

A stable, distributed code for cue value in mouse cortex during reward learning

David J. Ottenheimer^{1,2,3,5}, Madelyn M. Hjort^{1,2,5}, Anna J. Bowen^{3,5},
Nicholas A. Steinmetz^{1,3,6}, Garret D. Stuber^{1,2,4,6}

¹*Center for the Neurobiology of Addiction, Pain, and Emotion,*

²*Anesthesiology and Pain Medicine,* ³*Biological Structure,*

⁴*Pharmacology, University of Washington, Seattle.*

⁵These authors contributed equally

⁶Corresponding authors: nick.steinmetz@gmail.com, gstuber@uw.edu

Summary

The ability to associate reward-predicting stimuli with adaptive behavior is frequently attributed to the prefrontal cortex, but the stimulus-specificity, spatial distribution, and stability of prefrontal cue-reward associations are unresolved. We trained headfixed mice on an olfactory Pavlovian conditioning task and measured the coding properties of individual neurons across space (prefrontal, olfactory, and motor cortices) and time (multiple days). Neurons encoding cues or licks were most common in olfactory and motor cortex, respectively. By quantifying the responses of cue-encoding neurons to six cues with varying probabilities of reward, we unexpectedly found value coding, including coding of trial-by-trial reward history, in all regions we sampled. We further found that prefrontal cue and lick codes were preserved across days. Our results demonstrate that individual prefrontal neurons stably encode components of cue-reward learning within a larger spatial gradient of coding properties.

Introduction

Association of environmental stimuli with rewards and the subsequent orchestration of value-guided reward-seeking behavior are crucial functions of the nervous system linked to the pre-frontal cortex (PFC) (Klein-Flügge et al., 2022; Miller and Cohen, 2001). PFC is heterogeneous, with many studies noting subregional differences in both neural coding of (Hunt et al., 2018; Kennerley et al., 2009; Sul et al., 2010; Wang et al., 2020a) and functional impact on (Buckley et al., 2009; Dalley et al., 2004; Kesner and Churchwell, 2011; Rudebeck et al., 2008) value-based reward seeking in primates and rodents. Further, functional manipulations of PFC subregions exhibiting robust value signals do not always cause a discernible impact on reward-guided behavior (Chudasama and Robbins, 2003; Dalton et al., 2016; St. Onge and Floresco, 2010; Verharen et al., 2020; Wang et al., 2020a), encouraging investigation of differences between value signals across PFC. Within individual PFC subregions, multiple studies have observed evolving neural representations across time, calling into question the stability PFC signaling (Hyman et al., 2012; Malagon-Vina et al., 2018). A systematic comparison of coding properties across rodent PFC and related motor and sensory regions, as well as across days and stimulus sets, is necessary to provide a full context for the contributions of PFC subregions to reward processing.

Identifying neural signals for value requires a number of considerations. One issue is that other task features can vary either meaningfully or spuriously with value. In particular, action coding is difficult to parse from value signaling, given the high correlations between behavior and task events (Musall et al., 2019; Zagha et al., 2022) and widespread neural coding of reward-seeking actions (Steinmetz et al., 2019). Additionally, without a sufficiently rich value axis, it is possible to misidentify neurons as ‘value’ coding even though they do not generalize to valuations in other contexts (Hayden and Niv, 2021; Stalnaker et al., 2015; Zhou et al., 2021). Because reports of value have come from different experiments across different species, it is

difficult to compare the presence of value signaling even across regions within prefrontal cortex (Hayden and Niv, 2021; Hunt et al., 2018; Kennerley et al., 2009; Namboodiri et al., 2019; Otis et al., 2017; Stalnaker et al., 2015; Sul et al., 2010; Wang et al., 2020a; Zhou et al., 2021).

In this work, we sought to address the existing ambiguity in the distribution and stability of value signaling. We implemented an olfactory Pavlovian conditioning task that permitted identification of value correlates within the domain of reward probability across two separate stimulus sets. With acute *in vivo* electrophysiology recordings, we were able to assess coding of this task across 11 brain regions, including five PFC subregions, as well as olfactory and motor cortex, in a single group of mice, permitting a well-controlled comparison of coding patterns across a large group of task-relevant regions in the same subjects. Unexpectedly, in contrast to the graded cue and lick coding across these regions, there was a similar proportion of neurons encoding cue value in all regions. A subset of value coding neurons were sensitive to trial-by-trial fluctuations in value according to reward history. To assess coding stability, we performed 2-photon imaging of neurons in PFC for multiple days and determined that the cue and lick codes we identified were stable over time. Our data demonstrate universality and stability of cue-reward coding in mouse cortex.

Results

Distributed neural activity during an olfactory Pavlovian conditioning task.

We trained mice on an olfactory Pavlovian conditioning task with three cue (conditioned stimulus) types that predicted reward on 100% ('CS+'), 50% ('CS50'), or 0% ('CS-') of trials (Fig. 1A). Each mouse learned two odor sets, trained and imaged on separate days and then, for electrophysiology experiments, presented in six alternating blocks of 51 trials during recording sessions (Fig. 1B). Mice developed anticipatory licking (Fig. 1C-D), and the rate of this licking correlated with reward probability (Fig. S1), indicating that subjects successfully learned the

meaning of all six odors.

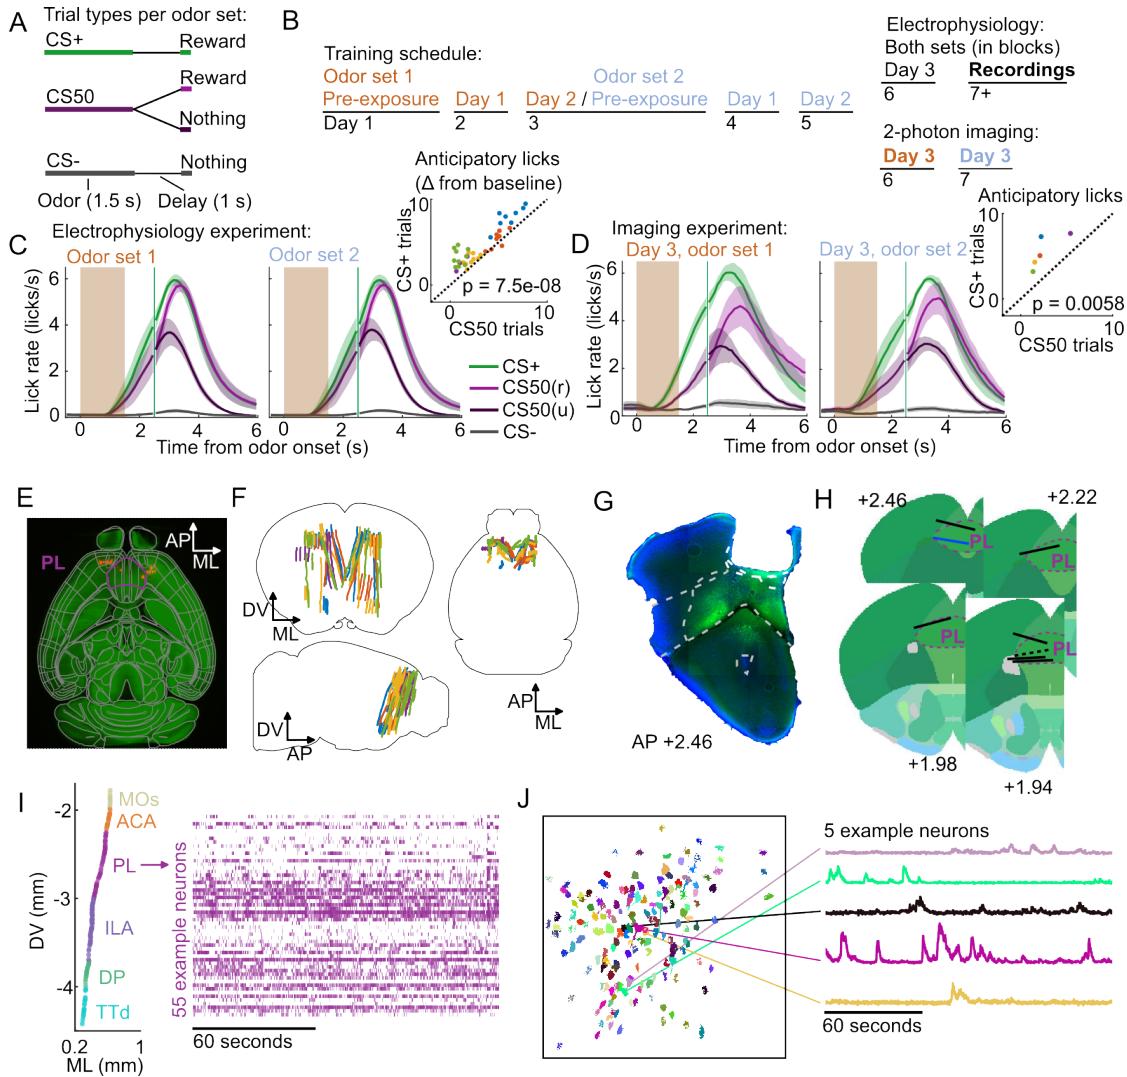


Figure 1: Electrophysiology and imaging during olfactory Pavlovian conditioning.

Figure 1. Electrophysiology and imaging during olfactory Pavlovian conditioning.

- (A) Trial structure in Pavlovian conditioning task.
- (B) Timeline for mouse training.
- (C) Mean (+/- standard error of the mean (SEM)) lick rate across mice ($n = 5$) on each trial type for each odor set during electrophysiology sessions. CS50(r) and CS50(u) are rewarded and unrewarded trials, respectively. Inset: mean anticipatory licks (change from baseline) for the CS+ and CS50 cues for every session, color-coded by mouse. $F(1, 66) = 36.6$ for a main effect of cue in a two-way ANOVA including an effect of subject.
- (D) Same as (C), for Day 3 imaging sessions ($n = 5$ mice). $t(4) = -5.4$ for a t-test comparing anticipatory licks on CS+ and CS50 trials.
- (E) Neuropixels probe tracks labeled with fluorescent dye (red) in cleared brain (autofluorescence, green). AP, anterior/posterior; ML, medial/lateral; DV, dorsal/ventral. Allen CCF regions delineated in gray. Outline of prelimbic area in purple.
- (F) Reconstructed recording sites from all tracked probe insertions ($n = 44$ insertions, $n = 5$ mice), colored by mouse.
- (G) Sample histology image of lens placement. Visualization includes DAPI (blue) and GCaMP (green) signal with lines indicating cortical regions from Allen Mouse Brain Common Coordinate Framework.
- (H) Location of all lenses from experimental animals registered to Allen Mouse Brain Common Coordinate Framework. Blue line indicates location of lens in (A). The dotted black line represents approximate location of tissue that was too damaged to reconstruct an accurate lens track. The white dotted line indicates PL borders.
- (I) ML and DV coordinates of all neurons recorded in one example session, colored by region, and spike raster from example PL neurons.
- (J) ROI masks for identified neurons and fluorescence traces from 5 example neurons.

Using Neuropixels 1.0 and 2.0 probes (Jun et al., 2017; Steinmetz et al., 2021), we recorded the activity of individual neurons in **PFC**, including anterior cingulate area (ACA), frontal pole (FRP), prelimbic area (PL), infralimbic area (ILA), and orbital area (ORB) (Laubach et al., 2018; Wang et al., 2020b). We also recorded from: secondary **motor cortex** (MOs), including anterolateral motor cortex (ALM), which has a well-characterized role in licking (Chen et al., 2017); **olfactory cortex** (OLF), including dorsal peduncular area (DP), dorsal taenia tecta (TTd), and anterior olfactory nucleus (AON), which receive input from the olfactory bulb

(Igarashi et al., 2012; Mori and Sakano, 2021); and **striatum**, including caudoputamen (CP) and nucleus accumbens (ACB) (Fig. 1E-F), which are major outputs of PFC (Heilbronner et al., 2016). In a separate group of mice, we performed longitudinal 2-photon imaging through a GRIN lens to track the activity of individual neurons in PL across several days of behavioral training (Fig. 1G-H). Both techniques permitted robust measurement of the activity of neurons of interest and generated complementary results (Figs. 1I-J, S2).

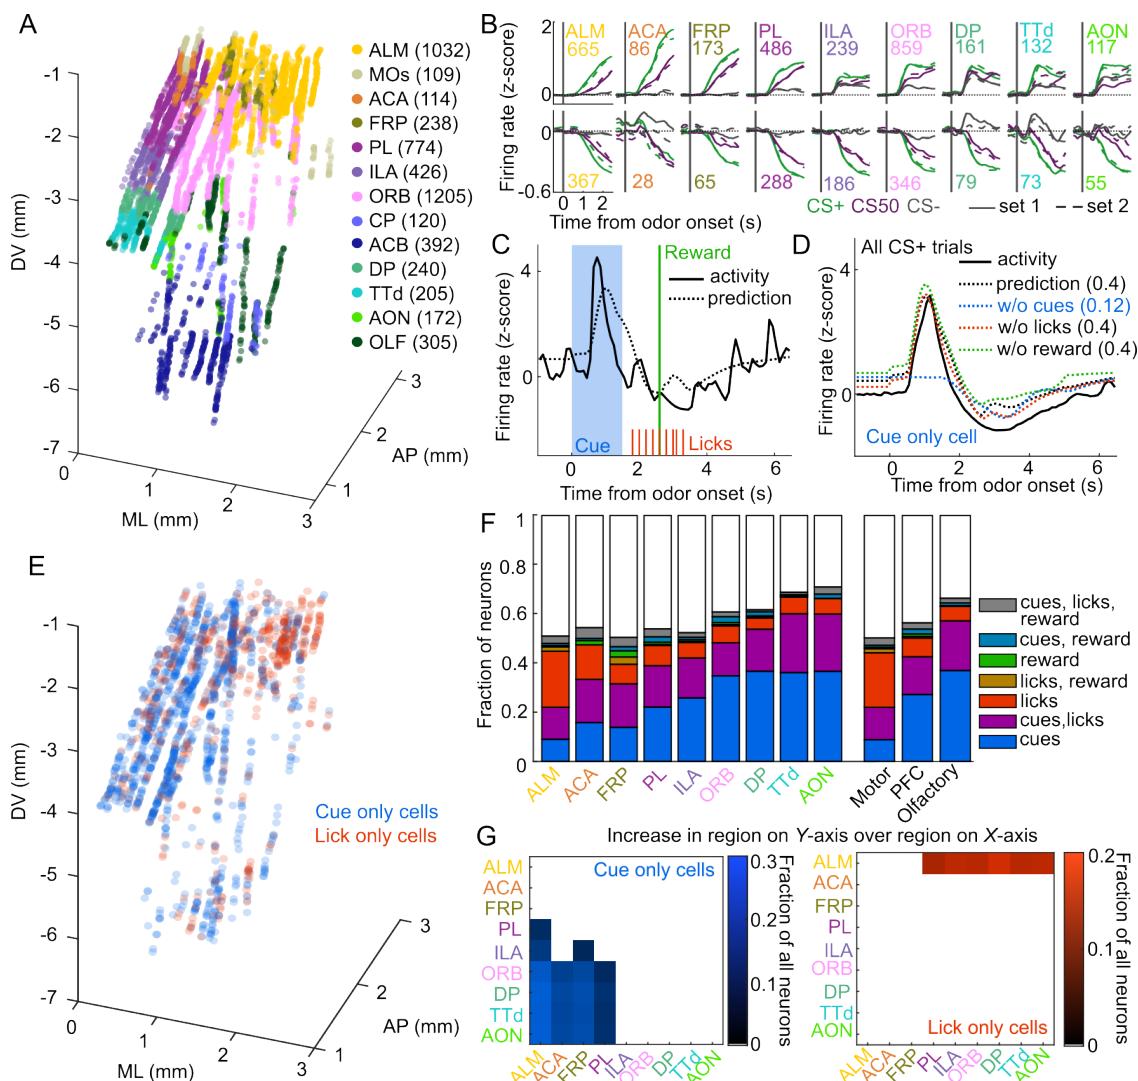


Figure 2: Graded cue and lick coding across the recorded regions.

