein, G. L. Role of mammalian auditory cortex in the
536 perception of elementary sound properties. *Journal of Neurophysiology* **85**, 2350–2358
537 (2001).

538 [18] Boatman, J. A. & Kim, J. J. A thalamo-cortico-amygdala pathway mediates auditory
539 fear conditioning in the intact brain. *European Journal of Neuroscience* **24**, 894–900
540 (2006).

541 [19] Hartwigsen, G. Flexible Redistribution in Cognitive Networks. *Trends in Cognitive*
542 *Sciences* **22**, 687–698 (2018). URL <http://dx.doi.org/10.1016/j.tics.2018.05.008>.

544 [20] Slonina, Z. A., Poole, K. C. & Bizley, J. K. What can we learn from inactivation
545 studies? Lessons from auditory cortex. *Trends in Neurosciences* **45**, 64–77 (2022). URL
546 <https://doi.org/10.1016/j.tins.2021.10.005>.

547 [21] Li, N. *et al.* Spatiotemporal constraints on optogenetic inactivation in cortical circuits.
548 *eLife* 1–31 (2019).

549 [22] Jaramillo, S. & Zador, A. M. The auditory cortex mediates the perceptual effects of
550 acoustic temporal expectation. *Nature Neuroscience* **14**, 246–253 (2011). URL <http://dx.doi.org/10.1038/nn.2688>.

552 [23] Kuchibhotla, K. V. *et al.* Parallel processing by cortical inhibition enables context-
553 dependent behavior. *Nature Neuroscience* **20**, 62–71 (2016). [15334406](#).

554 [24] Gimenez, T. L., Lorenc, M. & Jaramillo, S. Adaptive categorization of sound frequency
555 does not require the auditory cortex in rats. *Journal of Neurophysiology* **114**, 1137–1145
556 (2015).

557 [25] Stoilova, V. V. *et al.* Auditory cortex reflects goal-directed movement but is not necessary
558 for behavioral adaptation in sound-cued reward tracking. *Journal of Neurophysiology*
559 **124**, 1056–1071 (2020).

560 [26] Williams, A. H. *et al.* Unsupervised Discovery of Demixed, Low-Dimensional Neural
561 Dynamics across Multiple Timescales through Tensor Component Analysis. *Neuron* **98**,
562 1099–1115.e8 (2018). URL <https://doi.org/10.1016/j.neuron.2018.05.015>.

563 [27] Acar, E., Dunlavy, D. M., Kolda, T. G. & Mørup, M. Scalable tensor factorizations for
564 incomplete data. *Chemometrics and Intelligent Laboratory Systems* **106**, 41–56 (2011).
565 URL <http://dx.doi.org/10.1016/j.chemolab.2010.08.004>. **1005.2197**.

566 [28] Jenkins, W. M., Merzenich, M. M., Ochs, M. T., Allard, T. & Guic-Robles, E. Functional
567 reorganization of primary somatosensory cortex in adult owl monkeys after behaviorally
568 controlled tactile stimulation. *Journal of Neurophysiology* **63**, 82–104 (1990).

569 [29] Kilgard, M. P. & Merzenich, M. M. Cortical map reorganization enabled by nucleus
570 basalis activity. *Science* **279**, 1714–1718 (1998).

571 [30] Recanzone, G. H., Merzenich, M. M., Jenkins, W. M., Grajski, K. A. & Dinse, H. R.
572 Topographic reorganization of the hand representation in cortical area 3b of owl monkeys
573 trained in a frequency-discrimination task. *Journal of Neurophysiology* **67**, 1031–1056
574 (1992).

575 [31] Xerri, C., Stern, J. M. & Merzenich, M. M. Alterations of the cortical representation of
576 the rat ventrum induced by nursing behavior. *Journal of Neuroscience* **14**, 1710–1721
577 (1994).

578 [32] Xerri, C., Merzenich, M. M., Jenkins, W. & Santucci, S. Representational plasticity in
579 cortical area 3b paralleling tactal- motor skill acquisition in adult monkeys. *Cerebral
580 Cortex* **9**, 264–276 (1999).

581 [33] Polley, D. B. Perceptual Learning Directs Auditory Cortical Map Reorganization
582 through Top-Down Influences. *Journal of Neuroscience* **26**, 4970–4982 (2006).

583 [34] Rutkowski, R. G. & Weinberger, N. M. Encoding of learned importance of sound by
584 magnitude of representational area in primary auditory cortex. *Proceedings of the Na-
585 tional Academy of Sciences of the United States of America* **102**, 13664–13669 (2005).

586 [35] Blake, D. T., Heiser, M. A., Caywood, M. & Merzenich, M. M. Experience-Dependent
587 Adult Cortical Plasticity Requires Cognitive Association between Sensation and Reward.
588 *Neuron* **52**, 371–381 (2006).

589 [36] Elias, G. A., Biesczad, K. M. & Weinberger, N. M. Learning strategy refinement
590 reverses early sensory cortical map expansion but not behavior: Support for a theory
591 of directed cortical substrates of learning and memory. *Neurobiology of Learning and
592 Memory* **126**, 39–55 (2015).

593 [37] Weinberger, N. M. Dynamic regulation of receptive fields and maps in the adult sensory
594 cortex. *Annual Review of Neuroscience* **18**, 129–158 (1995).

595 [38] Buonomano, D. V. & Merzenich, M. M. Cortical plasticity: From synapses to maps.
596 *Annual Review of Neuroscience* **21**, 149–186 (1998).

597 [39] Weinberger, N. M. Specific long-term memory traces in primary auditory cortex. *Nature*
598 *Reviews Neuroscience* **5**, 279–290 (2004).

599 [40] Maor, I. *et al.* Neural Correlates of Learning Pure Tones or Natural Sounds in the
600 Auditory Cortex. *Frontiers in Neural Circuits* **13**, 1–21 (2020).

601 [41] Shepard, K. N., Chong, K. K. & Liu, R. C. Contrast enhancement without transient
602 map expansion for species-specific vocalizations in core auditory cortex during learning.
603 *eNeuro* **3** (2016).

604 [42] Hubel, D. H. & Wiesel, T. N. Receptive fields of single neurones in the
605 cat's striate cortex. *The Journal of Physiology* **148**, 574–591 (1959). URL
606 <http://www.ncbi.nlm.nih.gov/article/1363130>
607 7B%25%7D7B%7B&%7D%7B%25%7D7Dtool=pmcentrez%7B%25%7D7B%7B&%7D%7B%
608 25%7D7Drendertype=abstract%7B%25%7D7B%7B%25%7D25%7B%25%7D7D5Cnhttp:
609 //doi.wiley.com/10.1113/jphysiol.2009.174151.

610 [43] Olshausen, B. A. & Field, D. J. What Is the Other 85 Percent of V1 Doing? In van
611 Hemmen, J. L. & Sejnowski, T. J. (eds.) *23 Problems in Systems Neuroscience*, 1–29
612 (Oxford University Press, 2009).

613 [44] Banerjee, A. *et al.* Value-guided remapping of sensory circuits by lateral orbitofrontal
614 cortex in reversal learning. *Nature* **585**, 245–250 (2020).

615 [45] Kato, H. K., Gillet, S. N. & Isaacson, J. S. Flexible Sensory Representations in Auditory
616 Cortex Driven by Behavioral Relevance. *Neuron* **88**, 1027–1039 (2015). URL <http://dx.doi.org/10.1016/j.neuron.2015.10.024>. 15334406.

617 [46] Schneider, D. M., Nelson, A. & Mooney, R. A synaptic and circuit basis for corollary
618 discharge in the auditory cortex. *Nature* **513**, 189–194 (2014). URL <http://dx.doi.org/10.1038/nature13724>. NIHMS150003.

619 [47] Audette, N. J., Zhou, W. X., La Chioma, A. & Schneider, D. M. Precise movement-based
620 predictions in the mouse auditory cortex. *Current Biology* **32**, 4925–4940.e6 (2022). URL
621 <https://doi.org/10.1016/j.cub.2022.09.064>.

622 [48] Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial
623 neural dynamics are dominated by richly varied movements. *Nature Neuroscience* **22**,
624 1677–1686 (2019). URL <http://dx.doi.org/10.1038/s41593-019-0502-4>.

625 [49] Niell, C. M. & Stryker, M. P. Modulation of Visual Responses by Behavioral State in
626 Mouse Visual Cortex. *Neuron* **65**, 472–479 (2010). URL <http://dx.doi.org/10.1016/j.neuron.2010.01.033>.

630 [50] Shuler, M. G. & Bear, M. F. Reward Timing in the Primary Visual Cortex. *Science*
631 **311**, 1606–1609 (2006).

632 [51] Sharma, J. *et al.* Spatial attention and temporal expectation under timed uncertainty
633 predictably modulate neuronal responses in Monkey V1. *Cerebral Cortex* **25**, 2894–2906
634 (2015).

635 [52] Namboodiri, V. M. K., Huertas, M. A., Monk, K. J., Shouval, H. Z. & Shuler, M. G.
636 Visually cued action timing in the primary visual cortex. *Neuron* **86**, 319–330 (2015).
637 URL <http://dx.doi.org/10.1016/j.neuron.2015.02.043>.

638 [53] Suri, H. *et al.* A Cortico-Striatal Circuit for Sound-Triggered Prediction of
639 Reward Timing. *bioRxiv : the preprint server for biology* (2023). URL
640 <http://www.ncbi.nlm.nih.gov/pubmed/38045246> %0A <http://www.ncbi.nlm.nih.gov/articlerender.fcgi?artid=PMC10690153>.

641 [54] Stanisor, L., Van Der Togt, C., Pennartz, C. M. & Roelfsema, P. R. A unified selection
642 signal for attention and reward in primary visual cortex. *Proceedings of the National
643 Academy of Sciences of the United States of America* **110**, 9136–9141 (2013).

644 [55] Guo, L., Weems, J. T., Walker, W. I., Levichev, A. & Jaramillo, S. Choice-selective
645 neurons in the auditory cortex and in its striatal target encode reward expectation.
646 *Journal of Neuroscience* **39**, 3687–3697 (2019).

647 [56] Carcea, I., Insanally, M. N. & Froemke, R. C. Dynamics of auditory cortical activity
648 during behavioural engagement and auditory perception. *Nature Communications* **8**,
649 1–12 (2017).

650 [57] Otazu, G. H., Tai, L. H., Yang, Y. & Zador, A. M. Engaging in an auditory task
651 suppresses responses in auditory cortex. *Nature Neuroscience* **12**, 646–654 (2009).
652 [NIHMS150003](https://doi.org/10.1038/nn.2340).

653 [58] Yin, P., Fritz, J. B. & Shamma, S. A. Rapid spectrotemporal plasticity in primary
654 auditory cortex during behavior. *Journal of Neuroscience* **34**, 4396–4408 (2014).

655 [59] David, S. V., Fritz, J. B. & Shamma, S. A. Task reward structure shapes rapid receptive
656 field plasticity in auditory cortex. *Proceedings of the National Academy of Sciences* **109**,
657 2144–2149 (2012).

658 [60] Fritz, J., Shamma, S., Elhilali, M. & Klein, D. Rapid task-related plasticity of spec-
659 trotemporal receptive fields in primary auditory cortex. *Nature Neuroscience* **6**, 1216–
660 1223 (2003).

661 [61] Atiani, S. *et al.* Emergent selectivity for task-relevant stimuli in higher-order auditory
662 cortex. *Neuron* **82**, 486–499 (2014). URL <http://dx.doi.org/10.1016/j.neuron.2014.02.029>.

663

664

665 [62] De Franceschi, G. & Barkat, T. R. Task-induced modulations of neuronal activity along
666 the auditory pathway. *Cell Reports* **37**, 110115 (2021). URL <https://doi.org/10.1016/j.celrep.2021.110115>.

668 [63] Chillale, R. K., Shamma, S., Ostojic, S. & Boubenec, Y. Dynamics and maintenance
669 of categorical responses in primary auditory cortex during task engagement. *eLife* **12**,
670 1–20 (2023).

671 [64] Nelson, A. *et al.* A Circuit for Motor Cortical Modulation of Auditory Cortical Activity.
672 *Journal of Neuroscience* **33**, 14342–14353 (2013). URL <http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.2275-13.2013>.

674 [65] Schneider, D. M. & Mooney, R. Motor-related signals in the auditory system for listening
675 and learning. *Current Opinion in Neurobiology* **33**, 78–84 (2015). URL <http://dx.doi.org/10.1016/j.conb.2015.03.004>.

677 [66] Yang, Y., Lee, J. & Kim, G. Integration of locomotion and auditory signals in the mouse
678 inferior colliculus. *eLife* **9**, 1–17 (2020).

679 [67] Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* (MIT Press,
680 Cambridge, MA, 2018), second edi edn.

681 [68] Fitzgerald, T. H., Friston, K. J. & Dolan, R. J. Characterising reward outcome signals
682 in sensory cortex. *NeuroImage* **83**, 329–334 (2013). URL <http://dx.doi.org/10.1016/j.neuroimage.2013.06.061>.

684 [69] Huang, Y., Heil, P. & Brosch, M. Associations between sounds and actions in early audi-
685 tory cortex of nonhuman primates. *eLife* 1–22 (2019). URL <https://elifesciences.org/articles/43281>.

687 [70] Brosch, M., Selezneva, E. & Scheich, H. Representation of reward feedback in primate
688 auditory cortex. *Frontiers in Systems Neuroscience* **5**, 1–12 (2011).

689 [71] Zhu, F., Elnozahy, S., Lawlor, J. & Kuchibhotla, K. V. The cholinergic basal forebrain
690 provides a parallel channel for state-dependent sensory signaling to auditory cortex.
691 *Nature Neuroscience* **26**, 810–819 (2023).

692 [72] Metherate, R. & Weinberger, N. M. Cholinergic Modulation of Responses to Single
693 Tones Produces Tone-Specific Receptive Field Alterations in Cat Auditory Cortex.
694 *Synapse* **6**, 133–145 (1990).

695 [73] Pinto, L. *et al.* Fast modulation of visual perception by basal forebrain cholinergic neu-
696 rons. *Nature neuroscience* 1–10 (2013). URL <http://www.ncbi.nlm.nih.gov/pubmed/24162654>.

698 [74] Guo, W., Robert, B. & Polley, D. B. The Cholinergic Basal Forebrain Links Auditory

699 Stimuli with Delayed Reinforcement to Support Learning. *Neuron* **1–14** (2019). URL
700 <https://linkinghub.elsevier.com/retrieve/pii/S089662731930594X>.

701 [75] Rho, H.-J., Kim, J.-H. & Lee, S.-H. Function of Selective Neuromodulatory Projections
702 in the Mammalian Cerebral Cortex: Comparison Between Cholinergic and No-
703 radrenergic Systems. *Frontiers in Neural Circuits* **12**, 1–13 (2018). URL <https://www.frontiersin.org/article/10.3389/fncir.2018.00047/full>.

705 [76] Liu, D. *et al.* Orbitofrontal control of visual cortex gain promotes visual associative
706 learning. *Nature Communications* **11**, 1–14 (2020). URL <http://dx.doi.org/10.1038/s41467-020-16609-7>.

708 [77] Winkowski, D. E. *et al.* Orbitofrontal Cortex Neurons Respond to Sound and Activate
709 Primary Auditory Cortex Neurons. *Cerebral Cortex* **28**, 868–879 (2018).

710 [78] Miller, K. J., Botvinick, M. M. & Brody, C. D. Value representations in the rodent
711 orbitofrontal cortex drive learning, not choice. *eLife* **11**, 1–27 (2022).

712 [79] Pachitariu, M. *et al.* Suite2p: beyond 10,000 neurons with standard two-photon mi-
713 croscopy. bioRxiv. *Bioarxiv* **20**, 2017 (2017).

714 [80] Vega-Villar, M., Horvitz, J. C. & Nicola, S. M. NMDA receptor-dependent plas-
715 ticity in the nucleus accumbens connects reward-predictive cues to approach re-
716 sponses. *Nature Communications* **10** (2019). URL <http://dx.doi.org/10.1038/s41467-019-12387-z>.

718 [81] Hitchcock, F. L. The expression of a tensor or a polyadic as a sum of products. *Journal
719 of Mathematics and Physics* **6**, 164–189 (1927).

720 [82] Hitchcock, F. L. Multilple invariants and generalized rank of a p-way matrix or tensor.
721 *Journal of Mathematics and Physics* **7**, 39–79 (1927).

722 [83] Kolda, T. G. & Bader, B. W. Tensor Decompositions and Review. *SIAM Review* **51**,
723 455–500 (2009).

724 [84] Harshman, R. a. Foundations of the PARAFAC procedure: Models and conditions for an
725 “explanatory” multimodal factor analysis. *UCLA Working Papers in Phonetics* **16**, 1–
726 84 (1970). URL <http://www.psychology.uwo.ca/faculty/harshman/wppffac0.pdf>.

727 [85] Carroll, J. D. & Chang, J. J. Analysis of individual differences in multidimensional
728 scaling via an n-way generalization of ”Eckart-Young” decomposition. *Psychometrika*
729 **35**, 283–319 (1970).

730 [86] Kruskal, J. B. Three-way arrays: rank and uniqueness of trilinear decompositions, with
731 application to arithmetic complexity and statistics. *Linear Algebra and Its Applications*
732 **18**, 95–138 (1977).

733 [87] Kolda, T. G. Multilinear operators for higher-order decompositions. *Sandia report*
734 (2006).

735 [88] Acar, E., Dunlavy, D. M. & Kolda, T. G. A Scalable Optimization Approach for Fitting
736 Canonical Tensor Decompositions. *Journal of Chemometrics* **25**, 67–86 (2011).

737 [89] Tomasi, G. & Bro, R. A comparison of algorithms for fitting the PARAFAC model.
738 *Computational Statistics & Data Analysis* **50**, 1700–1734 (2006).

739 Acknowledgments

740 We thank P. Janak, C. Honey, S. Mysore, D. Lee and the members of the Kuchibhotla
741 laboratory for comments on the manuscript. This work was supported by the Johns Hopkins
742 Kavli NDI Distinguished Postdoctoral Fellows Program (C.D.), National Institute of Health
743 K99 DA059024 (C.D.), Johns Hopkins Kavli NDI Distinguished Graduate Student Fellowship
744 (Z.Z.), National Institute of Health R00 DC015014 and R01 DC018650 (K.K.), National
745 Science Foundation CAREER 2145247 (K.K) and Brain & Behavior Research Foundation
746 NARSAD (K.K.).

747 Author contributions

748 C.D. and K.K. designed the experiments and data analyses; C.D. and Z.Z. performed the two-
749 photon experiments; C.D., Z.Z. and K.F. preprocessed calcium imaging data; C.D., Z.W.,
750 K.F., A.W., and S.E. performed the optogenetic experiments; C.D. analyzed the data; C.D.
751 and K.K. wrote the manuscript.

