Article

Neuron

Context-Dependent Decision Making in a Premotor Circuit

Highlights

- Mice performed an olfactory delayed match to sample (DMS) task by licking left or right
- Inactivation of premotor area ALM in the sample and delay periods impaired performance
- Signals in upstream areas can be decoded to solve the task, but ALM disregards them
- Neurons in layer 2 of ALM represent the identity of the sample odor through the delay

Authors

Zheng Wu, Ashok Litwin-Kumar, Philip Shamash, Alexei Taylor, Richard Axel, Michael N. Shadlen

Correspondence

ra27@columbia.edu (R.A.), shadlen@columbia.edu (M.N.S.)

In Brief

Wu et al. demonstrate that premotor area ALM is crucial to decide whether a test odor matches a preceding sample. ALM maintains a trace of the sample in layer 2 and establishes the appropriate association between test odor and behavioral response.

Wu et al., 2020, Neuron *106*, 316–328
April 22, 2020
https://doi.org/10.1016/j.neuron.2020.01.034

CellPress

CellPress

Neuron
Article

Context-Dependent Decision Making in a Premotor Circuit

Zheng Wu,[1,2] Ashok Litwin-Kumar,[2] Philip Shamash,[2,3] Alexei Taylor,[1,2] Richard Axel,[1,2,*] and Michael N. Shadlen[1,2,4,*]
[1]Howard Hughes Medical Institute, Columbia University, New York, NY 10027, USA
[2]Mortimer B. Zuckerman Mind Brain Behavior Institute and Department of Neuroscience, Columbia University, New York, NY 10027, USA
[3]Present address: Sainsbury Wellcome Centre, University College London (UCL), London, UK
[4]Lead Contact
*Correspondence: ra27@columbia.edu (R.A.), shadlen@columbia.edu (M.N.S.)
https://doi.org/10.1016/j.neuron.2020.01.034

SUMMARY

Cognitive capacities afford contingent associations between sensory information and behavioral responses. We studied this problem using an olfactory delayed match to sample task whereby a sample odor specifies the association between a subsequent test odor and rewarding action. Multineuron recordings revealed representations of the sample and test odors in olfactory sensory and association cortex, which were sufficient to identify the test odor as match or non-match. Yet, inactivation of a downstream premotor area (ALM), but not orbitofrontal cortex, confined to the epoch preceding the test odor led to gross impairment. Olfactory decisions that were not context-dependent were unimpaired. Therefore, ALM does not receive the outcome of a match/non-match decision from upstream areas. It receives contextual information—the identity of the sample—to establish the mapping between test odor and action. A novel population of pyramidal neurons in ALM layer 2 may mediate this process.

INTRODUCTION

Brain functions deemed cognitive exhibit complexities that extend an organism's behavioral repertoire beyond simple sensory-response associations, motor programs, and instructed actions. Cognitive functions exploit contingent, hierarchical processes of decision-making and executive control. For example, instead of choosing an action, a decision might lead an organism to choose a strategy, switch tasks, or make yet another decision. Cognitive functions may also transpire over flexible timescales without precipitating an immediate behavior, as when a decision is based on information acquired at some moment and combined with information acquired later.

Neural mechanisms of decision-making and executive control have been studied primarily in the parietal and prefrontal cortex of nonhuman primates. These studies focused mainly on neurons that exhibit responses over long timescales and implicated processes such as working memory (Funahashi et al., 1989; Fuster and Alexander, 1971), planned action (Cisek and Kalaska, 2005; Evarts and Tanji, 1976; Snyder et al., 1997), behavioral state (Harvey et al., 2012; Kadohisa et al., 2013), representation of stimulus qualities (Freedman and Assad, 2006; Romo et al., 1999), and reasoning (Gold and Shadlen, 2007; Yang and Shadlen, 2007). The study of the neural mechanisms of perceptual decision-making may therefore provide an initial logic for higher-order cognition. This will require an understanding of neural circuits at a level that is not yet possible to achieve in nonhuman primates. Thus, there has been growing interest in the pursuit of elementary cognitive functions in the mouse, for which genetic and viral tools for circuit manipulation abound (Carandini and Churchland, 2013; Luo et al., 2018). The challenge is to find simple behaviors within the mouse repertoire that have the potential to elucidate more complex cognitive functions.

We developed a simple task that allows us to study rudiments of executive control and decision-making in a mouse. The task is a variant of the logical *exclusive or* (XOR) problem, realized as an olfactory delayed match to sample task (DMS; Figure 1) (Liu et al., 2014). The mouse is exposed to a *sample odor*, either S_A or S_B, and, after a short delay, receives a *test odor*, T_A or T_B. To receive a reward, the mouse must lick to the left if the sample and test odors are the same and to the right if they are different. The XOR problem interests computational neuroscientists because it cannot be solved by a simple linear classifier. Our interest was in the possibility that it could be solved by a hierarchy of decisions. For example, the mouse might use the sample odor to decide on the appropriate association between the test odor and the correct behavioral response—a left or right lick.

We hypothesized that this hierarchical control might be solved by changing the configuration of cortical circuitry in the premotor cortex, in the area anterolateral motor cortex (ALM). Studies from Svoboda and colleagues have shown that ALM plays an essential role in behaviors in which a sensory cue serves as an instruction to lick to the left or right (Guo et al., 2014a; Li et al., 2015; Svoboda and Li, 2018). They showed that many neurons in ALM can hold such an instruction—or the plan to lick to the left or right—in persistent activity through a delay period. In the DMS task we study, the sample odor cannot provide an instruction to lick left or right, only a context to interpret the test odor. Thus, the mouse cannot plan a lick until the arrival of the test odor after the delay. We investigated three hypotheses for the

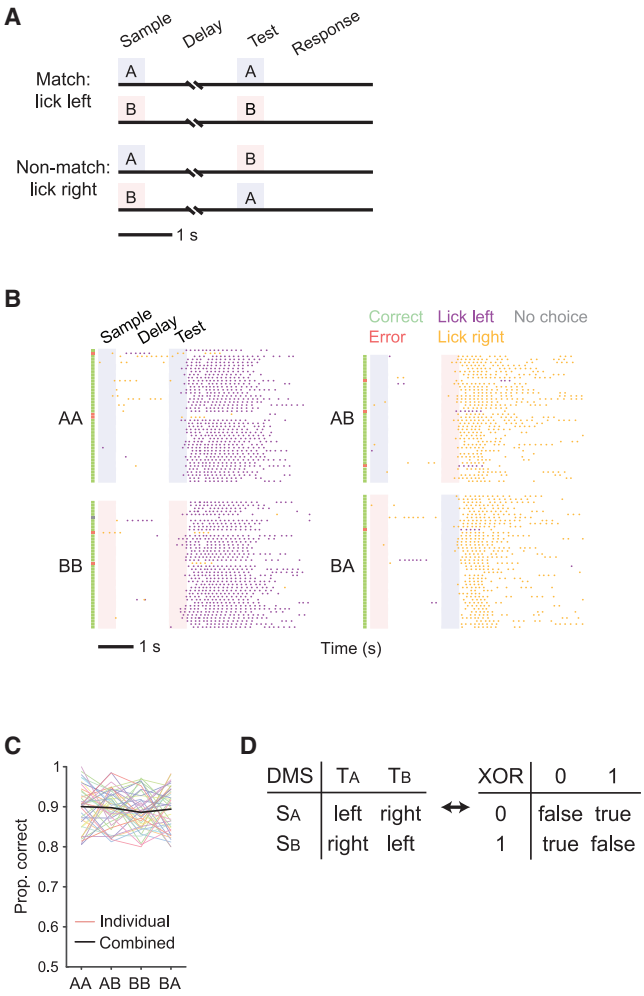

Figure 1. Lick-Left Lick-Right Olfactory Delayed Match to Sample Task

(A) Task structure. Mice are presented with two odorants (A, pinene; B, hexenol) separated by a delay and must decide whether they are the same (match) or different (non-match). The two odors create four unique pairs or trial types. Match trials are rewarded on the left port and non-match on the right (see STAR Methods).

(B) An example behavior session from a well-trained animal. The four trial types are randomly presented in the session but grouped here for plotting. Each row is a trial. The green, red, and gray markers on the left side of each trial denote, respectively, the outcome of correct, error, and no choice. The A and B odors are indicated by the light blue and light red shadings. The magenta and yellow tick marks are the left and right licks, respectively. Mice were trained to suppress their premature licks during the sample and delay epochs. The sample and test epochs are 0.5 s each and the delay is 1.5 s.

(C) The proportion of correct trials in well-trained animals. Each colored line represents data from one animal and the black line is the mean of all animals (n = 41).

(D) The DMS task is a variant of the logical XOR problem. The logical outcome and the lick response depend on whether two bits of information are the same or not ($T_{A|B}$, test odor A or B; $S_{A|B}$, sample odor A or B).

See also Figure S1.

involvement of ALM in this task. First, areas upstream to ALM could solve the match/non-match discrimination and project this solution to ALM, which organizes an appropriate lick response. Second, ALM could decode upstream representations that combine sample and test odors to select the appropriate response. Third, ALM could receive information about the identity of the sample odor to instantiate the appropriate response to the test odor. We provide experimental evidence for this third possibility. We show that the sample odor is represented in ALM itself, and that this representation allows ALM to associate the test odor with the appropriate lick response.

RESULTS

Mice were trained to compare a sample and test odor, separated by a delay, and to report their decisions of *match* or *non-match* by licking to the left or to the right, respectively (Figures 1A, S1A, and S1B). The lick-left, lick-right design requires distinct actions to report both match and non-match. Importantly, animals cannot solve the task until they smell the test odor. For most experiments, we used the same two odors, (+)-α-pinene and *cis*-3-hexen-1-ol (odors A and B, respectively) on all trials. Animals were trained to suppress their premature licks before test onset (Figure 1B) and do not appear to use licks to represent and remember the sample odors (Figures S1C–S1E).

Mice performed the task with a median accuracy of 90% (Figure 1C). When challenged with a novel pair of odorants after training with pinene and hexenol, mice performed at chance level with a high no-choice rate (Figure S1F). This suggests that the mice failed to generalize the match/non-match rule. Rather, they learned that *test odor* T_A instructs *lick left* if preceded by *sample odor* S_A (AA trial), whereas it instructs *lick right* if preceded by *sample odor* S_B (BA trial). The complementary logic holds for the *test odor* T_B (Figure 1D). This flexible association between test odor and licking response is contingent on the sample odor identity, which must be represented through the delay in order to affect the match/non-match decision.

We first characterize the neural representation of the sample and test odors by surveying neural responses in the Piriform cortex (Pir), the orbitofrontal cortex (OFC), and the ALM while the mouse performed the DMS task. We wished to determine whether each of these areas contain a persistent representation of the sample odor and whether they encode the test odor in a way that might inform the match/non-match decision. As detailed in the next section, the findings support the possibility that the match/non-match decision can be established within Pir and OFC and then transmitted to the ALM to render the lick response. We then test this model by inactivating ALM during the sample and delay epochs of the task. We find evidence against this model and instead establish a necessary role of ALM during the sample and delay epochs. Finally, using 2-photon Ca imaging, we expose a class of neurons in superficial ALM that could mediate the circuit changes required to allow ALM to associate the test odor with an appropriate lick response.

Representations of Sensory and Motor Signals in Pir, OFC, and ALM

We performed electrophysiological recordings in brain areas likely to be involved in the DMS task: Pir, OFC, and ALM. Pir is a primary olfactory sensory area (Giessel and Datta, 2014;

Sosulski et al., 2011) that encodes odor identity. Pir projects to OFC, a higher-order associative area that is thought to encode value, expectation, and working memory (Bechara et al., 2000; Mainen and Kepecs, 2009; Padoa-Schioppa and Assad, 2006; Ramus and Eichenbaum, 2000). Figures 2A–2C provide examples of neurons that respond selectively to (1) one or the other sample odor during the sample and delay periods (1st and 2nd columns); (2) one or the other test odor (3rd column); and (3) one or the other choice (4th column). They also show the strength of the selectivity across all the neurons recorded. The selectivity index quantifies the degree to which the distributions of firing rates to the two conditions (e.g., odor A or B) are non-overlapping (see STAR Methods). In what follows we refer to a neuron as selective if the difference in the mean responses is statistically reliable (indicated by shading in Figures 2D–2G).

More than a third of neurons (37%) in Pir responded selectively to either odor A or odor B during the sample epoch (Figures 2A and 2D). Of these, 71% exhibited the same odor preference during the sample and test epochs (Figure S2A), suggesting that these neurons represent sensory information and encode odor identity. Most of the sample-selective neurons responded transiently (e.g., Figure 2A, left most panel), but 28% exhibited persistent selectivity through the delay (e.g., Figure 2A, 2nd panel from left; Figure S2B). These neurons can therefore inform downstream neurons of the identity of the sample odor at the time of test. A subset of Pir neurons responded to the test odor in a way that depended on the identity of the sample odor (Figure S2C; STAR Methods). These neurons are characterized as trial-type-selective (6.2% of all Pir neurons). Finally, 4.5% of Pir neurons exhibited selective responses during the test epoch that reflect either a match or non-match between sample and test odors (Figures 2A and 2G). The information represented by these last two classes of neurons, trial type and match/non-match, would appear to be sufficient to guide motor output in downstream brain areas.