Figure 2. Graded cue and lick coding across the recorded regions.

- (A) Location of each recorded neuron relative to bregma, projected onto 1 hemisphere. Each neuron is colored by CCF region. Numbers indicate total neurons passing quality control from each region.
- (B) Mean normalized activity of all neurons from each region, aligned to odor onset, grouped by whether peak cue activity (0 - 2.5 s) was above (top) or below (bottom) baseline in held out trials. Number of neurons noted for each plot.
- (C) Example kernel regression prediction of an individual neuron's normalized activity on an example trial.
- (D) CS+ trial activity from an example neuron and predictions with full model and with cues, licks, and reward removed. Numbers in parentheses are model performance (fraction of variance explained).
- (E) Coordinates relative to bregma of every neuron encoding only cues or only licks, projected onto one hemisphere.
- (F) Fraction of neurons in each region and region group classified as coding cues, licks, reward, or all combinations of the three.
- (G) Additional cue (left) or lick (right) neurons in region on Y-axis compared to region on X-axis as a fraction of all neurons, for regions with non-overlapping 95% confidence intervals (see Methods).

Graded cue and lick coding across the recorded regions.

In the electrophysiology experiment, we isolated the spiking activity of 5332 individual neurons in regions of interest across 5 mice (449-1550 neurons per mouse, Figs. 2A, S3A). The activity of neurons in all regions exhibited varying degrees of modulation in response to the six trial types (Fig. 2B). Broadly, there was strong modulation on CS+ and CS50 trials that appeared to be common to both odor sets (Fig. S3B). Across regions, there was heterogeneity in both the magnitude and the timing of the neural modulation relative to odor onset (Fig. S3C).

To quantify the relative contribution of cues and conditioned responding (licking) to the activity of neurons in each region, we implemented reduced rank kernel regression (Steinmetz et al., 2019), using cues, licks, and rewards to predict neurons' activity on held out trials (Figs. 2C, S4A). To determine the contribution of cues, licks, and rewards to each neuron's activity, we

calculated unique variance explained by individually removing each predictor from the model and calculating the reduction in model performance (Fig. 2D).

We identified individual neurons encoding cues, licks, or rewards as those for which that predictor uniquely contributed to 2% or more of their variance (a cutoff with an estimated false positive rate of 0.02%, see Methods). Neurons encoding cues (24% of all neurons), licks (11%), or both (16%) were most common. Neurons with any response to reward (independent of licking) were rare (5%) (Horst and Laubach, 2013). Cue neurons were characterized by sharp responses aligned to odor onset; in contrast, lick neurons' responses were delayed and peaked around reward delivery (Fig. S4B-C), consistent with the timing of licks (Fig. 1C). The activity of cue neurons on rewarded and unrewarded CS50 trials validated our successful isolation of neurons with cue but not lick responses (Fig. S4D). The spatial distributions of cue and lick cells were noticeably different (Fig. 2E). The differences could be described as graded across regions, with the most lick neurons in ALM, and the most cue neurons in olfactory cortex and ORB, though each type of neuron was observed in every region (Fig. 2F-G, S5). Thus, our quantification of task encoding revealed varying prioritization of cue and lick signaling across all regions.

Cue value coding is present in all regions.

To expand upon our analysis identifying cue-responsive neurons, we next assessed the presence of cue value coding in this population. The three cue types (CS+, CS50, or CS-) in our behavioral task varied in relative value according to the predicted probability of reward (Eshel et al., 2016; Fiorillo et al., 2003; Winkelmeier et al., 2022). We reasoned that a neuron encoding cue value should have activity that scaled with the relative value of the cues (Fig. 3A). We modeled this relationship on a per-neuron basis by scaling a single cue kernel by the typical number of anticipatory licks for each cue (see Methods, Fig. 3B). This model describes cue

activity as similar across odors of the same value, and scaling in magnitude according to each odor's value. As a shuffle control, we also fit each neuron with 89 additional models containing all possible permutations of reassigning the original values across the six odors, as well as an 'untuned' model with the same value for all odors (Fig. S6). After removing neurons best fit by the untuned model, the remaining 90 models would be equally likely to fit each neuron best if cue responses were independent of cue association, as would be expected with a sensory olfactory code (Pashkovski et al., 2020; Stettler and Axel, 2009). We found, however, that the original model with cue responses that scale directly with value (CS+ > CS50 > CS-) was the best model for a large fraction of cue neurons (30%), far exceeding chance (1%) (Fig. 3C). We refer to these neurons as **value** cells (Fig. S7A). Interestingly, five additional models stood out as the best model for sizable fractions of cue neurons. These models corresponded to the five alternative rankings of cue types (for example, CS50 > CS- > CS+ rather than CS+ > CS50 > CS-) irrespective of odor set (Fig. 3D) and accounted for 22% of cue neurons. This population of cells, which we refer to as **trial type** cells (Fig. S7B), encoded each cue's reward probability independent of its particular odor, but with firing rates not proportional to value. The difference between value and trial type cells, therefore, is that value cells have both CS- activity closer to baseline and CS+ and CS50 activity occurring along the same dimension relative to CS- activity, evident in a smaller angle between CS+ and CS50 population response vectors for value cells (Fig. S7C-E).

The frequency of value cells was similar across the recorded regions (Fig. 3E). Indeed, despite the regional variability in number of cue cells broadly (Fig. 2F-G), there were no regions that statistically differed in their proportions of value cells (Figs. 3E, S8). Though there were fewest cue neurons in motor cortex, they were more likely than cue neurons in other regions to encode value, followed by PFC (Fig. S9). The frequency of trial type cells was more variable, peaking in DP (Fig. 3F) and decreasing from olfactory to PFC to motor cortex (Fig. S8).

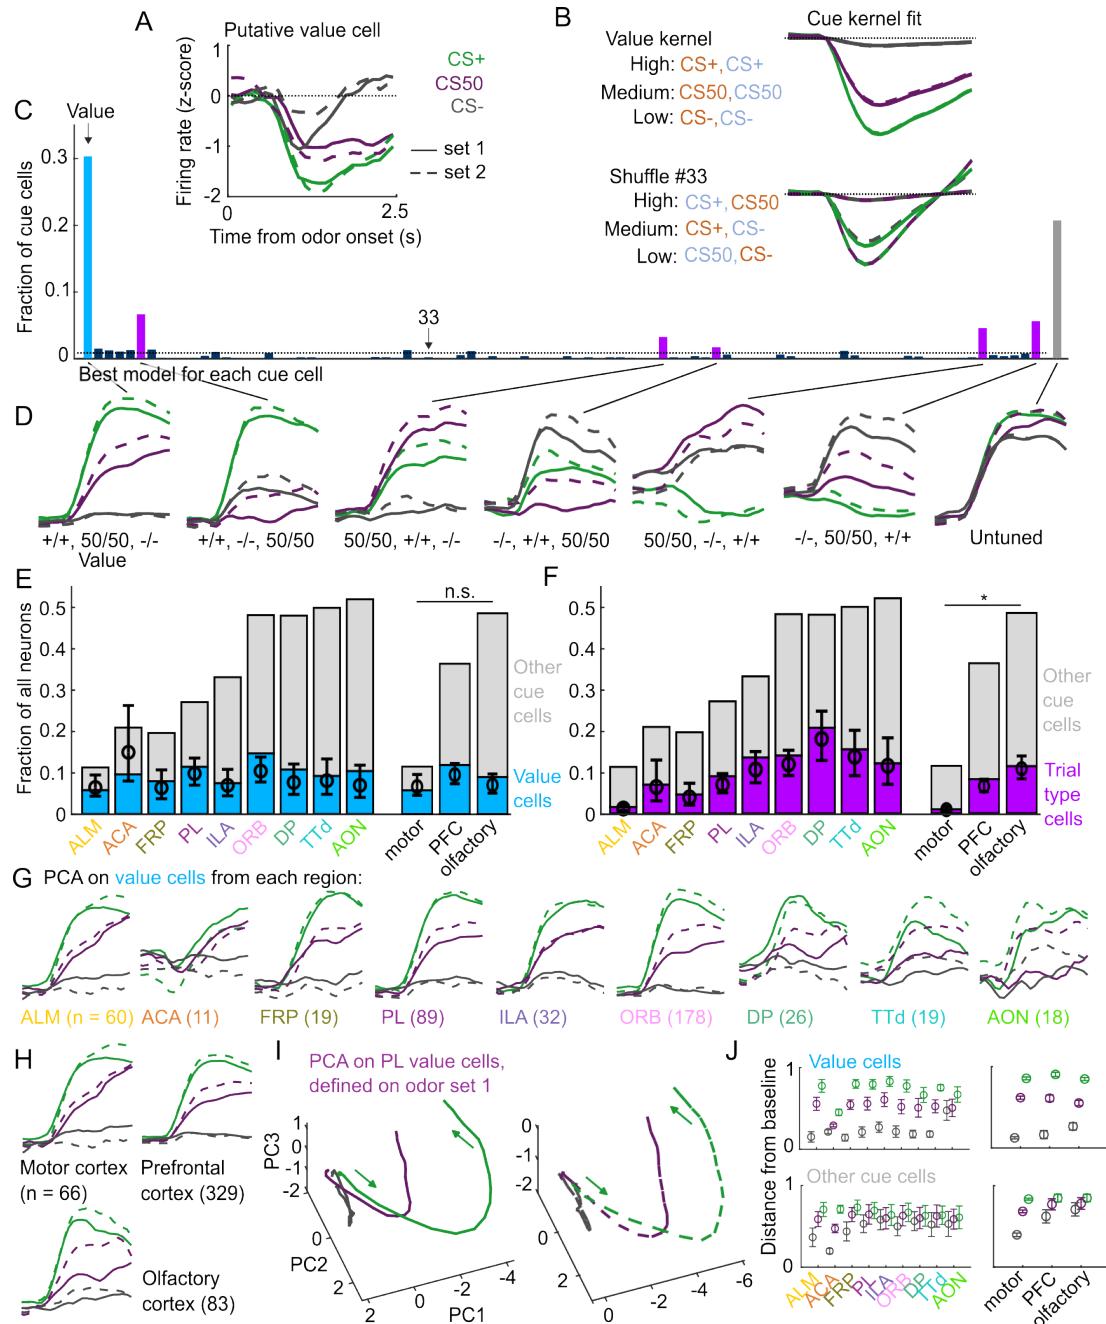


Figure 3: Widespread cue value coding.

Figure 3. Widespread cue value coding.

- (A) Normalized activity of an example value cell with increasing modulation for cues of higher value.
- (B) For the same neuron, model-fit cue kernel for the original value model and with one of the 89 alternatively-permuted value models.
- (C) Distribution of best model fits across all cue neurons. Light blue is value model, purple is trial type models, gray is untuned model, and the remaining models are dark blue.
- (D) First principal component of all neurons best fit by the original value model or other trial type and untuned models.
- (E) Fraction of neurons in each region and region group classified as value cells (blue) and other cue neurons (gray), as well as fraction (+/- 95% CI) estimated from a linear mixed effects model with random effect of session (see Methods). n.s. indicates overlapping 95% CI for all three region groups.
- (F) Additional value cells in region on Y-axis compared to region on X-axis as a fraction of all neurons, for regions with non-overlapping 95% confidence intervals. * indicates non-overlapping 95% CI for all three region groups.
- (G) Principal component most related to value for value cells from all regions (2nd component in ACA, 1st component in all others).
- (H) Same as (F), for region groups.
- (I) PL population activity 0 to 2.5 s from odor onset projected onto first 3 principal components, defined on odor set 1 activity.
- (J) Normalized distance (+/- SD) from baseline firing in odor set 1 PCA space (first 3 components) for odor set 2 trial types for value cells (top) and other cue cells (bottom) for 5000 selections of neurons (see Methods).

We next investigated the robustness of the value representation in each of our recorded regions. Principal component analysis on value cells from each region revealed similarly strong value-related dynamics across motor, prefrontal, and olfactory regions (Fig. 3G-H). Within the population space defined by the first 3 principal components of odor set 1 activity of value cells in each region, odor set 2 activity of this population showed strong value features; specifically, population activity deflected furthest from baseline for CS+ trials, less for CS50 trials, and least for CS- trials, with similar robustness in ALM, FRP, PL, ILA, ORB, and TTd (Fig. 3I-J). Taken together, these data illustrate that, in contrast to cue and lick coding broadly, and in contrast to

trial type cue coding, value coding is similarly represented across the regions we sampled. In fact, this observation extended to the striatal regions we sampled as well, indicating that such value coding is widespread even beyond cortex (Fig. S10).

Because cue valuations can be influenced by preceding reward outcomes, We next considered whether the cue value signaling we detected was sensitive to the history of reinforcement (Nakahara et al., 2004; Ottenheimer et al., 2020; Winkelmeier et al., 2022). To estimate the subjects' trial-by-trial cue valuation, we fit a linear model predicting the number of anticipatory licks on each trial from cue type, overall reward history, and cue type-specific reward history; we used the model prediction of licks per trial as our estimate of value. We found a strong influence of cue type-specific reward history and a more modest influence of overall reward history (Fig. 4A). These effects of reward history on lick rate are also visible when comparing the lick rate on trials grouped by the model-estimated value (Fig. 4B). Thus our behavioral model revealed that subjects more highly valued trials preceded by rewards.

We therefore investigated whether value cells showed similar trial-by-trial differences in their firing rates (Fig. 4C). To test this, we compared the fit of our original value model (Fig. 3B) with an alternative model in which the kernel scaled with the per-trial value estimates from our lick regression (Fig. 4D). Overall, 30% of value cells were better fit by the history-dependent value model than by the model without history dependence. To further evaluate the history component, we calculated these neurons' activity on CS50 trials of varying value and projected it onto the population dimension maximizing the separation between CS+ and CS-. We hypothesized that high value CS50 trials would be closer to CS+ activity while low value CS50 trials would be closer to CS- activity. Indeed, history value cells (and lick cells) demonstrated graded activity along this dimension, in contrast to non-history value cells and trial type cells (Fig. 4E-F). Finally, we examined the spatial distribution of history value cells and found similar, low numbers across all regions, but with higher prevalence overall in PFC

than in olfactory cortex (Fig. 4G). Also, as a fraction of all value cells, history cells were least common in olfactory cortex (Fig. 4H), providing evidence for less dynamic task-centric cue signaling there.