752 Competing interests

753 The authors declare no competing interests.

754 Materials & Correspondence

755 Correspondence and material requests should be addressed to cdrieu1@jhu.edu or kkuchib1@jhu.edu.

756 Methods

757 Animals

758 All procedures were approved by Johns Hopkins University Animal Care and Use Committee
759 (MO20A272). Male and female double (PV-ChR2; test mice) or single (PV-cre or flox-ChR2;
760 control mice) transgenic mice between 6 and 12 weeks at the start of experiments were used for
761 the optogenetic experiments. PV-cre (Jackson laboratory, strain #017320), flox-ChR2 (Ai32,
762 Jackson laboratory, strain #012569) and PV-ChR2 mice were bred in-house. PV-ChR2 mice
763 were obtained by crossing male PV-cre⁺⁻ mice with female flox-ChR2⁺⁺ or by crossing
764 male flox-ChR2⁺⁺ with female PV-cre⁺⁻. To obtain PV-cre⁺⁻ line, we bred female PV-
765 cre⁺⁺ with male C57BL/6J (Jackson laboratory, strain #000664). Offspring genotypes were
766 confirmed by PCR (Lucigen EconoTaq Plus GREEN 2X) and using two-photon imaging to
767 observe expression of the reporter protein (GFP, see subsection ‘Optogenetic experiments’).
768 Male C57BL/6J (Jackson laboratory, strain #000664) aged between 6 and 12 weeks at the
769 start of experiments were used for two-photon calcium imaging experiments. Animals were
770 group housed in standard plastic cages with food available *ad libitum* and maintained on a
771 12-hour reversed light-dark cycle at stable temperature (19.5-22°C) and humidity (35-38%).
772 Experiments took place during the dark phase. Mice were kept on a mild water restriction
773 diet (>85% of body weight) after surgery and throughout task training.

774 Surgical procedures

775 Mice were anesthetized with isoflurane (5% at induction and maintained at 2% during
776 surgery) and their body temperature was maintained at ~35°C throughout the surgery.

777 Calcium imaging experiments

778 Mice were injected (34 gauge, 25.4 mm, 12-degree bevel needle; Hamilton Company) with
779 1 μ l of AAV9-CaMKII-GCaMP6f (Addgene, #100834-AAV9, dilution 1/15) at 0.75 μ l.min⁻¹
780 (microinjection pump, Harvard Apparatus) in the left primary auditory cortex (centered at
781 1.75 mm anterior to the intersection of the lambdoid and interparietal-occipital sutures, DV:-
782 200 μ m). Above the injection coordinates, a cranial window was implanted replacing a circular
783 piece of skull by a 3-mm diameter cover glass slip (Warner Instruments) that was secured in
784 place using a mix of dental cement and Krazy Glue. A custom-made, three-point stainless
785 steel headpost was secured to the skull with C&B Metabond dental cement (Parkell). The
786 headpost consisted on a two-point kinematic fixation on the right side of the head, prolonged
787 by a rod encircling the cranial window and descending at ~45° ventrally on the left. Mice
788 were given a two-week recovery period to allow weight recovery and viral expression.

789 *Optogenetic experiments*

790 3-mm diameter cover glass slips were implanted bilaterally over the auditory cortex (centered
791 at 1.75 mm anterior to the intersection of the lambdoid and interparietal-occipital sutures,
792 on the ridge line of the temporal bone). Custom-made aluminum funnels were implanted
793 above each cranial window. The role of these funnels was threefold: 1) to precisely center the
794 end of the patch cord on the cranial windows, 2) to hold the patch cord perpendicular to the
795 cranial window (optimizing in-depth light diffusion), and 3) to fix the distance between the
796 patch cord and the cranial window to allow identical light delivery across days. A custom-
797 made, two-point stainless steel headpost was fixed onto the skull with C&B metabond dental
798 cement (Parkell) and dental cement. Mice were allowed to recover for at least one week
799 following surgery.

800 *Optogenetic silencing verification experiments*

801 For silencing verification experiments (Supplementary Figure 2, n=2), PV-ChR2 mice were
802 injected with 1 μ l of AAV-CaMK2-GCaMP6f (Addgene, #100834-AAV9, dilution 1/15) at
803 0.75 μ l.min⁻¹ in the left primary auditory cortex (centered at 1.75 mm anterior to the inter-
804 section of the lambdoid and interparietal-occipital sutures, DV:-200 μ m) and implanted with
805 a 3-mm cover glass slip and a custom-made, two-point stainless steel headpost. Mice were
806 given a two-week recovery period to allow weight recovery and viral expression.

807 **Auditory Go/No-go task**

808 All mice (optogenetic and two-photon imaging) underwent the same habituation and training
809 procedures. After recovery from surgery, mice were water restricted for at least 5 days so
810 that their weight stabilized at 85% of their *ad libitum* weight. During this period, mice were
811 handled daily. Mice were then head-fixed and placed in the experimental context, where they
812 were trained to lick from a lick tube or water cup to receive a drop of water (3 μ l). No tone
813 was presented during lick training. Lick training session ended after 30 min or when 1 ml
814 of water was consumed. After two days of lick training, mice were trained on the auditory
815 Go/No-go task for at least 15 days.

816 Mice were trained to lick to a target (S+) tone to receive a water drop (3 μ l) and with-
817 hold licking to the foil (S-) tone to avoid a timeout. Auditory stimuli were three quarter
818 octave-spaced pure tones. Target and foil tones were presented pseudo-randomly and coun-
819 terbalanced every 20 trials. Each trial consisted of a no lick period (1 s), tone presentation
820 (100 ms), dead period (200 ms), response period (2.5 s) and a delay period: hit: 4 s (to
821 enable full licking of the reward), miss and correct reject: 2 s, false alarm: 7 s (timeout).
822 In this learning context, called the ‘reinforced’ context, the lick-tube delivering water was
823 positioned within reach of the tongue. In contrast, in the ‘probe context’, the lick-tube was
824 moved out of tongue and whisker reach by an automated actuator. The blocks of probe trials
825 were interspersed between reinforced trials and no additional delay was introduced by lick-

826 tube movement. Importantly, we have shown that the performance gap observed between
827 probe and reinforced trials early in learning is not driven by the change in the sensory context
828 induced by the absence of the lick-tube in the probe context⁷.

829 Optogenetic experiments

830 Mice were trained in the Go/No-go task for 300 trials every day: 280 trials in the rein-
831 forced context interspersed with a short block of 20 non-reinforced (probe) trials starting
832 at trial #141. Head-fixation habituation, lick training and Go/No-go task training took
833 place in custom-made, sound-attenuated behavioral boxes (ambient noise level ~53 dB SPL)
834 controlled with custom-written MATLAB programs interfacing with Bpod State Machines
835 (Sanworks). Pure tones (4,757 and 8,000 Hz) were delivered through an electrostatic speaker
836 driver (TDT) to a free field electrostatic speaker (TDT) at an intensity of 70 dB SPL and
837 licks were detected through an infrared beam. Blue light (453nm, DPSS laser, Opto-Engine
838 LLC) was delivered in a 20-Hz sinewave generated by Arduino. The power recorded at the
839 end of the patch cord (splitter branching fiber-optic patch cords, Doric Lenses) was 6-8mW.
840 When dispersed over a diameter of 3mm, that yields a light intensity of 0.85-1.13 mw/mm²
841 at the cortical surface. Sound amplitude, water drop size, and laser power were calibrated at
842 the beginning of each experiment. To dissociate the effect of AC silencing on behavior from
843 its consequence on the learning process, we used a probabilistic approach whereby no light
844 was delivered during probe trials and a subset of reinforced trials. These light-off trials were
845 critical to assess behavior when the auditory cortex was available again.

846 Full trial experiment ($n = 8$ PV-ChR2, $n = 8$ control mice, $n = 8$ PV-ChR2 visual cortex): light
847 was turned on on 90% of reinforced trials pseudo-randomly (18 trials – 9 S+ and 9 S– –
848 every 20-trial block). In light-on trials, the light was turned on 100 ms before tone onset and
849 stayed on for ~2.5 s for all trial types (hit: 2.5 s post operant lick, CR and miss: stop at the
850 end of response window, FA: 2.5 s post first lick).

851 Expert only full trial experiment ($n = 4$ PV-ChR2 mice): Mice were trained for 18 days with-
852 out light. Afterward and for 5 days, from day 19 to 23, the light was turned on following the
853 ‘full trial experiment’ protocol or on 90% of reinforced trials consecutively.

854 Tone experiment ($n = 4$ PV-ChR2, $n = 3$ control mice): light was turned on on 90% of re-
855inforced trials pseudo-randomly (18 trials – 9 S+ and 9 S– – were light-on every 20-trial
856 block). In light-on trials, the light was turned on 100 ms before tone onset and turned off at
857 tone offset.

858 Post hit experiment ($n = 8$ PV-ChR2, $n = 5$ control mice, $n = 6$ PV-ChR2 visual cortex): we
859 used a closed-loop lick-triggered stimulation approach, whereby light was turned on after a
860 rewarded lick on 90% of reinforced trials pseudo-randomly (light could be turned on on 9
861 over 10 S+ trials every 20-trial block). In light-on trials, the light was turned on 70 ms after
862 the first lick detection (to allow the lick cycle to complete and the tongue to retract) and

863 100 ms before reward delivery and stayed on for 2.5 s.

864 Post false alarm experiment ($n = 8$ PV-ChR2, $n = 7$ control mice): we used a closed-loop
865 lick-triggered stimulation approach, whereby light was turned on after a non-rewarded lick
866 on 90% of reinforced trials pseudo-randomly (light could be turned on on 9 over 10 S- trials
867 every 20-trial block). In light-on trials, the light was turned on 70 ms after the first lick
868 detection (to allow the lick cycle to complete and the tongue to retract) for 2.5 s.

869 At the end of the experiments, mice were anesthetized (isoflurane 5% at induction and 2%
870 during surgery; body temperature maintained at $\sim 35^\circ\text{C}$) and the left funnel was drilled out.
871 Mice were then put under the two-photon microscope and the field of view was excited at
872 980nm. Green fluorescence was detected in test mice (ChR2-EYFP) but not in control mice.
873 This procedure allowed to confirm mice genotypes and to assess cell health. Z-stacks were
874 collected (unidirectional, 30.98 Hz; magnification 1.7 or 2.0X; range: 450 μm , step: 10 μm ,
875 50 frames per step; depth from brain surface 420-445 μm) to generate 3D reconstruction
876 (ImageJ).

877 Longitudinal two-photon calcium imaging during learning

878 Two-photon fluorescence of GCaMP6f was excited at 980nm using a mode locked Ti:Sapphire
879 laser (Spectra-Physics) and detected in the green channel (GFP emission). Imaging was
880 performed with a two-photon resonant-scanning microscope (Neurolabware) equipped with
881 a water immersion objective (16x, 0.8NA, Nikon) tilted to an angle of 40-50° to image the
882 auditory cortex. The arm of the microscope was enclosed in a custom-made sound-attenuated
883 box. An electronically tunable lens was used to record near-simultaneously two planes in
884 layer 2/3 (150-250 μm below dura, 50 μm spaced, 312x192 μm^2 , at 15.96Hz per plane, with a
885 laser power of ≤ 40 mW). Images were collected at 1.7x or 2x magnification using ScanBox
886 (Neurolabware) and task events (sounds, rewards, licks and frames) were recorded using a
887 digitizer (Digidata 1550b). Pure tones were delivered through an electrostatic speaker driver
888 (RZ6, TDT) to a free field electrostatic speaker (TDT) located at $\sim 5\text{cm}$ from the right ear
889 at intensity of 70dB SPL. Licks were detected through an infrared beam. Scanner noise
890 (8kHz) was attenuated using a custom-made foam sound enclosure directly surrounding the
891 animal and the resonant scanner was set to continuous throughout the recording session (to
892 avoid any scanning onset-related activity). Custom-written MATLAB program interfaced
893 with RPvdsEx to control task events. Mice were placed in a plastic tube and head-fixed
894 via a two-point pneumatic clamp on the right and a one-point, 360°-rotational clamp on the
895 left (at 45-50° in the horizontal plane). The whole behavioral platform was installed on a
896 rotation platform so that the field of views could be precisely retrieved one day to the next.
897 Imaging fields were retrieved every day before task training by visual inspection (see also
898 'Pre- and post- task tonotopic mappings'). Typically, mice were trained for three blocks
899 of 80-100 trials, with either two blocks of 10 probe trials interleaved in two of these three
900 blocks, or one block of 20 probe trials. The field of view was adjusted in between blocks

901 to compensate for z-drift, if necessary. An additional 10,000 frames of spontaneous activity
902 were recorded in a separate block at the end of each behavior session.

903 **Pre- and post- task tonotopic mapping**

904 One day before lick training, mice were placed under the microscope and were presented
905 with a set of 17 pure tones (duration 100ms), three-quarter octave spaced, in a pseudo-
906 random order ranging from 4 to 64 kHz at 70 dB SPL. Target and foil tones were selected
907 for the Go/No-go task as pure tones that were similarly represented in the recorded neuronal
908 population. The same mapping procedure took place immediately after or one day after the
909 last behavior session, and 7 and 14 days later.

910 **Two-photon calcium imaging and one-photon blue light stimulation for silencing
911 verification**

912 To validate our optogenetic silencing protocol and determine light power to use for efficient
913 and reliable silencing of cortical networks, we recorded calcium activity of layer 2/3 pyramidal
914 cells while stimulating ChR2-expressing PV interneurons with blue light (Supplementary
915 Figure 2a,b). Two-photon imaging was performed as indicated in ‘Longitudinal two-photon
916 calcium imaging during learning’, except that only one plane was recorded (15.49Hz, 150-
917 250 μ m below dura, 312x192 μ m², x1.7 or x2 magnification, laser power \leq 40 mW). A mounted
918 LED (490nm, M490L4, Thorlabs) and a LED driver (Thorlab, LEDD1B) were used to deliver
919 blue light at six different power levels over the AC. Pure tones (4-64kHz, 80dB SPL) and
920 complex sounds were played (100-ms duration each, 100-frame intervals) and blue light was
921 delivered in a counterbalanced manner. On a silencing trial, a trigger command is sent 100ms
922 before sound onset from Clampex to the Tower electronics (Scanbox) that generates control
923 signals for the LED and the PMT shutter (LED on for 1ms, PMTs off for 9ms, repeat for
924 5 frames; Supplementary Figure 2c). The first pulse was triggered 68ms before the onset of
925 the sound, and the stimulation continued for a total of 320ms (Supplementary Figure 2c).
926 To estimate the LED powers at the cortical surface (in mW/mm²), we measured the LED
927 power coming out of the objective and estimated the cortical surface illuminated to be 2 mm
928 (16X Nikon objective), leading to LED powers ranging from 0 to 3.15 mW/mm².

929 Non-rigid registration and cell segmentation were performed using suite2p⁷⁹ (<https://github.com/MouseLand/suite2p>). Fluorescence of each putative neuron ($n = 454$) was extracted
930 and converted into $\Delta F/F$ by taking the mean activity as the baseline. We aligned neural
931 responses to tone presentation, and quantified the effect of optogenetic silencing by compar-
932 ing the mean activity of each neuron across all repetitions of sound presentations at different
933 light powers (Supplementary Figure 2d,e). Only $\Delta F/F$ in frames immediately following light
934 presentation were considered for quantification to avoid light contamination of the signal.

936 **Calcium imaging preprocessing**

Upon acquisition, images were cropped (to remove artifact bands on plane 1 due to the electronically tunable lens) and converted to HDF5 files. Non-rigid registration (suite2p, ⁷⁹ <https://github.com/MouseLand/suite2p>) was run on the concatenated movie of all files recorded for a given mouse. All motion-corrected movies were visually inspected. Because recordings were made over weeks for a given dataset, our dataset could contain cells only weakly active overall. We, therefore, opted for manual detection of regions of interest (ROIs) rather than a semi-automatic one that uses cell activity to detect ROIs (e.g. suite2p cell registration). Manual ROI drawing was done in ImageJ using mean enhanced and maximum projection images. We identified 7,137 ROIs in 8 mice, with an average of 892 ± 109 ROIs per mouse. The stability of each ROI throughout the entire recording was then assessed using a custom-written GUI in Matlab (MathWorks, Natick, MA). Overall, 2,332/3,935 cells were tracked every day of the task training in Learning mice (mean proportion of $67.3 \pm 7.5\%$ of total ROIs per mouse), and 2,321/3,202 cells were tracked every day of passive exposure in Passive mice (mean proportion of $87.6 \pm 6.2\%$ of total ROIs per mouse). Fluorescence activity from the ROIs was extracted using custom functions (Matlab). Raw fluorescence of each cells was then normalized as:

$$\Delta F/F = F_{all} - \eta_{all}$$

where

$$F_{all} = F_1 \cap F_2 \cap \dots \cap F_n$$

937 where the symbol \cap represents a concatenation, n is the number of files, $F_i = \frac{F - F_0}{F_0}$ with F the
938 raw fluorescence extracted from recording file i and F_0 the median of this time series. η_{all} is
939 the median of F_{all} over a sliding window of ~ 3 minutes. To compare calcium activity across
940 trials, baseline fluorescence (activity during the inter-trial interval, before tone onset) was
941 subtracted from the trial activity, so the $\Delta F/F$ reflected changes of intensity to the original
942 intensity before trial onset.

943 **Data analysis**

944 *Statistics*

945 Analyses were performed in Matlab (MathWorks, Natick, MA), using custom written pro-
946 grams, FMAToolbox (M. Zugaro, <http://fmatoobox.sourceforge.net>), and Tensor Tool-
947 box for MATLAB (<https://www.tensorbox.org/>). Descriptive statistics are reported
948 as mean \pm standard error of the mean when the underlying distribution is Gaussian-shaped
949 (Jarque-Bera test) or median \pm standard error of the median otherwise. Unless indicated
950 otherwise, bars represent median \pm standard error of the median, box-plots represent median
951 (center line), upper and lower quartiles (box limits) and 1.5x interquartile range (whiskers),
952 and all statistical tests were two-sided. Student's t-test was used for two group comparisons
953 of Gaussian distributions, paired t-test for paired Gaussian distributions. For non-Gaussian

954 distributions of independent data, two group comparisons were made using Wilcoxon rank
955 sum tests. Wilcoxon sign rank tests were used for two group comparisons of non-Gaussian
956 paired data or to compare medians of non-Gaussian distributions to single values. Two-way
957 ANOVAs were performed to evaluate the effects of two independent variables on data and
958 their interaction. All ANOVA statistics are reported in Supplementary Table 1. Proportions
959 were compared using the binomial proportion test. Distributions were compared using the
960 Kolmogorov-Smirnov test. No statistical methods were used to pre-determine sample sizes,
961 but our sample sizes are similar to those generally employed in the field. Data collection and
962 analysis were not performed blind to the conditions of the experiments.

963 *Behavior analysis*

964 Rare non-learner mice were excluded and massive drops in performance after reaching high
965 performance (accuracy > 0.7) were not analyzed. Accuracy in probe and reinforced context
966 was computed as $(n_{HIT} + n_{CR})/(n_{S+} + n_{S-})$, where n_{HIT} , n_{CR} , n_{S+} , and n_{S-} are the number
967 of hit, correct reject, S+ and S- trials, respectively. To have trial-resolution assessment of
968 behavior, we also computed response index curves (Fig.1i), which reflected the latency to
969 respond to the cues compared to local, spontaneous licking rate^{6,80}. Response index curves
970 were computed for the two cues (S+ and S- trials) separately as the latency to lick in a
971 2.5s window before the cue onset minus the latency to lick in the response window (2.5s
972 after cue onset). If no lick was detected in either of these windows, the latency was set to
973 the window duration, i.e. 2.5s. Therefore, for a given trial, the response index ranges from
974 -2.5 to +2.5, with positive values indicating that the response to the cue was shorter than
975 the local spontaneous licking rate of the animal, negative values indicating a decrease of
976 licking in response to the cue, and values around 0 indicated that the cue did not impact the
977 response rate. Performance index (Fig.4s,t) was computed as the difference between S+ and
978 S- cumulative response index curves. From the S+ response index, we identified the 'change
979 point' (CP)^{6,80}, i.e. the trial after which there is a consistent expression of cued behavior
980 (Fig.1i). We used the method described here⁸⁰, itself a variation of the method used in⁶.
981 Briefly, a recursive algorithm successively run over each data point i of the cumulative S+
982 response index curve and performs the following steps: 1) draws a straight line from trial
983 i to trial 0 or the previous true CP, whatever is the closest to i and identifies the point
984 that deviates maximally from this line as a putative CP; 2) calculates the strength of the
985 evidence that it is a true CP, i.e. the log of the odds against the null hypothesis of no change
986 (the logit). If logit > 1.3^{6,80}, the putative CP becomes a true CP. As multiple CPs can be
987 identified on a single curve, we reported in Fig.1i only the first CP associated with a positive
988 change of the slope of the cumulative behavioral responses⁸⁰.