In OFC, 24.5% of the neurons were odor-selective during the sample epoch (Figures 2B and 2D) and 60% of these exhibited the same odor preference at test (Figure S2D), suggesting that they encode odor identity. Compared to Pir, fewer of the sample-selective neurons in OFC exhibited persistent spiking through the delay epoch (Figure S2E; 13% in OFC versus 28% in Pir). Only 3.7% of the neurons were selective to trial type (Figure S2F), but a greater fraction (11%) responded selectively to match/non-match trials (Figure 2G). Therefore, the OFC also appears to have integrated sensory inputs in a way that could guide appropriate motor output.

Unlike Pir and OFC, which largely represent odor identity, the dominant task-related activity in ALM was correlated with the lick response—the outcome of the match/non-match decision; 28% of ALM neurons distinguished match from non-match trials during the test epoch (Figures 2C and 2G). Analyses of error trials and premature licks show that the neural activity reflects lick direction rather than the identities of the sample and test odors or the true match/non-match condition, consistent with the role of ALM in motor planning (Figures S2J and S2K) (Guo et al., 2014a; Li et al., 2015; Svoboda and Li, 2018). A largely non-overlapping population in ALM (12%) exhibited weak odor selectivity during the sample epoch, and among these, 12% (less than 1.5% overall) maintained this selectivity through the end of the delay epoch (Figures 2D and S2H). Very few neurons exhibited a selective response to the test odor, and only 1.4% of ALM neurons were trial-type-selective (Figure S2I). Together, these recordings are consistent with the known role for ALM in the preparation of an appropriate motor response (Svoboda and Li, 2018) following the match/non-match decision.

This survey of a sensory, association, and premotor cortex reveals neurons in each area that exhibit one or more of the properties required to solve the DMS task: (1) a representation of the sample odor during the delay epoch, (2) a response to the test odor that is possibly modulated by the identity of the sample odor, and (3) the conversion of these sensory signals to choice-related activity (Figure 2H). Figure 3A shows the averaged difference in firing rate to each neuron's preferred versus nonpreferred odor during the trial. The assignment of preferred odor was derived from five randomly selected trials to each odor, which are excluded from the averages. The procedure ensured an unbiased estimate of the difference (see STAR Methods). The average difference is strongest in Pir and weakest in ALM. In the left panel, the magnitude of the difference reflects a combination of greater selectivity and the fraction of neurons that are selective, as nonselective neurons drive the average toward zero. In the right panel, the averages comprise only the selective cases from each area (filled histograms in Figure 2). All three areas contain signals that could convey the representation of the sample odor (Pir > OFC > ALM). However, this conclusion is based on averaged activity across trials.

To determine whether such task relevant information is available on single trials, we trained linear classifiers to decode these variables from simultaneously recorded neurons (from 14 to 120 simultaneously recorded neurons; see STAR Methods). As shown in Figure 3B, all three areas contain signals that support classification of the sample odor above chance levels throughout the sample and delay epochs. The performance of the classifier is lower for ALM recordings compared to Pir recordings, for which the sample odor identity can be decoded throughout the delay period and into the test epoch. This extended sample-selective response could potentially be used to perform a comparison between sample and test odors. Consistent with this, classification of all four trial types is possible using the Pir recordings (Figure 3C). In this same epoch (test odor onset to the first lick response), it is possible to decode the binary choice from all three areas, with ALM exhibiting the highest performance (Figure 3D).

The neural recordings and decoding analyses would appear to support a traditional hierarchical view of the mechanism responsible for resolving the match/non-match decision. That is, the expression of the decision by the left- and right-lick neurons in ALM could arise by reading out activity in OFC or Pir, perhaps via another intermediate area between Pir and ALM. On this view, ALM is only essential once the test odor arrives, when it must either convert the match/non-match decision to a lick response or perform a computation, similar to our decoder, that converts the population response from upstream areas to a lick response. Both mechanisms lead to the prediction that inactivation of ALM during the sample and delay epochs should

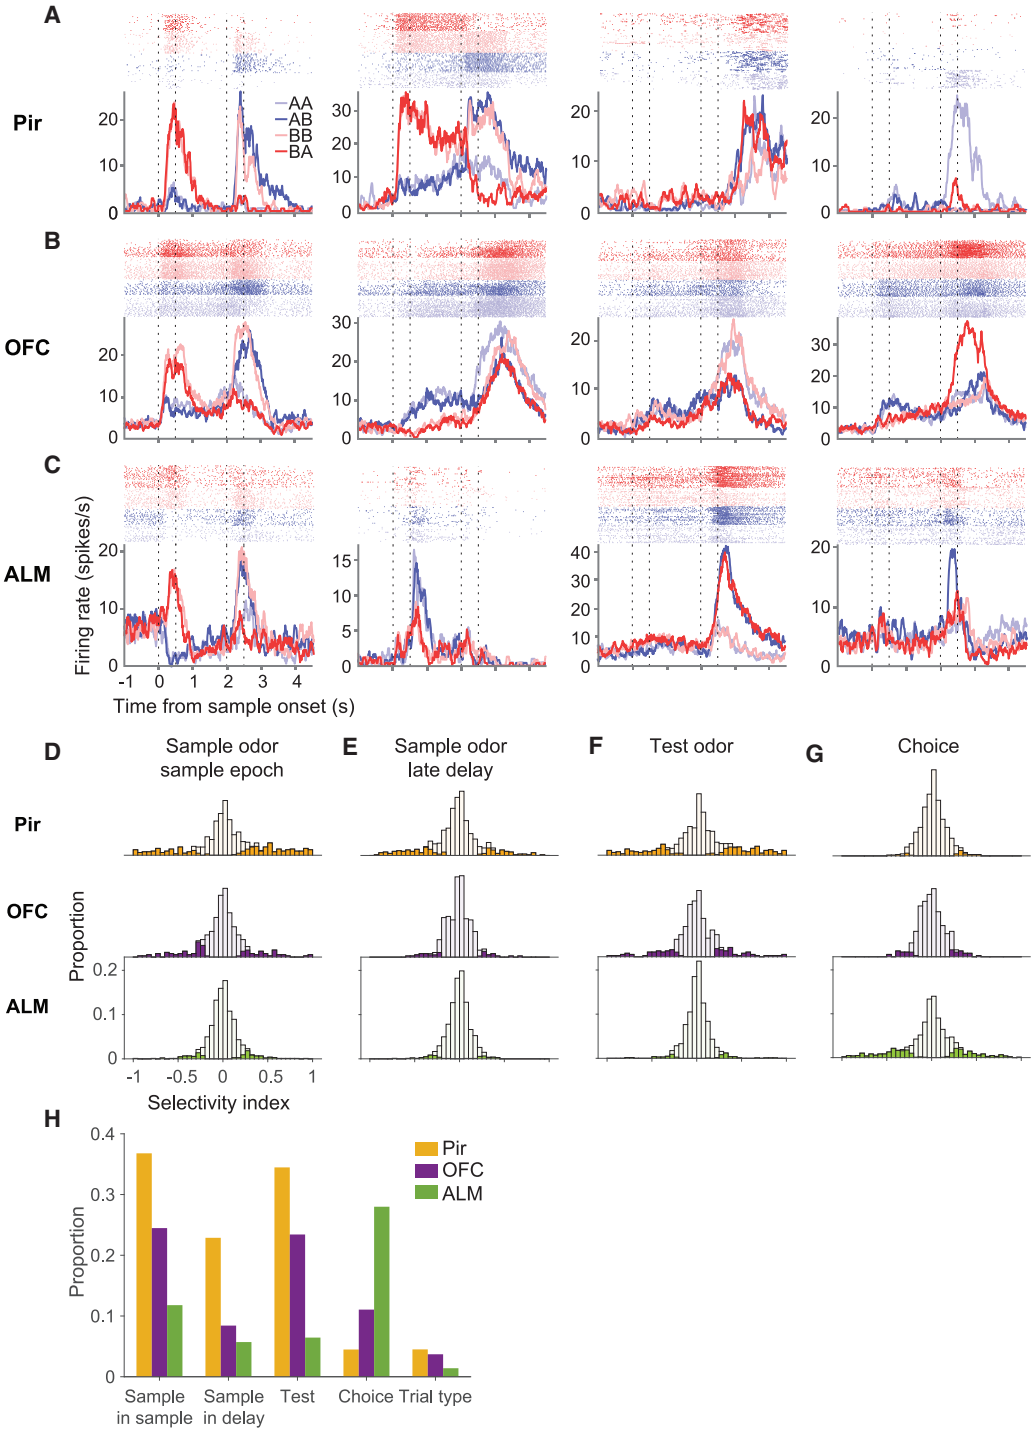

Figure 2. Neural Recordings of Pir, OFC, and ALM

(A–C) Firing rates of example neurons from Pir (A), OFC (B), and ALM (C) recordings. The rasters (top) show the times of action potentials in individual trials (rows). Traces are mean firing rates (100 ms bins) across trials. The dotted lines delineate the sample and test epochs.

(D–G) Selectivity indices of Pir, OFC, and ALM neurons for sample odor during the sample epoch (D), sample odor during the last 0.5 s of the delay epoch (E), test odor during the test epoch (F), and choice during the test epoch (G). n = 647, 380, and 1,086 neurons for Pir, OFC, and ALM, respectively. Shading indicates neurons that are selective for an odor (D)–(F) or behavioral choice (G) determined by Mann-Whitney U test, $p < 0.01$, not corrected for multiple comparisons.

(H) Summary of the feature-selective neurons in Pir, OFC, and ALM. Left to right: proportion of neurons selective for sample odor in the sample epoch, sample odor in the late delay, test odor in the test epoch, choice in the test epoch, and trial type in the test epoch.

See also Figures S2 and S7.

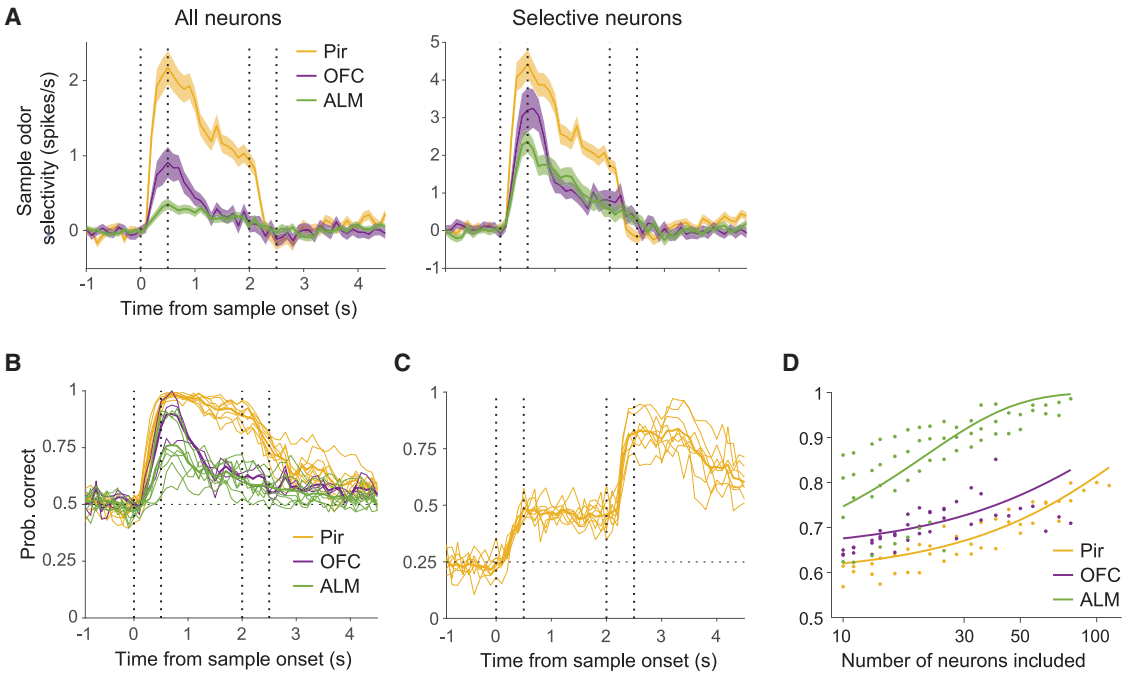

Figure 3. Odor Selectivity and Decoding Task-Relevant Features Using the Neural Responses

(A) Population sample odor selectivity of Pir, OFC, and ALM neurons. The ordinate is the difference in firing rate to the preferred and the nonpreferred odor (STAR Methods). The lines represent the means across neurons and the shading represents the SEM. Left graph, all recorded neurons from each area contribute to the averages. Right graph, only neurons that are selective for sample odor in the sample and delay epochs contribute to the averages (n = 319, 105 and 147 neurons for Pir, OFC, and ALM, respectively).

(B) Performance of a support vector machine trained to classify the sample odor identity using simultaneously recorded neurons from each of the three areas. Each thin line represents the probability of correct classifications of the held out trials using data from one session (see STAR Methods). Thick lines represent the mean of all sessions from an area (color).

(C) Performance of a classifier of trial type using all simultaneously recorded Pir neurons from each session. Same conventions as in (A).

(D) Performance of a classifier of match/non-match trials using the neural responses after test odor onset and before the animal's first lick. Each point represents the performance using N randomly selected neurons as input to the classifier, calculated using sessions in which at least N neurons were recorded simultaneously. The traces are fits to the classifier performance (logistic regression). For each area, the performance improves as more neurons are included.

not impair performance on the task. As we next show, this seemingly obvious prediction is incorrect.