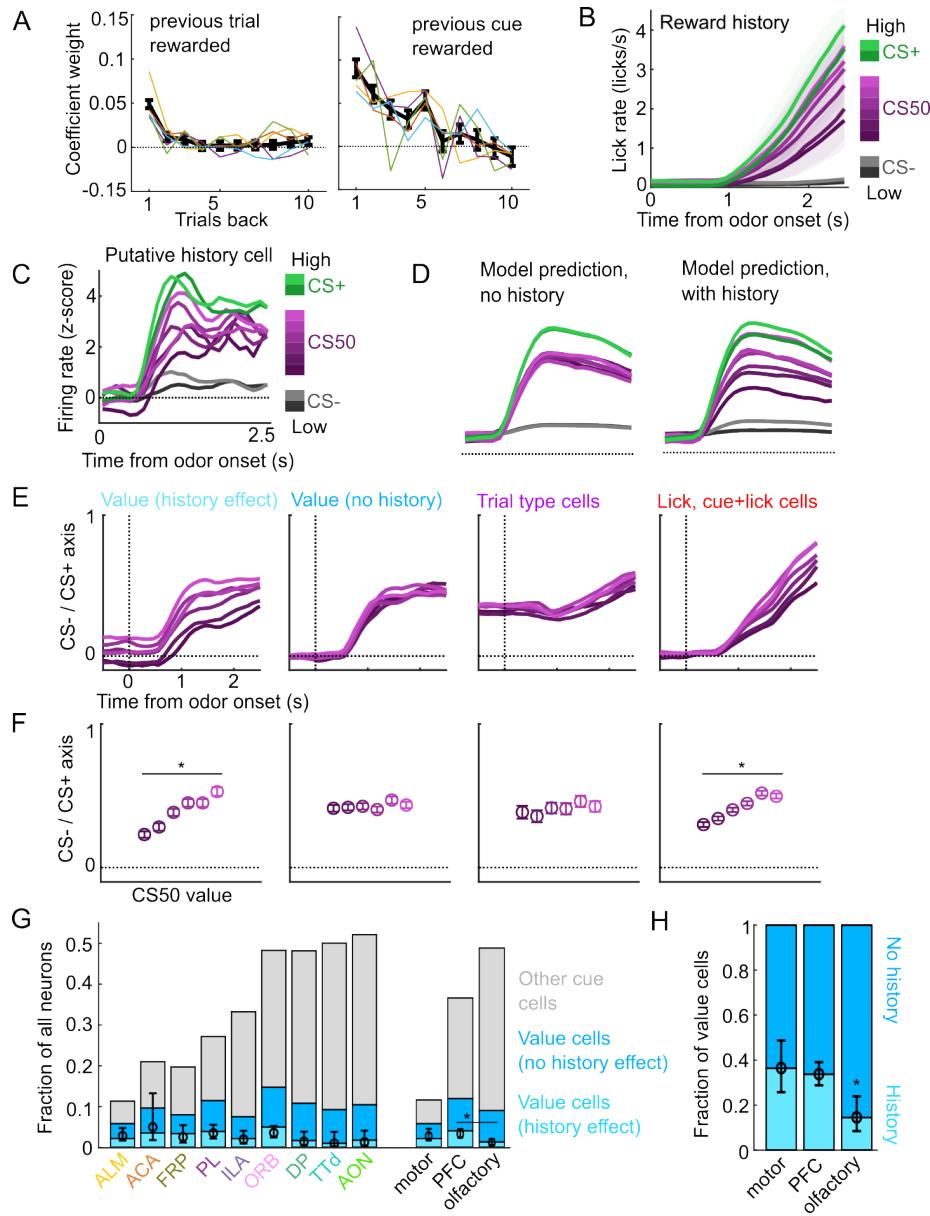


Figure 4: A subset of value cells incorporate reward history.

Figure 4. A subset of value cells incorporate reward history.

(A) Coefficient weight (+/- standard error from model fit) for reward outcome on the previous 10 trials of any type (left) and on the previous 10 trials of the same cue type (right) for a linear model predicting the number of anticipatory licks on every trial of every session. Lick rates were normalized so that the maximum lick rate for each session was equal to 1. Colored lines are models fit to each individual mouse.

- (B) Mean (+/- SEM) lick rate across mice ($n = 5$ mice) on trials binned according to value estimated from the lick model, incorporating recent reward history.
- (C) Normalized activity of an example history value cell with increasing modulation for cues of higher value.
- (D) For the same neuron, model-predicted activity with the original value model (left) and with trial-by-trial value estimates from the lick model (right).
- (E) The activity of all cells in each category projected onto the coding dimension maximally separating CS- and CS+ for CS50 trials binned by value estimated from the lick model.
- (F) The mean (+/- std across 5000 bootstrapped selections of neurons) activity (1 - 2.5 s from odor onset) along the coding dimension maximally separating CS- and CS+ for CS50 trials binned by value estimated from the lick model. * = $p < 10^{-7}$ comparing highest and lowest value CS50 trials (other categories $p > 0.23$).
- (G) Fraction of neurons in each region and region group classified as history value (light blue), non-history value (blue), and other cue neurons (gray), as well as estimated fraction (+/-95% CI) with random effect of session (see Methods). * indicates non-overlapping 95% CI for PFC and olfactory regions.
- (H) Fraction of value neurons in each region group with history effect and estimated fraction (+/- 95% CI) with random effect of session. * indicates non-overlapping 95% CI for olfactory compared to PFC and motor regions.

Cue coding emerges along with behavioral learning.

To determine the timescales over which these coding schemes emerged and persisted, we performed longitudinal 2-photon imaging and tracked the activity of individual neurons across several days of behavioral training (Fig. 5A). We targeted our GRIN lenses to PL, a location with robust cue and lick coding (Fig. 2F) and where cue responses predominantly encode cue trial type and value (Fig. 3E-F). Mice ($n = 8$) developed anticipatory licking on day 1 that differentiated CS+ trials from CS50 ($t(7) = 3.2, p = 0.015$) and CS- ($t(7) = 7.0, p = 0.0002$) trials and CS50 trials from CS- ($t(7) = 3.7, p = 0.008$) trials (Fig. 5B-C). Visualizing the normalized activity across the imaging plane following CS+ presentation early and late on day 1 revealed a pronounced increase in modulation in this first session (Fig. 5D-E). Individual neurons ($n = 705$, 41-165 per mouse) also displayed a notable increase in modulation in response to the CS+ after task learning (Fig. 5F).

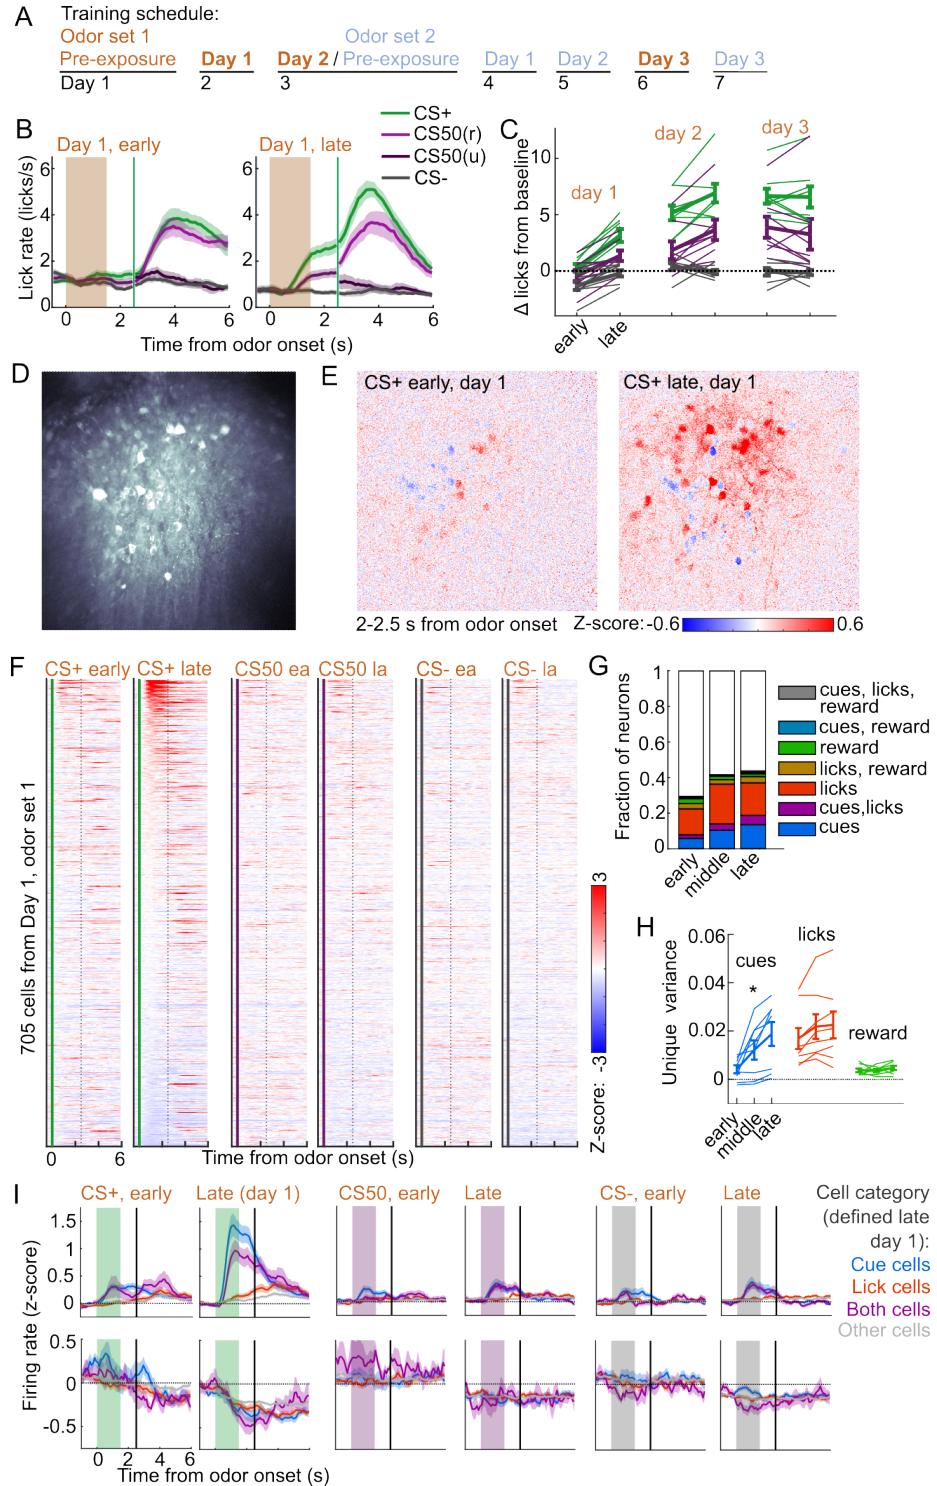


Figure 5: Acquisition of conditioned behavior and cue encoding in PFC.

Figure 5. Acquisition of conditioned behavior and cue encoding in PFC.

- (A) Training schedule for 5 of the mice in the imaging experiment. An additional 3 were trained only on odor set 1.
- (B) Mean (+/- SEM) licking on early (first 60) and late (last 60) trials from day 1 of odor set 1 ($n = 8$ mice).
- (C) Mean (+/- SEM) baseline-subtracted anticipatory licks for early and late trials from each day of odor set 1. Thin lines are individual mice ($n = 8$ mice).
- (D) Standard deviation of fluorescence from example imaging plane.
- (E) Normalized activity of each pixel following CS+ presentation in early and late day 1 trials.
- (F) Normalized deconvolved spike rate of all individual neurons for early and late trials on day 1.
- (G) Proportion of neurons classified as coding cues, licks, rewards, and all combinations for each third of day 1.
- (H) Mean(+/- SEM) unique variance explained by cues, licks, and rewards for neurons from each mouse. Thin lines are individual mice. Unique variance was significantly different across session thirds for cues ($F(2, 21) = 3.71, p = 0.04$) but not licks ($F(2, 21) = 0.37, p = 0.69$) or reward ($F(2, 21) = 0.65, p = 0.53, n = 8$ mice, one-way ANOVA).
- (I) Mean (+/- SEM) normalized deconvolved spike rate for cells coding cues, licks, both, or neither on early and late trials, sorted by whether peak cue activity (0 - 2.5 s) was above (top) or below (bottom) baseline for late trials.

To determine whether this increase in activity was best explained by a cue-evoked response, licking, or both, we again used kernel regression to fit and predict the activity of each neuron for early, middle, and late trials on day 1. The number of individual neurons encoding cues more than doubled from early to late day 1 trials (Fig. 5G). The unique variance cues increased across this first session, in contrast to licks and reward (Fig. 5H). This stark change in cue coding was also noticeable when plotting neurons encoding cues, licks, or both, as defined at the end of the session, on both early and late trials (Fig. 5I). These data indicated that PFC neural activity related to cues (but not licks) rapidly emerges during initial learning of the behavioral task.

Cue and lick coding is stable across days

We next assessed whether cue and lick coding were stable across days. By revisiting the same imaging plane on each day of training, we were able to identify neurons that were present on all three days of odor set 1 training ($n = 371$, 20-65 per mouse) (Fig. 6A-B). There was remarkable conservation of task responding across days, both on an individual neuron level (Fig. 6C) and across all imaged neurons (Fig. 6D). To quantify coding stability, we fit our kernel regression to the activity of each neuron on day 3 (Fig. 6E) and then used these models to predict activity on early, middle, and late trials on days 1-3 (Fig. 6F). Day 3 model predictions were most highly correlated with true activity on day 3, but they outperformed shuffle controls at all time points except early day 1, demonstrating preservation of a learned coding scheme. We then asked more specifically whether cells coding cues, licks, and both maintained their coding preferences across days. For each group of cells, we calculated their unique cue, lick, and reward variance at each time point. The preferred coding of each group, as defined on day 3, was preserved in earlier days (Fig. 6G). Thus, cue and lick coding are stable properties of PFC neurons across multiple days of behavioral training.

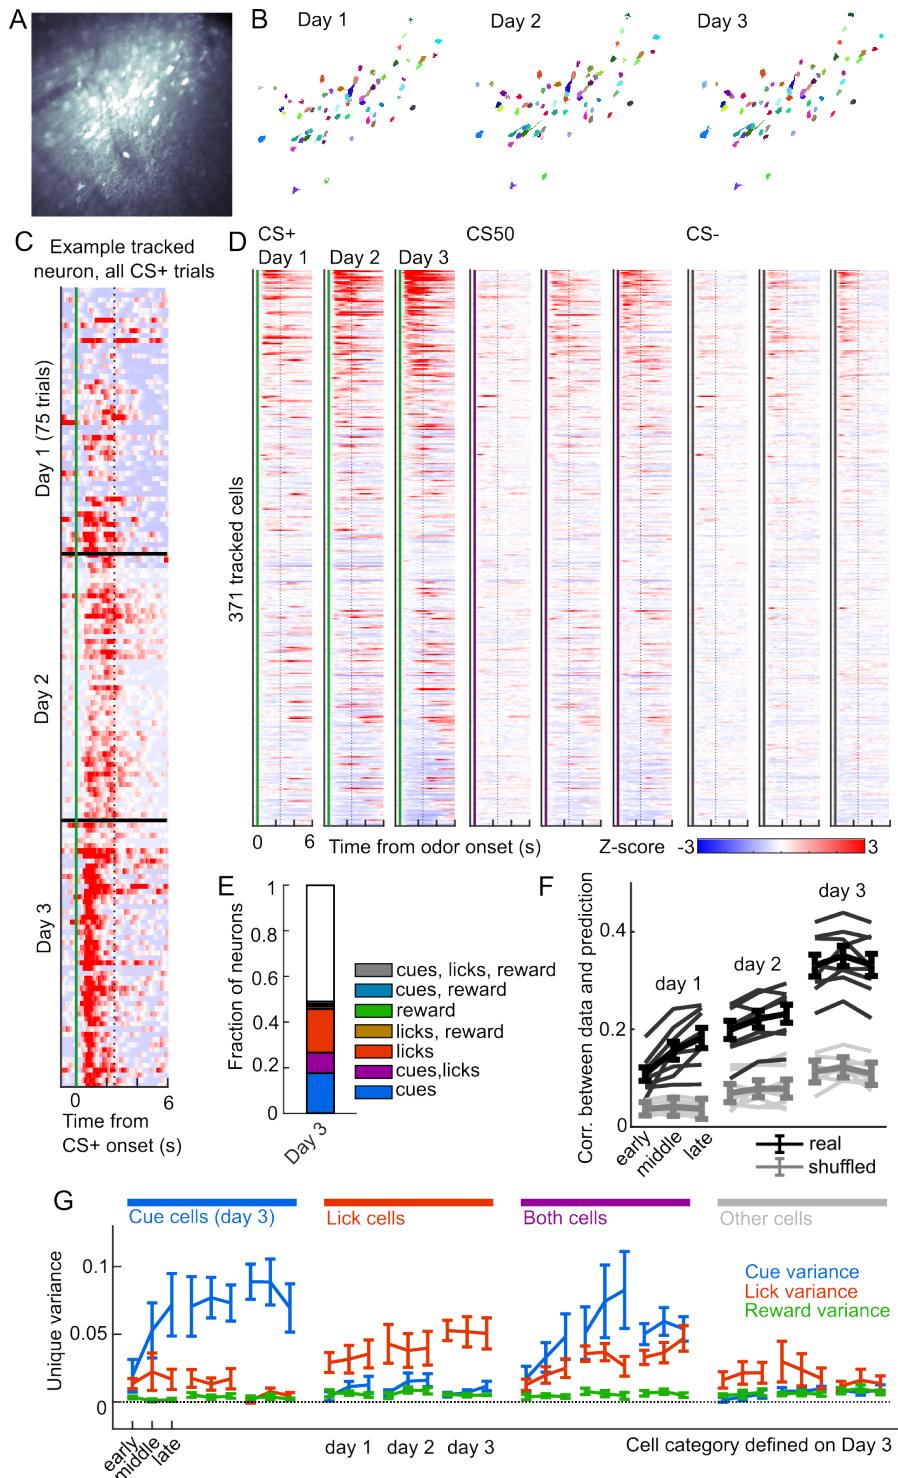


Figure 6: Cue and lick coding is stable across days.