989 *Best frequency*

990 Single cell responses to the 17 tones presented were evaluated with paired t-test comparing
991 pre- vs post-tone mean activity (over 10 frames, ~626ms). Bonferroni correction for the

992 number of sounds ($n = 17$) was applied. For each cell, the peak amplitude response to each
993 tone was determined as the maximum value of the averaged traces in the 10-frame post-tone
994 window. A neuron's best frequency was determined as the pure tone for which the peak
995 amplitude response was the highest among significant responses only.

996 *Tone-evoked responses across days*

997 Evolution of tone-evoked responses in the reinforced context was analyzed using all cells
998 recorded (Supplementary Figure 12 and Supplementary Figure 13) but the conclusions held
999 when restricted to cells tracked every day. Response to S+ and S-, or stimulus 1 (S1) and
1000 stimulus 2 (S2) for Passive mice, were analyzed separately with paired t-tests comparing
1001 pre- vs post-tone mean activity (in 11-frame windows, ~ 688 ms). A cell was considered tone-
1002 responsive in a given day if it significantly responded to either S+/S1 or S-/S2. Given that
1003 response profiles were identical to S1 and S2, responses to the two tones were sometimes
1004 represented together (Supplementary Figure 13).

1005 *Tone-evoked responses, responsiveness, response index and stimulus selectivity index*

Tone-evoked responses were defined as the mean $\Delta F/F$ in a 11-frame window (~ 688 ms) post tone onset. Responsiveness was defined as the proportion of cells exhibiting a significant tone response (paired t-tests; Supplementary Figure 13). To compute response indices (Fig.3k), the peak of the average $\Delta F/F$ for hit and S- trials (FA trials until mid-expression, CR trials after that) in 80-trial blocks was calculated, followed by the proportion of blocks with significant (peak $\Delta F/F > 2\%$ of baseline) response throughout learning. The response index of a neuron was computed as the average response probability in hit and S- trials over learning. Stimulus selectivity was computed for each neuron in 80-trial blocks over learning and defined as:

$$SI = \frac{|S^+ - S^-|}{|S^+| + |S^-|},$$

1006 where S^+ is the peak $\Delta F/F$ in the tone-evoked response window on hit trials, S^- is the peak
1007 $\Delta F/F$ in the tone-evoked response window on S- trials. SI could therefore ranged from 0
1008 to 1, with 1 indicating maximal selectivity for either the S+ or the S-. Values of S^+ and S^-
1009 $< 2\%$ were set to zero, and SI in blocks where S^+ and S^- were both equal to zero was set to
1010 zero. The selectivity index of a neuron was its average SI over learning (Fig.3l).

1011 *Stimulus decoding*

1012 For each mouse, cue identity was decoded across trial frames from activity of cells tracked
1013 across all days using linear discriminant analysis with 5-fold cross-validation (Supplementary
1014 Figure 1). Tone decoding accuracy in the tone-evoked window referred to the mean accuracy
1015 in the tone-evoked window (11 frames post-tone onset; Supplementary Figure 1e). Chance

accuracy level was estimated by decoding cue identity across trial after randomly shuffling cue identity across trials ($n = 20$ shuffles/day/mouse).

1018 Data organization and tensor decomposition

1019 To analyze our high dimensional dataset, we took advantage of tensor decomposition^{81–83}, a
1020 method that enables unbiased and interpretable descriptions of dynamic changes at multiple
1021 timescales, also referred as ‘tensor component analysis’ or TCA²⁶. Here we used it not only to
1022 reveal within and across trial dynamics²⁶, but also to identify shared and distinct variability
1023 in cell networks recorded from Learning and Passive mice. We organized calcium traces into
1024 a fourth-order tensor (or four-dimensional array) with four axes corresponding to individual
1025 neurons (recorded in Learning and Passive mice), time within trial, trials over time, and trial
1026 types. We then fit a tensor CANDECOMP/PARAFAC (CP) decomposition model^{83–85} to
1027 identify in an unsupervised way a set of low-dimensional components describing variability
1028 along each of these four axes (also referred here as factors; Supplementary Figure 9).

1029 Data organization. We first built two arrays for learning and passive data separately and
1030 combined them afterwards. Only data from the reinforced context was taken for Learning
1031 mice. We filtered out disengagement periods (hit rate ≤ 0.5 in a 20-trial block), sometimes
1032 occurring during the last dozens of trials of the day and associated with significant changes
1033 in neuronal dynamics compared to engaged state^{23,45,56–63}. For both Learning and Passive
1034 data, $\Delta F/F$ of each trial was selected from -1s to $+4\text{s}$ relative to tone onset (2nd tensor
1035 dimension). With 4,643 cells tracked all days, 75 frames/trials, ~ 300 trials/day over 15 days,
1036 our dataset approximated 1,567,000,000 data points. To reduce computation time, trials
1037 of identical types (hit, miss, FA or CR) within 20-trial blocks were averaged together. In
1038 other words, from a given 20-trial block, up to four trial traces could be obtained (4th tensor
1039 dimension). Because of the exclusion of disengaged periods and the tendency of the animals
1040 to lick, miss trials were too rare in the Learning group to be considered without adding sig-
1041 nificant noise and were excluded. As a result, the 4th tensor dimension dissociated S+ (hit
1042 trials for Learning data, miss trials for Passive data), FA and CR trials. Finally, a crucial
1043 goal of this analysis was to be able to identify neural dynamics associated with task learn-
1044 ing, and more precisely to isolate any dynamics associated with task contingency acquisition
1045 (measured in the probe context) or performance improvement (measured in the reinforced
1046 context). To this end, we aligned the trial traces to learning phases (3rd tensor dimension).
1047 First, we identified Acquisition, Expression and Expert phases in our 5 learning mice (see
1048 Supplementary Figure 8). The Acquisition phase started at the first trial of training and
1049 continued until maximum accuracy was reached in probe or when accuracy was ≥ 0.65 in
1050 probe and ≤ 0.70 in reinforced trials. This marked the beginning of Expression phase, which
1051 continued until Expert phase started at the second day of high and stable performance. Data
1052 in between Acquisition and Expert phases was part of the Expression phase. Evolution of
1053 individual mouse performance per identified phases is quantified in Supplementary Figure 8f.
1054 Resultant mega-mouse performance (i.e. pooled performance in 20-trial block across mice) is

1055 shown in Supplementary Figure 8d,e. Second, because these phases varied in duration across
 1056 animals, we identified the mouse with the minimum number of trial traces in a given phases
 1057 and downsampled the number of trial traces of the other mice to match this number. Down-
 1058 sampling was performed by preserving the duration/performance range in each mouse (i.e.
 1059 keeping first and last trial traces) and removing trial traces at consistent intervals in-between,
 1060 such as the overall learning evolution of the phase was preserved. Third, each Passive mouse
 1061 was assigned with the learning phases of a Learning mouse, and the same downsampling
 1062 procedure was used. Finally, the two four-dimensional arrays containing Learning and Pas-
 1063 sive data, respectively, were concatenated in the first (neurons) dimension (referred as the
 1064 ‘mega-mouse’ tensor) and $\Delta F/F$ traces were z-scored. Because Passive mice essentially did
 1065 not lick, any data for FA trials for Passive cells were zeroed out. Any missing entries of the
 1066 mega-mouse tensor were also zeroed out.

Tensor decomposition. To deal with incomplete data (absence of FA trials in Passive mice and possible missing CR early in learning or missing FA at expert level for Learning mice), we fitted an R-component weighted CP model²⁷ to our mega-mouse tensor. Briefly, CP decomposition decomposes a tensor into a sum of rank-one tensors. For a third-order tensor $X \in \mathbb{R}^{I \times J \times K}$, we wish to write it as:

$$X \approx \sum_{r=1}^R a_r \otimes b_r \otimes c_r,$$

where \otimes represents the vector outer product, $a_r \in \mathbb{R}^I$, $b_r \in \mathbb{R}^J$ and $c_r \in \mathbb{R}^K$ for $R = 1, \dots, R$, and $a_r \otimes b_r \otimes c_r$ is a rank-one tensor. With perfect data we would obtain equality; however, in practice the presence of noise prevents it. We can use the Kruskal operator to simplify the previous expression^{86,87}:

$$\sum_{r=1}^R a_r \otimes b_r \otimes c_r \equiv \llbracket A, B, C \rrbracket,$$

where factor matrices $A \in \mathbb{R}^{I \times R}$, $B \in \mathbb{R}^{J \times R}$ and $C \in \mathbb{R}^{K \times R}$, with

$$A^{(n)} = [a_1^{(n)} \dots a_R^{(n)}] \text{ for } n = 1, \dots, 3$$

To fit the CP decomposition model to data, we used the CP-WOPT (CP Weighted OPTimization) algorithm²⁷ that uses a first-order optimization approach to solve the weighted least squares problem, i.e. minimize the error function

$$f(A, B, C) = \frac{1}{2} \|W(X - \llbracket A, B, C \rrbracket)\|^2,$$

where W is a nonnegative weight tensor with same size as X defined as

$$w_{ijk} = \begin{cases} 1 & \text{if } x_{ijk} \text{ is known,} \\ 0 & \text{if } x_{ijk} \text{ is missing,} \end{cases} \quad \text{for all } i = 1, \dots, I, j = 1, \dots, J, k = 1, \dots, K$$

The weighted least squares objective function is solved over all the factor matrices simultaneously.

In practice, the rank R of a tensor is generally not known and is not easily determined⁸⁸. To fit the CP models and choose the number of components, we closely followed the pipeline detailed in²⁶. Briefly, we ran models 20 times with different random initializations for different numbers of low-dimensional components $R = 1, \dots, 6$. We used two metrics to compare and assess models: 1) the (normalized) weighted squared reconstruction error, computed for each fitted model, defined as:

$$\frac{1}{2} \left(\frac{\|W(X - [\![A, B, C]\!])\|^2}{\|WX\|^2} \right)$$

and 2) a similarity score^{26,89}, quantifying the match between two fitted models i.e. how similar are the components resulting from two different runs. Let's consider the Kruskal form of the tensor X (or ktensor)

$$X = \sum_{r=1}^R \lambda_r a_r \otimes b_r \otimes c_r,$$

where λ_r is the scaling factor after rescaling a_r , b_r and c_r to be unit length. Considering two tensors $[\![A, B, C]\!]$ and $[\![D, E, F]\!]$,

$$\max_{\omega \in \Omega} \frac{1}{R} \sum_{r=1}^R \text{penalty} \times (a_r' d_{\omega(r)} \otimes b_r' e_{\omega(r)} \otimes c_r' f_{\omega(r)}),$$

with

$$\text{penalty} = 1 - \frac{|\lambda_r - \lambda_{\omega(r)}|}{\max(\lambda_r, \lambda_{\omega(r)})}$$

where Ω is the set of all permutations of the R components, and ω a particular permutation. With increasing number of components R , considering all possible matches is exponentially expensive and can be computationally prohibitive and factors were matched in a greedy fashion to identify good alignment (although not necessarily optimal). Similarity for each model fit was computed with respect to the best-fit model with the same number of components. Adding more components caused models to be less reliably identified (lower similarity score). For a given number of components R , the model fits were also visually inspected and compared. With our dataset, models with similarity scores above 0.8 were qualitatively similar while consistency dropped for values closed to 0.5. Therefore, a decomposition into 4 components was chosen for our dataset. The output of our decomposition was therefore a set of four components, each composed of four factors (i.e. weight vectors): 1) neuron factor (W_N), reflecting cell ensembles, 2) within trial factor (W_W), indicating when the activity occur in the trial, 3) across trial factor (W_A), reflecting the evolution profile over learning/time at trial resolution, and 4) outcomes factor (W_O), reflecting contribution of sensory, motor and cognitive variables. When R is small, increasing number of components demixed the activity

1082 until providing redundant information (when $R > 4$ for this tensor). Importantly, other types
1083 of decomposition were run, and other tensors (individual mouse, Passive and Learning data
1084 separately) were decomposed, and they all converged into the same description of the data.

1085 *Identification of learning-related dynamics*

1086 Quantification. To determine whether the low-dimensional dynamics described by the tensor
1087 decomposition were selectively attributed to the cells from Learning or Passive mice, we
1088 analyzed the neuronal factor, i.e. the neuronal weights (W_N) of the four components. We
1089 first compared the contribution of Learning and Passive networks to the highest (absolute)
1090 neuronal weights across components (Fig.10c, Supplementary Figure 11c). Given that no
1091 constraint was applied on the sign of the weights, a given component could describe up to
1092 two distinct dynamics. We therefore also analyzed positive and negative neuronal weights
1093 separately (Supplementary Figure 10d,e, Supplementary Figure 11d) and obtained the same
1094 results: components 1 and 2 described dynamics largely driven by the passive network while
1095 components 3 and 4 described neural dynamics driven by the learning network. Importantly,
1096 we verified that this effect was not driven only by one mouse: for each component, we
1097 compared the neuronal weights of cell populations recorded in each mouse of a group (e.g.
1098 passive) and compared it to the other group (e.g. learning) (Supplementary Figure 10e).
1099 Because the components described different neuronal dynamics, this result therefore implied
1100 that learning and passive networks contained different low-dimensional dynamics.

1101 Visualization. To visualize how the revealed neural dynamics maps onto our two experimental
1102 groups (learning and passive), we used two different dimensionality reduction approaches to
1103 project the data into a two- or three-dimensional space. First, we used t-distributed stochastic
1104 neighbor embedding (t-SNE) on the neuronal weight matrix W_N of size $N \times R$, where N is
1105 the number of cells in tensor and R the number of components (Fig.2l,m, Supplementary
1106 Figure 10f). Second, we used principal component analysis (PCA) on different combinations
1107 of factors: $W_N \otimes W_W$ (Supplementary Figure 10g), $W_N \otimes W_W \otimes W_A$ (Fig.2k), $W_N \otimes W_W \otimes W_O$
1108 (Supplementary Figure 10h), and $W_N \otimes W_W \otimes W_A \otimes W_O$ (Supplementary Figure 10i), and
1109 projected learning and passive data separately into the same principal component subspace.

1110 *Unique participation: defining cell ensembles*

1111 For visualization and quantification purposes, we attributed each neural dynamic to unique
1112 cell ensembles based on neurons' weights (Supplementary Figure 11a). As indicated earlier,
1113 factor weights could be positive or negative and therefore up to two distinct dynamics could
1114 be represented per component. With this in mind, each neuron i was associated with a two
1115 digit code [componentID sign], i.e. a unique dynamic, where componentID is the component
1116 where the $|W_N|$ of the neuron i was maximal. This approach therefore filtered out non-
1117 participating (i.e. low weighted) neurons in describing neuronal dynamics, as illustrated in
1118 Supplementary Figure 11b. Finally, in order to assess the nature of encoding of these cell

1119 ensembles, cell ensembles 1 and 2 were restricted to cells recorded in the passive mice, while
1120 cell ensembles 3 to 6, describing dynamics of components 3 and 4, were restricted to cells
1121 recorded in learning mice (Fig.2m).

1122 *Comparison of calcium responses between trial outcomes with a time-changing signal*

1123 For each $\Delta F/F$ comparison between different trial types, both the number of trials taken
1124 ('how many') and the trial numbers ('when') were matched between group to control for
1125 time/learning effect and power/noise difference (Figs. 4d-h,m-n, 5f,g).

1126 *Analysis of licks outside task events*

1127 Lick bouts outside task events were defined as lick bouts that preceded the first tone presen-
1128 tation at the beginning of each behavioral block. The analysis was restricted to the first day
1129 of training, to remove learning confound as much as possible (Supplementary Figure 15b,c).
1130 A lick bout was defined as a succession of at least 3 licks with less than 1s interval in-between
1131 each lick. In addition, it had to be preceded by a 1s no lick period, used to z-score the traces.

1132 *Classification of false alarm trials based on reward prediction activity*

1133 For each learning mouse, we trained a two-class support vector machine (SVM) algorithm
1134 to decode trial identity (matched hit and CR trials) from late-in-trial activity (single trial
1135 AUCs) of neurons part of cell ensemble 5. This decoding gave us access to a misclassification
1136 rate (for each class and global), representing the noise level in the data (Supplementary
1137 Figure 16a,b,e). We then used this trained SVM to classify FA trials, reasoning that if a
1138 reward prediction signal is present during an FA trial, it will be decoded as a hit trial. In
1139 each mouse, the proportion of FA trials with a RP signal was higher than the misclassification
1140 rate of the decoder (Supplementary Figure 16e).

1141 *Isolating brief disengagement periods during behavior*

1142 Once mice acquire task contingencies and start increasing their correct rejection in the re-
1143 inforced context, they generally stop behaving in the probe context (hit rate close to zero;
1144 e.g. Supplementary Figure 8)⁷. We therefore found these periods by looking for probe blocks
1145 with hit rate < 0.4 (Fig.5g).

1146 *Pre- vs post-behavior changes in tonopy*

To assess how learning and passive exposure affected the cortical tonotopic map, we compared best frequency surfaces from tuning curve recording sessions before and after learning (see 'Pre- and post- task tonotopic mapping'). We first split the field of views in 30×30 pixels ($\sim 41 \times 41 \mu\text{m}$) and computed the best frequency mode of the local neuronal population in

each of those pixel blocks (Fig.3o). We estimated the change in surface before and after behavior as:

$$\Delta_{surface} = \left[\left(\frac{n_T}{n_{any}} \right)^{post} - \left(\frac{n_T}{n_{any}} \right)^{pre} \right] \times 100,$$

1147 where n_T is the number of pixel block with T best frequency mode and n_{any} the number
1148 of pixel block with any best frequency. In our analysis, T could be the S+, S-, tones in
1149 between S+ and S-, and tones with lower or higher frequency than S+ or S- (Fig.3n).
1150 We also evaluated best frequency mode differences before and after behavior in pixel blocks
1151 (Fig.3q).

1152 *Spatial clustering of contingency-related cell ensembles*

1153 To assess the spatial distribution of reward prediction and action suppression cell ensembles
1154 (referred to here as ‘clusters’), we compared the distance between the two ensembles to a
1155 random spatial organization (Fig.6a,b). To do so, we computed the median of between-
1156 cluster cell distances and compared it to a median distribution obtained with cell ensemble
1157 identity shuffles ($n = 500$). This allowed us to assess the clustered nature of these two cell
1158 ensembles while preserving the spatial cell distribution in the fields of view. We considered
1159 the cell ensembles significantly clustered if the median distance of the cell ensembles was
1160 $> 97.5\%$ of the shuffle distribution. Because of the different statistics of cell distribution
1161 inside a field of view for each mouse, comparing raw cell ensembles distances between mice
1162 was prohibited. Instead, we computed a z-scored distance for each mouse by subtracting the
1163 mean and dividing by the standard deviation of the shuffle distribution to the data median
1164 distance (Fig.6c).