ALM Is Required for the Match/Non-match Computation

We inactivated ALM bilaterally on a fraction of the trials by photostimulation of GABAergic interneurons expressing channelrhodopsin-2 (ChR2) (Figures 4A and S3; STAR Methods) (Guo et al., 2014a; Zhao et al., 2011). Inactivation was confined to the sample and delay epochs and was tapered gradually at the end of the delay (see STAR Methods) such that by the onset of the test odor, ALM should have been capable of receiving information from upstream sensory and association areas. We reasoned that if ALM reads out the match/non-match decision computed in upstream areas, then inactivation of ALM before the arrival of the test odor should not impair performance.

Contrary to this prediction, bilateral inactivation of ALM markedly impaired performance on the DMS task (Figure 4B). The proportion of correct trials decreased from 0.86 on the control trials to 0.68 (0.5 reflects chance performance). In contrast, bilateral inactivation of OFC induced no impairment on the same task (proportion correct: 0.82 and 0.84 with and without inactivation; Figure 4C). Photostimulation of ALM in animals that did not express ChR2 did not diminish performance (Figure 4D). Inactivation of ALM caused less impairment when it was restricted to only a portion of the sample plus delay epochs (Figure 4E). Inactivation during the late delay reduced the proportion of correct trials from 0.86 to 0.79, whereas inactivation during the sample and early delay reduced the proportion to 0.83. These more modest effects suggest that the impairment can be ameliorated when there is a time window in which ALM can receive task related information.

The behavioral impairment induced by ALM inactivation is a consequence of altered activity within ALM itself during the sample and delay. Photoinactivation of ALM did not affect odor or match/non-match selectivity in upstream areas, Pir and OFC. During inactivation, the population odor selectivity in Pir or OFC did not change appreciably (Figures 4I, 4J, 4L, and 4M), and our ability to decode match versus non-match from activity was similarly unchanged (Figures 4O and 4P). Neural recordings from ALM during silencing demonstrated that optogenetic suppression eliminated the representation of the sample odor, but neural activity returned after the laser was ramped down (Figures 4K, 4N, 4Q, and 4R). Despite their diminished

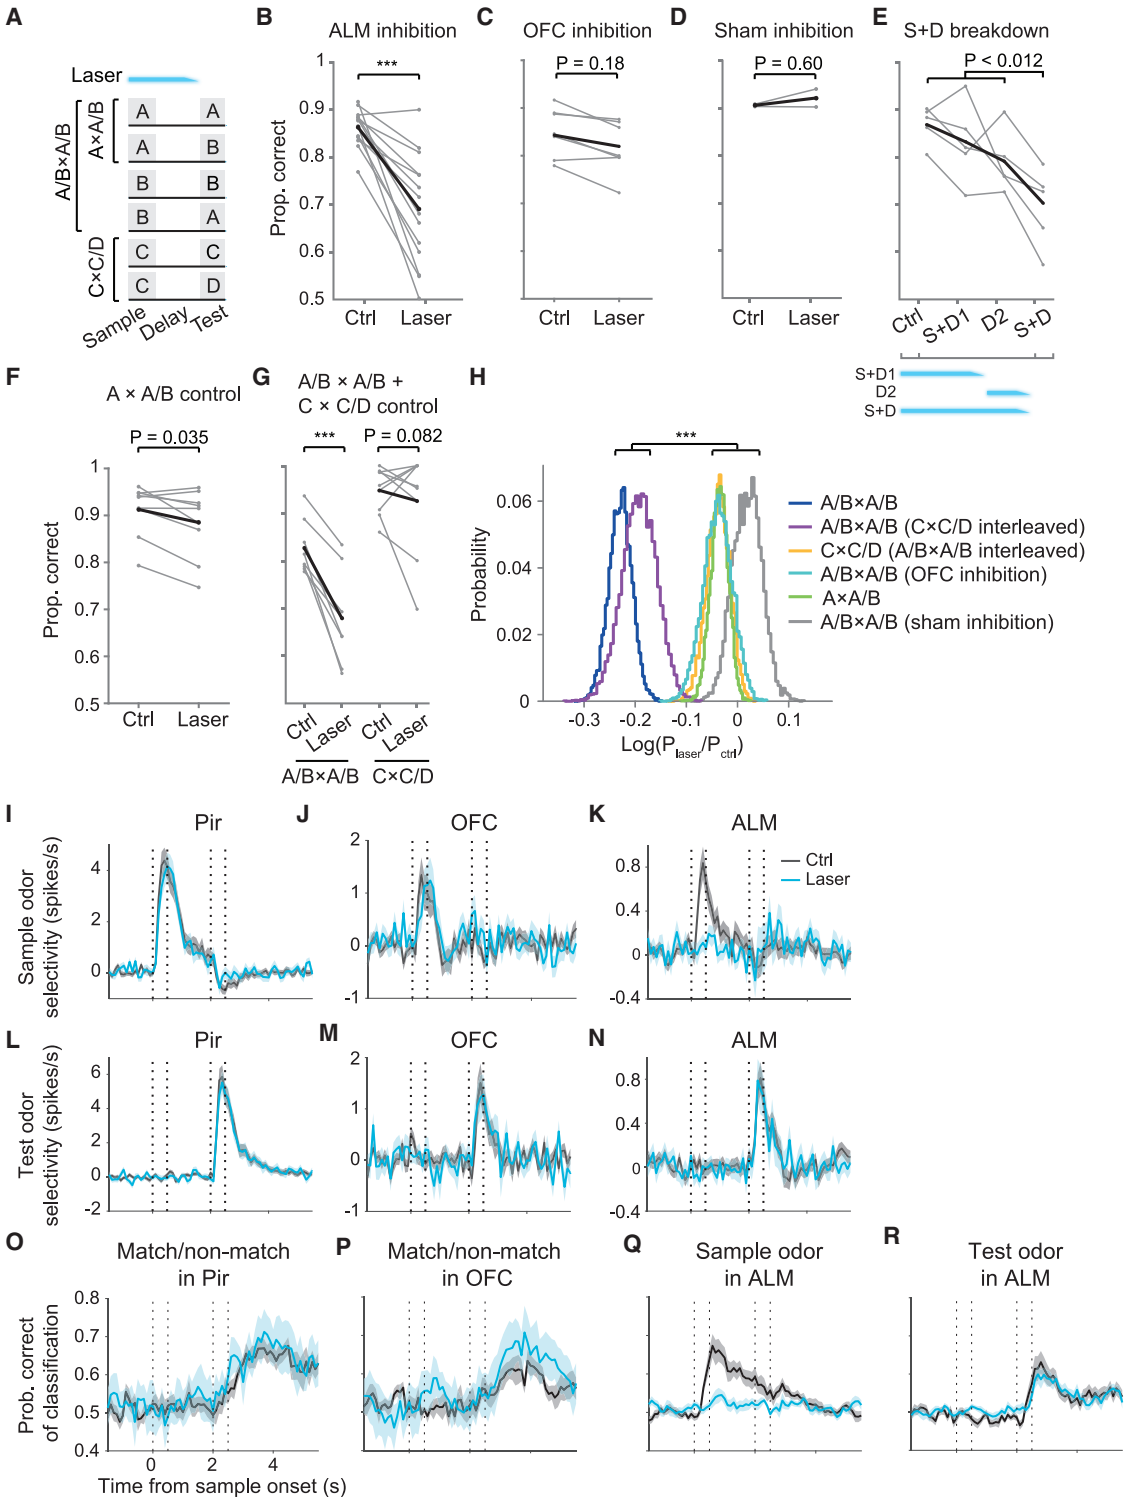

Figure 4. Optogenetic Inactivation Experiments

(A) Schematic of the DMS and control tasks employed in the inhibition experiments (time not drawn to scale). A/B × A/B design is the same DMS task as in the recording experiments. On inactivation trials, unless otherwise indicated, the laser was on for 2 s including a 0.25 s power ramp-down at the end (tapered blue bar). In the interleaved A/B × A/B and C × C/D task, the delay epoch is 4 s (STAR Methods).

(legend continued on next page)

performance, animals invariably licked to one of the two ports in inactivation trials, and there was no consistent change in the reaction time (Figures S4A and S4B). From these observations, we conclude that the impairment was not explained by a loss of processing capacity in Pir or OFC or from physiological sequelae of photoinactivation of ALM that might affect its function during the test epoch. This last conclusion deserves further scrutiny.

A possible concern is that the inactivation of ALM during sample and delay epochs disrupts the ALM circuitry such that it is unable to process information about the test odor or receive information about the decision from an upstream area. We evaluated this possibility with two control experiments that required the mice to lick to the left or right based on the identity of a test odor but did not require a comparison of the test odor to the sample odor. In the first control, an uninformative sample odor A was presented on all trials, but the mouse was rewarded for licking left or right for test odor A or B, respectively (AA and AB trials; Figure 4A). Inactivation during the sample and delay epochs produced minimal impairment (0.91 to 0.88 correct; Figure 4F).

In the second control, we incorporated two additional trial types into the AB × AB design using two new odors (C and D). In these CC and CD trials, as in the previous control, the correct lick behavior was determined only by the test odor. The six trial types were randomly interleaved (Figure 4A). Inactivation of ALM during the sample and delay epochs led to minimal impairment in CC and CD trials (0.94 to 0.91 correct), whereas significant impairment was replicated in the interleaved DMS trials (0.82 to 0.68 correct; Figure 4G). The degree of impairment on each type of experiment is captured by a simple statistic: the log probability ratio of correct choices with and without inactivation (Equation 4). We used Monte Carlo methods to estimate the magnitude and uncertainty of this statistic for each of the experimental and control experiments. As is clear from Figure 4H, inactivation impaired performance selectively when it was targeted to ALM and when the sample odor was informative (Figure 4H, sample informative tasks versus sample uninformative tasks, $p < 0.0001$; STAR Methods). Further, the laser introduced no side bias in the CC/CD control ($\beta_2 = -0.46$, SEM = 0.40, $p = 0.25$, Equation 8; see STAR Methods), demonstrating that the bias observed in some of the DMS trials is not due to unbalanced inactivation of the two hemispheres (see also Figure S4C).

These control experiments demonstrate that after recovery from inactivation, ALM is capable of processing information that instructs a licking response via a simple association between two odors and two actions. The impairment on the DMS task must therefore arise by interfering with the process that allows the sample odor to establish, on each trial, the appropriate association between test odor and lick response. Moreover, it implies that this process occurs in ALM.

Enrichment of Sample-Selective Neurons in ALM Layer 2

Based on the neural recordings, the necessity of ALM during the sample and delay seems highly perplexing. At the end of the delay period, less than 1.5% of neurons in ALM had activity that was informative about the identity of the sample. We considered that we might have missed neurons, especially from superficial cortical layers (Figure S2L; see also Figure 3D in Guo et al., 2017). We therefore examined neural responses in ALM with 2-photon calcium imaging while mice performed the DMS task. Imaging was performed in mice expressing the calcium indicator GCaMP6f in pyramidal cells (Chen et al., 2013; Madisen et al., 2015) (STAR Methods). Consistent with our electrical recordings, calcium imaging revealed that ~37% of the neurons were choice-selective across all cortical depths examined (Figures S5A and S5B).

However, we also observed a large subset of neurons with striking odor selectivity during the sample and delay epochs (Figures 5A, S5C, and S5D). Interestingly, the odor-selective calcium activity was heterogeneous in its timing, showing a variety of latencies and timescales (Figures 5B–5E). Some of the signals spanned the sample and part of the delay period, peaking at consistent times across trials (long duration, low scatter; e.g., Figure 5C). Other neurons exhibited more transient responses that occurred at different times across trials (short duration, high scatter; e.g., Figure 5D). Although they are ordered on the graph for visualization, these brief, scattered activations appeared random. For example, their timing is uncorrelated in

(B) Proportion of correct trials in the A/B × A/B task with and without bilateral ALM inactivation. Each thin line represents data from one animal. The thick line is the combined performance from 4,092 trials (14 animals). Statistical reliability of effects in panels (B)–(G) is based on permutation tests (STAR Methods) (*, **, and *** denote $P < 0.01$, 0.001, and 0.0001).
(C) Proportion of correct trials in the A/B × A/B task with and without bilateral OFC inactivation (7 animals, 2,278 trials).
(D) Proportion of correct trials in the A/B × A/B task by ChR2 non-carriers with and without sham ALM inactivation (2 animals, 437 trials).
(E) Effect of inactivation during portions of the sample and delay epochs (5 animals, 3,732 trials). Same conventions as in (B)–(D). Ctrl, no inactivation; S+D1, sample and first 1.5 s of the delay; D2, last 1.0 s of the delay; S+D, sample and delay.
(F) Proportion of correct trials in the A × A/B control task with and without bilateral ALM bilateral inactivation (9 animals, 2,198 trials).
(G) Proportion of correct trials in the interleaved A/B × A/B and C × C/D task with and without bilateral ALM inactivation (9 animals, 2,728 trials).
(H) ALM inactivation caused greater impairment on the A/B × A/B task than on the control tasks and OFC inactivation on the same A/B × A/B task. The effect of inactivation is expressed as the difference in log probability of a correct choice with and without inactivation (STAR Methods; Equation 4). Distributions were estimated by a bootstrap procedure. The largest effects were evident when ALM was inactivated and the sample was informative (t test, d.f. determined by size of dataset). Table S1 provides additional identifying information about the experiments displayed in this panel.
(I–K) Population sample odor selectivity of Pir (I), OFC (J), and ALM (K) with and without ALM inactivation. The ordinate is the difference in firing rate to the preferred and the nonpreferred odor across neurons (see STAR Methods). Note the different ordinate scales. Shading is SEM. n = 188, 104, and 470 neurons for Pir, OFC, and ALM, respectively. The same neurons are used in the analyses below.
(L–N) Population test odor selectivity of Pir (L), OFC (M), and ALM (N) with and without ALM inactivation.
(O–R) Performance of the SVM decoder to classify match/non-match using Pir response (O), match/non-match using OFC responses (P), sample odor identity using ALM responses (Q), and test odor identity using ALM responses (R) with and without ALM inactivation.
See also Figures S3 and S4.