Figure 6. Cue and lick coding is stable across days.

- (A) Standard deviation fluorescence from example imaging plane.
- (B) Masks (randomly colored) for all tracked neurons from this imaging plane.
- (C) Deconvolved spike rate on every CS+ trial from all three days of odor set 1 for an example neuron. Vertical dashed line is reward delivery. Color axis as in (D).
- (D) Normalized deconvolved spike rate for all tracked neurons on all three days of odor set 1.
- (E) Fraction of tracked neurons coding cues, licks, rewards, and their combinations on day 3.
- (F) Model performance when using day 3 models to predict the activity of individual neurons across odor set 1 training, plotted as mean (+/- SEM) correlation between true and predicted activity across mice. Thin lines are individual mice. Performance was greater than shuffled data at all time points ($p < 0.0001$) except early day 1 ($p = 0.21$, Bonferroni-corrected, $n = 8$ mice).
- (G) Mean (+/- SEM) unique cue, lick, and reward variance for cells classified as coding cues, licks, both, or neither on day 3. Day 3 cue cells had increased cue variance on day 2 ($p < 10^{-7}$, see Methods) and 1 ($p < 0.03$) relative to lick and reward variance. Same pattern for lick cells on day 2 ($p < 0.0001$) and day 1 ($p < 0.01$).

A subset of mice ($n = 5$) also learned a second odor set, presented on separate days. Activity was very similar for both odor sets, evident across the entire imaging plane (Fig. 7A), for individual tracked neurons ($n = 594$, 81-153 per mouse) (Fig. S2B), and for kernel regression classification of these neurons (Fig. 7B). Notably, odor set 1 models performed similarly well at predicting both odor set 1 and odor set 2 activity (Fig. 7C). Moreover, cue, lick and both neurons maintained their unique variance preference across odor sets (Fig. 7D). Finally, to investigate the presence of value coding across odor sets over separate days, we fit tracked cue neurons with the value model and its shuffles. Even with odor sets imaged on separate days, we again found that the value model was the best model for a sizable fraction (27%) of cue neurons, demonstrating that value coding is conserved across stimulus sets and days (Fig. 7E).

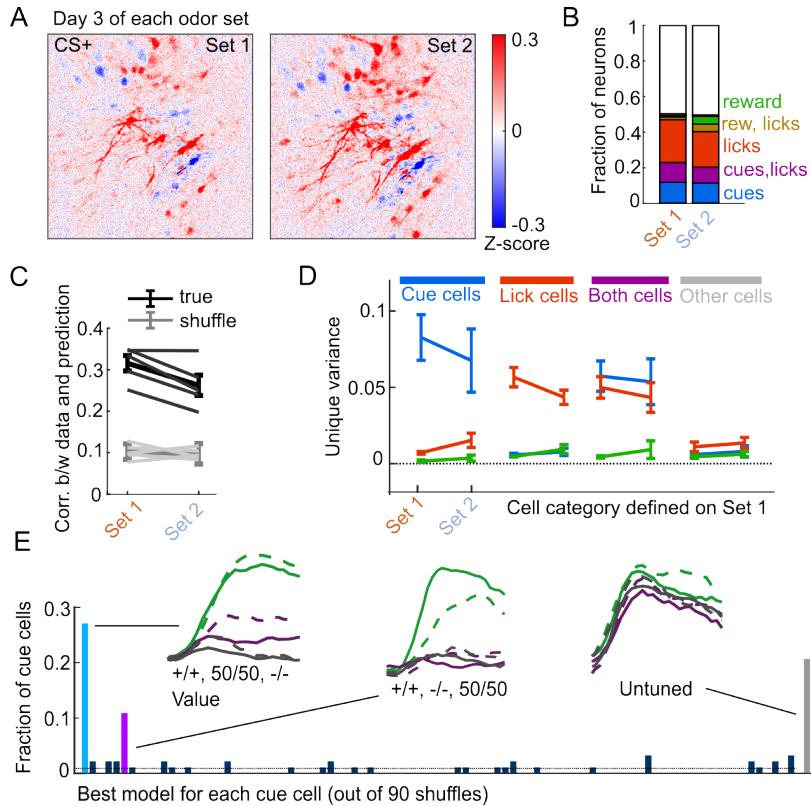


Figure 7: Cue and lick coding in separately trained odor sets.

Figure 7. Cue and lick coding in separately trained odor sets.

- (A) Normalized activity of all pixels in the imaging plane following CS+ presentation on day 3 of each odor set.
- (B) Fraction of neurons coding for cues, licks, rewards, and their combinations for day 3 of each odor set.
- (C) Mean(+/- SEM, across mice) correlation between activity predicted by odor set 1 models and true data, for real (black) and trial shuffled (gray) activity. Thin lines are individual mice. $F(1, 16) = 3.2, p = 0.09$ for main effect of odor set, $F(1, 16) = 135, p < 10^{-8}$ for main effect of shuffle, $F(1, 16) = 2.2, p = 0.16$ for interaction, $n = 5$ mice, two-way ANOVA.
- (D) Mean(+/- SEM, across mice) unique cue, lick, and reward variance for cells classified as coding cues, licks, both, or neither for odor set 1. For each category, odor set 1 unique variance preference was maintained for odor set 2 ($p < 0.04$) except for both cells, for which lick and reward variance were not different in odor set 2 ($p = 0.22$, Bonferroni-corrected, $n = 5$ mice).
- (E) Distribution of best model fits across all cue neurons, and the first principal component of the activity of all neurons best fit by the top 3 models.

Discussion

Our experiments assessed how coding for reward-predicting cues and reward-seeking actions differed across brain regions and across multiple days of training. We found coding for cues and licks in all regions we sampled, but their proportions varied in a graded way across those regions. In contrast to regional differences in the proportion of cue-responsive cells, cue value coding was similarly represented in all regions. Coding for cue value was robust, occurring with far greater frequency than chance, and, in a subset of neurons, incorporated the recent reward history. Cue coding was established within the first day of training and neurons encoding cues or licks maintained their coding preference across multiple days of the task. These results demonstrate a lack of regional specialization in value coding and the stability of cue and lick codes in PFC.

Graded cue and lick coding across regions.

We found robust and separable coding for licks and cues (and combined coding of both) in all regions using electrophysiology and in PL using calcium imaging. The widespread presence of lick coding is consistent with recent reports of distributed movement and action coding (Musall et al., 2019; Steinmetz et al., 2019; Stringer et al., 2019); however, we saw sizable differences in the amount of lick coding across recorded regions. Notably, ALM had the greatest number of lick neurons, as well as the fewest cue neurons, perhaps reflecting its specialized role in the preparation and execution of licking behavior (Chen et al., 2017). Conversely, the olfactory cortical regions DP, TTd, and AON had the most cue neurons (especially non-value coding cue neurons), suggesting a role in early odor identification and processing (Mori and Sakano, 2021). PFC subregions balanced lick and cue coding, consistent with their proposed roles as association areas (Klein-Flügge et al., 2022; Miller and Cohen, 2001), but there was variability within PFC as well. In particular, ORB had a greater fraction of cue cells than any other subregions,

consistent with its known dense inputs from the olfactory system (Ekstrand et al., 2001; Price, 1985; Price et al., 1991). Thus, our results establish that the neural correlates of this Pavlovian conditioned behavior consist of a gradient of cue and response coding rather than segmentation of sensory and motor responses.

Widespread value signaling.

Value signals can take on many forms and occur throughout task epochs. In our experiments, we focused on the predicted value associated with each conditioned stimulus, which is crucial for understanding how predictive stimuli produce motivated behavior (Berridge, 2004). Most comparisons of single neuron stimulus-value signaling across PFC have been conducted in primates. These studies have found neurons correlated with stimulus-predicted value in many subregions, with the strongest representations typically in ORB (Hunt et al., 2018; Kennerley et al., 2009; Roesch and Olson, 2004; Sallet et al., 2007). In rodents, there is also a rich history of studying value signaling in ORB (Kuwabara et al., 2020; Namboodiri et al., 2019; Schoenbaum et al., 2003; Stalnaker et al., 2014; Sul et al., 2010; van Duuren et al., 2009; Wang et al., 2020a), but there have been many reports of value-like signals in frontal cortical regions beyond ORB, as well (Allen et al., 2019; Kondo and Matsuzaki, 2021; Otis et al., 2017; Wang et al., 2020a). In our present experiment, we sought to expand upon the results from the rodent literature by separating cue activity from licking, which can track with value and may confound interpretation of the signal, by including more than two cue types, which provided a rich space to assess value coding, and by sampling from many frontal regions in the same experiment.

When considering the number of neurons responsive to cues rather than licks, our data confirmed the importance of ORB, which has more cue-responsive neurons than motor and other prefrontal regions. However, by analyzing the activity of cue-responsive neurons across all 6 odors predicting varying probabilities of reward, we were able to separate out neurons coding

value from other neurons, which included a population that had consistent responses for odors with the same associated reward probability (trial type) but activity that did not scale according to probability, consistent with the possibility of nonlinear value coding (Enel et al., 2021). When only considering cue neurons with linear coding of value, the distribution was even across regions. One consequence of a widely distributed value signal is that manipulating only one subregion would be less likely to fully disrupt value representations, which is consistent with the results of studies comparing functional manipulations of different PFC subregions (Chudasama and Robbins, 2003; Dalton et al., 2016; St. Onge and Floresco, 2010; Verharen et al., 2020; Wang et al., 2020a). Different subregional impacts on behavior may reveal biases in how the value signal in each region contributes to reward-related behaviors, for instance during learning or expression of a reward association (Namboodiri et al., 2019; Otis et al., 2017; Wang et al., 2020a). A related interpretation is that, in this task, there may be other properties that correlate with cue value, and the homogeneous value representation we observed across regions masks regional differences in tuning to these other correlated features, such as motivation (Roesch and Olson, 2004) and a host of related concepts, including salience, uncertainty, vigor, and arousal (Hayden and Niv, 2021; Stalnaker et al., 2015; Zhou et al., 2021), which can have different contributions to behavior. This interpretation is consistent with broader views that observations of ‘value’ signals are often misconstrued (Zhou et al., 2021) and that pure abstract value may not be encoded in the brain at all (Hayden and Niv, 2021). Although the identification of value in our task was robust to three levels of reward probability across two stimulus sets, the fact that this signal was widespread contributes to the case for revisiting the definition and interpretation of value to better understand regional specialization.

In our analysis, we uncovered a distinction between neurons encoding the overall value of cues and those with value representations that incorporated the recent reward history. Neurons with history effects were rarer but also widespread. These neurons may have a more direct

impact on behavioral output in this task, because the lick rate also incorporated recent reward history. Notably, the impact of reward history on these neurons was noticeable even prior to cue onset, consistent with a previously proposed mechanism for persistent value representations encoded in the baseline firing rates of PFC neurons (Bari et al., 2019).

Given the presence of value coding in olfactory cortex, the question remains of where odor information is first transformed into a value signal. In fact, there have been multiple reports of some association-related modification of odor representations as early as the olfactory bulb (Chu et al., 2016; Doucette et al., 2011; Koldaeva et al., 2019; Li et al., 2015). Considering the prevalence of value and non-value trial type coding we observed in AON, DP, and TTd, perhaps these regions are a crucial first step in processing and amplifying task-related input from the olfactory bulb. Because they provide input to PFC (Bhattarai et al., 2021; Igarashi et al., 2012), they may be an important source of the odor coding we observed there. Previous recordings in AON, DP, and TTd were in anesthetized rodents (Cousens, 2020; Kikuta et al., 2008; Lei et al., 2006; Tsuji et al., 2019); as the first recordings in awake behaving animals, our results bring these regions into focus for future work on the transformation of odor information into task-relevant coding.

Stability of PFC codes.

Previous reports have observed drifting representations in PFC across time (Hyman et al., 2012; Malagon-Vina et al., 2018), and there is compelling evidence that odor representations in piriform drift over weeks when odors are experienced infrequently (Schoonover et al., 2021). On the other hand, it has been shown that coding for odor association is stable in ORB and PL, and that coding for odor identity is stable in piriform (Wang et al., 2020a), with similar findings for auditory Pavlovian cue encoding in PL (Grant et al., 2021; Otis et al., 2017) and ORB (Namboodiri et al., 2019). We were able to expand upon these data in PL by identifying both cue

and lick coding and showing separable, stable coding of cues and licks across days and across sets of odors trained on separate days. We were also able to detect value coding common to two stimulus sets presented on separate days. This consistency in cue and lick representations indicates that PL serves as a reliable source of information about cue associations and licking during reward seeking tasks, perhaps contrasting with other representations in PFC (Hyman et al., 2012; Malagon-Vina et al., 2018). Interestingly, the presence of lick, but not cue coding at the very beginning of day 1 of training suggests that lick cells in PL are not specific to the task but that cue cells are specific to the learned cue-reward associations.

Overall, our work emphasizes the importance of evaluating regional specialization of neural encoding with systematic recordings in many regions using the same task. Future work will clarify whether cue value is similarly widely represented in other reward-seeking settings and whether there are regional differences in the function of the value signal.

Acknowledgments

Thank you to Vijay Namboodiri and Charles Zhou for assistance with the imaging. Thank you to Noam Roth for the spike sorting quality control metrics. This work was supported by National Institutes of Health grants F32DA053714 (D.O.), F31DA053706 (M.H.), T32DK007247 (A.B.), R37DA032750 (G.S.), and P30DA048736 (G.S.), a UW Center for the Neurobiology of Addiction, Pain, and Emotion 2-photon pilot project grant (D.O.), a Klingenstein-Simons Fellowship in Neuroscience (N.S.), and the Pew Biomedical Scholars Program (N.S.).

Author contributions

Conceptualization: D.O., M.H., N.S., G.S.; data collection: D.O., M.H., A.B.; data curation: D.O., A.B.; data interpretation: D.O., M.H., N.S., A.B., G.S.; formal analysis: D.O.; visualization: D.O., M.H., A.B.; writing - original draft: D.O., M.H., A.B.; writing - review & editing:

D.O., M.H., A.B., N.S., G.S.; funding acquisition: N.S., G.S..

Declaration of interests

The authors declare no competing interests.