1165 *Pre-task stimulus selectivity index*

For cells with positive tone-evoked responses to both S+ and S- in pre-task tuning curve session, pre-task stimulus index (SI, (Fig.6d,e) was computed as:

$$SI = \frac{S^+ - S^-}{S^+ + S^-},$$

1166 where S^+ is the peak $\Delta F/F$ in the tone-evoked response window to the S+ tone and S^- is the
1167 peak $\Delta F/F$ in the tone-evoked response window to the S- tone. SI could therefore range
1168 from -1 to 1 , with 1 indicating total selectivity for the S+, -1 indicating total selectivity
1169 for the S-, and zero an absence of selectivity (similar response to both tone).

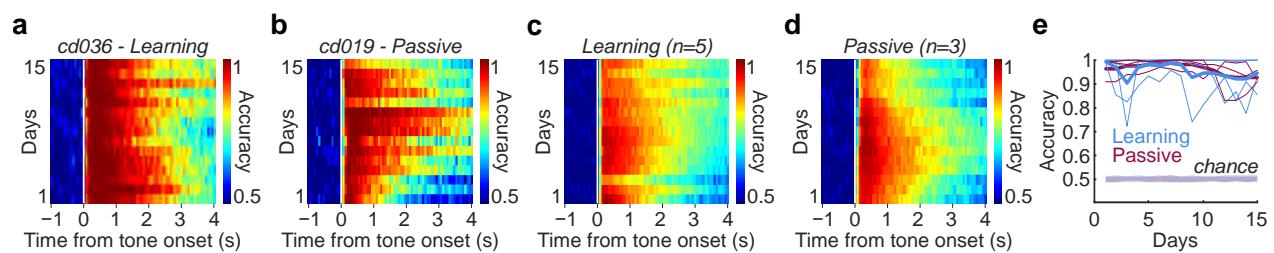
1170 *Assessing the relationship between tonotopic map and contingency organization*

1171 To assess whether reward prediction cells were S+ preferring cells and action suppression
1172 cells were S- preferring cells before training started, we generated two separate statistical
1173 tests (Fig.6g,h). First, we tested the hypothesis that the reward prediction cell ensemble

1174 emerged from S+ preferring cells. We constructed a distribution of best-frequency distance
1175 to S+ if H0 was true, i.e. if reward prediction cells were to have a best frequency the closest
1176 to S+ given the field of view statistics (Fig.6g). Separately, we tested the hypothesis that
1177 the action suppression cell ensemble emerged from S- preferring cells. We constructed a
1178 distribution of best-frequency distance to S- if H0 was true, i.e. if action suppression cells
1179 were to have a best frequency the closest to S- given the field of view statistics (Fig.6h).
1180 Finally, we compared the proportion of S+ and S- preferring cells among reward prediction
1181 and action suppression cell ensembles and observed no differences (Fig.6i).

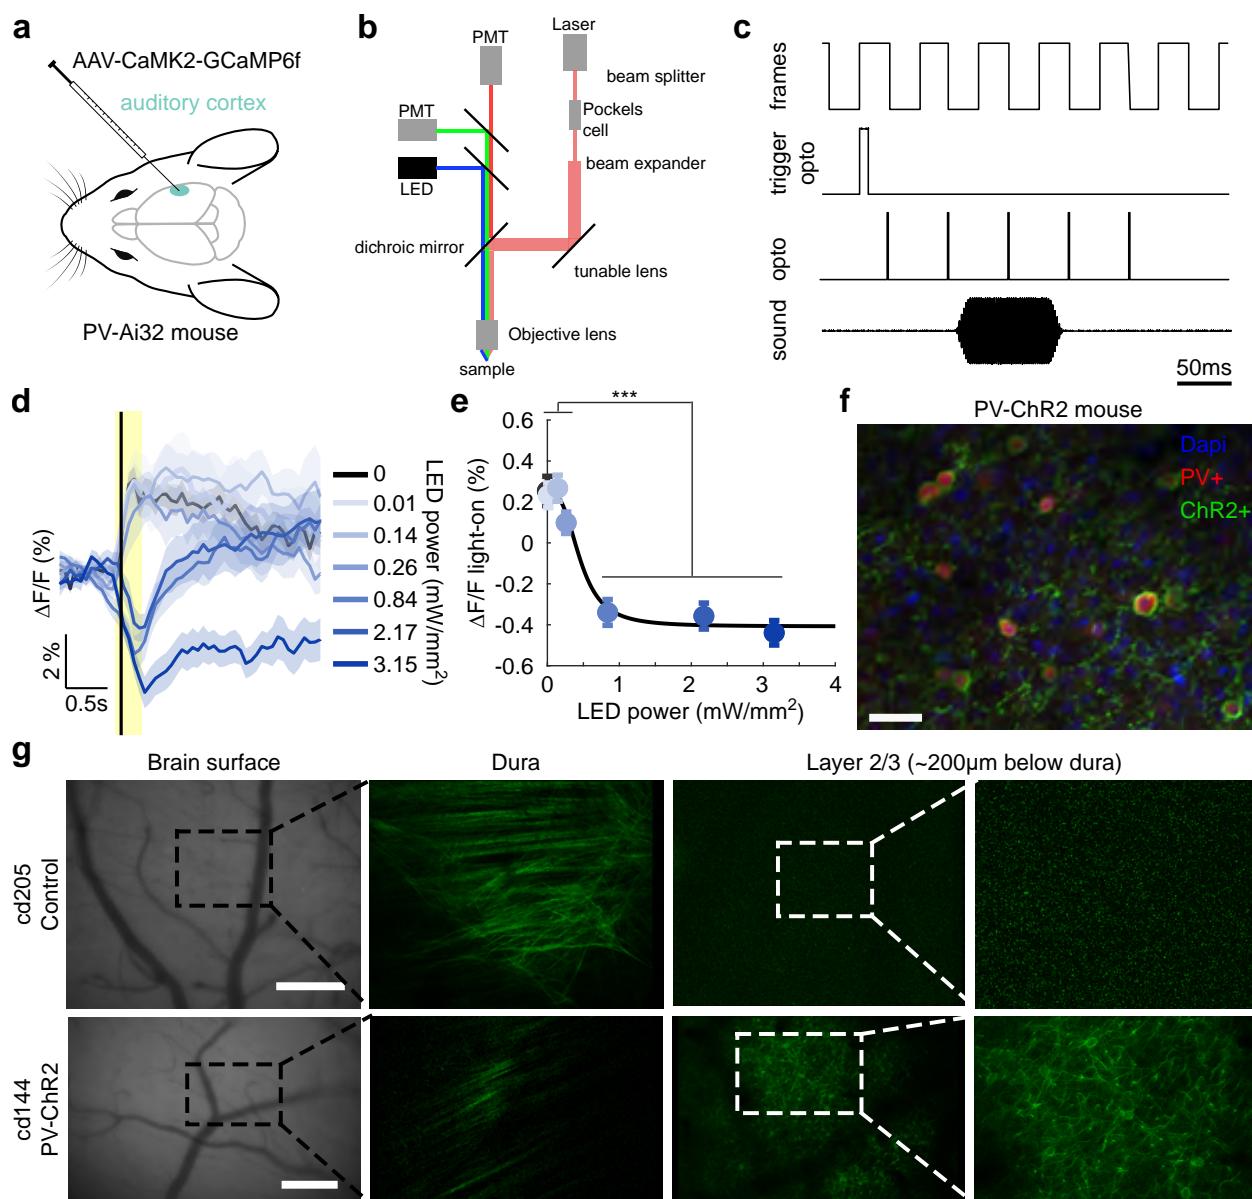
1182 **Data availability**

1183 The data that support the findings of this study are available from the corresponding authors
1184 upon request.



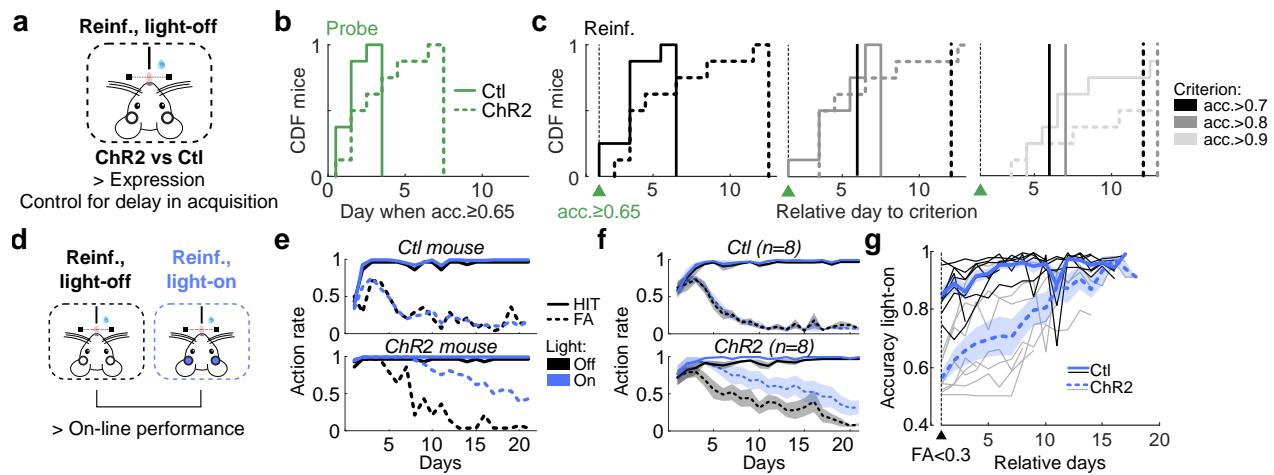
Supplementary Figure 1. Stimulus decoding in the auditory cortex is at ceiling from Day 1 of learning.

a, Stimulus decoding is at ceiling on Day 1 and remains high throughout learning (example mouse). Only the cells tracked across all days were used to decode tone identity. **b**, Stimulus decoding is at ceiling on day 1 and remains high throughout passive exposure over 15 days (example mouse). **c**, Average decoding accuracy for all Learning mice ($n = 5$). **d**, Average decoding accuracy for all Passive mice ($n = 3$). **e**, Evolution of tone decoding accuracy in the tone-evoked window across days for Learning and Passive mice compared to chance level (trial shuffle, see Methods).



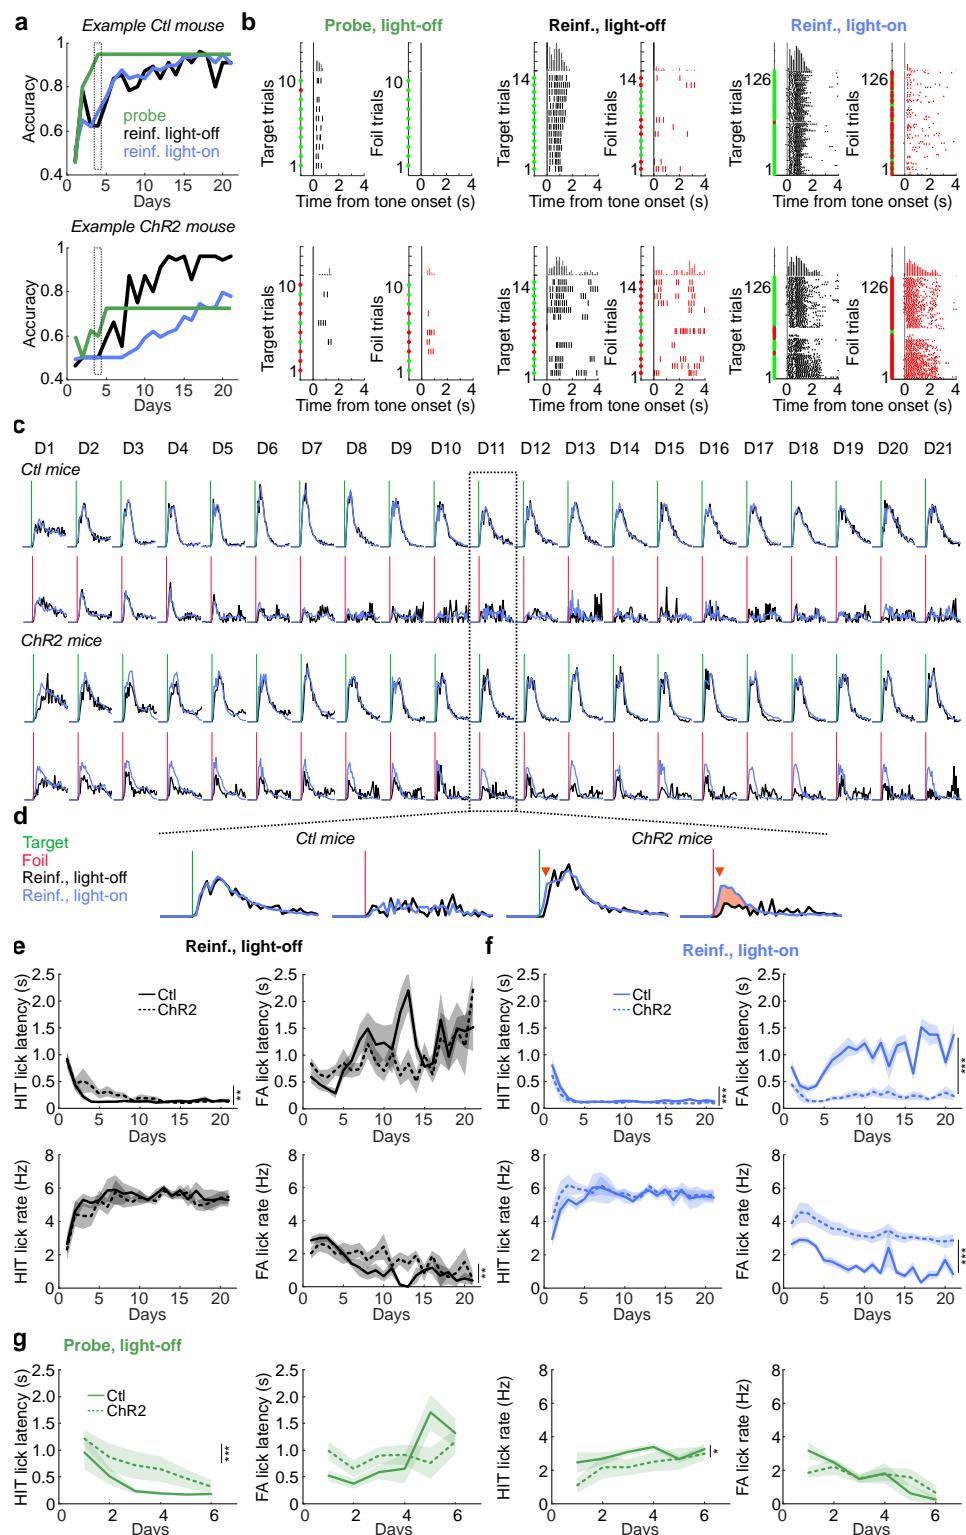
Supplementary Figure 2. Activating PV+ neurons in the auditory cortex robustly suppresses stimulus-evoked activity of excitatory neurons.

a, PV-ChR2 mice ($n = 2$) were injected with AAV-CaMKII-GCaMP6f to allow simultaneous one-photon excitation of PV cells and two-photon recordings of pyramidal cell population. **b**, Schematic of simultaneous widefield optogenetics and two-photon imaging. **c**, Optogenetic activation was locked to frame acquisition. **d**, Trial-averaged $\Delta F/F$ aligned to tone onset (black vertical line) of an example neuron at different intensity of LED power (blue scale). Yellow rectangle indicates period of light delivery. mean \pm s.e.m. **e**, Effect of optogenetic silencing as a function of LED power ($n = 454$ neurons; Friedman test, $p \sim 0$). $\Delta F/F$ at powers 0-0.26 mW/mm^2 are all significantly different from $\Delta F/F$ at powers 0.84-3.15 mW/mm^2 (post hoc comparisons with Tukey-Kramer test, *** $p < 0.001$). Black line is the logistic fit. median \pm s.e.m. **f**, Immunostaining of PV-ChR2 mice auditory cortex showing ChR2 expression in PV cells (PV+ and ChR2+ colocalization). **g**, Post-task imaging of a representative control (top) and a representative test (PV-ChR2, bottom) mouse used in AC silencing experiments. Note that no fluorescence below the dura is detected in control mice.



Supplementary Figure 3. AC full trial silencing impairs expression and on-line performance

a, Assessment of the impact of AC full trial silencing over learning on Expression by controlling for the delay in Acquisition. **b**, Cumulative distribution function (CDF) of mice as function of the day to reach an accuracy ≥ 0.65 in probe trials. **c**, Cumulative distribution function (CDF) of mice as function of the relative number of days to reach accuracy (acc.) criteria of >0.7 (left), >0.8 (middle), and >0.9 (right) in reinforced light-off trials after reaching an accuracy ≥ 0.65 in probe trials. Black and dark gray vertical lines correspond to when CDF was reach for acc. >0.7 and >0.8 , respectively. **d**, Comparing action rate and accuracy between reinforced light-off versus reinforced light-on trials to assess the impact of AC silencing on on-line performance. **e**, Hit (solid line) and FA (dashed line) of an example control mouse (top) and an example PV-ChR2 mouse (bottom) in reinforced light-off (black) and reinforced light-on (blue) trials across learning. **f**, Averaged action rate in reinforced light-off (black) and reinforced light-on (blue) trials per day for control (top) and PV-ChR2 (bottom) groups. **g**, Accuracy in light-on reinforced trials from the day when FA <0.3 in light-off reinforced trials. Note how PV-ChR2 mice (gray lines) increase accuracy (positive slopes) with light-on, showing that performance impairment fades away.