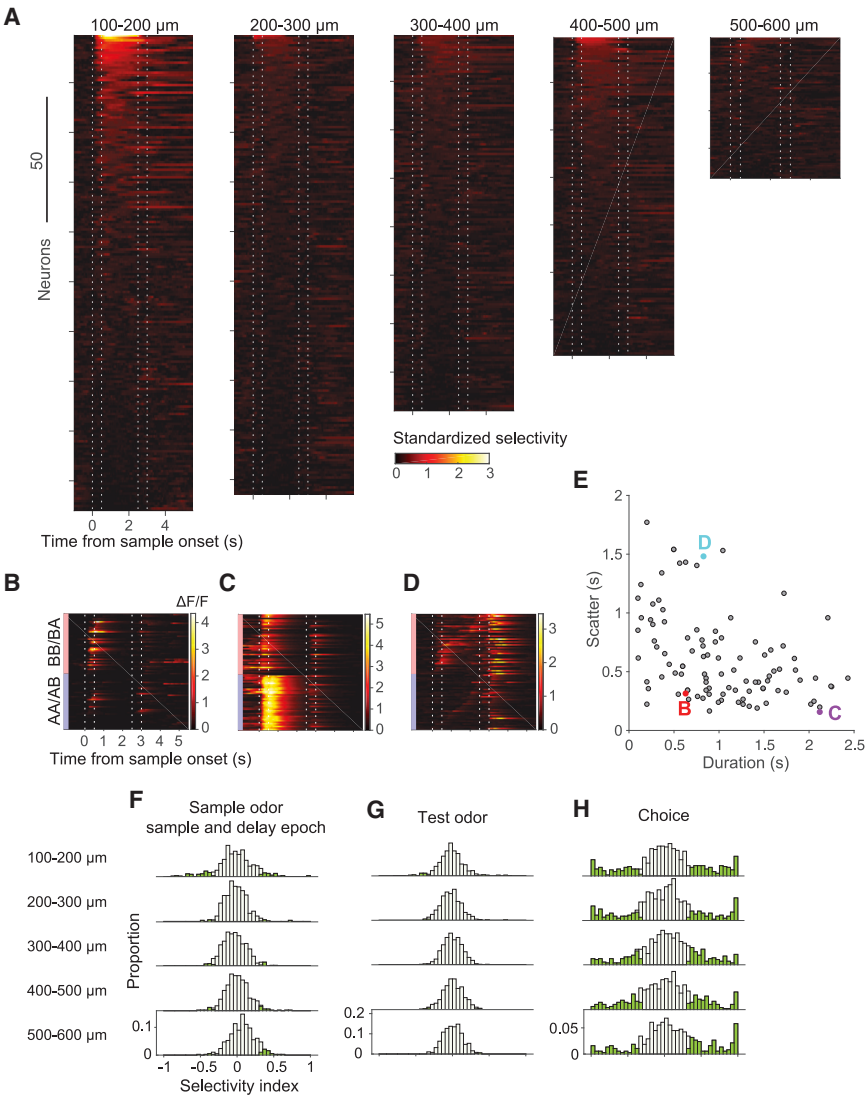

Figure 5. Two-Photon Calcium Imaging of ALM

(A) Sample odor selectivity of ALM neurons in an example animal. The columns show all task-related neurons at five cortical depths. Each row is a neuron, sorted by the mean standardized selectivity during the sample and delay epochs. The standardized odor selectivity is the difference between the ΔF/F to odor A and B, normalized by the standard deviation (STAR Methods).

(B–D) ΔF/F traces of example sample-selective neurons. Each row is a trial, grouped by sample odor identity and sorted by peak response time to sample odor. The three examples illustrate variation in response duration (B, D, brief; C, long) and consistency in timing (B, C, consistent across trials; D, scattered across trials).

(E) Time course of sample-selective responses. Each point shows the duration of the calcium transient (abscissa) against the scatter of the peak response (ordinate). Response duration is the median duration of each neuron's response. Scatter is defined as the interquartile range of the peak response times across trials (STAR Methods). The three colored points correspond to the neurons shown in panels (B)–(D).

(F–H) Selectivity indices of neurons across five cortical depths for sample odor, test odor, and choice, respectively. n = 614, 514, 608, 512, and 379 neurons for the five cortical depths, respectively, from superficial to deep layers. Shading denotes statistical significance ($p < 0.01$, Mann-Whitney U test, two-tailed). (F) Selecivity for sample odor during the sample and delay epochs; (G) Selectivity for test odor; (H) Selectivity for choice. See also Figures S5 and S6.

simultaneously recorded neurons (Figures S6A and S6B). These characterizations also hold for the estimated spike rates, achieved through deconvolution of the raw Ca signals (Figures S6C–S6F), despite a revision of the estimates of signal duration. From the imaging data, we educe that a subpopulation of ALM neurons represents the identity of the sample odor through the sample and delay period.

These sample selective neurons appear to constitute a distinct cell type. They were concentrated in the superficial layer 2 (L2; 100–200 µm from the pial surface), where 18% responded selectively to one or the other sample odor (Figures 5A and 5F). They were only rarely encountered between 200 µm and 600 µm below the pial surface ($p < 0.05$, Kruskal-Wallis test with Tukey-Kramer multiple comparison). This layer specificity is not explained by an inability to measure calcium signals in deeper layers, as comparable fractions of choice-selective cells were found at all depths examined ($p = 0.49$, Kruskal-Wallis test; Figures 5H and S5A). Indeed, a decoder, trained to classify the sample odor, performs much better using L2 neurons than using neurons from the other layers we sampled (Figure 6A). In contrast, the capacity to decode choice is similar across all layers (Figure 6B). Most of these sample-selective neurons did not respond to the test odor (87 of 108, 81%; Figure 6C), and only a few were choice-selective or active during licking (Figure 6D). Their time course distinguishes them from the rare odor-selective responses encountered in our electrical recordings, and the absence of lick and choice activity distinguishes them from the dominant cell type in ALM.

The L2 neurons might resolve the perplexing inactivation result. ALM contains a representation of the sample odor throughout the delay period. It is therefore possible that the match/non-match decision is made within ALM, based on two external inputs that convey (1) the identity of the sample odor during the sample or delay period and (2) the identity of the test odor after the delay. Indeed, the activity of the L2 neurons during the sample and delay is informative about whether the mouse will ultimately succeed in the trial by making the correct choice. This assertion is supported by the logistic regression analysis summarized in Figure 6E. The concentration of errors in the upper left corner of the graph demonstrates that errors were more likely to occur when neurons selective for sample odor A, say, responded relatively weakly to sample odor A or

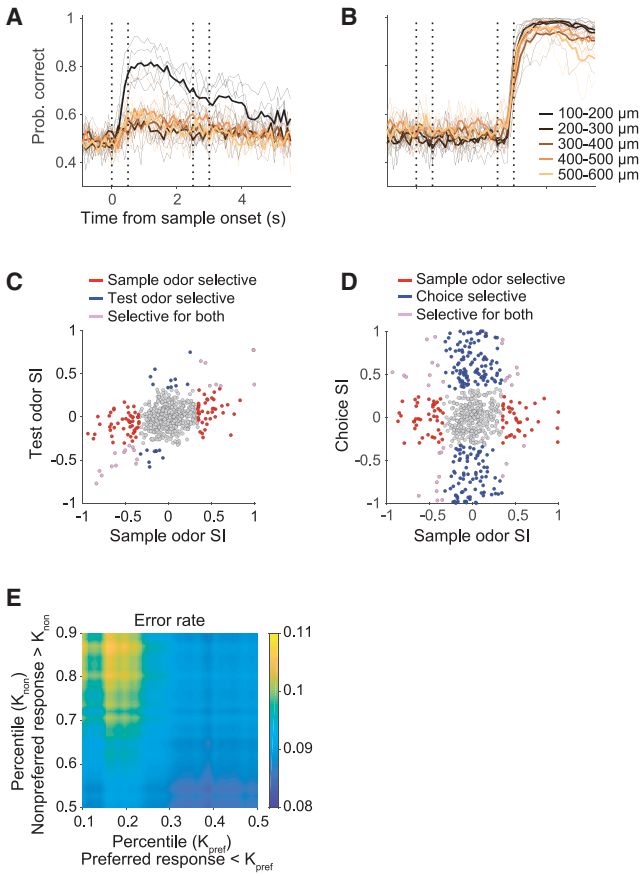

Figure 6. Sample-Selective Cells in ALM L2

(A and B) Performance of a support vector machine trained to classify the sample odor identity (A) or choice (B) using a subset of simultaneously imaged neurons from each of the five cortical depths. Each thin line represents the probability of correct classifications of the held out trials using data from one session (see STAR Methods). Thick lines represent the mean of all sessions from a depth.

(C) Selectivity indices of ALM L2 neurons for sample and test odors. Red, blue, or purple denotes neurons selective for the sample odor, test odor, or both, respectively. Non-selective neurons are shown in gray.

(D) Selectivity indices of ALM L2 neurons for sample odor and choice. Red, blue, or purple denotes neurons selective for the sample odor, choice, or both, respectively. Non-selective neurons are shown in gray.

(E) Trial-by-trial association between neural response and choice accuracy. The heatmap shows the fraction of errors when L2 neurons responded weakly to the preferred odor (*pref*) or strongly to the nonpreferred odor (*non*). The criteria for weak and strong are varied parametrically up to the median for pref (abscissa) and above the median for non (ordinate). The graph shows an increased probability of an error when the *pref* response is in the lower 20th percentile or the *non* response is in the upper 30th percentile. Logistic regression demonstrates that the effect is reliable across the dataset ($p < 0.001$; STAR Methods).

respond more strongly than usual to sample odor B. Such single trial correlations between neural response and choice, termed choice probability, have been exploited in perceptual decisions to support the proposal that a neuron's response contributes to the decision process, either directly or via correlation with other neurons with similar response selectivity (Britten et al., 1996). In the present case, the contribution is not to the match/non-match or lick-left lick-right decision. Rather, it is a decision about which mapping—from test odors to lick responses—to apply on each trial (Figure 1D).

Potential Mechanisms

We do not yet know the mechanism by which a representation of the sample odor in ALM affects the match/non-match decision, but it appears that some process requiring ALM in the sample and delay period must establish a state at the time of test that can implement the correct mapping between externally derived signals about the identity of the test odor and activation of the appropriate lick neurons. Several possible mechanisms could underlie the formation of such a state. One possibility is that the state is defined by the persistent firing of sample-selective neurons, as in standard attractor models of working memory (Figure 7A) (Amit and Brunel, 1997; Goldman-Rakic, 1995). In this model, the requirement of ALM during the sample and delay epochs (Figure 4) demands that ALM itself maintain a persistent representation of the sample odor.

We determined whether ALM and/or earlier stages form persistent representations by controlling whether or not recurrent connections at each stage are modified during network training (Figure 7, curved arrows; see STAR Methods). If the persistent sample representation is maintained in upstream areas, but not ALM, the network model predicts that inactivating ALM would cause little behavioral impairment (Figure 7B, top trace). This is inconsistent with our experimental observation. If the persistent sample representation is only maintained in ALM, the network model predicts that inactivation during the sample and early delay epochs should produce gross impairment (Figure 7B, bottom trace). This too is inconsistent with the data (Figure 4E). The weaker effects of inactivation when it was restricted to portions of the sample and delay epochs suggest that ALM's representation of the sample odor can be partially recovered by activity in upstream areas. Thus, our experimental results argue for models in which multiple brain regions, including both ALM and additional upstream areas, each maintain a persistent representation of the sample odor (Figure 7B, middle trace).