Data and code availability

The data and code for this manuscript are publicly available at

<https://doi.org/10.5281/zenodo.6686927> (Ottenheimer et al., 2022).

Methods

Subjects.

Subjects ($n = 5$ for electrophysiology, $n = 8$ for imaging) were male and female C57BL/6 mice single-housed on a 12hr light/dark cycle and aged 12-28 weeks at the time of recordings. Imaging experiments were performed during the dark cycle, electrophysiology during light cycle. Mice were given free access to food in their home cages for the duration of the experiment. Mice were water restricted for the duration of the experiments and maintained around 85% of their baseline weight (Guo et al., 2014a). All experimental procedures were performed in strict accordance with protocols approved by the Animal Care and Use Committee at the University of Washington.

Surgical procedures.

Mice were anesthetized with isoflurane (5%) and maintained under anesthesia for the duration of the surgery (1-2%). Mice received injections of carprofen (5 mg/kg) prior to incision.

Electrophysiology. A brief (1 h) initial surgery was performed, as described in (Guo et al., 2014b; Steinmetz et al., 2017, 2019), to implant a steel headplate (approximately $15 \times 3 \times 0.5$ mm, 1 g) and a 3D-printed recording chamber that exposed the skull for subsequent craniotomies. Briefly, an incision was made around the circumference of the dorsal surface of the skull, from the interparietal bone to the frontonasal suture. The skin and periosteum were removed to expose the dorsal surface of the skull. Skull yaw, pitch, and roll were leveled, and exposed bone was texturized with a brief application of green activator (Super-Bond C&B, Sun Medical). The incised skin was secured around the circumference of exposed skull with application of cyanoacrylate (VetBond; World Precision Instruments), and the chamber was attached to the skull with L-type radiopaque polymer (Super-Bond C&B). A thin layer of cyanoacrylate was applied to the skull inside the chamber and allowed to dry. Multiple (2-4) thin layers of UV-curing optical glue (Norland Optical Adhesives #81, Norland Products) were applied inside the chamber to cover the entire exposed surface of the skull and cured with UV light. The headplate was attached to the skull over the interparietal bone posterior to the chamber with Super-Bond polymer, and more polymer was applied around the headplate and chamber. A second brief (15-30 min) surgery was conducted to perform craniotomies for probe insertion. Briefly, following induction of anesthesia a small (2×1.5 mm (w \times h) craniotomy was made over frontal cortex ($+2.5 - 1$ mm AP, $\pm 2.5 - 0.3$ mm ML) with a handheld dental drill. The craniotomy was covered with silicone gel (DOWSIL 3-4680) and the recording chamber was covered with a 3D-printed lid sealed with Kwik-Cast elastomer for protection.

Imaging. A GRIN lens and metal headcap were implanted following previously described procedures (Namboodiri et al., 2019) with the following modifications. In most mice, once the dura was removed from the craniotomy, we injected, $0.5 \mu\text{L}$ of virus containing the GCaMP gene construct (AAVDJ-CamKIIa-GCaMP6s, 5.3×10^{12} viral particles/mL from UNC Vector core lot AV6364) using a glass pipette microinjector (Nanoject II) targeted at Bregma $+1.94$ mm AP, 0.3 and 1.2 mm ML, -2 mm DV. Ten minutes elapsed before microinjector withdrawal to allow virus to diffuse away from the infusion site. Then, mice were implanted with a 1×4 mm GRIN lens (Inscopix) aimed at $+1.94$ mm AP, 0.3 and 1.2 mm ML, -1.8 mm DV. A subset of mice did not receive viral injections; instead, a lens with the imaging face coated $1 \mu\text{L}$ of the GCaMP6s virus mixed with 5 percent aqueous silk fibroin solution (Jackman et al., 2018) was implanted at the same coordinate. GCaMP expression and transients were similar in both preparations. Mice were allowed to recover for at least 5 weeks before experiments began.

Behavioral training.

Mice were headfixed during training and recording sessions using either a headring (imaging experiments) or headbar (electrophysiology experiments). After initial habituation to head fixation, mice were first trained to lick for $2.5\mu L$ rewards of 10% sucrose solution, delivered every 8 - 12 s through a miniature inert liquid valve (Parker 003-0257-900). After 4 - 5 days of lick training, mice experienced their first odor exposure (without reward delivery). Odors were delivered for a total of 1.5 s using a 4-channel olfactometer (Aurora 206A) with 10% odor flow rate and 800 SCCM overall flow rate of medical air. Odors were randomly assigned to sets and cue identities, counterbalanced across mice. Odors were -carvone, -limonene, alpha-pinene, butanol, benzaldehyde, and geranyl acetate (Sigma Aldrich 124931, 218367, 147524, 281549, 418099, 173495, respectively), selected because of they are of neutral valence to naive mice (Devore et al., 2013; Saraiva et al., 2016). Odors were diluted 1:10 in mineral oil and 10 μL was pipetted onto filter paper within the odor delivery vials (Thermo Fisher SS246-0040) prior to each session. Airflow was constant onto the mouse's nose throughout the session and switched from clean air to scented air for the duration of odor delivery on each trial.

On days 1 - 2 of Pavlovian conditioning, mice received 50-75 trials each of 3 odor cues, followed by reward on 100% (CS+), 50% (CS50), or 0% (CS-) of trials, 2.5 s following odor onset, with 8 - 12 s between odor presentations. Mice then received training days 1 - 2 with a second odor set with three new odors. For electrophysiology experiments, the odors were subsequently presented in the same sessions in 6 blocks of 51 trials. Odor set order alternated and was counterbalanced across days. For imaging experiments, mice received day 3 of odor set 1 and then day 3 of odor set 2. An additional 3 imaging mice were only trained on one odor set.

Electrophysiological recording and spike sorting.

During recording sessions, mice were headfixed. Recordings were made using either Neuropixels 1.0 or Neuropixels 2.0 electrode arrays (Jun et al., 2017; Steinmetz et al., 2021), which have 384 selectable recording sites. Recordings were made with either 1.0 (1 shank, 960 sites, 2.1 (1 shank, 1280 sites) or 2.4 (4 shanks, 5120 sites) probes, depending on the regions of interest. Probes were mounted to a dovetail and affixed to a steel rod held by a micromanipulator (uMP-4, Sensapex Inc.). To allow later track localization, probes were coated with a solution of DiI (ThermoFisher Vybrant V22888) by holding 2 μl in a droplet on the end of a micropipette and painting the probe shank. In each session, one or two probes were advanced through the Duragel covering the craniotomy over frontal cortex, then advanced to their final position at approximately $3 \mu m s^{-1}$. Electrodes were allowed to settle for around 15 min before starting recording. Recordings were made in internal reference mode using the 'tip' reference site, with a 30 kHz sampling rate. Recordings were repeated at different locations on each of multiple subsequent days, performing an additional craniotomy over contralateral frontal cortex. The resulting data were automatically spike sorted with Kilosort2.5 and Kilosort3 (<https://github.com/MouseLand/Kilosort>). Extracellular voltage traces were preprocessed with common-average referencing by subtracting each channel's median to remove baseline offsets, then subtracting the median across all channels at each time point to remove artifacts. Sorted units were curated using automated quality control (International-Brain-Laboratory et al., 2022): exclusions were based on spike floor violations (the estimated proportion of spikes that were missed because they fell below the noise level of the recording, i.e. esatimed false negative rate), and refractory period violations (the estimated proportion of spikes which did not arise from the primary neuron, i/e/ the estimate false positive rate due

to contamination, with a 10% cutoff). Quality control accuracy was assessed by manually reviewing a subset of the data using the phy GUI (<https://github.com/kwikteam/phy>). Because Kilosort2.5 and Kilosort3 use different clustering algorithms that can be advantageous for different types of recordings (stability, region, number of channels), for each session, we used units sorted with either Kilosort2.5 or Kilosort3 depending on which yielded the greatest number of high quality units for that session. Brain regions were only included for subsequent analysis if there were recordings from at least three subjects and in total over 100 neurons in the region. When we analyzed all of motor cortex together, we included ALM and MOs neurons. When we analyzed all of olfactory cortex, we included DP, TTd, AON, and other neurons in PIR, EPd, and OLF. We relabeled PIR and EPd as OLF because there were not enough neurons to analyze them as separate regions.

Imaging and ROI extraction.

During imaging sessions, mice were headfixed and positioned under the 2-photon microscope (Bruker Ultima2P Plus) using a 20x air objective (Olympus LCPLN20XIR). A Spectra-Physics InSight X3 tuned to 920nm was used to excite GCaMP6s through the GRIN lens. Synchronization of odor and 10 percent sucrose delivery, lick behavior recordings, and 2-photon recordings was achieved with custom Arduino code. After recording, raw TIF files were imported into suite2p (<https://github.com/MouseLand/suite2p>). We used their registration, region-of-interest (ROI) extraction, and spike deconvolution algorithms, inputting a decay factor of $\tau = 1.3$ to reflect the dynamics of GCaMP6s, and manually reviewed putative neuron ROIs for appropriate morphology and dynamics. To find changes in activity across the entire imaging plane, found the mean pixel intensity for frames in the time of interest (2 to 2.5 s from CS+), subtracted the mean intensity of each pixel prior to cue onset (-2 to 0 s from all cues), and divided by the standard deviation for each pixel across those frames prior to cue onset.

Histology.

Animals were anesthetized with pentobarbital or isoflurane. Mice were perfused intracardially with 0.9% saline followed by 4% paraformaldehyde (PFA).

Electrophysiology. Brains were extracted immediately following perfusion and post-fixed in 4% paraformaldehyde for 24 h. In preparation for light sheet imaging brains were cleared using organic solvents following the 3DISCO protocol (Ertürk et al., 2012) (<https://idisco.info/>), with some modification. Briefly, on day 1 brains were washed 3× in PBS, then dehydrated in a series of increasing MeOH concentrations (20%, 40%, 60%, 80%, 100%, 100%; 1-h each) and incubated overnight for lipid extraction in 66% dichloromethane (DCM) in MeOH. On day 2 the brains were washed twice for 1-h each in 100% MeOH, then bleached in 5% H₂O₂ in MeOH at 4°C overnight. On day 3 brains were washed 2× in 100% MeOH, then final lipid extraction was accomplished in a series of DCM incubations (3-h in 66% DCM in MeOH, 2× 100% DCM for 15 min each) before immersion in dibenzyl ether (DBE) for refractive index matching. Brains were imaged on a light sheet microscope (LaVision Biotec UltraScope II) 2-7 days after clearing. Brains were immersed in DBE in the imaging well secured in the horizontal position, and illuminated by a single light sheet (100% width, 4 μ m thick) from the right. Images were collected through the 2X objective at 1X magnification, from the dorsal surface of the brain to the ventral surface in 10 μ m steps in 488 nm (autofluorescence, 30% power) and 594 nm (DiI, 2-10% power) excitation channels. The approximately 1000 raw TIF images were compiled into a single multi-image file with 10 μ m

voxels, then spatially downsampled to $25 \mu\text{m}$ voxels for transformation to the Allen common-coordinate framework (CCF) volume (Wang et al., 2020b) using the Elastix algorithm (Shamonin et al., 2014). Transformed volumes were used to track fluorescent probe tract locations in CCF using Lasagna (<https://github.com/SainsburyWellcomeCentre/lasagna>), generating a series of CCF pixel coordinates for points along each probe tract. CCF pixel coordinates (origin front, top, left) were transformed to bregma coordinates ($x == \text{ML}$, $y == \text{AP}$ and $z == \text{DV}$) in preparation for final integration with electrophysiology recordings using the International Brain Lab electrophysiology GUI (Faulkner M, Ephys Atlas GUI; 2020. <https://github.com/int-brain-lab/iblapps/tree/master/atlaselectrophysiology>). For recording alignment, sorted spikes and RMS voltage on each channel were displayed spatially in relation to the estimated channel locations in Atlas space from the tracked probe. The recording sites were then aligned to the Atlas by manually identifying a warping such that recording sites were best fit to the electrophysiological characteristics of the brain regions (e.g. matching location of ventricles or white matter tracts with low firing activity bands). This procedure has been estimated to have $70 \mu\text{m}$ error (Liu et al., 2021; Steinmetz et al., 2019). Brain regions were then ascribed to each unit based on location of the recording site with maximum waveform. We additionally assigned MOs neurons to anterolateral motor cortex (ALM) if they were within a 0.75 mm radius of 2.5 mm AP , 1.5 mm ML (Chen et al., 2017).

Imaging. Following perfusion, intact heads were left in PFA for an additional week before brain extraction. Brains were then sliced on a Leica Vibratome (VT1000S) at $70 \mu\text{m}$ before mounting and nuclear staining via Fluoroshield with DAPI (Sigma-Aldrich F6057-20ML). Slices with GRIN lens tracks were then imaged on a Zeiss Axio Imager M2 Upright Trinocular Phase Contrast Fluorescence Microscope with ApoTome. The resulting images were manually aligned to the Allen Brain Atlas to reconstruct the location of each GRIN lens.

Neuron tracking.

To identify the same neurons across imaging sessions, we used two approaches. To track neurons across the two odor sets on day 3, we concatenated the TIF files from each session and extracted ROIs simultaneously. To track neurons across days 1 - 3, we manually identified ROIs from the ROI masks outputted by suite2p. We linked the ROIs using a custom Python script that permitted selection of the same ROI across the 3 imaging planes using OpenCV and saved the coordinates on each day. The tracking results across days 1 - 3 from one of the mice is displayed in Fig. 6B.

Behavioral analysis.

For electrophysiology experiments, eye and face movements were monitored by illuminating the subject with infrared light (850 nm, CMVision IR30). The right eye was monitored with a camera (FLIR CM3-U3-13Y3M-CS) fitted with a zoom lens (Thorlabs MVL7000) and long-pass filter (Thorlabs FEL0750), recording at 70 Hz. Face movements were monitored with another camera (same model with a different lens, Thorlabs MVL16M23) directed at a $2 \times 2 \text{ cm}$ mirror reflecting the left side of the face, recording at 70 Hz. Licks were detected by thresholding the average intensity of an ROI centered between the lips and the lick spout, calculated for every frame. For imaging experiments, licks were detected with a capacitance sensor (MPR121, Adafruit Industries) connected to an Arduino board. To determine the impact of cues and previous outcomes on anticipatory licking, we fit a linear model on all electrophysiology sessions simultaneously (and for each mouse). We predicted the number of licks 0 to 2.5 s from odor onset using cue identity, outcomes on previous 10 trials, outcomes on previous 10 of that

cue type, and total number of presentations of that cue type so far (to account for cue-specific satiety) using ‘fitlm’ in MATLAB. When dividing sessions into ‘early’ and ‘late’, we used the first 60 and last 60 trials of the session. When dividing sessions into thirds for the GLM (‘early’, ‘middle’, ‘late’), we used even splits of trials into thirds.

PSTH creation.

Peri-stimulus time histograms (PSTHs) were constructed using 0.1s bins surrounding cue onset.

Electrophysiology. Neuron spike times were first binned in to 0.02s bins and smoothed with a half-normal causal filter ($\sigma = 300$ ms) across 50 bins. PSTHs were then constructed in 0.1s bins surrounding each cue onset. Each bin of the PSTH was z-scored by subtracting the mean firing rate and dividing the standard deviation across the 0.1s bins in the 2s before all trials. When splitting responses by polarity (above/below baseline, Figs. 2B, S4B), we used half of trials to determine polarity and plotted the mean across the other half of trials.

Imaging. Frames were collected at 30Hz with 2-frame averaging, so the fluorescence for each neuron and the estimated deconvolved spiking were collected at 15Hz. We interpolated the smoothing filter from the electrophysiology analysis (which was calculated at 50Hz) and applied it to the deconvolved spiking traces. We then constructed PSTHs in 0.1s bins surrounding each cue onset and z-scored (same as electrophysiology).