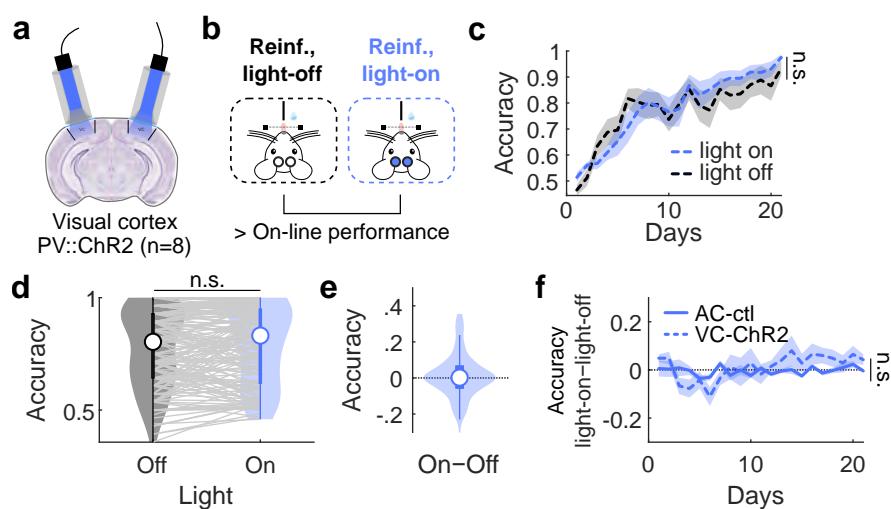
Supplementary Figure 4. Effect of AC full trial silencing on lick patterns

a, Example control (top) and ChR2 (bottom) mice accuracy in probe light-off, reinforced light-off and reinforced light-on trials across day. Dashed rectangle indicates day where licks in **b** are extracted from.

continued →

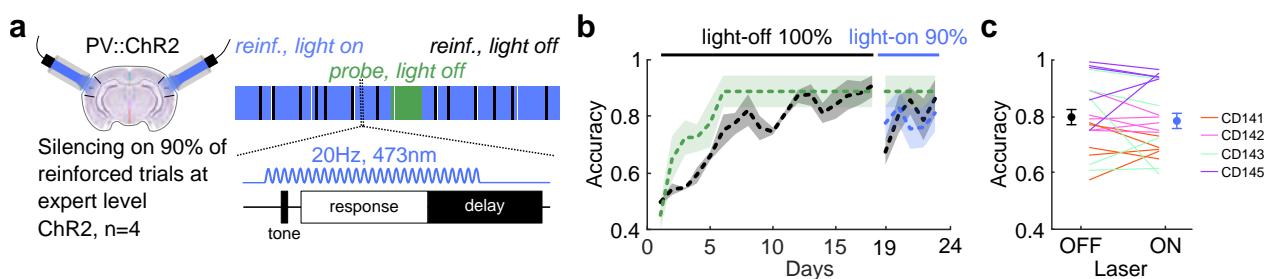
Supplementary Figure 4 (continued).

b, Lick raster plots from day 4 from the example mouse from **A** in probe light-off (left), reinforced light-off (middle) and reinforced light-on (right) trials, split into target (black, left) and foil (red, right) trials. Green and red dots indicates correct and incorrect trials, respectively. Note the difference in discrimination in all contexts between control and PV-ChR2 mice. **c**, Average lick probability across training days for control ($n = 8$) and ChR2 ($n = 8$) mice in response to target (vertical green line) and foil (vertical red line) tones, in reinforced light-off (black) and light-on (blue) trials. **d**, Insets showing faster lick latencies (red arrow heads) in response to both tones and higher lick probability in response to the foil (incorrect licking) in reinforced light-on compared to light-off in ChR2 mice (right). Light has no effect on lick structure in control mice (left). **e**, Lick latencies (top) and lick rate (bottom) in response to target (HIT trials; left) and foil (false alarm (FA) trials; right) tones in reinforced light-off trials (HIT lick latencies, Days: $F(20, 256) = 8.2738$, $p < 10^{-17}$, Groups: $F(1, 256) = 8.1568$, $p = 0.0046$, Days*Groups: $F(20, 256) = 0.9176$, $p = 0.56$; FA Lick latencies, Days: $F(20, 190) = 2.2393$, $p = 0.0027$, Groups: $F(1, 190) = 1.8422$, $p = 0.18$, Days*Group: $F(20, 190) = 1.5563$, $p = 0.067$; HIT lick rate, Days: $F(20, 256) = 4.3619$, $p < 10^{-8}$, Groups: $F(1, 256) = 2.9549$, $p = 0.087$, Days*Groups: $F(20, 256) = 0.2927$, $p = 0.99$; FA lick rate, Days: $F(20, 190) = 4.04477$, $p < 10^{-6}$, Groups: $F(1, 190) = 7.4070$, $p = 0.0071$, Days*Groups: $F(20, 190) = 1.1944$, $p = 0.26$). **f**, Lick latencies (top) and lick rate (bottom) in response to target (HIT trials; left) and foil (false alarm (FA) trials; right) tones in reinforced light-on trials (HIT lick latencies, Days: $F(20, 256) = 10.5303$, $p < 10^{-22}$, Groups: $F(1, 256) = 11.2328$, $p < 10^{-3}$, Days*Groups: $F(20, 256) = 0.6211$, $p = 0.90$; FA Lick latencies, Days: $F(20, 254) = 3.9111$, $p < 10^{-6}$, Groups: $F(1, 254) = 450.4358$, $p < 10^{-57}$, Days*Group: $F(20, 254) = 2.1947$, $p = 0.0029$; HIT lick rate, Days: $F(20, 256) = 2.6372$, $p < 10^{-3}$, Groups: $F(1, 256) = 3.7748$, $p = 0.0531$, Days*Groups: $F(20, 256) = 0.4520$, $p = 0.98$; FA lick rate, Days: $F(20, 254) = 6.4469$, $p < 10^{-13}$, Groups: $F(1, 254) = 301.2679$, $p < 10^{-44}$, Days*Groups: $F(20, 254) = 0.6326$, $p = 0.89$). **g**, Lick latencies (left) and lick rate (right) in response to target (HIT) and foil (FA) tones in probe light-off trials (HIT lick latencies, Days: $F(5, 83) = 6.4522$, $p < 10^{-4}$, Groups: $F(1, 83) = 11.7734$, $p < 10^{-3}$, Days*Groups: $F(5, 83) = 0.2878$, $p = 0.92$; FA Lick latencies, Days: $F(5, 58) = 2.9217$, $p = 0.020$, Groups: $F(1, 58) = 0.9337$, $p = 0.338$, Days*Group: $F(5, 58) = 2.1909$, $p = 0.068$; HIT lick rate, Days: $F(5, 83) = 2.0103$, $p = 0.086$, Groups: $F(1, 83) = 5.9422$, $p = 0.017$, Days*Groups: $F(5, 83) = 0.5721$, $p = 0.72$; FA lick rate, Days: $F(5, 58) = 5.6386$, $p < 10^{-3}$, Groups: $F(1, 58) = 0.0192$, $p = 0.89$, Days*Groups: $F(5, 58) = 1.6182$, $p = 0.17$).



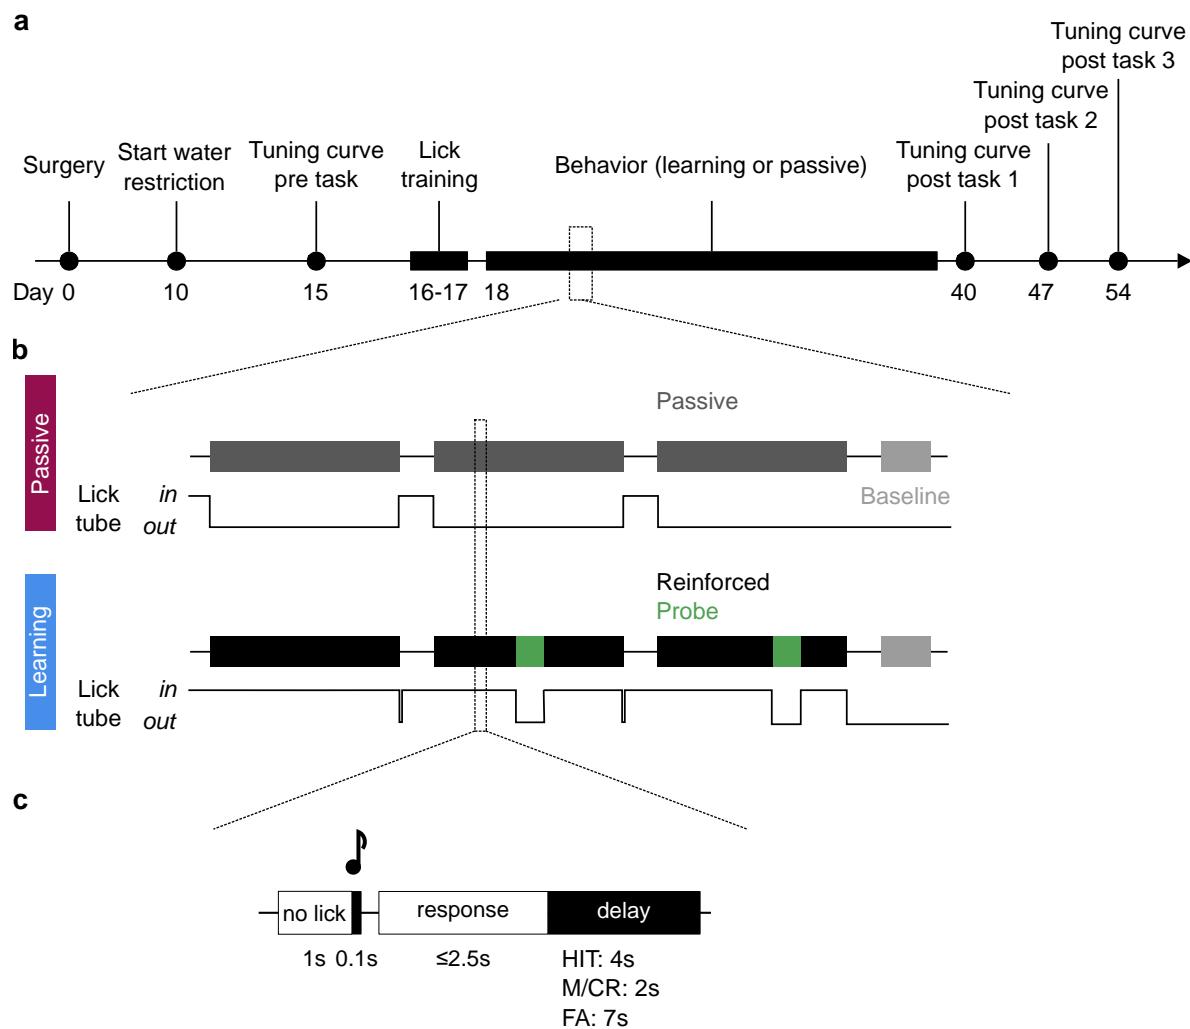
Supplementary Figure 5. Silencing of the visual cortex does not impair performance throughout learning

a, Silencing of the visual cortex in 90% of the reinforced trials throughout learning ($n = 8$ PV-ChR2 mice). **b**, Comparison of reinforced light-off versus light-on trials shows no deficit when silencing the VC demonstrating the specificity of the effects of AC silencing. **c**, Accuracy in reinforced light-off and light-on trials across days (two-way repeated measures ANOVA, Group: $F(1, 140) = 0.5093$, $p = 0.50$). **d**, Accuracy in reinforced light-off and light-on trials ($n = 168$ sessions; Wilcoxon signed rank, $p = 0.41$). **e**, Difference in accuracy in reinforced light-on versus light-off trials per session. **f**, Difference in accuracy in reinforced light-on versus light-off trials across days in visual cortex PV-ChR2 mice (dashed line) versus auditory cortex control mice (solid line) (two-way ANOVA, Days: $F(20, 271) = 1.5547$, $p = 0.06$, Groups: $F(1, 271) = 2.3072$, $p = 0.13$, Days*Groups: $F(20, 271) = 1.1540$, $p = 0.2950$).



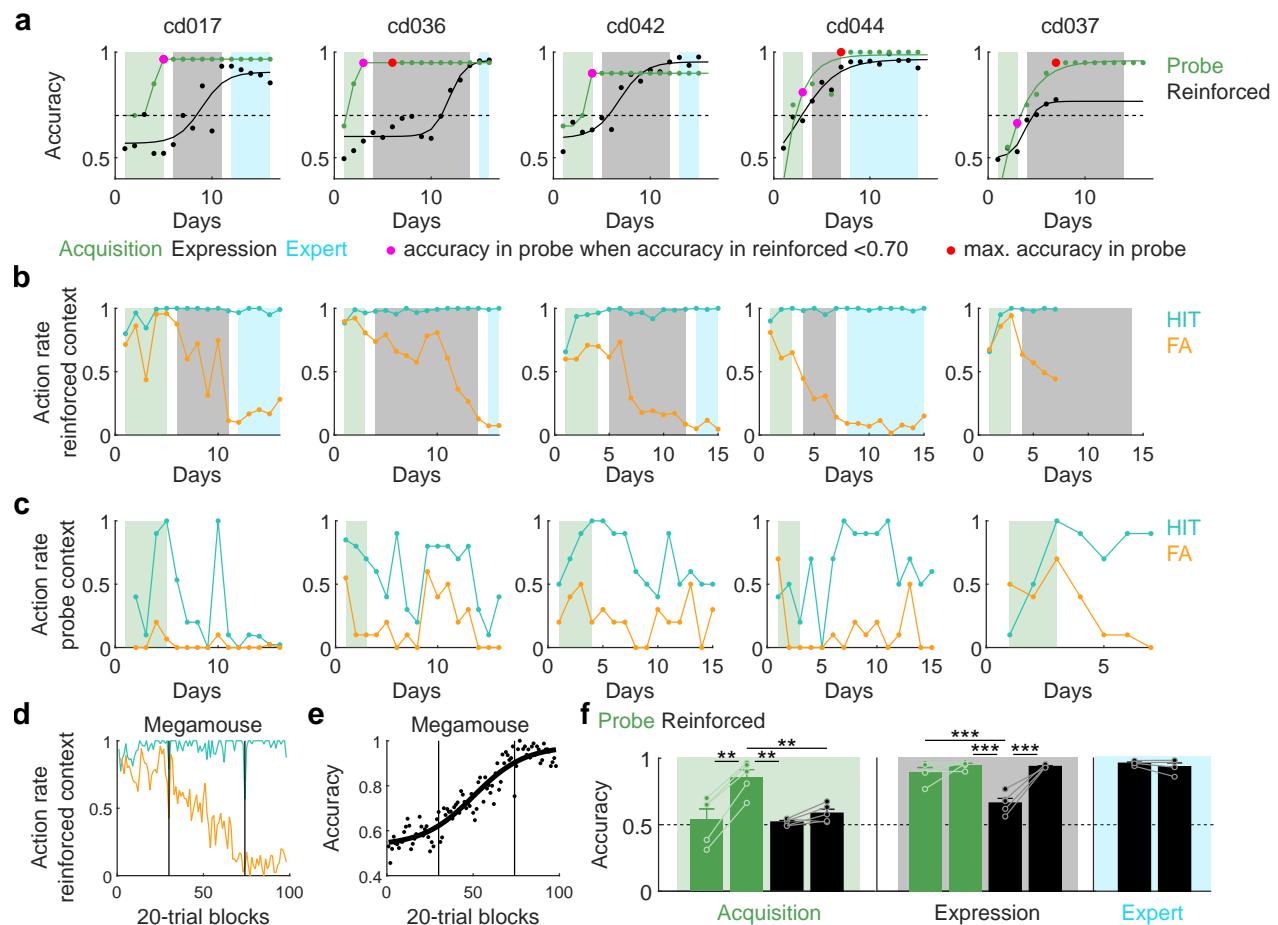
Supplementary Figure 6. AC full trial silencing at expert level

a, Probabilistic optogenetic silencing of the auditory cortex at expert level. Silencing starts once stable performance is reached. **b**, Accuracy in probe light-off (green), reinforced light-off (black) and reinforced light-on (blue) trials. Silencing is performed from day 19 to 23. **c**, Accuracy in reinforced light-off and light-on trials (paired t-test, $p = 0.602$).



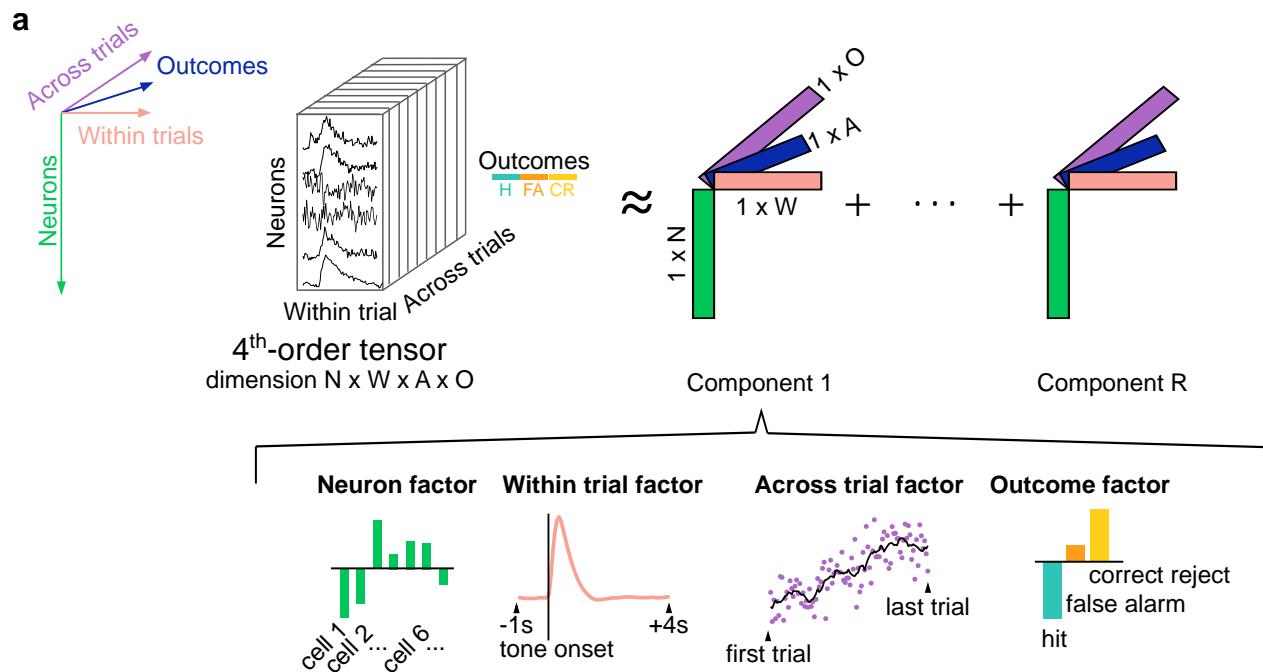
Supplementary Figure 7. Experimental design and timeline of imaging experiments.

a, After surgery, animals underwent a 10-day recovery period after which water restriction started. Tonotopic mapping (tuning curve session) of the auditory cortex took place 5 days later under the two-photon microscope, followed by two days of lick training under the two-photon microscope. These two sessions also allowed for habituation to head fixation and context. Behavior sessions started the following day for 15 or 16 days, after which tonotopic mapping sessions took place at day +1, +7 and +15 post learning. **b**, One behavioral session consisted of three blocks of 80 or 100 trials, and a baseline session (no tone presented). Two groups of mice were imaged under the two-photon microscope: the Passive group (*n* = 3) was presented with two pure tones but was never rewarded (lick tube out), and the Learning group (*n* = 5) was rewarded (3 μ l water drop) if licking in the response window after the S+ tone. Two probe blocks of 10 trials each were introduced in two of the three reinforced blocks. **c**, Trial structure. After a no-lick period of 1s, a 100-ms tone was played, followed by a 200-ms dead period and a ≤ 2.5 s response period. The length of the delay period was 2s after a miss (M, no lick after S+) or a correct reject (CR, no lick after S-), 4s after a hit (H, lick after S+) and 7s after a false alarm (FA, lick after S-).