The attractor models posit that a representation of the sample odor is maintained through the persistent firing of neuronal assemblies. However, this memory might be maintained in other ways. We observed that impairment was more profound when ALM was silenced during the entire sample and delay epochs, compared with inhibition only late in the delay (Figure 4E). This suggests that a trace of the sample odor identity can persist in ALM during silencing. This trace may be maintained by facilitation or depression of synapses formed by axons of odor-selective neurons onto lick-left or lick-right neurons (Mongillo et al., 2008) (Figure 7C). It could also be achieved through dendritic gating mechanisms that selectively route test odor information depending on the sample odor identity (Yang et al., 2016) (Figure 7D). In this model, neurons selective for a particular sample odor enhance or suppress dendritic branches of lick-selective neurons where either A or B inputs concentrate. This synaptic modulation could be controlled by sample-selective neurons in ALM layer 2. Indeed, the observation of sample-selective

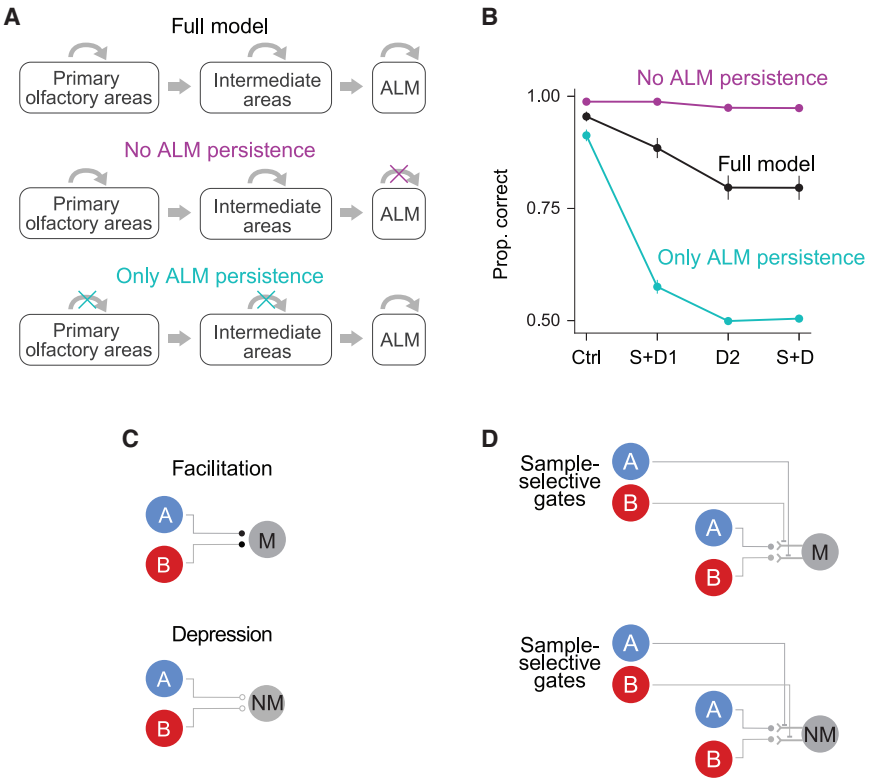

Figure 7. Neural Models of the DMS Task

(A) Three variants of the recurrent neural network model. Each variant comprises three stages, corresponding roughly to primary olfactory sensory areas, intermediate association areas, and ALM. The stages are connected by fixed, random, feed-forward projections (horizontal straight arrows). Top: all recurrent connections within each stage are trained. Middle: recurrent connections within Pir and OFC, but not ALM, are trained. Bottom: only recurrent connections within ALM are trained. Note that only those recurrent connections that are trained are able to support persistent activity that is not inherited from an upstream area.

(B) Simulated effects of inactivating ALM during portions of the sample and delay epochs. Same conventions for the behavioral epochs as in Figure 4E.

(C) Synaptic facilitation and depression model. Facilitation leads to enhanced response to repeated presentation of the same stimulus in match-selective neurons, whereas depression results in suppression in non-match-selective neurons. We assume that the soma is inhibited during the sample epoch so that action potentials only occur after the test odor.

(D) Dendritic gating model. Top: circuit supporting match (M). The sample odor suppresses excitatory inputs that convey the identity of the opposite test odor. Bottom: circuit supporting non-match (NM). A similar result could be achieved with presynaptic gating. This circuit also requires a mechanism to suppress spiking before the test epoch and a mechanism to reset the gate.

responses that occur throughout the sample and delay (Figure 5E) could explain why impairment is reduced when inactivation is limited to only a part of the sample and delay periods. If instead of participating in a dynamical computation to sustain the representation of the sample odor, these neurons alter the state of the neurons that will ultimately process the test odor, then full impairment could be ameliorated by an interruption to the inactivation at any time (cf. Kopec et al., 2015).

DISCUSSION

Decisions that engage cognitive capacities entail at least three elemental processes: (1) the maintenance of information over timescales that can extend for seconds prior to a response; (2) the association of the same information with different responses; and (3) the association of the same response with different information. The first process requires planning and working memory (Fuster, 1973; Goldman-Rakic, 1995; Romo et al., 1999), whereas the latter two involve flexible, context-dependent routing of information to appropriate outputs.

Our findings suggest an unexpected role of ALM in processing perceptual information. Neurons in ALM are known to play an essential role in the planning of licking (Svoboda and Li, 2018). In primates, and more recently in rodents, preparatory activity in premotor and parietal circuits has been shown to play a role in perceptual decisions, in which attributes of sensory stimuli instruct the selection of a movement (Hanks et al., 2015; Kopec et al., 2015; Shadlen and Newsome, 2001). In some cases, this preparatory activity evolves gradually as evidence accrues, thereby reflecting the deliberation process leading to a decision (Cisek, 2011; Gold and Shadlen, 2003; Schall and Hanes, 1993; Selen et al., 2012; Spivey et al., 2005; Yang and Shadlen, 2007). These observations might lead one to anticipate a role for premotor area ALM in an olfactory decision that is reported by licking to the left or right. But there is an essential difference.

In the DMS task, a plan to lick left or right cannot be prepared until the test odor arrives. Before this moment, all that can be prepared is the sensory-response mapping between the possible test odors and the appropriate lick direction. Therefore, the impairment in performance arising from the inactivation of ALM does not reflect a disruption of movement preparation, but rather a disruption of the state of the circuit that readies it to process information about the test odor to select the appropriate lick response. This is a high-level decision about how to make a subordinate decision, that is, the type of control process associated with higher prefrontal circuits in primates (Wallis and Miller, 2003; cf. Ebbesen et al., 2018). In our DMS task, this control process appears to require a normally functioning ALM during at least some stretch of the sample plus delay periods. The requirement was only demonstrable when the sample odor established the sensory-response mapping between test odor and lick direction. Inactivation of ALM during the sample and delay did not impair performance on the A × A/B and C × C/D control tasks. These controls require only simple, learned associations between the test odor and a lick response.

The absence of impairment on these control tasks also demonstrates that ALM can receive information from upstream neurons about the identity of the test odor or its associated lick direction. In contrast, upstream areas do not supply ALM with information about match versus non-match. This assertion holds for any upstream area—not just Pir and OFC, but the many areas that connect the main olfactory bulb to ALM. We did not record from the entire network of areas; hence, we cannot be certain that they do not contain critical information that helps the mouse solve the DMS task. However, they do not supply the match/non-match (or lick-left lick-right) solution to ALM or its downstream targets. We cannot rule out the possibility that inactivation of ALM during the sample and delay periods induce a disruption in whatever essential computation such upstream areas might contribute. However, this seems unlikely because task-relevant information in OFC and Pir was unaffected by ALM inactivation (Figures 4I, 4J, 4L, 4M, 4O, and 4P).

It is nonetheless remarkable that match versus non-match can be decoded from small populations of neurons in Pir and OFC. It serves as a reminder that information that can be decoded by the experimenter is not necessarily decoded by the brain. We assume that Pir and OFC supply information, directly or indirectly, to ALM during the test epoch as well as the sample and delay epochs. The absence of impairment associated with bilateral OFC inactivation does not imply otherwise, as it could be explained by the multiplicity of pathways that connect Pir to ALM. Also, although the OFC is thought to play an essential role in reversal learning (Schoenbaum et al., 2002), our mice had already learned the sample-instructed mappings—effectively trial-by-trial reversals—by the time of the inactivation experiments.

Previous studies have demonstrated persistent activity in ALM in association with perceptual decision making. For example, elevated activity was reported in direction-selective lick neurons when mice associated different haptic stimuli (applied to a whisker) with a lick to the left or right (Chen et al., 2017; Guo et al., 2014a; Li et al., 2015; Svoboda and Li, 2018). This persistent activity is involved in planning and driving movement based on a sensory-response association. In our study, on the other hand, persistent activity in sample-selective neurons represents a sensory category rather than a plan of action, because the action cannot be selected until after the delay period upon receipt of the test odor. Hence, unlike in a sensory-instructed delayed response task (Guo et al., 2014a; Li et al., 2015) in which persistent activity represents decision outcome, the activity of our L2 neurons is not the outcome of the match/non-match decision because this decision has yet to be made. This activity could be viewed as representing an intermediate decision about which mapping from test odor to lick response to deploy (Figure 1D). It is possible that our sample-selective neurons overlap those reported previously, but we suspect the overlap is minimal. The persistent activity we observed was concentrated primarily in superficial L2 neurons that did not exhibit lick responses, whereas the persistent activity reported in the somatosensory instructed delay task was found in all cortical layers and in many neurons that responded to error responses in a way that agrees with their role in motor planning.

We reason that the sample-selective neurons play a critical role in the DMS task. First, to perform the task, it is essential that neural activity in ALM is normal during the delay period. Second, the only task-relevant activity in ALM that spans the delay period is a representation of the sample odor by these neurons concentrated in L2. Third, the activity of these sample-selective neurons is predictive of performance on a trial-by-trial basis (Figure 6E). We conclude that these neurons are likely to affect the local circuitry so that the two test odors are associated with the appropriate lick response. More definitive evidence for this interpretation would require selective inactivation of this population of neurons. This is not possible presently, but if these neurons represent a distinct subtype, as we suspect (e.g., Luo et al., 2017), it may be possible in the near future. We also do not know if the sample-selective neurons affect the lick cells directly or influence the local circuit.

We considered several alternative explanations of the impairment produced by ALM inactivation. The control experiments rule out a simple motor impairment or motor bias, allowing us to focus on the delay period activity. Our interpretation is that the representation of the sample odor in ALM establishes a state of the ALM circuit such that it can perform a simple sensory-response association at the time of test. An alternative is that ALM must prepare both potential motor plans during the delay period (Cisek and Kalaska, 2005) in order for one to be selected by an instruction from another brain area, and it is solely the disruption of this preparatory activity that explains the impairment. This account predicts incorrectly that inactivation would also impair performance on the A × A/B and C × C/D tasks. It is also inconsistent with the neural recordings, which failed to reveal simultaneous preparation of the two lick directions during the delay period.

Another alternative explanation is that inactivation before the test period disrupts a general state of readiness in ALM to respond to any signal instructing an action to be performed. We find this unlikely, because to account for our behavioral observations, this state would need to be specific to the DMS task and not control tasks with identical temporal structures and actions. Finally, we note that DMS is a more difficult task than the controls. Therefore, it may be argued that DMS is more susceptible to perturbations. However, it is not the motor action that makes DMS more difficult than the controls, because all the tasks require the identical motor responses. If the only role of ALM is motor preparation, then inactivating ALM should not result in differential impairment of the DMS and control tasks (Figure 4H). More importantly, our perturbation affected ALM before a decision could be prepared, because the decision cannot be made before the test odor arrives. Previous studies of decisions to lick to the left or right examined only the preparatory phase after the decision was instructed (Guo et al., 2014a; Li et al., 2015).

Neurons in superficial layers of the neocortex receive input from controlling structures such as intralaminar and accessory thalamic nuclei (Jones, 1998) as well as long range feedback, e.g., perirhinal cortex to ALM (Zingg et al., 2014). These layers contain the distal dendrites of projection neurons in deeper layers, and activation or suppression of these dendrites can alter the functional properties of the projection neurons (Bittner et al., 2017; Larkum et al., 1999, 2009; Polsky et al., 2004). Thus, it is possible that the selectivity to the sample odor that we observed

in L2 neurons might reflect a general mechanism that alters the state of ALM circuitry to flexibly implement different sensorimotor mappings.

It may seem counterintuitive that a premotor area would play a critical role in a decision about the relationship between two sensory stimuli. However, the ultimate goal of both sensory and motor systems is not to identify stimulus features or categories, but to determine whether to act one way or another. A potential action might require support from different sources of sensory input. Thus, decisions to act—even if only provisionally—rest on a capacity of the motor system to establish functional connections that allow it to integrate different sources of sensory information, and this information infiltrates the responses of neurons in cortical areas that are appropriately designated premotor or associative (Cisek, 2011; Shadlen and Kiani, 2013). These functional connections must be rapidly modifiable as the demands of an organism and its context change. This capacity for dynamic circuit configuration is likely to support a wide range of cognitive phenomena involving flexible routing, decision making, and perceptual inference (Shadlen and Shohamy, 2016).

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- LEAD CONTACT AND MATERIALS AVAILABILITY
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Behavior training
  - Electrophysiology
  - Two-photon calcium imaging
  - Photostimulation
  - Simultaneous photostimulation and recording
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Electrophysiology data analysis
  - Imaging data analysis
  - Behavior/inhibition data analysis
  - Attractor network models
- DATA AND CODE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j.neuron.2020.01.034.

ACKNOWLEDGMENTS

We thank T. Harris for the acute H3 probes; T. Tabachnik, D. Peterka, G. Johnson, J. Wang, J. Goldbas, A. Baker, and M. Chin for technical assistance; K. Svoboda, Z. Guo, L. Abbott, A. Losonczy, R. Bruno, C. Rodgers, C. Lacefield, K. Shin, and B. Hu for advice on experiments and analysis; D. Wolpert, D. Shohamy, J.D. Murray, and members of the Axel and Shadlen labs for helpful discussions; C. Eccard for assistance in the preparation of the manuscript; and M. Myles for figure artwork. Z.W. is a Simons Junior Fellow. A.L.-K. is supported by NSF NeuroNex Award DBI-1707398, the Burroughs Wellcome Foundation, and the Gatsby Charitable Foundation. R.A. and M.N.S. are supported by the Howard Hughes Medical Institute (HHMI). R.A. is also supported by the Simons Foundation.

AUTHOR CONTRIBUTIONS

Z.W., R.A., and M.N.S. conceived and designed the study. Z.W. performed all the experiments with assistance from P.S. and A.T. A.L.-K. implemented the attractor network models. Z.W., A.L.-K., R.A., and M.N.S. interpreted the results and wrote the paper.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: September 3, 2019
Revised: December 12, 2019
Accepted: January 24, 2020
Published: February 26, 2020

REFERENCES

Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., et al. (2016). TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. arXiv, arXiv:1603.04467 https://arxiv.org/abs/1603.04467.

Amit, D.J., and Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb. Cortex *7*, 237–252.