Licks. Licking PSTHS were constructed in 0.1s bins surrounding cue onset. Each trial was then smoothed with a half-normal filter ($\sigma = 800$ ms) across the previous (but not upcoming) 10 bins. For the GLM, the lick rate was calculated across the whole session by first counting licks in either the 0.02s (electrophysiology) or 15Hz (imaging) bins, smoothed with a half-normal filter over 25 previous bins, and then converted to 0.1s bins relative to each cue.

Kernel regression.

To identify coding for cues, licks, and rewards in individual neurons, we fit reduced rank kernel-based linear model (Steinmetz et al., 2019).

Data preparation. The discretized firing rates $f_n(t)$ for each neuron n were calculated as described above for PSTH creation. We used the activity -1 to 6.5 s from each cue onset for our GLM analysis.

Predictor matrix. The model included predictor kernels for cues (CS+, CS50, and CS- for each odor set, as relevant), licks (individual licks, lick bout start, and lick rate), and reward (initiation of consummatory bout). The cue kernels were supported over the window 0 to 5s relative to stimulus onset. The lick predictor kernels were supported from -0.3 to 0.3 s relative to each lick, from -0.3 to 2 s relative to lick bout start, and lick rate was shifted from -0.4 to 0.6 s in 0.2 s increments from original rate. The reward kernel was supported 0 to 4s relative to first lick following reward delivery. For electrophysiology experiments, the model also included 6 constants that identified the block number, accounting for changes in firing rate across blocks. For each kernel to be fit we constructed a Toeplitz predictor matrix of size $T \times l$, in which T is the total number of time bins and l is the number of lags required for the kernel. The predictor matrix contains diagonal stripes starting each time an event occurs and 0 otherwise. The predictor matrices were horizontally concatenated to yield a global prediction matrix \mathbf{P} of size $T \times L$ containing all predictor kernels. Rate vectors of all N neurons were horizontally concatenated to form \mathbf{F} , a $T \times N$ matrix.

Reduced-rank regression. To prevent noisy and overfit kernels we implemented reduced-rank regression (Steinmetz et al., 2019), which allows regularized estimation by factorizing the kernel matrix \mathbf{K} into the product of a $L \times r$ matrix \mathbf{B} and a $r \times N$ matrix \mathbf{W} , minimizing the

total error: $E = \|\mathbf{F} - \mathbf{PBW}\|^2$. The $T \times r$ matrix \mathbf{PB} may be considered as a set of ordered temporal basis functions, which can be linearly combined to estimate the neuron's firing rate over the whole training set, resulting in the best possible prediction from any rank r matrix. To estimate each neuron's kernel functions we generated the reduced rank predictor matrix \mathbf{PB} for $r = 20$, estimated the weights \mathbf{w}_n to minimize the squared error $E_n = |\mathbf{f}_n - \mathbf{PBw}_n|^2$ with lasso regularization (using the MATLAB function 'lassoglm') with parameters $\alpha = 0.5$ and $\lambda = [0.001, 0.005, 0.01, 0.02, 0.05, 0.1, 0.2, 0.5]$, using 4-fold cross-validation to determine the optimal value for λ for each neuron. The kernel functions for neuron n were then unpacked from the L-length vector obtained by multiplying the first $r = 20$ columns of \mathbf{B} by \mathbf{w}_n .

Predictor unique contributions. To assess the importance of each group of kernels for predicting a neuron's activity we first fit the activity of each neuron using the reduced-rank regression procedure, then fit the model (with 4-fold cross-validation) again excluding the kernels belonging to the predictor to be tested (cues, licks, rewards). If the difference in variance explained between the full and held-out model was $> 2\%$, and the total variance explained by the full model was $> 2\%$, the neuron was deemed selective for those predictors (Steinmetz et al., 2019). We validated this cutoff by randomly shuffling the onset time of each trial, which resulted in only 1/5332 neurons with $> 2\%$ unique variance explained by any variable.

Training and testing on other time points. In the imaging experiments, we also fit the models independently to each session third (early, middle, late) of days 1-3 with odor set 1 to determine how fits and unique contributions evolved over time. To assess coding stability of individual neurons in the imaging experiment, we used the kernels resulting from fitting the full model on day 3 and the predictors from each session third to predict neural activity at those time points. We assessed the the accuracy of the prediction by correlating it with the true activity and comparing that to the correlation with trial-shuffled data. We also did this with models trained on day 3 odor set 1 and tested on day 3 odor set 1 and 2.

Cue coding models. To assess cue coding schemes, we fit a new set of models focusing on a more restricted time window (-1 to 2.5 s from cue onset) using only cues and licks as predictors. Cue and lick neurons were identified as before, and subsequent cue characterization was performed on neurons with only a unique contribution of cues. To identify value coding among cue neurons, we fit an additional version of the kernel model with only one cue kernel that scaled according to cue value. We estimated cue value for each cue type in each session by finding the mean value predicted by the lick linear model (described in section 'Behavioral analysis'), using cue type, 10 previous outcomes, and 10 previous cue outcomes as predictors. These values were used to scale the height of cue kernel for each trial type and were approximately 0.05, 0.35, and 0.5 for CS-, CS50, and CS+, respectively. We also fit 89 additional versions of this model consisting of all permutations of cue value assignment to the 6 odors. We further fit each neuron with an 'untuned' model that had the same value for all cues. These models are visualized in Fig. S6. After removing neurons best fit with the untuned model (signifying non-specific odor responses), we classified the remaining neurons according to which of the 90 models best fit an individual neuron. For neurons where the value model was the best, we also fit another version of the model where instead of scaling the cue kernel by mean value, we scaled it by the trial-by-trial prediction of value from the lick linear model, which we called the history value model. For neurons better fit by the history model, we also fit 1000 additional models with shuffled trial values within each cue. The median percentage of shuffles that the true history model improved upon was 96.5% across all neurons for which the history model improved over the mean model. All value neurons best fit by the history model and improving over $> 65\%$ shuffles ($< 8\%$ of value neurons better fit by the history model than the mean model were below 65%) were classified as history coding. We also fit the value model and its 89 shuffles to the neurons imaged on separate days (Fig. 7E), concatenating the data from each

odor set and adding a constant for each day to account for day differences.

Principal component analysis.

To visualize the dominant firing pattern of PL neurons (Fig. S2), and of value and trial type cells (Fig. 3), irrespective of direction (excitation or inhibition), we performed principal component analysis ('PCA' in MATLAB) on the concatenated PSTHs across all 6 cues for the neurons of interest, with each neuron's activity normalized by peak modulation so that each neuron's concatenated PSTH peaked at -1 or 1. We then plotted the score of the first components. To visualize the value dynamics of value neurons from each region (Fig. 3), we performed PCA as before on this subset and plotted the component most related to value (first component for all regions except ACA, which was second). To quantify value dynamics within PCA space for each region, we randomly selected 10 neurons from each region (or 60 neurons from region group) and performed PCA as before but using only the 3 cues from odor set 1. We then projected the activity of those neurons during odor set 2 cues onto the first 3 components. We then calculated the distance between baseline (0 to 0.5 s) and cue (1 to 2.5 s) activity for each odor set 2 cue in this 3-dimensional space as a fraction of maximum total distance spanned across all time points. We performed this 5000 times with different selections of neurons to estimate the distribution of differences between cue distances.

Cue coding dimension.

To project population activity onto the coding dimensions separating CS- activity from CS+ and CS50 activity, respectively, we adapted an approach from Li et al. (2016). We first normalized the odor set 1 PSTH activity of each neuron by dividing all of that neurons' activity by its peak activity across odor set 1 PSTHs. This prevented neurons with particularly large z-score values from dominating the dimension. We then used half of odor set 1 trials to define the dimensions. For each neuron, we found the 0.5s bin in the range 0 to 2.5s from cue onset that maximally separated CS- activity from CS+ or CS50 activity, respectively. The difference between the normalized CS+ and CS- activity for all neurons in their preferred bins comprised a vector defining the coding dimension for each group of neurons of interest. We then multiplied that difference vector by the original z-score values of each neuron in their peak bins to find the values of peak CS+ and CS- coding; we used these values to transform the data onto a 0 to 1 scale for CS- to CS+ activity. To find population activity along that dimension at each moment in CS+, CS50, and CS- trials, we multiplied the activity of all neurons in each 0.1s bin of the other half of odor set 1 trials (z-score) by the difference vector and used the same conversion to 0 to 1 scale ('same odor set'). We also multiplied the activity of neurons for cues in the other odor set by the difference vector ('other odor set'). We repeated the same process for CS- and CS50 activity. To find the baseline distance from CS-, we bootstrapped (5000 iterations of selecting which neurons to include in the analysis, with replacement) the euclidean distance of the average baseline values for all 3 cues from odor set 1 from the CS- position (0, 0). To find the angle between the CS+ and CS50 projections, we bootstrapped the vectors that connected baseline activity to peak activity of CS50 and CS+ along the CS- / CS+ and CS- / CS50 axes and found the angle between these vectors. To find population activity along the CS+ / CS- dimension at each moment for CS50 trials of various values, we multiplied the activity of all neurons in each 0.1s bin of the CS50 PSTH from each value level (z-score) by the difference vector and used the same conversion to 0 to 1 scale. To estimate the distribution of values along the CS+ / CS- dimension for each CS50 value condition, we bootstrapped (5000 iterations, with replacement) the population projection and took the mean 1 - 2.5 s from odor onset. We

calculated a p-value by finding the bootstrapped distribution of differences between CS50 high value projections and CS50 low value projections and calculating the fraction of the distribution that was less than 0 (supporting the null hypothesis that CS50 high value activity is not greater than CS50 low value activity).

Statistics.

All statistical tests were performed in MATLAB (MathWorks). To compare the fraction of neurons of a specific coding type across regions, we fit a generalized linear mixed-effects model ('fitglme' in MATLAB) with logit link function and with fixed effects of intercept and region and a random effect of session and then found the estimated mean and 95% confidence interval for each region. Regions with non-overlapping CIs were considered to have significantly different fractions of neurons of that coding type. To compare the number of anticipatory licks on different trial types, we found the mean number of anticipatory licks for each cue in each session and then performed a two-way ANOVA with effects of cue and subject and session as our n (Fig. 1C). To compare variance explained during each third of the first session, we found the mean value across neurons from each mouse and then performed a one-way ANOVA on those means with mouse as our n (Fig. 5H). To compare day 3 model performance on true and shuffled data across each time point (Fig. 6F), we found the mean value across neurons from each mouse at each time point and then performed a two-way ANOVA with main effects of shuffle and time point, with mouse as our n. We then calculated pairwise statistics using 'multcompare' in MATLAB with Bonferroni correction. To compare cue, lick, and reward unique variance at each time point for each cell category (determined on day 3, Fig. 6G), we found the mean from the cells in that category in each mouse at each time point and performed a two-way ANOVA with main effects of variable and day, with mouse as our n. We then calculated pairwise statistics using 'multcompare' in MATLAB with Bonferroni correction.

References

Allen WE, Chen MZ, Pichamoorthy N, Tien RH, Pachitariu M, Luo L, Deisseroth K. Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science* 364: 253–253, 2019.

Bari BA, Grossman CD, Lubin EE, Rajagopalan AE, Cressy JI, Cohen JY. Stable representations of decision variables for flexible behavior. *Neuron* 103: 922–933, 2019.

Berridge KC. Motivation concepts in behavioral neuroscience. *Physiology & behavior* 81: 179–209, 2004.

Bhattarai JP, Etyemez S, Jaaro-Peled H, Janke E, Tolosa UDL, Kamiya A, Gottfried JA, Sawa A, Ma M. Olfactory modulation of the medial prefrontal cortex circuitry: Implications for social cognition. In *Seminars in Cell & Developmental Biology*. Elsevier, 2021.

Buckley MJ, Mansouri FA, Hoda H, Mahboubi M, Browning PG, Kwok SC, Phillips A, Tanaka K. Dissociable components of rule-guided behavior depend on distinct medial and prefrontal regions. *Science* 325: 52–58, 2009.

Chen TW, Li N, Daie K, Svoboda K. A map of anticipatory activity in mouse motor cortex. *Neuron* 94: 866–879, 2017.

Chu MW, Li WL, Komiyama T. Balancing the robustness and efficiency of odor representations during learning. *Neuron* 92: 174–186, 2016.

Chudasama Y, Robbins TW. Dissociable contributions of the orbitofrontal and infralimbic cortex to pavlovian autoshaping and discrimination reversal learning: further evidence for the functional heterogeneity of the rodent frontal cortex. *Journal of Neuroscience* 23: 8771–8780, 2003.

Cousens GA. Characterization of odor-evoked neural activity in the olfactory peduncle. *IBRO reports* 9: 157–163, 2020.

Dalley JW, Cardinal RN, Robbins TW. Prefrontal executive and cognitive functions in rodents: neural and neurochemical substrates. *Neuroscience & Biobehavioral Reviews* 28: 771–784, 2004.

Dalton GL, Wang NY, Phillips AG, Floresco SB. Multifaceted contributions by different regions of the orbitofrontal and medial prefrontal cortex to probabilistic reversal learning. *Journal of Neuroscience* 36: 1996–2006, 2016.

Devore S, Lee J, Linster C. Odor preferences shape discrimination learning in rats. *Behavioral neuroscience* 127: 498, 2013.

Doucette W, Gire DH, Whitesell J, Carmean V, Lucero MT, Restrepo D. Associative cortex features in the first olfactory brain relay station. *Neuron* 69: 1176–1187, 2011.

Ekstrand JJ, Domroese ME, Johnson DM, Feig SL, Knodel SM, Behan M, Haberly LB. A new subdivision of anterior piriform cortex and associated deep nucleus with novel features of interest for olfaction and epilepsy. *Journal of Comparative Neurology* 434: 289–307, 2001.

Enel P, Perkins AQ, Rich EL. Heterogeneous value coding in orbitofrontal populations. *Behavioral neuroscience* 135: 245, 2021.

Ertürk A, Becker K, Jährling N, Mauch C, Hojer C, Egen J, Hellal F, Bradke F, Sheng M, Dodt HU. Three-dimensional imaging of solvent-cleared organs using 3DISCO. *Nature protocols* 7: 1983–95, 2012.

Eshel N, Tian J, Bukwich M, Uchida N. Dopamine neurons share common response function for reward prediction error. *Nature neuroscience* 19: 479, 2016.

Fiorillo CD, Tobler PN, Schultz W. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299: 1898–1902, 2003.

Grant RI, Doncheck EM, Vollmer KM, Winston KT, Romanova EV, Siegler PN, Holman H, Bowen CW, Otis JM. Specialized coding patterns among dorsomedial prefrontal neuronal ensembles predict conditioned reward seeking. *Elife* 10: e65764, 2021.

Guo ZV, Hires SA, Li N, O'Connor DH, Komiyama T, Ophir E, Huber D, Bonardi C, Morandell K, Gutnisky D, et al. Procedures for behavioral experiments in head-fixed mice. *PloS one* 9: e88678, 2014a.

Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K. Flow of cortical activity underlying a tactile decision in mice. *Neuron* 81: 179–194, 2014b.

Hayden BY, Niv Y. The case against economic values in the orbitofrontal cortex (or anywhere else in the brain). *Behavioral Neuroscience* 135: 192, 2021.

Heilbronner SR, Rodriguez-Romaguera J, Quirk GJ, Groenewegen HJ, Haber SN. Circuit-based corticostriatal homologies between rat and primate. *Biological psychiatry* 80: 509–521, 2016.

Horst NK, Laubach M. Reward-related activity in the medial prefrontal cortex is driven by consumption. *Frontiers in neuroscience* 7: 56, 2013.