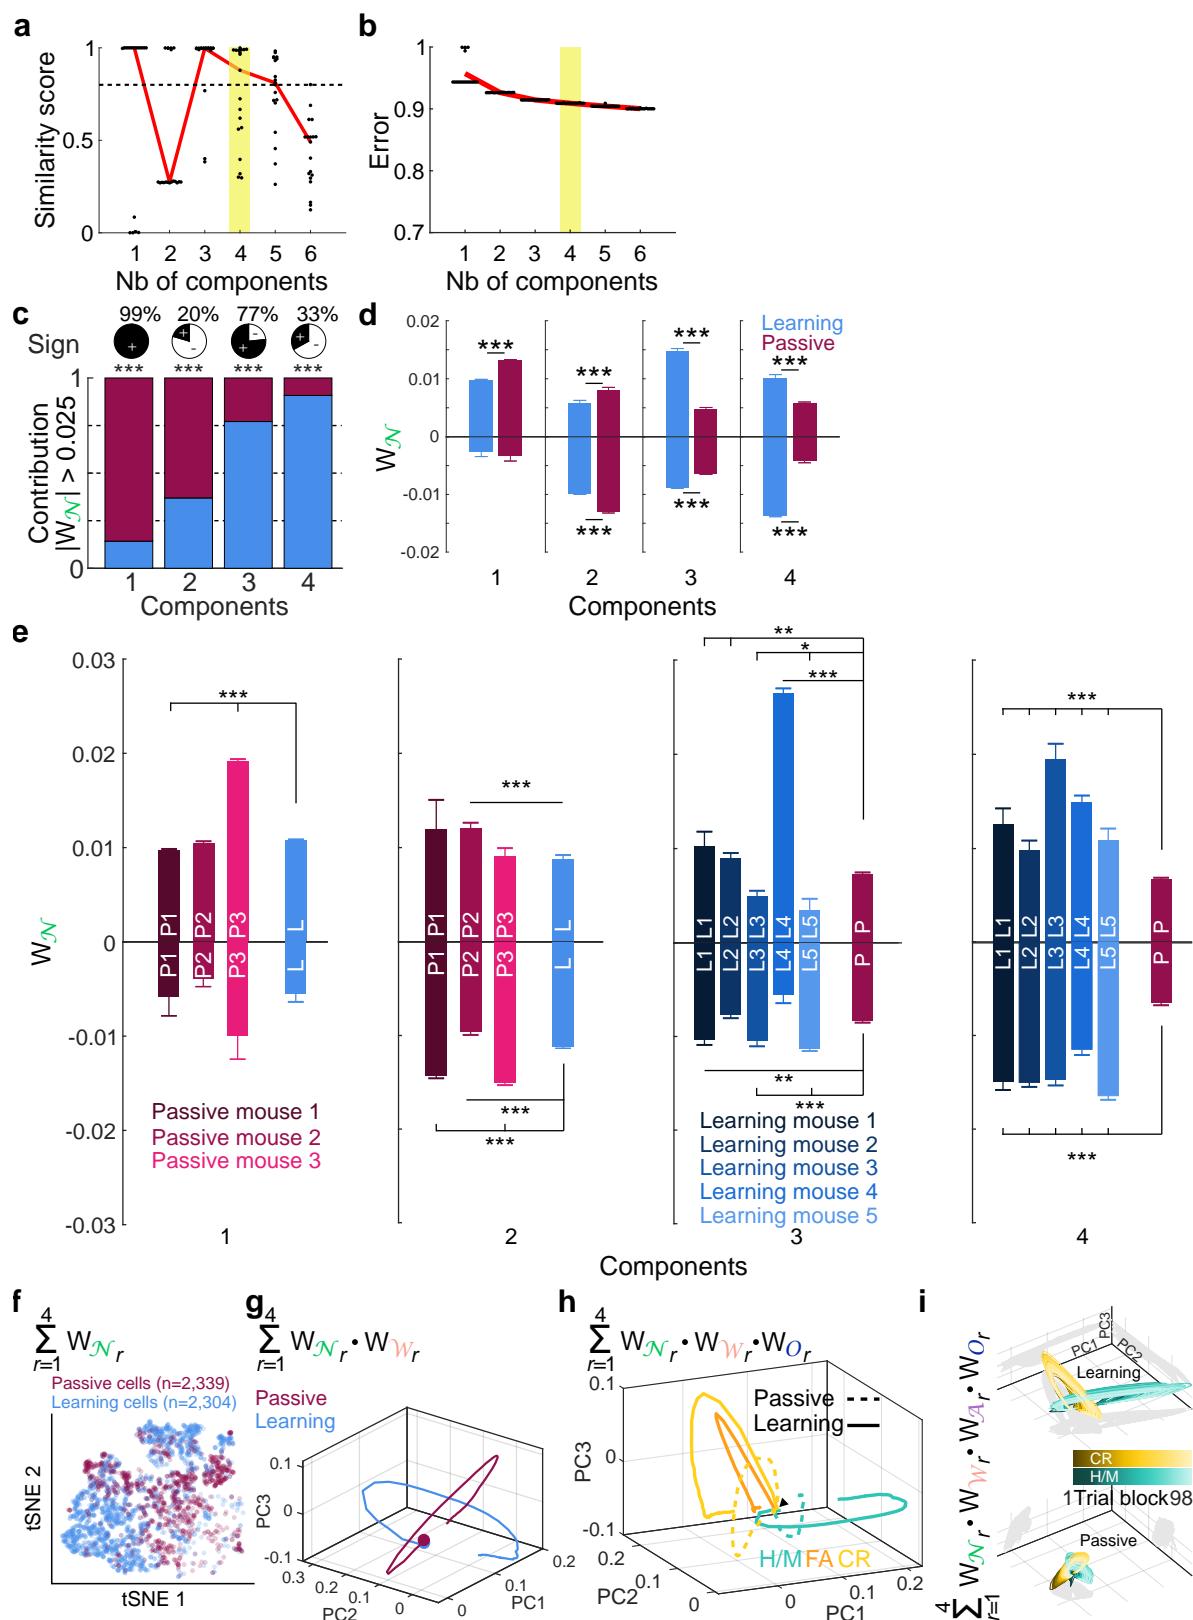
Supplementary Figure 8. Inter-subject performance alignment for megamouse tensor.

a, Accuracy in probe and reinforced contexts across days of all Learning mice. **b**, Action rate in reinforced context across days of all Learning mice. **c**, Action rate in probe context across days of all Learning mice. Please note that we fixed the probe performance at the maximum discrimination that was followed by a decrease in hit rate due to extinction. **d**, After the alignment procedure, action rate from the megamouse (all learning mice pooled) in reinforced context across learning phases. **e**, Megamouse accuracy in reinforced context across learning phases. **f**, Accuracy difference between the start and the end of the three learning phases in probe (green) and reinforced (black) contexts. Acquisition is characterized by an increase of accuracy in probe trials (paired t-test, $p = 5.47 \cdot 10^{-4}$) but not in reinforced trials (paired t-test, $p = 0.07$), Expression corresponds to an increase of accuracy in reinforced trials (paired t-test, $p = 0.008$) and Expert is when accuracy in reinforced trials is high and stable (paired t-test, $p = 0.27$).



Supplementary Figure 9. Tensor representation of neural data.

a, Data are organized into a fourth-order tensor with dimensions $N \times W \times A \times O$. Tensor decomposition approximates the data as a sum of outer products of four vectors. Each outer product contains a neuron factor (green rectangles), within trial factor (pink rectangles), across trial factor (blue rectangles) and outcome factor (purple rectangles). Each set of low-dimensional factors (i.e. component) describes the activity of group of neurons within and across trials according to trial outcomes.



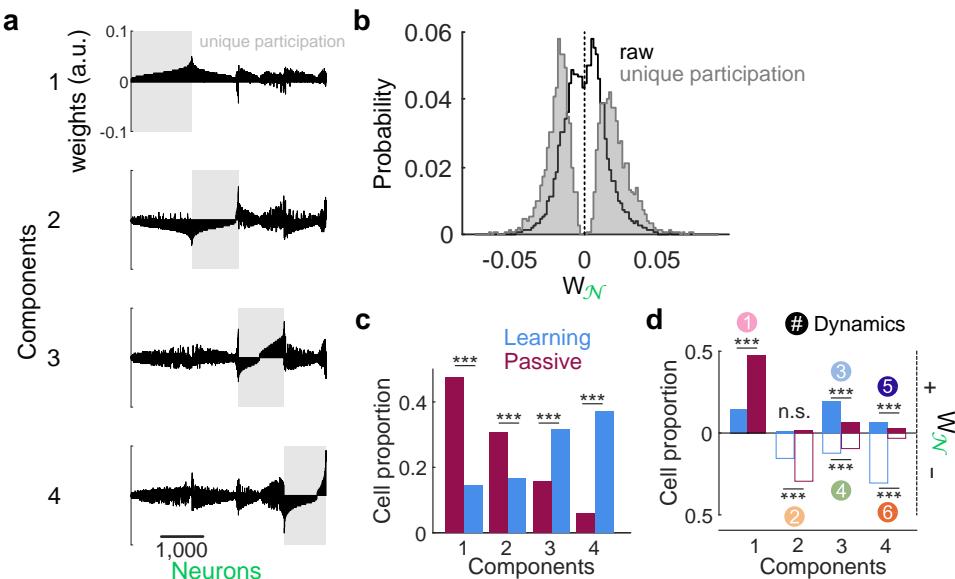
Supplementary Figure 10. Low-rank tensor decomposition.

a, Similarity score as a function of model components. Each dot shows the similarity of a single optimization run compared to the best-fit model within each category.

continued →

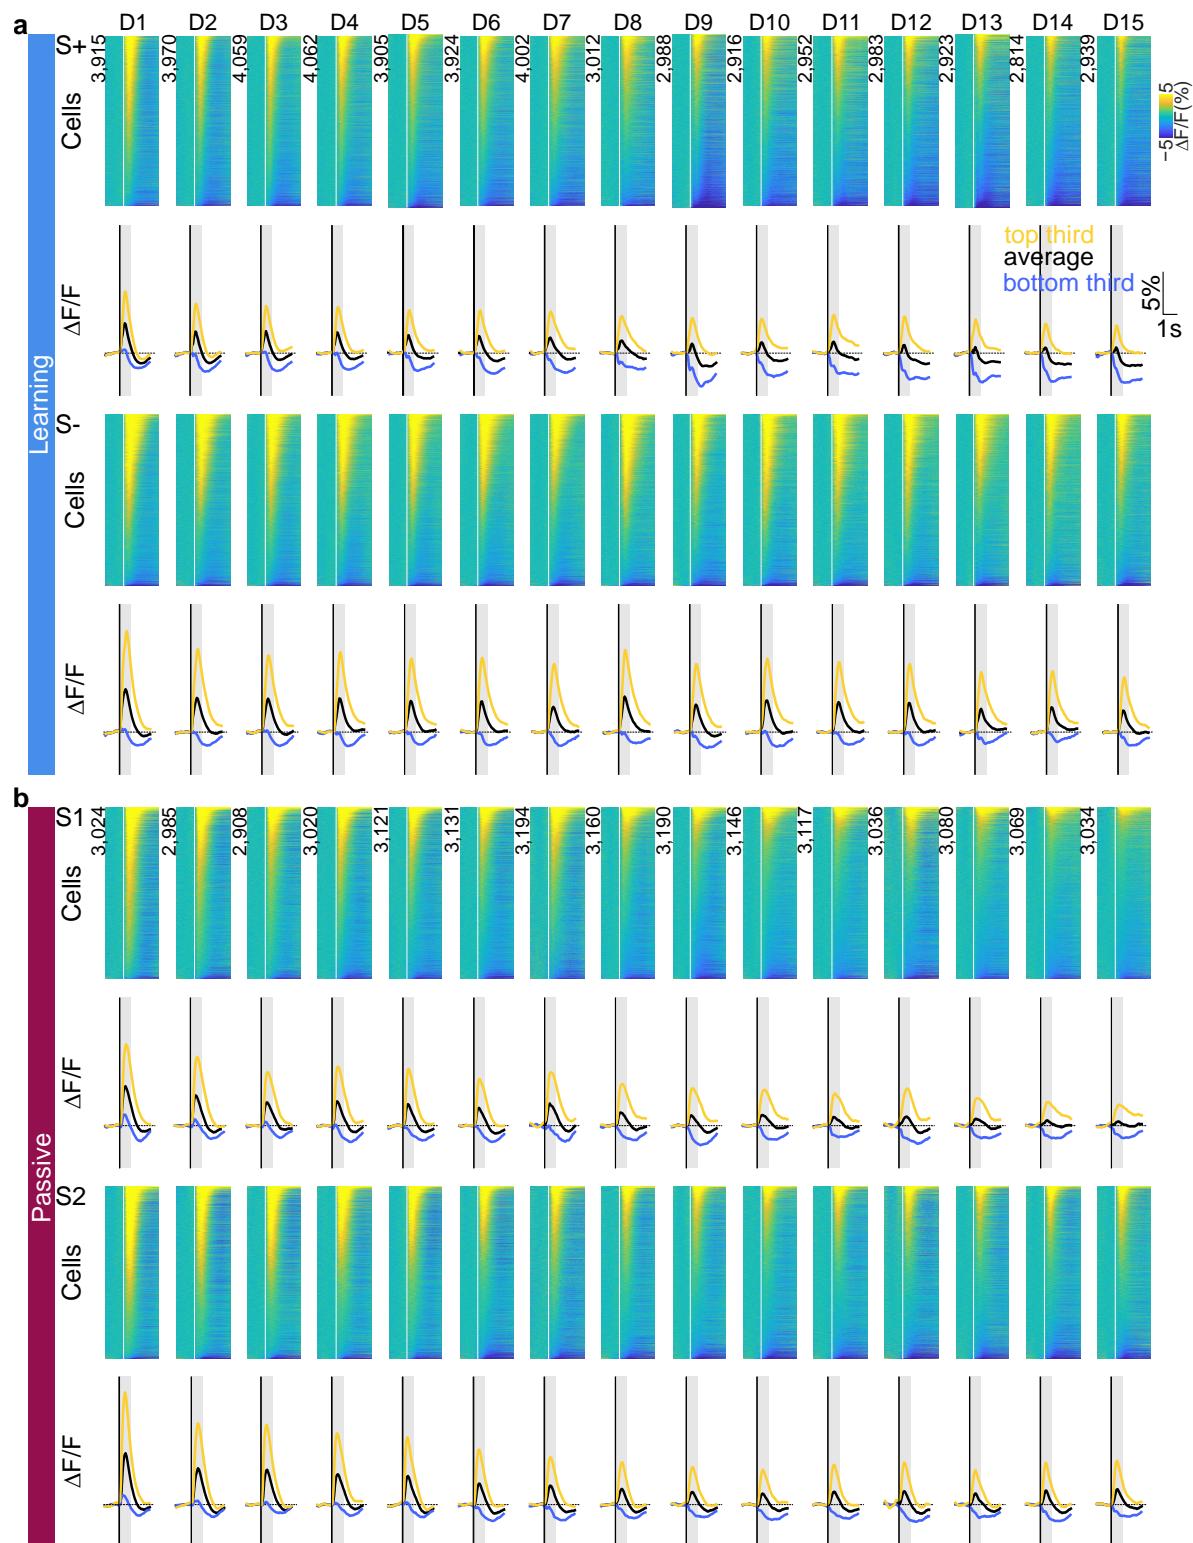
Supplementary Figure 10 (continued).

b, Model reconstruction error as a function of the number of components, where each dot corresponds to a different optimization run. **c**, Neuronal contribution (Learning vs Passive cells) per components (binomial proportion tests, all $p < 0.001$). **d**, Positive and negative neuronal weights across components in cell population recorded in learning mice (Learning) or in passive mice (Passive) (Wilcoxon tests). **e**, Positive and negative neuronal weights across components and individual mice. **f**, t-SNE of neuronal weights. Note how Learning and Passive cell populations are largely non-overlapping. **g**, Projection of neuronal \times within trial weights of Learning and Passive network activity into principal component space. **h**, Projection of neuronal \times within trial \times trial outcome weights of Learning and Passive network activity into principal component space. **i**, Projection of neuronal \times within trial \times across trials \times trial outcome (H/M and CR only) weights of Learning and Passive network activity into principal component space.



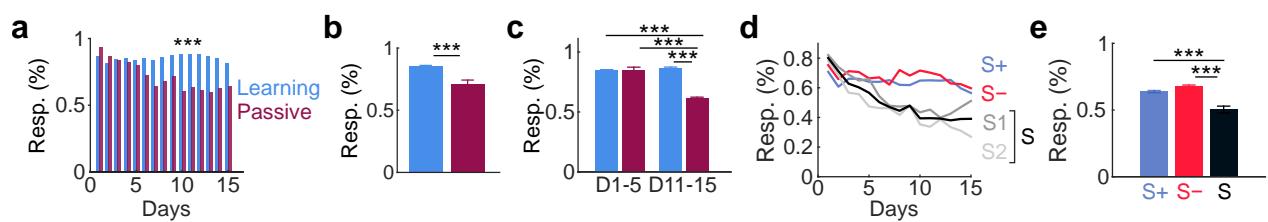
Supplementary Figure 11. Defining unique cell ensembles based on neuronal weights.

a, Neuronal weights in the four components. Each neuron is attributed to a given dynamic according to its highest absolute weights, i.e. highest contribution. As a result, each dynamic is attributed to a unique cell ensemble (gray rectangles). **b**, Neuronal weights distribution before (raw, black) and after unique contribution attribution (gray). **c**, Learning and Passive cell proportion among components after unique attribution (binomial proportion tests). **d**, Learning and Passive cell proportion among components and given neuronal weight sign after unique attribution. In other words, proportion of cells from Learning and Passive networks describing the tensor-revealed neuronal dynamics (binomial proportion tests). *** $p < 0.001$, n.s.: not significant.