Bechara, A., Damasio, H., and Damasio, A.R. (2000). Emotion, decision making and the orbitofrontal cortex. Cereb. Cortex *10*, 295–307.

Bittner, K.C., Milstein, A.D., Grienberger, C., Romani, S., and Magee, J.C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. Science *357*, 1033–1036.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis. Neurosci. *13*, 87–100.

Carandini, M., and Churchland, A.K. (2013). Probing perceptual decisions in rodents. Nat. Neurosci. *16*, 824–831.

Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature *499*, 295–300.

Chen, T.W., Li, N., Daie, K., and Svoboda, K. (2017). A Map of Anticipatory Activity in Mouse Motor Cortex. Neuron *94*, 866–879.

Cisek, P. (2011). Cortical mechanisms of action selection: the affordance competition hypothesis. Model. Nat. Action Sel. 208–238.

Cisek, P., and Kalaska, J.F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. Neuron *45*, 801–814.

Ebbesen, C.L., Insanally, M.N., Kopec, C.D., Murakami, M., Saiki, A., and Erlich, J.C. (2018). More than just a "motor": Recent surprises from the frontal cortex. J. Neurosci. *38*, 9402–9413.

Evarts, E.V., and Tanji, J. (1976). Reflex and intended responses in motor cortex pyramidal tract neurons of monkey. J. Neurophysiol. *39*, 1069–1080.

Fan, R.-E., Chang, K.-W., Hsieh, C.-J., Wang, X.-R., and Lin, C.-J. (2008). LIBLINEAR: A library for large linear classification. J. Mach. Learn. Res. *9*, 1871–1874.

Freedman, D.J., and Assad, J.A. (2006). Experience-dependent representation of visual categories in parietal cortex. Nature *443*, 85–88.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J. Neurophysiol. *61*, 331–349.

Fuster, J.M. (1973). Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. J. Neurophysiol. *36*, 61–78.

Fuster, J.M., and Alexander, G.E. (1971). Neuron activity related to short-term memory. Science *173*, 652–654.

Giessel, A.J., and Datta, S.R. (2014). Olfactory maps, circuits and computations. Curr. Opin. Neurobiol. *24*, 120–132.

Gold, J.I., and Shadlen, M.N. (2003). The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. J. Neurosci. *23*, 632–651.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.

Goldman-Rakic, P.S. (1995). Cellular basis of working memory. Neuron *14*, 477–485.

Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014a). Flow of cortical activity underlying a tactile decision in mice. Neuron *81*, 179–194.

Guo, Z.V., Hires, S.A., Li, N., O'Connor, D.H., Komiyama, T., Ophir, E., Huber, D., Bonardi, C., Morandell, K., Gutnisky, D., et al. (2014b). Procedures for behavioral experiments in head-fixed mice. PLoS ONE *9*, e88678.

Guo, Z.V., Inagaki, H.K., Daie, K., Druckmann, S., Gerfen, C.R., and Svoboda, K. (2017). Maintenance of persistent activity in a frontal thalamocortical loop. Nature *545*, 181–186.

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220–223.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62–68.

Jones, E.G. (1998). Viewpoint : the Core and Matrix of Thalamic. Neuroscience *85*, 331–345.

Kadohisa, M., Petrov, P., Stokes, M., Sigala, N., Buckley, M., Gaffan, D., Kusunoki, M., and Duncan, J. (2013). Dynamic construction of a coherent attentional state in a prefrontal cell population. Neuron *80*, 235–246.

Kaifosh, P., Zaremba, J.D., Danielson, N.B., and Losonczy, A. (2014). SIMA: Python software for analysis of dynamic fluorescence imaging data. Front. Neuroinform. *8*, 80.

Kingma, D.P., and Ba, J.L. (2014). Adam: A Method for Stochastic Optimization. arXiv, arXiv:1412.6980 https://arxiv.org/abs/1412.6980.

Kopec, C.D., Erlich, J.C., Brunton, B.W., Deisseroth, K., and Brody, C.D. (2015). Cortical and Subcortical Contributions to Short-Term Memory for Orienting Movements. Neuron *88*, 367–377.

Larkum, M.E., Zhu, J.J., and Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. Nature *398*, 338–341.

Larkum, M.E., Nevian, T., Sandler, M., Polsky, A., and Schiller, J. (2009). Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. Science *325*, 756–760.

Li, N., Chen, T.-W., Guo, Z.V., Gerfen, C.R., and Svoboda, K. (2015). A motor cortex circuit for motor planning and movement. Nature *519*, 51–56.

Liu, D., Gu, X., Zhu, J., Zhang, X., Han, Z., Yan, W., Cheng, Q., Hao, J., Fan, H., Hou, R., et al. (2014). Medial prefrontal activity during delay period contributes to learning of a working memory task. Science *346*, 458–463.

Luo, H., Hasegawa, K., Liu, M., and Song, W.J. (2017). Comparison of the upper marginal neurons of cortical layer 2 with layer 2/3 pyramidal neurons in mouse temporal cortex. Front. Neuroanat. *11*, 115.

Luo, L., Callaway, E.M., and Svoboda, K. (2018). Genetic Dissection of Neural Circuits: A Decade of Progress. Neuron *98*, 256–281.

Madisen, L., Garner, A.R., Carandini, M., Zeng, H., Madisen, L., Garner, A.R., Shimaoka, D., Chuong, A.S., Klapoetke, N.C., Li, L., et al. (2015). Transgenic Mice for Intersectional Targeting of Neural Sensors and Effectors with High Specificity and Performance. Neuron *85*, 942–958.

Mainen, Z.F., and Kepecs, A. (2009). Neural representation of behavioral outcomes in the orbitofrontal cortex. Curr. Opin. Neurobiol. *19*, 84–91.

Mongillo, G., Barak, O., and Tsodyks, M. (2008). Synaptic Theory of Working Memory. Science *319*, 1543–1546.

Pachitariu, M., Steinmetz, N.A., Kadir, S.N., Carandini, M., and Harris, K.D. (2016). Fast and accurate spike sorting of high-channel count probes with KiloSort. NIPS Proc. 4448–4456.

Padoa-Schioppa, C., and Assad, J.A. (2006). Neurons in the orbitofrontal cortex encode economic value. Nature *441*, 223–226.

Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., et al. (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. Neuron *89*, 285–299.

Polsky, A., Mel, B.W., and Schiller, J. (2004). Computational subunits in thin dendrites of pyramidal cells. Nat. Neurosci. *7*, 621–627.

Ramus, S.J., and Eichenbaum, H. (2000). Neural correlates of olfactory recognition memory in the rat orbitofrontal cortex. J. Neurosci. *20*, 8199–8208.

Romo, R., Brody, C.D., Hernández, A., and Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. Nature *399*, 470–473.

Root, C.M., Denny, C.A., Hen, R., and Axel, R. (2014). The participation of cortical amygdala in innate, odour-driven behaviour. Nature *515*, 269–273.

Rossant, C., Kadir, S.N., Goodman, D.F.M., Schulman, J., Hunter, M.L.D., Saleem, A.B., Grosmark, A., Belluscio, M., Denfield, G.H., Ecker, A.S., et al. (2016). Spike sorting for large, dense electrode arrays. Nat. Neurosci. *19*, 634–641.

Schall, J.D., and Hanes, D.P. (1993). Neural basis of saccade target selection in frontal eye field during visual search. Nature *366*, 467–469.

Schoenbaum, G., Nugent, S.L., Saddoris, M.P., and Setlow, B. (2002). Orbitofrontal lesions in rats impair reversal but not acquisition of go, no-go odor discriminations. Neuroreport *13*, 885–890.

Selen, L.P.J., Shadlen, M.N., and Wolpert, D.M. (2012). Deliberation in the motor system: reflex gains track evolving evidence leading to a decision. J. Neurosci. *32*, 2276–2286.

Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. Neuron *80*, 791–806.

Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J. Neurophysiol. *86*, 1916–1936.

Shadlen, M.N., and Shohamy, D. (2016). Decision Making and Sequential Sampling from Memory. Neuron *90*, 927–939.

Snyder, L.H., Batista, A.P., and Andersen, R.A. (1997). Coding of intention in the posterior parietal cortex. Nature *386*, 167–170.

Sosulski, D.L., Bloom, M.L., Cutforth, T., Axel, R., and Datta, S.R. (2011). Distinct representations of olfactory information in different cortical centres. Nature *472*, 213–216.

Spivey, M.J., Grosjean, M., and Knoblich, G. (2005). Continuous attraction toward phonological competitors. Proc. Natl. Acad. Sci. USA *102*, 10393–10398.

Svoboda, K., and Li, N. (2018). Neural mechanisms of movement planning: motor cortex and beyond. Curr. Opin. Neurobiol. *49*, 33–41.

Wallis, J.D., and Miller, E.K. (2003). From rule to response: neuronal processes in the premotor and prefrontal cortex. J. Neurophysiol. *90*, 1790–1806.

Yang, T., and Shadlen, M.N. (2007). Probabilistic reasoning by neurons. Nature *447*, 1075–1080.

Yang, G.R., Murray, J.D., and Wang, X.-J. (2016). A dendritic disinhibitory circuit mechanism for pathway-specific gating. Nat. Commun. *7*, 12815.

Zhao, S., Ting, J.T., Atallah, H.E., Qiu, L., Tan, J., Gloss, B., Augustine, G.J., Deisseroth, K., Luo, M., Graybiel, A.M., and Feng, G. (2011). Cell type–specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. Nat. Methods *8*, 745–752.

Zingg, B., Hintiryan, H., Gou, L., Song, M.Y., Bay, M., Bienkowski, M.S., Foster, N.N., Yamashita, S., Bowman, I., Toga, A.W., and Dong, H.W. (2014). Neural networks of the mouse neocortex. Cell *156*, 1096–1111.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT OR RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Mouse: wild type C57BL/6J	Jackson Laboratory	000664
Mouse: VGAT-ChR2-EYFP	Jackson Laboratory	014548
Mouse: Ai93(TITL-GCaMP6f)-D;CaMK2a-tTA	Jackson Laboratory	024108
Mouse: Emx1-IRES-cre+/+	Jackson Laboratory	005628
Electrophysiology		
32 channel, acute Buzsaki32 probe	Neuronexus	Buzsaki32-A32
32 channel, chronic Buzsaki32 probe	Neuronexus	Buzsaki32-H32_21mm
32 channel, acute Poly3 probe	Neuronexus	A1x32-Poly3-5mm-25 s-177-A32
64 channel, acute 1x64 H3 probe	Tim Harris	https://www.cambridgeneurotech.com/neural-probes
Software		
MATLAB	Mathworks	https://www.mathworks.com
LabVIEW	National Instrument	https://www.ni.com/en-us/shop/labview/labview-details.html
SIMA	Kaifosh et al., 2014	https://www.losonczylab.org/software
CNMF	Pnevmatikakis et al., 2016	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4881387/
Python	Python	https://www.python.org
Tensorflow	Tensorflow	https://www.tensorflow.org

LEAD CONTACT AND MATERIALS AVAILABILITY

This study did not generate new unique reagents. Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Michael Shadlen (shadlen@columbia.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Animals were maintained on 12 h: 12 h light/dark cycle with food and water available *ad libitum*. Mice were water-restricted during the training and testing phases. Experimental sessions were 1-2 h, during which mice received 0.5-1.5 mL of water. Animals received supplemental water as necessary to maintain their body weights. Aseptic surgeries were carried out under ketamine (100 mg/kg)/xylazine (10 mg/kg) or 1%–3% isoflurane anesthesia. Buprenorphine (0.05-0.1 mg/kg) and carprofen (5 mg/kg) were administered for postoperative analgesia. Animal care and experiments were carried out in accordance with the NIH guidelines and approved by the Columbia University Institutional Animal Care and Use Committee (IACUC).

This study is based on data from 41 mice (both males and females, 2-8 months old). Five C57BL/6J and 6 VGAT-ChR2-EYFP (Jackson Laboratory, JAX 014548) mice were used for electrophysiology recording. Two untrained VGAT-ChR2-EYFP mice were used to characterize the inhibition at different laser powers. Twenty-nine VGAT-ChR2-EYFP mice were used for inhibition experiments, 6 of which were also used for simultaneous recording during inhibition. Two-photon imaging data were collected from 5 Emx1-cre+/−;TITL-GCaMP6f+/−;CaMK2a-tTA+ mice. These mice were created by crossing Ai93(TITL-GCaMP6f)-D;CaMK2a-tTA (JAX 024108) to Emx1-IRES-cre+/+ (JAX 005628) line.

METHOD DETAILS

Behavior training

Before training, mice were implanted with a custom-made titanium head post (Guo et al., 2014b). The scalp and periosteum over the dorsal surface of the skull were removed, and a head post was placed on the skull, aligned with the lambda suture and cemented in place with C&B metabond (Parkwell, S380). After at least 2-3 days of recovery, animals were water-restricted and accustomed to

head-fixation following procedures described in Guo et al., 2014 (Guo et al., 2014b) and then trained on a custom-assembled apparatus.

Odorants were delivered with a custom-made olfactometer and custom-written LabVIEW programs (National Instruments). (+)-α-pinene (odor A, Sigma-Aldrich, 268070), cis-3-hexen-1-ol (odor B, Sigma-Aldrich, W256307), (R)-(+)-limonene (odor C, Sigma-Aldrich, 183164), and methyl butyrate (odor D, Sigma-Aldrich, 246093) were chosen for their lack of innate valence to mice (Root et al., 2014) and their low adhesion to the surface of the olfactometer. The odorants were diluted 100-fold in mineral oil (Fisher Scientific, O121-1) and then loaded on syringe filters (GE healthcare, 6888-2527 or 6823-1327). The air flow was maintained at 1.0 L/min. We confirmed the rapid kinetics (Figure S1G) of these odorants with photo-ionization detector (Aurora Scientific, 200B). Over the course of a behavioral session, the odorants on the syringe filters gradually deplete. Thus animals were unlikely to rely on the absolute concentration of the odorants which changed constantly, but rather the identity of the odors, as evidenced by the stable sensory responses in Pir over a session (Figure 2A).