Hunt LT, Malalasekera W, de Berker AO, Miranda B, Farmer SF, Behrens TE, Kennerley SW. Triple dissociation of attention and decision computations across prefrontal cortex. *Nature neuroscience* 21: 1471–1481, 2018.

Hyman JM, Ma L, Balaguer-Ballester E, Durstewitz D, Seamans JK. Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences* 109: 5086–5091, 2012.

Igarashi KM, Ieki N, An M, Yamaguchi Y, Nagayama S, Kobayakawa K, Kobayakawa R, Tanifuji M, Sakano H, Chen WR, et al. Parallel mitral and tufted cell pathways route distinct odor information to different targets in the olfactory cortex. *Journal of Neuroscience* 32: 7970–7985, 2012.

International-Brain-Laboratory, Banga K, Benson J, Bonacchi N, Bruijns SA, Campbell R, Chappuis GA, Churchland AK, Davatolhagh MF, Lee HD, Faulkner M, et al. Reproducibility of in-vivo electrophysiological measurements in mice. *bioRxiv* , 2022.

Jackman SL, Chen CH, Chettih SN, Neufeld SQ, Drew IR, Agba CK, Flaquer I, Stefano AN, Kennedy TJ, Belinsky JE, et al. Silk fibroin films facilitate single-step targeted expression of optogenetic proteins. *Cell reports* 22: 3351–3361, 2018.

Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, Lee AK, Anastassiou CA, Andrei A, Aydin Ç, et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551: 232–236, 2017.

Kennerley SW, Dahmubed AF, Lara AH, Wallis JD. Neurons in the frontal lobe encode the value of multiple decision variables. *Journal of cognitive neuroscience* 21: 1162–1178, 2009.

Kesner RP, Churchwell JC. An analysis of rat prefrontal cortex in mediating executive function. *Neurobiology of learning and memory* 96: 417–431, 2011.

Kikuta S, Kashiwadani H, Mori K. Compensatory rapid switching of binasal inputs in the olfactory cortex. *Journal of Neuroscience* 28: 11989–11997, 2008.

Klein-Flügge MC, Bongioanni A, Rushworth MF. Medial and orbital frontal cortex in decision-making and flexible behavior. *Neuron* , 2022.

Koldaeva A, Schaefer AT, Fukunaga I. Rapid task-dependent tuning of the mouse olfactory bulb. *Elife* 8: e43558, 2019.

Kondo M, Matsuzaki M. Neuronal representations of reward-predicting cues and outcome history with movement in the frontal cortex. *Cell Reports* 34: 108704, 2021.

Kuwabara M, Kang N, Holy TE, Padoa-Schioppa C. Neural mechanisms of economic choices in mice. *Elife* 9: e49669, 2020.

Laubach M, Amarante LM, Swanson K, White SR. What, if anything, is rodent prefrontal cortex? *eneuro* 5, 2018.

Lei H, Mooney R, Katz LC. Synaptic integration of olfactory information in mouse anterior olfactory nucleus. *Journal of Neuroscience* 26: 12023–12032, 2006.

Li A, Gire DH, Restrepo D. Υ spike-field coherence in a population of olfactory bulb neurons differentiates between odors irrespective of associated outcome. *Journal of Neuroscience* 35: 5808–5822, 2015.

Li N, Daie K, Svoboda K, Druckmann S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 532: 459–464, 2016.

Liu LD, Chen S, Hou H, West SJ, Faulkner M, Laboratory TIB, Economo MN, Li N, Svoboda K. Accurate localization of linear probe electrode arrays across multiple brains. *eneuro* 8, 2021.

Malagon-Vina H, Ciocchi S, Passecker J, Dorffner G, Klausberger T. Fluid network dynamics in the prefrontal cortex during multiple strategy switching. *Nature communications* 9: 1–13, 2018.

Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. *Annual review of neuroscience* 24: 167–202, 2001.

Mori K, Sakano H. Olfactory circuitry and behavioral decisions. *Annual Review of Physiology* 83: 231–256, 2021.

Musall S, Kaufman MT, Juavinett AL, Gluf S, Churchland AK. Single-trial neural dynamics are dominated by richly varied movements. *Nature neuroscience* 22: 1677–1686, 2019.

Nakahara H, Itoh H, Kawagoe R, Takikawa Y, Hikosaka O. Dopamine neurons can represent context-dependent prediction error. *Neuron* 41: 269–280, 2004.

Namboodiri VMK, Otis JM, van Heeswijk K, Voets ES, Alghorazi RA, Rodriguez-Romaguera J, Mihalas S, Stuber GD. Single-cell activity tracking reveals that orbitofrontal neurons acquire and maintain a long-term memory to guide behavioral adaptation. *Nature neuroscience* 22: 1110–1121, 2019.

Otis JM, Namboodiri VM, Matan AM, Voets ES, Mohorn EP, Kosyk O, McHenry JA, Robinson JE, Resendez SL, Rossi MA, et al. Prefrontal cortex output circuits guide reward seeking through divergent cue encoding. *Nature* 543: 103–107, 2017.

Ottenheimer DJ, Bari BA, Sutlief E, Fraser KM, Kim TH, Richard JM, Cohen JY, Janak PH. A quantitative reward prediction error signal in the ventral pallidum. *Nature neuroscience* 23: 1267–1276, 2020.

Ottenheimer DJ, Hjort MM, Bowen AJ, Steinmetz NA, Stuber GD. Data and code for Ottenheimer et al 2022. *Zenodo* DOI: 10.5281/zenodo.6686927, 2022.

Pashkovski SL, Iurilli G, Brann D, Chicharro D, Drumme K, Franks KM, Panzeri S, Datta SR. Structure and flexibility in cortical representations of odour space. *Nature* 583: 253–258, 2020.

Price JL. Beyond the primary olfactory cortex: olfactory-related areas in the neocortex, thalamus and hypothalamus. *Chemical Senses* 10: 239–258, 1985.

Price JL, Carmichael ST, Carnes K, Clugnet M, Kuroda M, Ray J. Olfactory input to the prefrontal cortex. *Olfaction: A model system for computational neuroscience* pp. 101–20, 1991.

Roesch MR, Olson CR. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* 304: 307–310, 2004.

Rudebeck PH, Behrens TE, Kennerley SW, Baxter MG, Buckley MJ, Walton ME, Rushworth MF. Frontal cortex subregions play distinct roles in choices between actions and stimuli. *Journal of Neuroscience* 28: 13775–13785, 2008.

Sallet J, Quilodran R, Rothé M, Vezoli J, Joseph JP, Procyk E. Expectations, gains, and losses in the anterior cingulate cortex. *Cognitive, Affective, & Behavioral Neuroscience* 7: 327–336, 2007.

Saraiva LR, Kondoh K, Ye X, Yoon Kh, Hernandez M, Buck LB. Combinatorial effects of odorants on mouse behavior. *Proceedings of the National Academy of Sciences* 113: E3300–E3306, 2016.

Schoenbaum G, Setlow B, Saddoris MP, Gallagher M. Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala. *Neuron* 39: 855–867, 2003.

Schoonover CE, Ohashi SN, Axel R, Fink AJ. Representational drift in primary olfactory cortex. *Nature* 594: 541–546, 2021.

Shamonin DP, Bron EE, Lelieveldt BP, Smits M, Klein S, Staring M, Initiative ADN. Fast parallel image registration on CPU and GPU for diagnostic classification of Alzheimer's disease. *Frontiers in neuroinformatics* 7: 50, 2014.

St Onge JR, Floresco SB. Prefrontal cortical contribution to risk-based decision making. *Cerebral cortex* 20: 1816–1828, 2010.

Stalnaker TA, Cooch NK, McDannald MA, Liu TL, Wied H, Schoenbaum G. Orbitofrontal neurons infer the value and identity of predicted outcomes. *Nature communications* 5: 1–13, 2014.

Stalnaker TA, Cooch NK, Schoenbaum G. What the orbitofrontal cortex does not do. *Nature neuroscience* 18: 620–627, 2015.

Steinmetz NA, Aydin C, Lebedeva A, Okun M, Pachitariu M, Bauza M, Beau M, Bhagat J, Böhm C, Broux M, et al. Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* 372: eabf4588, 2021.

Steinmetz NA, Buetfering C, Lecoq J, Lee CR, Peters AJ, Jacobs EA, Coen P, Ollerenshaw DR, Valley MT, De Vries SE, et al. Aberrant cortical activity in multiple GCaMP6-expressing transgenic mouse lines. *eneuro* 4, 2017.

Steinmetz NA, Zatka-Haas P, Carandini M, Harris KD. Distributed coding of choice, action and engagement across the mouse brain. *Nature* 576: 266–273, 2019.

Stettler DD, Axel R. Representations of odor in the piriform cortex. *Neuron* 63: 854–864, 2009.

Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* 364: eaav7893, 2019.

Sul JH, Kim H, Huh N, Lee D, Jung MW. Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron* 66: 449–460, 2010.

Tsuji T, Tsuji C, Lozic M, Ludwig M, Leng G. Coding of odors in the anterior olfactory nucleus. *Physiological reports* 7: e14284, 2019.

van Duuren E, van der Plasse G, Lankelma J, Joosten RN, Feenstra MG, Pennartz CM. Single-cell and population coding of expected reward probability in the orbitofrontal cortex of the rat. *Journal of Neuroscience* 29: 8965–8976, 2009.

Verharen JP, den Ouden HE, Adan RA, Vanderschuren LJ. Modulation of value-based decision making behavior by subregions of the rat prefrontal cortex. *Psychopharmacology* 237: 1267–1280, 2020.

Wang PY, Boboila C, Chin M, Higashi-Howard A, Shamash P, Wu Z, Stein NP, Abbott L, Axel R. Transient and persistent representations of odor value in prefrontal cortex. *Neuron* 108: 209–224, 2020a.

Wang Q, Ding SL, Li Y, Royall J, Feng D, Lesnar P, Graddis N, Naeemi M, Facer B, Ho A, et al. The Allen mouse brain common coordinate framework: a 3D reference atlas. *Cell* 181: 936–953, 2020b.

Winkelmeier L, Filosa C, Hartig R, Scheller M, Sack M, Reinwald JR, Becker R, Wolf D, Gerchen MF, Sartorius A, et al. Striatal hub of dynamic and stabilized prediction coding in forebrain networks for olfactory reinforcement learning. *Nature Communications* 13: 1–21, 2022.

Zagha E, Erlich JC, Lee S, Lur G, O'Connor DH, Steinmetz NA, Stringer C, Yang H. The importance of accounting for movement when relating neuronal activity to sensory and cognitive processes. *Journal of Neuroscience* 42: 1375–1382, 2022.

Zhou J, Gardner MP, Schoenbaum G. Is the core function of orbitofrontal cortex to signal values or make predictions? *Current Opinion in Behavioral Sciences* 41: 1–9, 2021.

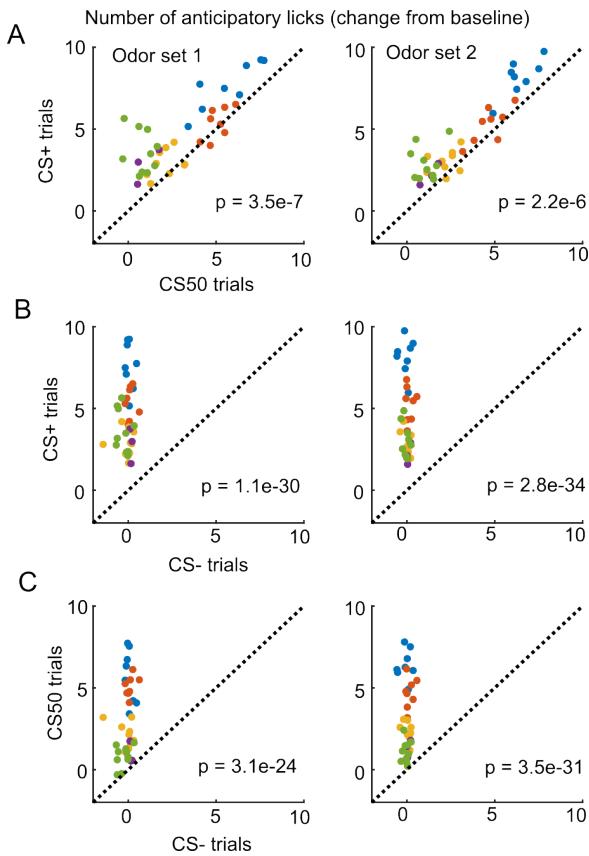


Figure S 1: Anticipatory licking during the electrophysiology sessions.

Figure S1. Anticipatory licking during the electrophysiology sessions.

- (A) Mean anticipatory licks (change from baseline) for the CS+ and CS50 from odor set 1 (left) and 2 (right) for every session, color-coded by mouse. $F(1, 66) = 32.07$ and $F(1, 66) = 26.93$ in each odor set for a main effect of cue in a two-way ANOVA including an effect of subject.
- (B) As above, for the CS+ and CS- from odor set 1 (left) and 2 (right). $F(1, 66) = 433.1$ and $F(1, 66) = 574.6$ in each odor set for a main effect of cue in a two-way ANOVA including an effect of subject.
- (C) As above, for the CS50 and CS- from odor set 1 (left) and 2 (right). $F(1, 66) = 252.3$ and $F(1, 66) = 450.1$ in each odor set for a main effect of cue in a two-way ANOVA including an effect of subject.

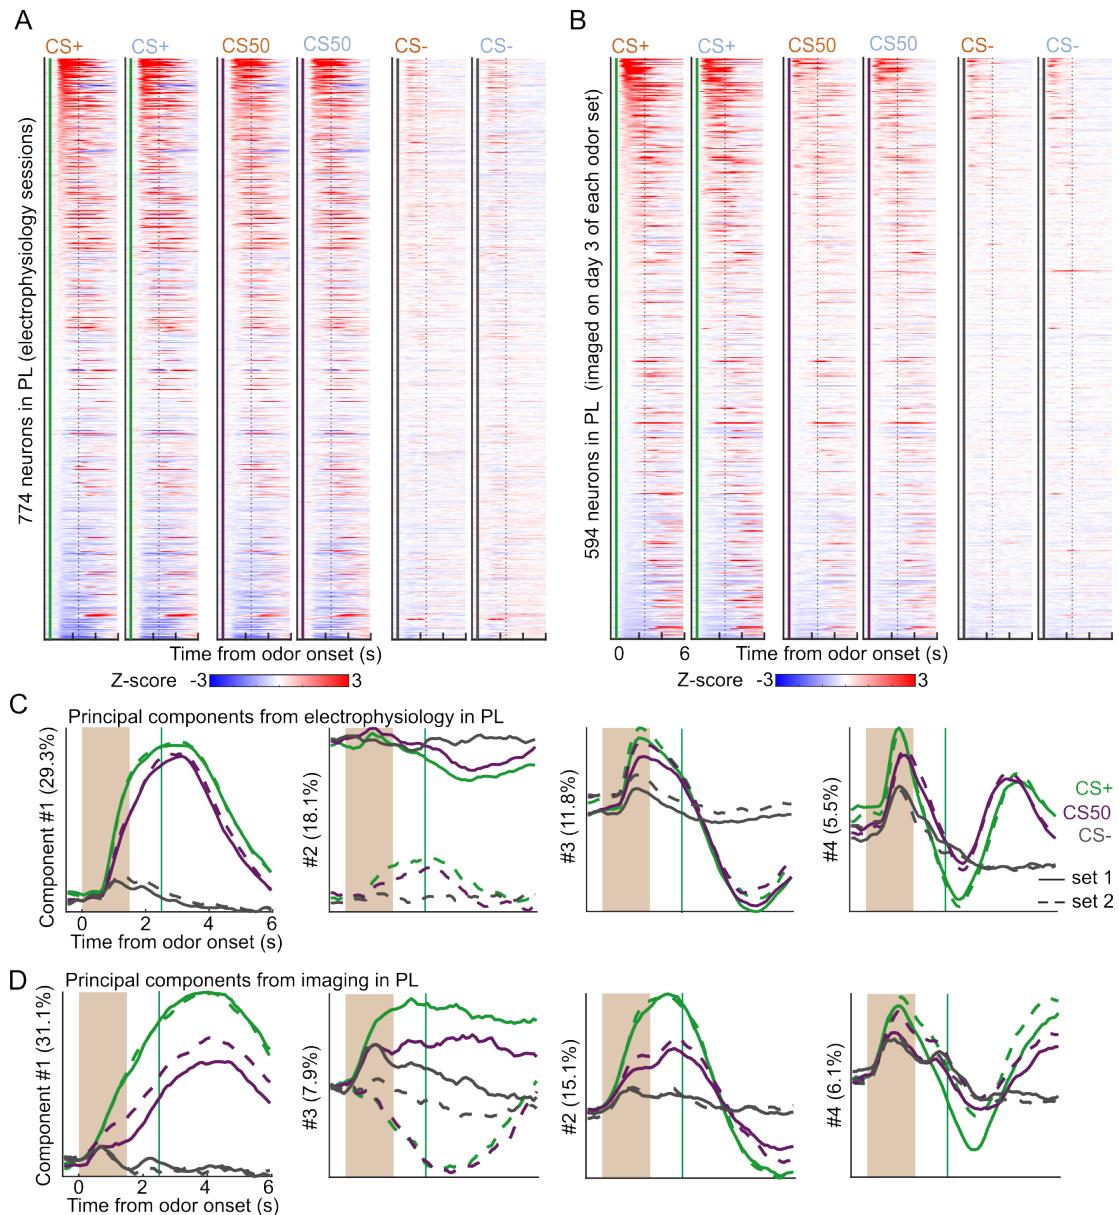


Figure S 2: Similar neural activity in prelimbic area using electrophysiology and imaging.