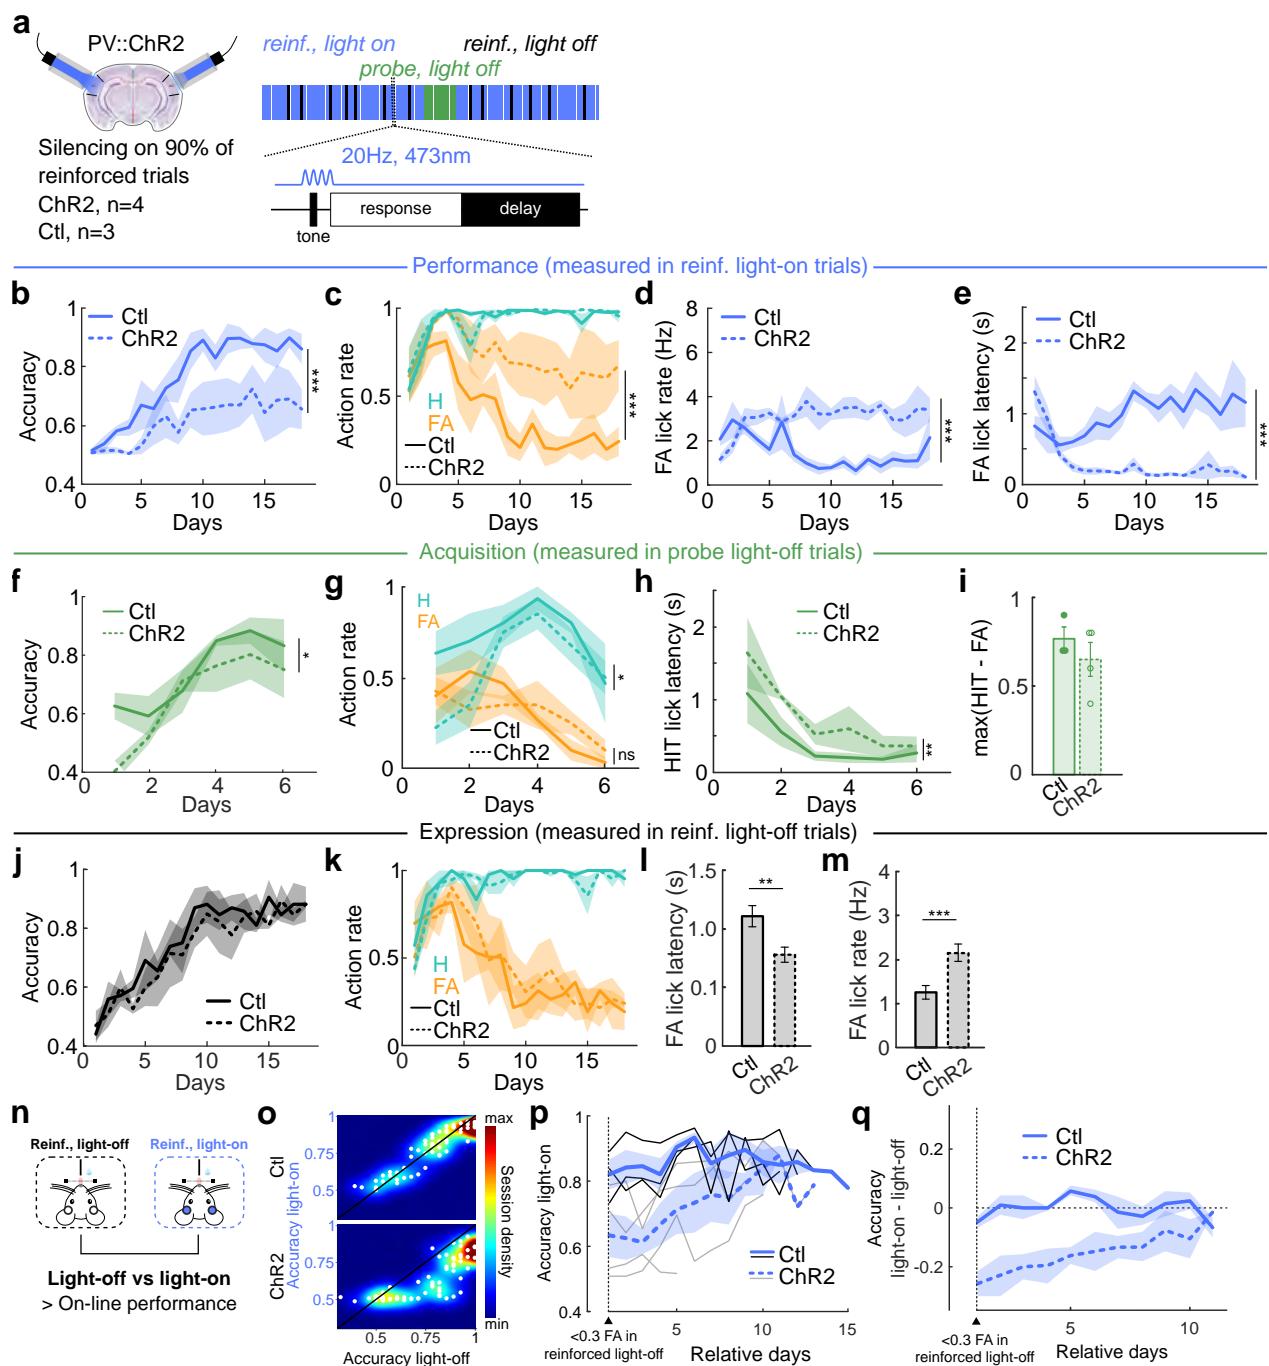
Supplementary Figure 12. Evolution of tone-evoked responses across days.

a, Tone-evoked responses to S+ and S- in Learning mice across days for all cells recorded. **b**, Tone-evoked responses to S1 and S2 in Passive mice across days for all cells recorded.



Supplementary Figure 13. Learning counteracts tone-evoked habituation.

a, Proportion of tone-responsive cells across days among Passive and Learning cells. **b**, Averaged proportion of tone-responsive cells in Passive and Learning networks (mean \pm s.e.m.; t-test, $p = 3.89 \cdot 10^{-5}$). **c**, Proportion of tone-responsive cells in days 1-5 versus days 11-15 in Learning and Passive networks (mean \pm s.e.m.; two-way ANOVA, Time \times Group, $p = 1.73 \cdot 10^{-7}$). **d**, Proportion of cells responsive to S+ and S- in Learning network and S1, S2 or S1 or S2 (S) in Passive network. **e**, Averaged proportion of cells responsive to S+, S- or S (mean \pm s.e.m.; ANOVA, $p = 1.93 \cdot 10^{-6}$).

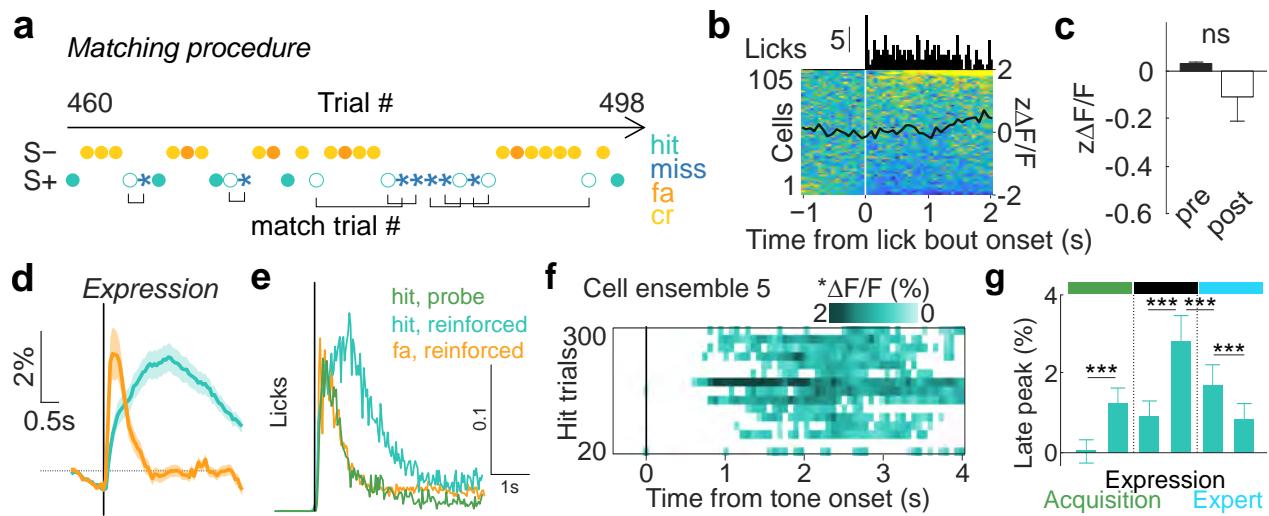


Supplementary Figure 14. AC silencing restricted to sound presentation impairs audiometer learning and on-line performance during learning.

a, Probabilistic optogenetic silencing of the auditory cortex during learning. Light-on periods were restricted to sound presentation only (see Methods). **b**, Accuracy in reinforced light-on trials (two-way ANOVA, Days: $F(17, 86) = 5.4950, p < 10^{-7}$; Groups: $F(1, 86) = 50.5343, p < 10^{-9}$; Days*Groups: $F(17, 86) = 0.70700, p = 0.79$). **c**, Action rate in reinforced light-on trials (HIT, two-way ANOVAs, HIT, Days: $F(17, 86) = 10.68010, p < 10^{-14}$; Groups: $F(1, 86) = 0.0200, p = 0.89$; Days*Groups: $F(17, 86) = 1.0647, p = 0.40$; FA, Days: $F(17, 86) = 2.7330, p = 0.0012$; Groups: $F(1, 86) = 41.5010, p < 10^{-8}$; Days*Groups: $F(17, 86) = 0.7255, p = 0.77$). *continued →*

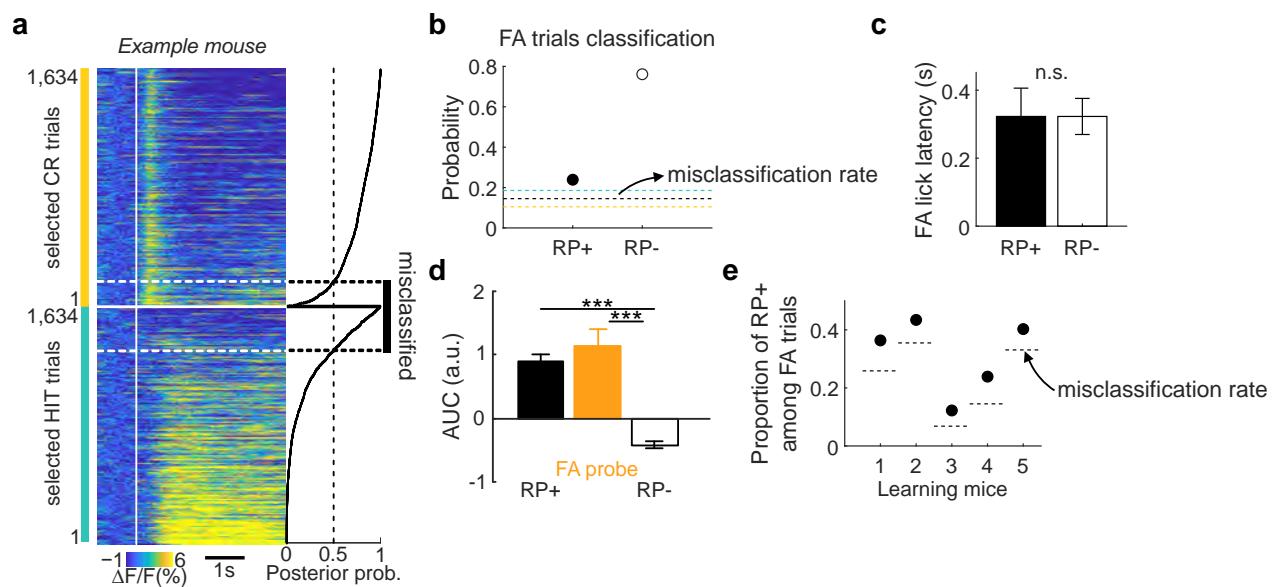
Supplementary Figure 14 (continued).

d, False alarm lick rate in reinforced light-on trials (two-way ANOVA, Days: $F(17, 86) = 0.8663, p = 0.6140$; Groups: $F(1, 86) = 89.3004, p < 10^{-14}$; Days*Groups: $F(17, 86) = 3.2285, p < 10^{-3}$).
e, False alarm lick latency in reinforced light-on trials (two-way ANOVA, Days: $F(17, 86) = 2.0216, p = 0.018$; Groups: $F(1, 86) = 251.7387, p < 10^{-26}$; Days*Groups: $F(17, 86) = 4.8600, p < 10^{-6}$).
f, Accuracy in probe light-off trials (two-way ANOVA, Days: $F(5, 30) = 8.3041, p < 10^{-4}$; Groups: $F(1, 30) = 4.7288, p = 0.038$; Days*Groups: $F(5, 30) = 0.7288, p = 0.619$). **g**, Action rate in probe light-off trials (two-way ANOVAs, HIT, Days: $F(5, 30) = 5.4632, p = 0.0011$; Groups: $F(1, 30) = 6.3510, p = 0.017$; Days*Groups: $F(5, 30) = 1.2158, p = 0.33$; FA, Days: $F(5, 30) = 5.5019, p = 0.0010$; Groups: $F(1, 30) = 0, p = 1$; Days*Groups: $F(5, 30) = 1.1320, p = 0.37$). **h**, HIT lick latency in probe light-off trials (two-way ANOVA, Days: $F(5, 29) = 6.0308, p < 10^{-3}$; Groups: $F(1, 29) = 10.3058, p = 0.0032$; Days*Groups: $F(5, 29) = 0.1542, p = 0.98$). **i**, Maximal difference between hit and false alarm rates in probe light-off trials over the first 6 days (t-test, $p = 0.40$).
j, Accuracy in reinforced light-off trials (two-way ANOVA, Days: $F(17, 86) = 8.3579, p < 10^{-11}$; Groups: $F(1, 86) = 1.6832, p = 0.20$; Days*Groups: $F(17, 86) = 0.2356, p = 1$).
k, Action rate in reinforced light-off trials (two-way ANOVAs, HIT, Days: $F(17, 86) = 11.1314, p < 10^{-14}$; Groups: $F(1, 86) = 2.1423, p = 0.15$; Days*Groups: $F(17, 86) = 0.9107, p = 0.56$; FA, Days: $F(17, 86) = 4.2760, p < 10^{-5}$; Groups: $F(1, 86) = 0.5043, p = 0.48$; Days*Groups: $F(17, 86) = 0.3026, p = 1$). **l**, FA lick latency in reinforced light-off trials (two-way ANOVA, Days: $F(17, 78) = 1.7364, p = 0.053$; Groups: $F(1, 78) = 9.0848, p = 0.0035$; Days*Groups: $F(17, 78) = 1.3749, p = 0.17$). **m**, FA lick rate in reinforced light-off trials (two-way ANOVA, Days: $F(17, 78) = 0.7983, p = 0.69$; Groups: $F(1, 78) = 13.4564, p < 10^{-3}$; Days*Groups: $F(17, 78) = 1.4494, p = 0.14$). **n**, Comparison of light-off versus light-on trials to measure auditory cortex silencing effect on on-line performance. **o**, Session density plot of accuracy in reinforced light-on against light-off. Top, control; bottom, PV-ChR2. **p**, Accuracy in light-on reinforced trials from day where FA < 0.3 in light-off reinforced trials. Note the general trend for ChR2 mice (gray lines) to increase accuracy (positive slopes), i.e. performance impairment fades away. **q**, Within subject difference between accuracy in reinforced light-on and light-off aligned to the day where false alarm rate < 0.3 in reinforced light-off.



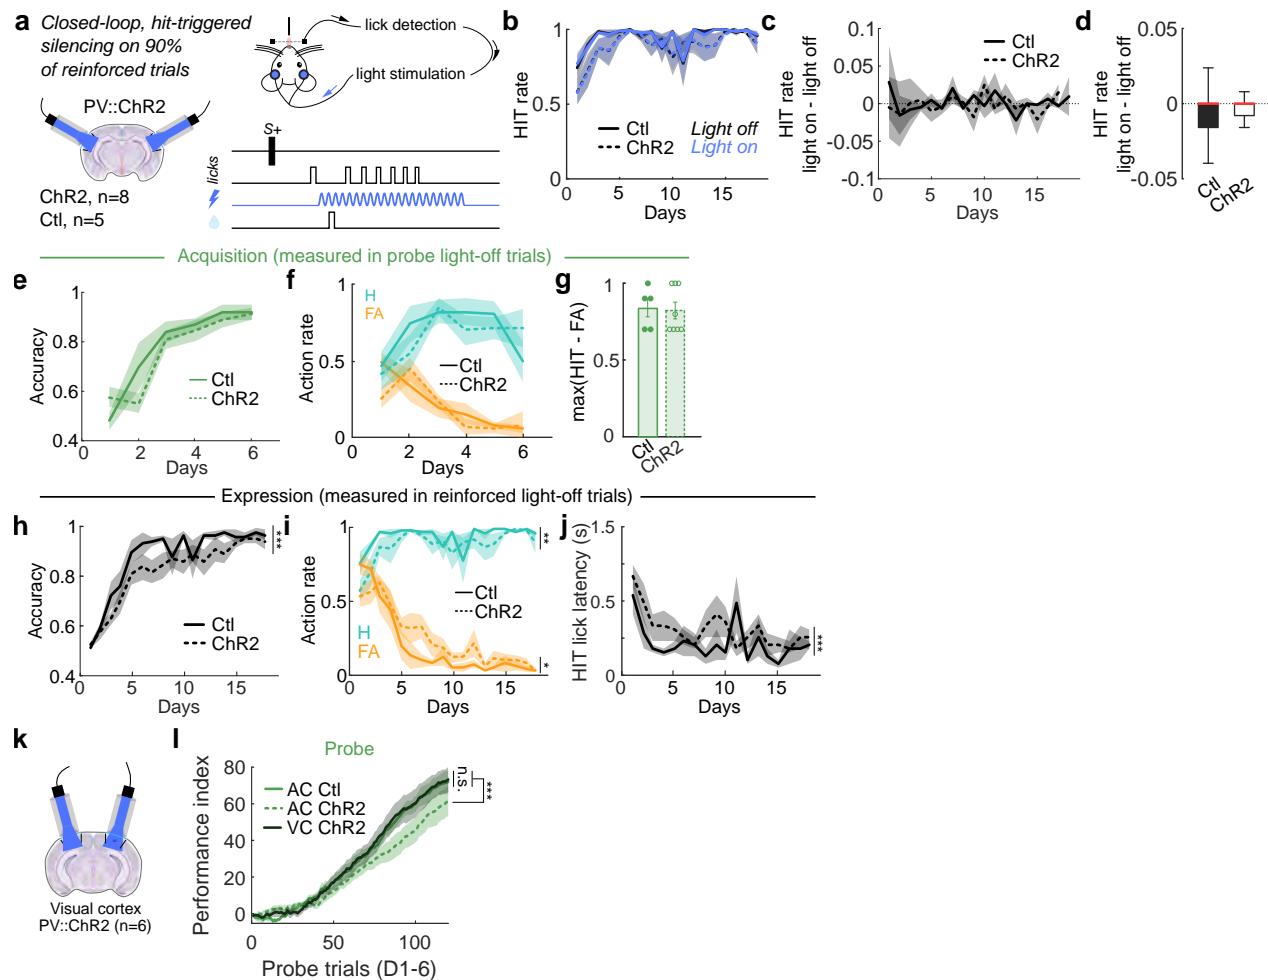
Supplementary Figure 15. Emergence of reward prediction signal.

a, Procedure of hit and miss trial matching. **b**, Heat map of members of cell ensemble 5 ($n = 105$) activity aligned to lick bout onset outside task events in day 1 of training. Lick PSTH is represented above. **c**, Quantification of z-scored calcium activity 1s pre- vs 1s post- lick bout onset (Wilcoxon test, $p = 0.11$). **d**, Average cell ensemble 5 activity in reinforced hit (green) and FA (orange) trials over Expression phase. **e**, Lick PSTHs aligned to tone onset of FA trials in Expression and hit trials in probe context. **f**, Cell ensemble 5 activity over the first 300 hit trials (20-trial blocks). Only significant activity (and higher than null population, see Methods) is represented. Note the emergence of a stable late-on-trial signal after 40 hit trials onwards. **g**, Quantification of Fig.4l, i.e. evolution of late-in-trial signal of cell ensemble 5 across learning, taking first and last two 40-hit trial blocks (KW test, $p = 1.05 \cdot 10^{-23}$). $*p < 0.05$, $**p < 0.01$, $***p < 0.001$, n.s.: not significant.