During training, mice were presented with a sample odor (1.0 s duration) and a test odor (1.0 s) separated by a delay epoch (1.0 s). After hearing an auditory ''go'' cue (0.1 s, 5 KHz pure tone), they were free to report their decision by licking to one of the two syringe ports positioned in front of their mouth, and collect a water reward at the same port if they were correct. Many animals started licking in the test epoch, which was permitted, but only the licks during the 2 s ''response window'' following the ''go'' cue were considered their choice. To prevent mice from ''probing'' for the correct port by rapidly switching between the two ports, we required mice to commit to their choices. A choice is scored as correct only when the first two licks were on the correct port. If they first licked on the incorrect port, that trial was scored as an error. If they licked once on the correct port and then on the incorrect port, that trial was also scored as an ''error.'' If they did not lick during the ''response window'' or only licked once on the correct port, that trial was scored as ''no choice.''

The two odorants give rise to four unique pairs of sample and test odors (AA, AB, BB, and BA), or trial types, and were randomly presented in each session. The match trials (AA, BB) were rewarded on the left port and non-match trials (AB, BA) on the right. Mice were punished by a brief timeout (3-8 s) when they made an error. ''No choice'' trials were rare and typically occurred in the very early training stage or at the end of a session when the mice were sated. Animals completed this training stage when they achieved a criterion of at least 80% correct for each trial type in a single session. They then underwent additional training to suppress ''premature'' licking before the test epoch. Such early licks were punished by an immediate 0.1-0.2 s siren (RadioShack, 273-079), followed by a 0.5-1.0 s pause in that trial, and a longer inter-trial interval. We required the proportion of trials with premature licking to be less than ~20%. After achieving these milestones (median 25 days; IQR 21-31 days) the sample and test durations were reduced from 1.0 s to 0.5 s, and the delay was increased from 1.0 s to 1.5-4.0 s depending on the experiment (1.5 s in extracellular recording, 2.0 s in imaging experiments, 1.5, 2.5 or 4.0 s in optogenetic inactivation). Mice performed at 90% correct (interquartile range 88%–92%) when they entered the testing phase.

Electrophysiology

Extracellular recordings were made acutely or chronically in head-fixed animals using 32- or 64-channel silicon probes (Buzsaki32 and Poly3, NeuroNexus; 1 × 64 acute H3 probe, HHMI). The probes were targeted stereotaxically to Pir, OFC and ALM using Bregma coordinates as follows: Pir: AP 1.9-2.4 mm, ML 1.5-2.1 mm, DV 2.3-2.9 mm; OFC: AP 2.2-2.8 mm, ML 0.7-1.2 mm, DV 1.2-1.8 mm; ALM: AP 2.5 mm, ML 1.5 mm, DV 0-1.1 mm. Four animals were used for Pir recordings, six for OFC recordings, and five for ALM recordings.

In acute recordings, a small craniotomy (0.3-0.8 mm in diameter) was made over the targeted area before the recording session. Recording depth from the pial surface was inferred from micromanipulator reading. After each recording session, the brain surface was covered with silicone gel (3-4680, Dow Corning) and Kwik-Sil (World Precision Instruments). The recording sites were confirmed by painting the recording probes with Vybrant DiO Cell-Labeling Solution (Invitrogen). Representative histology sections are shown in Figure S7A-C. In chronic recordings, the probe (Buzsaki32-H32_21mm, NeuroNexus) was attached to a custom-made microdrive which allows for advancement of the shanks. The microdrive was then implanted and cemented with C&B metabond and dental acrylic (Lang Dental Jet Repair Acrylic, 1223CLR). We advanced the shanks by 50 μm per day at the end of each recording session.

Two-photon calcium imaging

Calcium imaging was performed on Emx1-cre+/−;TITL-GCaMP6f+/−;CaMK2a-tTA+ mice. A square craniotomy (2mm side) was made above left or right ALM, along the superior sagittal sinus and the inferior cerebral vein. The imaging window was constructed from three stacked layers of custom-cut coverglass (CS-3S, Warner Instruments) and cemented with C&B metabond. Animals were allowed 1-2 weeks of recovery before the imaging sessions began. Images were acquired with a Bruker Ultima two-photon microscope under resonant galvo scanning mode. The light source was a femtosecond pulsed laser (Chameleon Vision II, Coherent). The objective was a 16X water immersion lens (Nikon, 0.8 NA, 3mm working distance). GCaMP6f was excited at 920nm and images (512 × 512 pixels, ~820 μm x 820 μm field of view) were acquired at ~30 Hz.

Photostimulation

Animals were prepared with a clear skullcap to achieve optical access to ALM (Guo et al., 2014a). Briefly, after removing the scalp and periosteum over the dorsal surface of the skull, a layer of cyanoacrylate adhesive (Krazy glue, Elmer's Products Inc.) was directly

applied to the intact skull. The entire skull was then covered with a thin layer of the clear dental acrylic (Lang Dental) with a head post cemented over the lambda suture. Before photostimulation sessions, the dental acrylic was polished (0321B, Shofu Dental Corporation) and covered with a thin layer of clear nail polish to reduce glare (Part No. 72180, Electron Microscopy Sciences). Light from a 473nm laser (MLL-FN-473-50mW, Ultralasers, Inc.) was directed to an optic fiber and split into two paths (FCMH2-FCL, Thorlabs, Ø200 μm Core, 0.39 NA). The two optic fibers were positioned over ALM on each hemisphere. The light transmission through the skullcap is ~50% in average power, as measured directly with a light meter (PM100D, Thorlabs) with freshly dissected skullcap, consistent with previous measurements (Guo et al., 2014a).

We used 40 Hz photostimulation with a sinusoidal temporal profile (3 mW average power as light reaches the skull, ~1.5 mW on the brain surface). The photoinhibition inactivated a cortical area of ~1mm radius, as the population firing rate drops to ~50% at 1mm away from the center of the laser beam (Figure S3A). To reduce rebound excitation after laser offset, we included a 250 ms linear power ramp-down at the end of the photostimulation (Figure 4A) unless otherwise indicated. In the interleaved A/B × A/B & C × C/D experiment, the delay epoch was extended to 4 s while sample and test epochs were kept at 0.5 s each. Here we used a 500 ms ramp-down at the end of the 4 s photostimulation, which terminates 500 ms before the test odor onset to allow more recovery time. To prevent the mice from distinguishing photostimulation trials from control trials, a masking flash (40 Hz sinusoidal profile) was delivered with 470 nm LED (Luxeon Star) and LED driver (SLA-1200-2, Mightex Systems) in front of the animals' eyes on all trials. The masking flash was not phased locked to the photostimulation, began at sample onset and lasted until the end of test, covering the entire stimulus and delay epochs in which photostimulation could occur.

For the experiments in which we varied the duration of inactivation (Figure 4E), photostimulation was either limited to (*i*) the 0.5 s sample epoch and the first 1.5 s of the 2.5 s delay epoch, or (*ii*) the last 1.0 s of the delay epoch, or (*iii*) the entire sample plus delay epoch. These inactivation trials were randomly interleaved to constitute 25% of all trials. For each animal, multiple (2-3) behavioral sessions were pooled to collect at least 10 trials for each of the four trial types and the three inactivation durations.

Simultaneous photostimulation and recording

We calibrated the laser power for ALM and OFC inactivation by recording from these two areas during photoinhibition in awake animals. For ALM inactivation, we positioned an acute 1x64 H3 probe at various distances from the optic fiber over ALM (0-2.0 mm in 0.5 mm increments). A range of laser powers (0.5 mW, 1.5 mW, 5.0 mW, and 10.0 mW) as well as controls were examined at each location. We chose 1.5 mW power on the brain surface as it inactivates a cortical area of ~1mm radius and produces minimal rebound activity after laser offset (Figure S3). For OFC inactivation, an optic fiber was targeted to Bregma AP 2.5 mm, ML 1.0 mm, DV 1.4 mm. The recording probe was then positioned at AP 2.5 mm, ML 1.5 mm, DV 1.0-2.3 mm, where it is close to the border of OFC with agranular insular cortex. We recorded the neural responses similarly at a range of laser powers and chose 1.0 mW to silence OFC while minimizing the impact on neighboring brain areas.

For simultaneous recording and inactivation when animals performed the DMS task, we used the clear skullcap preparation for photoinhibition and made a small craniotomy lateral to the optic fiber to insert the probe. The probe was advanced at 60-70° angle from the horizontal plane at Bregma AP 2.3-2.5 mm, ML 2.0-2.5 mm to record from ALM, OFC, and Pir at different depth.

QUANTIFICATION AND STATISTICAL ANALYSIS

Electrophysiology data analysis

The 32- and 64-channel recording data were digitized at 40 KHz and acquired with OmniPlex D system (Plexon Inc.) The voltage signals were high-pass filtered (200 Hz, Bessel) and sorted automatically with Kilosort (Pachitariu et al., 2016). The clusters were then manually curated with Phy GUI (Rossant et al., 2016) to merge spikes from the same units and to remove noises and units that were not well isolated. Recording depths were inferred from micromanipulator readings in acute recordings or microdrive turns in chronic recordings.

We determined sample odor selectivity for each neuron by comparing the spike counts during the sample epoch (0.1-0.6 s from sample onset) or late delay (1.5-2.0 s from sample onset) between sample odor A group (AA, AB) and sample odor B group (BA, BB). The 0.1 s time offset from the sample epoch accounts for the olfactometer valve time. A neuron was considered odor selective if its responses to sample odor A and B were significantly different by a two-tailed Mann-Whitney U test ($p < 0.01$, not corrected for multiple comparisons). The selectivity index (SI) was computed as follows for each unit: First, the trial-by-trial spike counts from the responses to odor A and B were used to construct a Receiver Operating Characteristic (ROC). The area under the ROC (AuROC) is the probability that a randomly sampled response associated with odor A is larger than a randomly sampled response associated with odor B. The selectivity index is

$$SI = 2\left(AuROC - \frac{1}{2}\right) \quad (1)$$

such that ± 1 represents perfect discriminability (i.e., no overlap of the response counts) and 0 indicates chance-level discriminability (i.e., complete overlap). Positive SI connotes a preference for sample odor A. We determined selectivity for the test odor and choice in a similar fashion by comparing spike counts during the test epoch (2.1-2.6 s from sample onset), with appropriate trial grouping. Positive SI for choice connotes a preference for licking to the left spout.

The trial-type selectivity in Figure S2C, F, I was computed as follows. Neurons from each of the three areas were first selected by their test odor selectivity. Only neurons deemed test-odor selective as determined by a Mann-Whitney U test ($p < 0.01$, not corrected for multiple comparisons) were considered further. These test-odor selective neurons were then treated separately based on their preference for test-odor A or B. If a neuron preferred test odor A, its selectivity index for trial type was computed based on its responses to AA versus BA trials, in a similar fashion described above. Positive SI connotes a preference for AA. If a neuron preferred test odor B, its trial type selectivity was computed based on its response to AB versus BB, and positive SI connotes a preference for AB.

Graphs depicting population odor selectivity of a brain area show the difference in firing rates associated with the preferred and nonpreferred odors (Figure 3A, 4I-N). The preferred odor of each neuron was designated using a subset of the trials (n = 5 per odor, chosen randomly), in the epoch under consideration (e.g., sample), based on the sign of the difference in the means, irrespective of statistical significance. These trials were then excluded and the odor selectivity was computed as the spike rate to preferred odor minus the nonpreferred odor on the remaining data in a sliding bin of 100 ms. All neurons from each of the three areas are included in Figure 3A (left graph). In Figure 4I-N, all neurons from each of the three areas are included.

For the decoding analysis (Figure 3B-D, 4O-R and 6A, B), we trained a support vector machine (SVM) (Fan et al., 2008) with neural responses recorded simultaneously from Pir, OFC, or ALM to classify sample odor identity, trial type, or match/nonmatch. 48-120 neurons were simultaneously recorded from Pir, 16-82 neurons from OFC and 14-82 neurons from ALM. In Figure 3B, we used the spike counts in a sliding bin of 500 ms from 40 randomly selected neurons from each of the three areas at 100 ms steps. Sessions with insufficient number of simultaneously recorded neurons were excluded. In Figure 3C, we used a 500 ms sliding bin with neural responses from all Pir neurons recorded in a session. In Figure 3D, we used the firing rates in the 500 ms time window before animal's first lick. The decoding capability of each area was estimated by using varying numbers of randomly selected neurons that are recorded simultaneously in a session. As we included more neurons, sessions with insufficient number of neurons dropped out of the analysis. The classifier was trained on randomly selected 90% of the trials in each session and then tested on the remaining 10% of the trials. Only correct trials were used. The training/testing was repeated 50 times for every given number of neurons and for all the sessions that may be included. The correct rates from the 50 repetitions were then averaged. When comparing performance in control and inactivation conditions (Figure 4O-R), the classifiers were trained on correct control trials and tested on correct laser trials and held-out correct control trials.