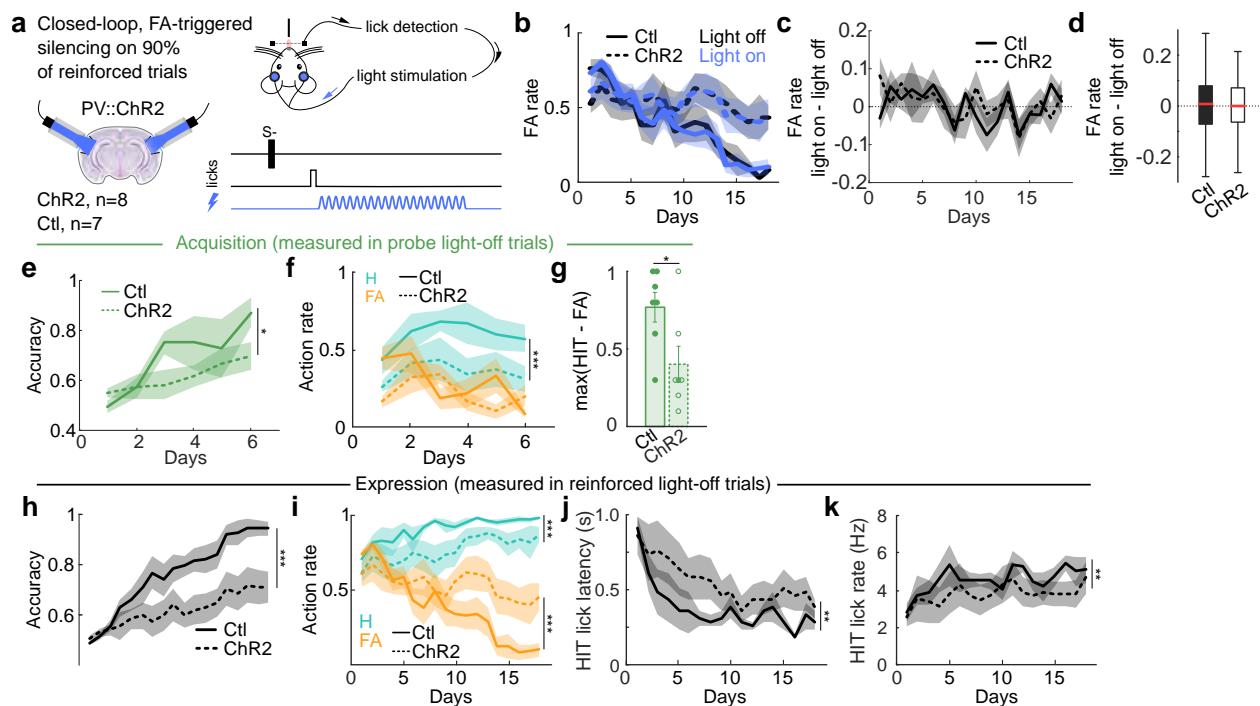
Supplementary Figure 16. Reward prediction signal on error trials.

a, Classification of hit versus CR trials in the reinforced context from the AUC post-tone of a fraction of cell ensemble 5 ($n = 51$) recorded in the example mouse showed in Fig.4p,q. Right: posterior probability of being part of CR class. **b**, Proportion of RP+ and RP– FA trials from the example mouse showed in Fig.4o,p. **c**, No difference in lick latency was observed between RP+ and RP– FA trials (Wilcoxon test, $p = 0.83$). **d**, AUC quantification of RP+, RP– and probe FA trials (KW, $p = 9.76 \cdot 10^{-28}$). **e**, Proportion of RP+ among all FA trials and misclassification rate in each learning mice. $*p < 0.05$, $**p < 0.01$, $***p < 0.001$, n.s.: not significant.



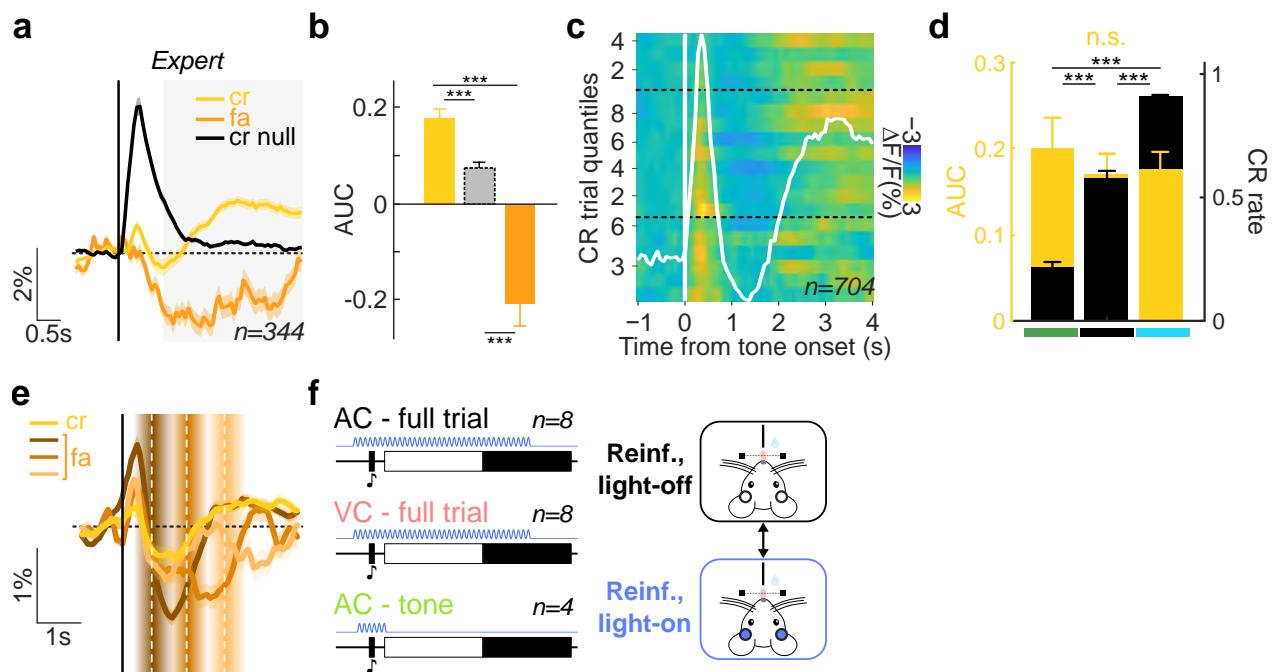
Supplementary Figure 17. Post-hit silencing over learning.

a, Experimental design of optogenetic silencing of AC activity throughout learning post hit only. **b**, Hit rate across days in control (Ctl) and test (ChR2) mice in reinforced light-on or light-off trials across days. **c**, Difference in hit rate in reinforced light-on versus light-off trials across days. **d**, Difference in hit rate in reinforced light-on versus light-off trials (Wilcoxon test, $p = 0.13$). **e**, Accuracy in probe light-off trials. **f**, Action rate (hit, H; false alarm, FA) in probe light-off trials. **g**, Maximum difference between hit and false alarm trials over the first 6 days in probe light-off trials. **h**, Accuracy in reinforced light-off trials. **i**, Action rate in reinforced light-off trials. **j**, Hit lick latency in reinforced light-off trials. **k**, Silencing of visual cortex (VC) activity throughout learning post hit only. **l**, Performance index in probe trials for AC control ($n = 5$), AC PV-ChR2 ($n = 8$) and VC PV-ChR2 ($n = 6$) (two-way ANOVA, $p = 1.90 \cdot 10^{-32}$).



Supplementary Figure 18. Post-FA silencing over learning.

a, Experimental design of optogenetic silencing of AC activity throughout learning post false alarm (FA) only. **b**, False alarm rate across days in control (Ctl) and test (ChR2) mice in reinforced light-on or light-off trials across days. **c**, Difference in false alarm rate in reinforced light-on versus light-off trials across days. **d**, Difference in false alarm rate in reinforced light-on versus light-off trials (t-test, $p = 0.76$). **e**, Accuracy in probe light-off trials. **f**, Action rate (hit, H; false alarm, FA) in probe light-off trials. **g**, Maximum difference between hit and false alarm trials over the first 6 days in probe light-off trials. **h**, Accuracy in reinforced light-off trials. **i**, Action rate in reinforced light-off trials. **j**, Hit lick latency in reinforced light-off trials. **k**, Hit lick rate in reinforced light-off trials.



Supplementary Figure 19. A signal for action suppression in Learning network.

a, Average activity of cell ensemble 6 or low weighted cells (null, black) in CR and FA trials in Expert phase. **b**, Quantification of late-in-trial activity (KW test, $p = 4.76 \cdot 10^{-44}$). **c**, Average cell ensemble 6 activity across learning phases. CR trials were split into 6, 9 and 4 quantiles over Acquisition, Expression and Expert phases, respectively. **d**, Quantification of late-in-trial activity (left axis) and CR rate (right axis) over learning phases. **e**, Averaged ensemble 6 activity in FA and CR trials. FA trials are split according to lick latencies (white dashed line, mean latency; graded rectangles, latency range extrema). **f**, Silencing protocols compared in Fig. 5h,i.

Test	Figure	Variable	Factor 1			Factor 2			Interaction	
			Name	F statistic	p-value	Name	F statistic	p-value	F statistic	p-value
2-way ANOVA	1f	Accuracy in reinf. Light-on trials	Days	F (20, 256) = 17.5118	< 10 ⁻³⁶	Groups	F (1, 256) = 195.4231	< 10 ⁻³²	F (20, 256) = 1.8215	0.02
2-way ANOVA	1g	HIT reinf. Light-on trials	Days	F (20, 256) = 6.9034	< 10 ⁻¹⁴	Groups	F (1, 256) = 3.3036	0.07	F (20, 256) = 0.7914	0.72
		FA reinf. Light-on trials	Days	F (20, 256) = 10.2452	< 10 ⁻²²	Groups	F (1, 256) = 197.5210	< 10 ⁻³³	F (20, 256) = 1.1138	0.33
2-way ANOVA	1h	Accuracy probe	Days	F (5, 84) = 17.5637	< 10 ⁻¹¹	Groups	F (1, 84) = 20.7994	< 10 ⁻⁴	F (5, 84) = 2.1360	0.07
2-way ANOVA	1i	S+ response index	Days	F (119, 1680) = 8.8	< 10 ⁻¹⁰⁸	Groups	F (1, 1680) = 532.07	< 10 ⁻¹⁰¹	F (119, 1680) = 1.85	< 10 ⁻⁵
2-way ANOVA	1l	Accuracy in reinf. Light-off trials	Days	F (20, 256) = 21.8381	< 10 ⁻⁴³	Groups	F (1, 256) = 39.9729	< 10 ⁻⁸	F (20, 256) = 1.1202	0.33
2-way ANOVA	1m	HIT reinf. Light-off trials	Days	F (20, 256) = 5.6985	< 10 ⁻¹¹	Groups	F (1, 256) = 0.3266	0.57	F (20, 256) = 0.4733	0.97
		FA reinf. Light-off trials	Days	F (20, 256) = 14.1390	< 10 ⁻³⁰	Groups	F (1, 256) = 38.8122	< 10 ⁻⁸	F (20, 256) = 0.8034	0.71
2-way ANOVA	1n	Aligned accuracy in reinf. Light-off trials	Days	F (20, 232) = 23.00	< 10 ⁻⁴³	Groups	F (1, 232) = 20.43	< 10 ⁻³	F (20, 232) = 1.4462	0.1
2-way ANOVA	1q	Aligned light-on-light-off reinf. Accuracy	Days	F (17, 181) = 1.91	0.02	Groups	F (1, 181) = 80.72	< 10 ⁻¹⁵		
rm ANOVA	4m	Accuracy	Groups	F (2,66)=9.13	3.16.10 ⁻⁴					
2-way ANOVA	4s	Performance index	Days	F(119,1320)=47.7299	0	Groups	F(1,1320)=93.4275	2.11.10 ⁻²¹	F(119,1320)=0.9718	0.57
		HIT lick latency	Days	F(8,95)=0.5799	0.79	Groups	F(1,95)=6.4473	0.013	F(8,95)=0.4567	0.8833
2-way ANOVA	4t	Performance index	Days	F(119,1140)=3.8374	5.51.10 ⁻³⁴	Groups	F(1,1440)=16.0877	6.36.10 ⁻⁵	F(119,1440)=0.5461	1
2-way ANOVA		HIT lick latency	Days	F(6,104)=0.9029	0.5172	Groups	F(1,104)=11.1571	0.0012	F(6,104)=0.8896	0.53
2-way ANOVA	5h	FA rate	Days	F(16,279)=1.4149	0.13	Groups	F(2,279)=39.6122	7.20.10 ⁻¹⁶	F(32,279)=1.1953	0.22
2-way ANOVA	5i	Av. Lick probability	Days	F(16,253)=1.1327	0.32	Groups	F(2,253)=11.8720	1.18.10 ⁻⁵	F(32,253)=1.4887	0.05
rm 2-way ANOVA	Extended Data Fig.5c	Accuracy	Days	F (20, 140) = 15.6714	1.47.10 ⁻²⁶	Groups	F (1, 140) = 0.5093	0.51		
2-way ANOVA	Extended Data Fig.5f	Accuracy light-on - light-off	Days	F (20, 271) = 1.5547	0.06	Groups	F (1, 271) = 2.3072	0.13	F (20, 271) = 1.154	0.295
2-way ANOVA	Extended Data Fig.13c	Proportion responsive cells	Time	F(1,16)=57.9347	1.05.10 ⁻⁶	Groups	F(1,16)=68.3328	3.62.10 ⁻⁷	F(1,16)=76.3993	1.73.10 ⁻⁷
1-way ANOVA	Extended Data Fig.13e	Proportion responsive cells	Groups	F(2,57)=16.7209	1.93.10 ⁻⁶					
2-way ANOVA	Extended Data Fig.14b	Accuracy light-on trials	Days	F (17, 86) = 5.4950	< 10 ⁻⁷	Groups	F (1, 86) = 50.5343	< 10 ⁻⁹	F (17, 86) = 0.70700	0.79
2-way ANOVA	Extended Data Fig.14c	HIT light-on trials	Days	F (17, 86)=10.68010	< 10 ⁻¹⁴	Groups	F (1, 86) = 0.0200	0.89	F (17, 86) = 1.0647	0.4
2-way ANOVA	Fig.14d	FA light-on trials	Days	F (17, 86) = 2.7330	0.0012	Groups	F (1, 86) = 41.5010	< 10 ⁻⁸	F (17, 86) = 0.7255	0.77
		FA lick rate light-on trials	Days	F (17, 86) = 0.8663	0.61	Groups	F (1, 86) = 89.3004	< 10 ⁻¹⁴	F (17, 86) = 3.2285	< 10 ⁻³
2-way ANOVA	Extended Data Fig.14e	FA lick latency light-on trials	Days	F (17, 86) = 2.0216	0.018	Groups	F (1, 86) = 251.7387	< 10 ⁻²⁶	F (17, 86) = 4.8600	< 10 ⁻⁶
2-way ANOVA	Extended Data Fig.14f	Accuracy probe	Days	F (5, 30) = 8.3041	< 10 ⁻⁴	Groups	F (1, 30) = 4.7288	0.038	F (5, 30) = 0.7288	0.62
2-way ANOVA	Extended Data Fig.14g	HIT probe	Days	F (5, 30)=5.4632	0.0011	Groups	F (1, 30)=6.351	0.017	F (5, 30)=1.2158	0.33
2-way ANOVA	Extended Data Fig.14g	FA probe	Days	F (5, 30)=5.5019	0.001	Groups	F (1, 30)=0	1	F (5, 30)=1.1320	0.37
2-way ANOVA	Extended Data Fig.14h	HIT lick latency	Days	F (5, 29) = 6.0308	< 10 ⁻³	Groups	F (1, 29) = 10.3058	0.0032	F (5, 29) = 0.1542	0.98
2-way ANOVA	Extended Data Fig.14j	Accuracy light-off trials	Days	F (17, 86) = 8.3579	< 10 ⁻¹¹	Groups	F (1, 86) = 1.6832	0.2	F (17, 86) = 0.2356	1
2-way ANOVA	Extended Data Fig.14k	HIT light-off trials	Days	F (17, 86) = 11.1314	< 10 ⁻¹⁴	Groups	F (1, 86) = 2.1423	0.15	F (17, 86) = 0.9107	0.56
2-way ANOVA	Extended Data Fig.14k	FA light-off trials	Days	F (17, 86) = 4.2760	< 10 ⁻⁵	Groups	F (1, 86) = 0.5043	0.48	F (17, 86) = 0.3026	1
2-way ANOVA	Extended Data Fig.14l	FA lick latency light-off trials	Days	F (17, 78) = 1.7364	0.053	Groups	F (1, 78) = 9.0848	0.0035	F (17, 78) = 1.3749	0.17
2-way ANOVA	Extended Data Fig.14m	FA lick rate light-off trials	Days	F (17, 78) = 0.7983	0.69	Groups	F (1, 78) = 13.4564	< 10 ⁻³	F (17, 78)=1.4494	0.14
2-way ANOVA	Extended Data Fig.17e	Accuracy probe	Days	F(5,66)=28.4499	2.96.10 ⁻¹⁵	Groups	F(1,66)=1.1059	0.3	F(5,66)=1.6693	0.15
2-way ANOVA	Extended Data Fig.17f	HIT probe	Days	F(5,66)=3.7418	0.0048	Groups	F(1,66)=0.3151	0.58	F(5,66)=0.9120	0.48
2-way ANOVA	Extended Data Fig.17h	FA probe	Days	F(5,66)=10.0172	3.66.10 ⁻⁷	Groups	F(1,66)=0.5165	0.47	F(5,66)=1.6770	0.1524
2-way ANOVA	Extended Data Fig.17h	Accuracy reinf. trials	Days	F(17,196)=21.3723	8.07.10 ⁻³⁶	Groups	F(1,196)=16.4149	7.33.10 ⁻⁵	F(17,196)=0.6588	0.84
2-way ANOVA	Extended Data Fig.17i	HIT reinf. trials	Days	F(17,196)=3.0906	7.78.10 ⁻⁵	Groups	F(1,196)=7.4877	0.0068	F(17,196)=0.7175	0.78
2-way ANOVA	Extended Data Fig.17i	FA reinf. trials	Days	F(17,196)=19.3955	2.70.10 ⁻³³	Groups	F(1,196)=5.2646	0.023	F(17,196)=1.2621	0.22
2-way ANOVA	Extended Data Fig.17j	HIT lick latency	Days	F(17,196)=1.0782	0.38	Groups	F(1,196)=9.9595	0.0019	F(17,196)=0.9034	0.57
2-way ANOVA	Extended Data Fig.17l	Performance index	Days	F(119,1920)=73.4277	0	Groups	F(2,1920)=75.8901	1.90.10 ⁻³²	F(238,1920)=0.7307	0.999
2-way ANOVA	Extended Data Fig.18e	Accuracy probe	Days	F(5,78)=4.7464	7.74.10 ⁻⁴	Groups	F(1,78)=4.9199	0.03	F(5,78)=1.2012	0.32
2-way ANOVA	Extended Data Fig.18f	HIT probe	Days	F(5,78)=1.0833	0.38	Groups	F(1,78)=16.6632	1.07.10 ⁻⁴	F(5,78)=0.1594	0.98
2-way ANOVA	Extended Data Fig.18f	FA probe	Days	F(5,78)=2.7704	0.02	Groups	F(1,78)=2.7111	0.1	F(5,78)=2.6530	0.029
2-way ANOVA	Extended Data Fig.18h	Accuracy reinf. trials	Days	F(17,199)=7.7483	1.19.10 ⁻¹⁴	Groups	F(1,199)=68.0233	2.17.10 ⁻¹⁴	F(17,199)=1.3683	0.16
2-way ANOVA	Extended Data Fig.18i	HIT reinf. trials	Days	F(17,199)=2.0458	0.011	Groups	F(1,199)=32.4866	4.28.10 ⁻⁸	F(17,199)=0.2519	0.1
2-way ANOVA	Extended Data Fig.18i	FA reinf. trials	Days	F(17,199)=4.1289	4.34.10 ⁻⁷	Groups	F(1,199)=19.9782	1.31.10 ⁻⁵	F(17,199)=1.5640	0.077
2-way ANOVA	Extended Data Fig.18j	HIT lick latency	Days	F(17,199)=1.8158	0.028	Groups	F(1,199)=11.5978	7.99.10 ⁻⁴	F(17,199)=0.2495	0.1
2-way ANOVA	Extended Data Fig.18k	HIT lick rate	Days	F(17,199)=0.9297	0.54	Groups	F(1,199)=8.1477	0.005	F(17,199)=0.2253	0.1

Supplementary Table 1. Report of ANOVA statistics.