Imaging data analysis

The raw images were first motion corrected with SIMA package (Kaifosh et al., 2014) (Release 1.3.2) and verified manually. Regions of interest (ROIs) were selected automatically with constrained nonnegative matrix factorization (CNMF) (Pnevmatikakis et al., 2016). The CNMF algorithm infers the time-varying background and extracts smoothed $\Delta F/F$ signals, which were used for plotting only. For data analysis, we manually computed unfiltered $\Delta F/F$ traces as follows. We obtained the raw fluorescent trace of each ROI by applying the spatial component (ROI filter) on the image sequence. We then smoothed the raw trace in each trial with a 1 s averaging window (boxcar) and take the minimal fluorescence value in the inter-trial interval as the baseline. The $\Delta F/F$ signals were calculated by subtracting the baseline from the raw trace and dividing the difference by the baseline. We used the constrained deconvolution spike inference algorithm (FOOPSI) in the CNMF package to infer the spikes (decay time constant 0.7 s). The deconvolved activity was then smoothed using a Gaussian filter over a five-element sliding window.

We compared the means of the response to odors A and B during the sample and delay epochs using a Mann-Whitney U test. Each trial contributed a scalar value: the average $\Delta F/F$ signal from t = 0-2.5 s from onset of the sample odor. Neurons were classified as sample odor-selective if $p < 0.01$ (two-tailed, not corrected for multiple comparisons). Test odor and choice selectivity were determined similarly using time bins of 2.5-4.0 s and 2.5-5.0 s from sample onset, respectively. Selectivity indices (SIs, Equation 1) were computed from the same scalar values. The distributions of the sample odor SIs of ALM L2 neurons acquired with calcium imaging and those of ALM neurons sampled by electrical recording were compared using a two-sample Kolmogorov-Smirnov test.

The standardized odor and choice selectivity (Figure 5A and S5A) were computed by dividing the absolute value of the difference between the mean $\Delta F/F$ responses to odor A and B (or match and non-match) by the common standard deviation: the standard deviation is the square root of the sum of the variances of the $\Delta F/F$ responses to odors A and B (or match and non-match).

To determine the duration of the calcium response, we first computed the mean and standard deviation of the baseline (the epoch before stimulus onset). Calcium transients were then identified as any response greater than 2 standard deviations away from the baseline and lasting for at least 100 ms. The peak time is the time of the maximum calcium response. Only the calcium transients during the sample and delay epochs were considered in Figure 5E. We used identical methods to determine the duration and peak time of the inferred spiking activity except that a minimal response duration was not required.

To ascertain whether the transient responses from L2 neurons occur at random times during the delay, we measured the pairwise correlation in the peak times of the responses (Figure S6A). Statistical significance of each r-value was established using Fisher z-transformed values and their s.e., without correction for multiple comparisons (gray shaded histograms, Figure S6B). We also estimated the distribution of correlation coefficients expected under the null hypothesis, using a shuffle control. The calculations are identical except each ordered pair from the two neurons comprises peak times from non-corresponding trials. The red dashed distribution (Figure S6B) was estimated using 1000 iterations of this procedure.

For the decoding analysis (Figure 6A-B), we trained a support vector machine (SVM) (Fan et al., 2008) using the calcium responses to classify sample odor identity (Figure 6A) or match/nonmatch (Figure 6B). The calcium responses were computed as the average ΔF/F in a sliding bin of 100 ms at 100 ms steps using 40 randomly selected neurons from each of the five cortical depths in ALM.

We characterized a trial-by-trial association between the Ca responses of L2 neurons and the likelihood that the mouse would make an error. The signals themselves are indirect measurements of neural activity and highly skewed. We therefore applied a variant of a choice probability measure, which is based on ordinal statistics (Britten et al., 1996). We included 109 neurons that had statistically reliable preference for sample odor A or B using the integrated Ca signal over the sample and delay epochs. We included both correct and error trials in determining the sample odor selectivity, to avoid biases favoring neurons with larger responses in correct trials due to random fluctuation. For each neuron, the integrated signal was ranked and scored as a percentile ranking on (0,1] using the trials in which the sample odor was the preferred odor to form the vector of trial-by-trial responses for each neuron, r_n^{pref}, where the subscript identifies the neuron. The percentiles were assigned for each neuron independently (not the population). The percentiles from all the neurons were then pooled. We used the same procedure using the responses to each neuron's nonpreferred odor to form r_n^{non}. The heatmap in Figure 6E was formed by parametrically varying the two criteria, κ_{pref} and κ_{non}, to compute the proportion of errors when $r^{pref} < \kappa_{pref}$ and $r^{non} > \kappa_{non}$, using the combined data from all 109 neurons (i.e., concatenating across all n). To evaluated the null hypothesis that these responses have only a random association with the behavioral outcome, we conducted a simple logistic regression using the percentiles themselves:

$$P_{err} = \{1 + \exp(-\beta_0 - \beta_1 \mathbf{X})\}^{-1} \quad (2)$$

Where X is the vector of the transformed percentiles:

$$\mathbf{X} = \begin{pmatrix} -r^{pref} + 1 \\ r^{non} \end{pmatrix} \quad (3)$$

Note that the percentiles are simply reversed for r^{pref} so that the larger percentiles correspond to the weakest responses. We report the p value associated with $\{H_0: \beta_1 = 0\}$.

Behavior/inhibition data analysis

Mouse performance (P) was reported as the fraction of correct responses in all trials. Animals may perform below chance (50%) due to "no choice" trials (e.g., when they are challenged with novel C/D × C/D pairs; Figure S1F). To assess the statistical reliability of photoinhibition on P we generated the distribution of log probability ratio $(\mathscr{L})$ under the null hypothesis

$$H_0 : \mathscr{L} \equiv \log \frac{P_{laser}}{P_{ctrl}} = 0 \quad (4)$$

We randomly permuted the designations, correct/incorrect, among the laser and control trials within each trial type (AA/AB/BB/BA), repeating the process 10,000 times. From this distribution, we obtain the two-tailed probability of obtaining $\mathscr{L}$ under H_0.

To compare the effect of inactivation across tasks, we generated a distribution of $\mathscr{L}$ for each task by bootstrapping. The laser and control trials were re-sampled respectively with replacement within each trial type (AA/AB/BB/BA). Repeating this process 10,000 times, we established a t-statistic from the means and variances of these distributions (degree of freedom based on the number of experimental trials). The reported significance reflects one-tailed comparisons.

We used the following logistic model to characterize an animal's bias. "No choice" trials in which animals did not respond were excluded in this analysis (< 1% of trials).

$$P_{left} = \{1 + \exp(-Q)\}^{-1} \quad (5)$$

$$Q = \beta_0 + \beta_1 S_M + \beta_2 I_{laser} + \beta_3 S_M I_{laser} \quad (6)$$

Where P_{left} is the probably that the animal licks to the left port, S_M is +1 or −1 if the trial is a match or non-match, respectively, and I_{laser} is 1 or 0 for laser on or off, respectively. The beta terms are fitted coefficients: β_0 quantifies the bias in favor of left on control trials in units of log odds; $\beta_0 + \beta_2$ quantifies the bias in laser-on trials; β_1 quantifies how well the animal uses the match/non-match information to determine the direction of licking (i.e., sensitivity to condition) on the control trials; and $\beta_1 + \beta_3$ quantifies the sensitivity in laser-on trials. In a well-trained animal, β_1 is always positive. Thus β_2 is an estimate of the side bias induced by inactivation and β_3 is an estimate of the impairment on sensitivity by inactivation.

When trials from all the animals are combined in this analysis, the bias coefficients could be underestimated (β_0, β_2) because the bias for left or right is different across experiments. It is theoretically possible that underestimation of the laser-induced bias, β_2, could lead to misattribution of this effect to a laser-induced change in sensitivity, β_3. To address this possibility, we first determined the bias of each animal in the laser condition by comparing the rate of correct match and non-match trials. Based on this bias, we designated left or right as the *preferred* lick port and match or non-match as the preferred trial for each animal. Then we combined the trials from all the animals to fit the following modified model:

$$P_{pref} = \{1 + \exp(-Q)\}^{-1} \quad (7)$$

$$Q = \beta_0 + \beta_1 S_{pref} + \beta_2 I_{laser} + \beta_3 S_{pref} I_{laser} \quad (8)$$

Where P_{pref} is the probability that the animal licks to the preferred port, and S_{pref} is +1 or −1 if the trial is preferred or non-preferred, respectively. In this model, β_2 estimates the side bias induced by inactivation across all sessions. Importantly, if inactivation only biased the mouse to lick more to one port or the other—a different bias on each experiment—this procedure would fail to reject the null hypothesis, H_0: $\beta_3 = 0$. We complemented this analysis using Monte Carlo methods to estimate the magnitude of impairment on the task that is not accounted for by a bias to the left or right lick port. We simulated datasets using the estimated β coefficients and their standard error, while setting $\beta_2 = 0$. This recovers the proportion correct on the control trials and models the proportion correct on the inactivation trials, were no bias induced by the laser. In Figure S4C, we show the distribution of impairments from 10,000 repetitions.

Attractor network models

We constructed recurrent neural network models consisting of three stages (Figure 7A): (1) primary sensory areas (e.g., Pir); (2) intermediate areas; (3) ALM. Each stage $a = 1, 2, 3$ contains $N = 80$ units whose activities are represented by an N-dimensional vector $\mathbf{x}_a$ and follow the dynamics

$$\tau \frac{d\mathbf{x}_a}{dt} = -\mathbf{x}_a + f\left(\mathbf{J}_a \mathbf{x}_a + \mathbf{J}_a^{FF} \mathbf{x}_{a-1} + \sigma \boldsymbol{\eta}_a(t) + \mathbf{b}_a\right), a = 1, 2, 3. \quad (9)$$

The matrices $\mathbf{J}_a$ and $\mathbf{J}_a^{FF}$ represent recurrent input and feedforward input from the previous area, respectively. The vector $\mathbf{x}_0$ is two-dimensional, and its two elements are indicator variables (with value 1 or 0) representing the presence or absence of odors A and B. The term $\sigma \boldsymbol{\eta}_a(t)$ represents independent white noise input with standard deviation σ. The vector $\mathbf{b}_a$ represents the bias inputs to each unit. The function f is rectified-linear and the time constant τ (which represents combined membrane and synaptic time constants) equals 100 ms.

In each trial, the network receives sample and test odors (A or B) for 500 ms each, beginning at times $t = 1$ and 3 s. Each trial is drawn randomly from one of the four trial types. A readout of the network must classify the trial as match or non-match during the response period, from $t = 3.5$ s to $t = 4$ s. The readout is a softmax function of the ALM activity $\mathbf{x}_3$. At the beginning of each trial, the initial values of the $\mathbf{x}_a$ vectors are taken to be independent rectified Gaussian random variables with standard deviation 0.05. Networks are simulated with a timestep of 20 ms.

The networks are trained through gradient descent with TensorFlow and the Adam optimizer (Agarwal et al., 2016; Kingma and Ba, 2014) to minimize a loss $L = L_{classifier} + L_{activity}$ determined by the classifier and an activity regularization term. Specifically, $L_{classifier}$ in a given training epoch equals the summed cross-entropy loss between the classifier's output and the desired output (match or non-match) during the response period, averaged over a batch size of 50 trials. The regularization term $L_{activity} = 10^{-4} \cdot \langle \|\mathbf{r}_a\|^2 \rangle$ is proportional to the ℓ_2 norm of the activities averaged across units, time, areas, and batches. Every 50 epochs of training, the network is tested on 1000 trials to determine its classification accuracy, and training ceases when this accuracy exceeds 95%. The noise σ equals 0.05 during training and 0.1 during this testing phase. The learning rate of the optimizer decreases logarithmically from 10^{-3} to 10^{-4} over 1000 epochs. Networks which do not reach the 95% criterion accuracy after these 1000 epochs are discarded.

At the beginning of training, all recurrent and feedforward weights $\mathbf{J}_a$, $\mathbf{J}_a^{FF}$ are initialized as independent random Gaussian variables with standard deviation $1/\sqrt{N}$ (except for $\mathbf{J}_0^{FF}$, which has standard devation $1/\sqrt{2}$ because of the dimension of $\mathbf{x}_0$). Half of the feedforward weights are then randomly set equal to zero, representing sparser connections across areas versus within areas. The softmax classifier weights are initialized with standard deviation $2/\sqrt{N}$, and the biases $\mathbf{b}_a$ are initialized to zero. We assume that only $\mathbf{J}_a$ and $\mathbf{b}_a$ are learned. All other variables are fixed during training.

To generate the performance curves in Figure 7B, 60 trained networks were tested with $\sigma = 0.2$ under four conditions. In the control condition, the dynamics were identical to those described above. In the other conditions, which mimic ALM inactivation during different epochs of the task, an inhibitory input was applied to ALM units during the sample and early delay epochs (the 1 s following sample onset), the late delay period (the final 1 s of the delay epoch), or the entire sample and delay epochs. The inhibitory input was equivalent to reducing the biases of all ALM units $\mathbf{b}_3$ by 5. To generate the curve for networks without ALM persistence, modifications of $\mathbf{J}_a$ were restricted to only $\mathbf{J}_1$ and $\mathbf{J}_2$ while $\mathbf{J}_3$ was held fixed, so that the training algorithm could not learn to implement attractor dynamics in ALM $(a = 3)$. Conversely, for networks with only ALM persistence, training was restricted to only $\mathbf{J}_3$, while $\mathbf{J}_1$ and $\mathbf{J}_2$ were held fixed.

DATA AND CODE AVAILABILITY

The datasets and code are available upon request to the Lead Contact.