Figure S2. Similar neural activity in prelimbic area using electrophysiology and imaging.

- (A) Heatmap of the normalized activity of each neuron recorded with electrophysiology in PL, aligned to each of the 6 odors. All columns sorted by mean firing 0 - 1.5 s following odor onset for odor set 1 CS+ trials.
- (B) As in (A), for all neurons imaged in PL on day 3 of each odor set.

- (C) The score from the first 4 principal components of the normalized activity presented in (A), with variance explained in parentheses.
- (D) As in (C), for the activity in (B).

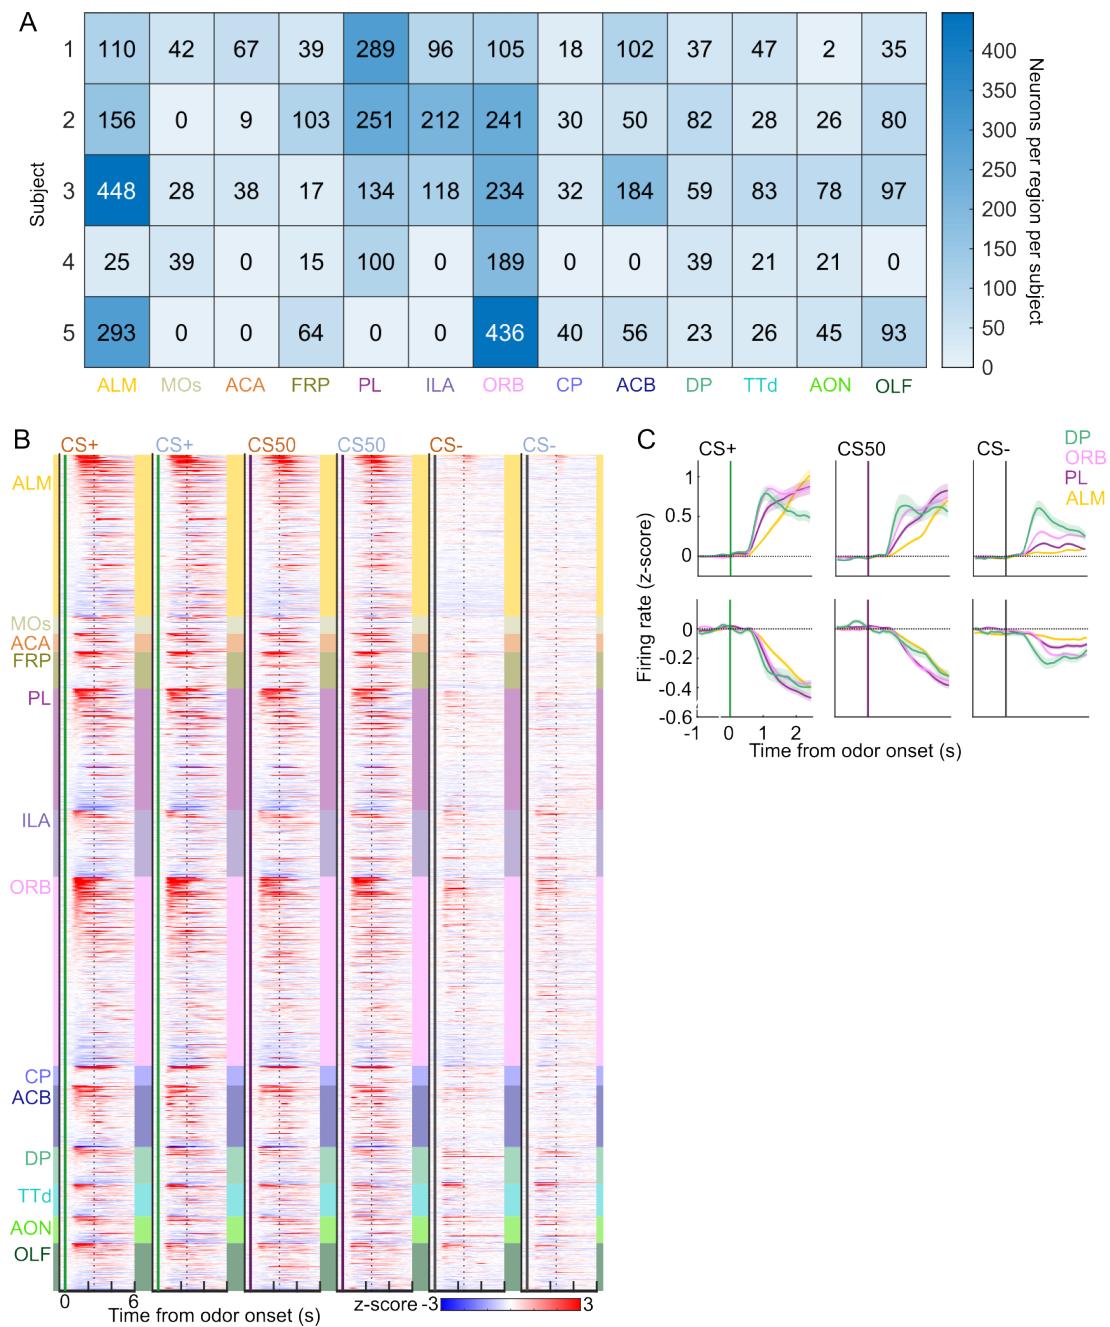


Figure S 3: Task-related neural activity across brain regions.

Figure S3. Task-related neural activity across brain regions.

(A) For each of the 5 mice in the electrophysiology experiment, the number of neurons recorded in each region.

- (B) Heatmap of the normalized activity of each neuron ($n = 51$ trials per cue). All columns sorted by region and then by mean firing 0 - 1.5 s following odor onset for odor set 1 CS+ trials.
- (C) Mean (+/- SEM) activity of neurons from 4 regions aligned to each cue type, grouped by whether peak cue activity (0 - 2.5 s) was above (top) or below (bottom) baseline in held out trials.

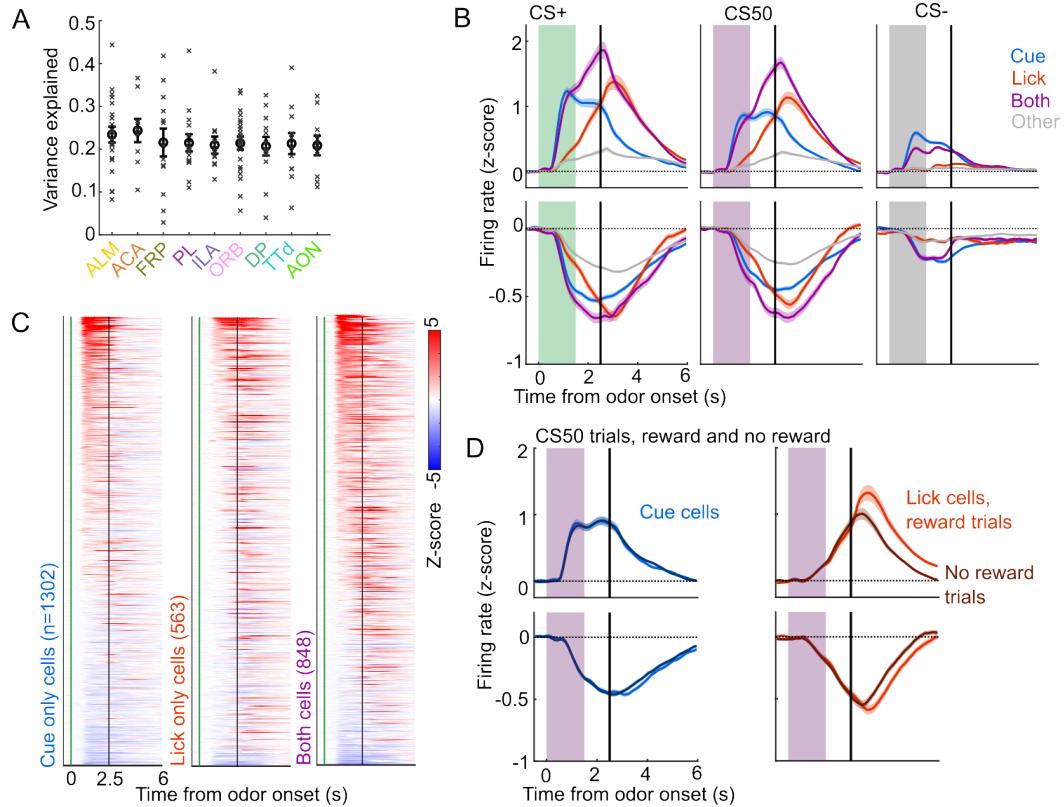


Figure S 4: Identification of cue and lick cells with GLM.

Figure S4. Identification of cue and lick cells with GLM.

- Mean variance explained (fraction) by linear models in each region for each session (x) and the mean (+/- SEM) across those sessions.
- Mean (+/- SEM) activity of neurons encoding cues, licks, both, or neither aligned to each cue type, grouped by whether peak cue activity (0 - 2.5 s) was above (top) or below (bottom) baseline in held out trials.
- Normalized activity of every neuron encoding cues, licks, or both, aligned to CS+ onset, sorted by mean firing 0 - 1.5 s following odor onset.
- Mean (+/- SEM) activity of neurons encoding cues or licks, grouped as in (B), on CS50 trials, divided into rewarded (lighter colors) or unrewarded (darker colors) trials.

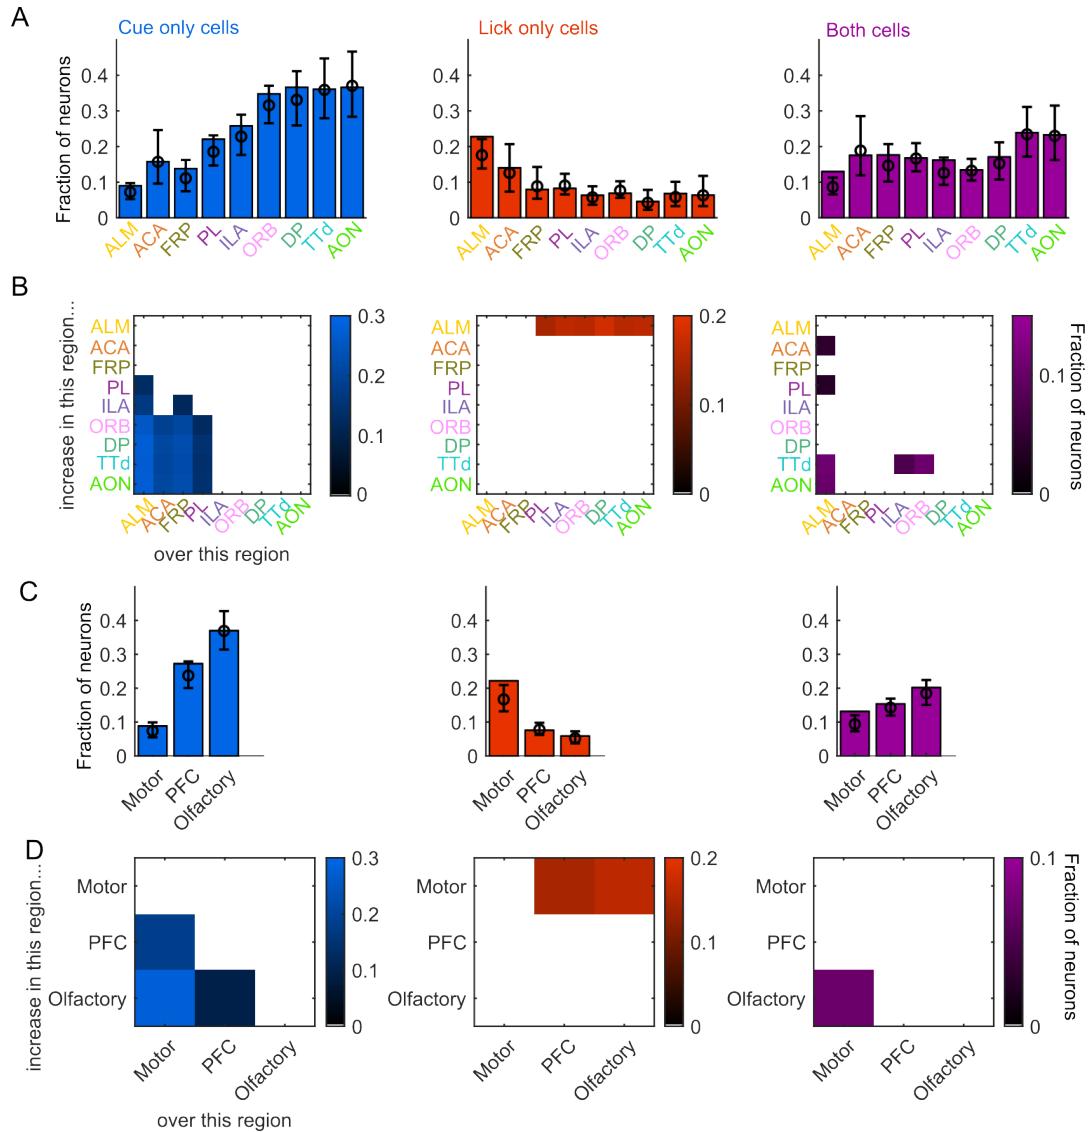


Figure S 5: Comparing proportions of cue and lick neurons across regions.

Figure S5. Comparing proportions of cue and lick neurons across regions.

- Fraction of neurons in each region classified as coding cues (left), licks (middle), or both (right), as well as estimated fraction(+/-95% CI) with random effect of session (see Methods).
- Additional cue/lick/both cells in region on Y-axis compared to region on X-axis as a fraction of all neurons, for regions with non-overlapping 95% confidence intervals.
- As in (A), for region groups.
- As in (B), for region groups.

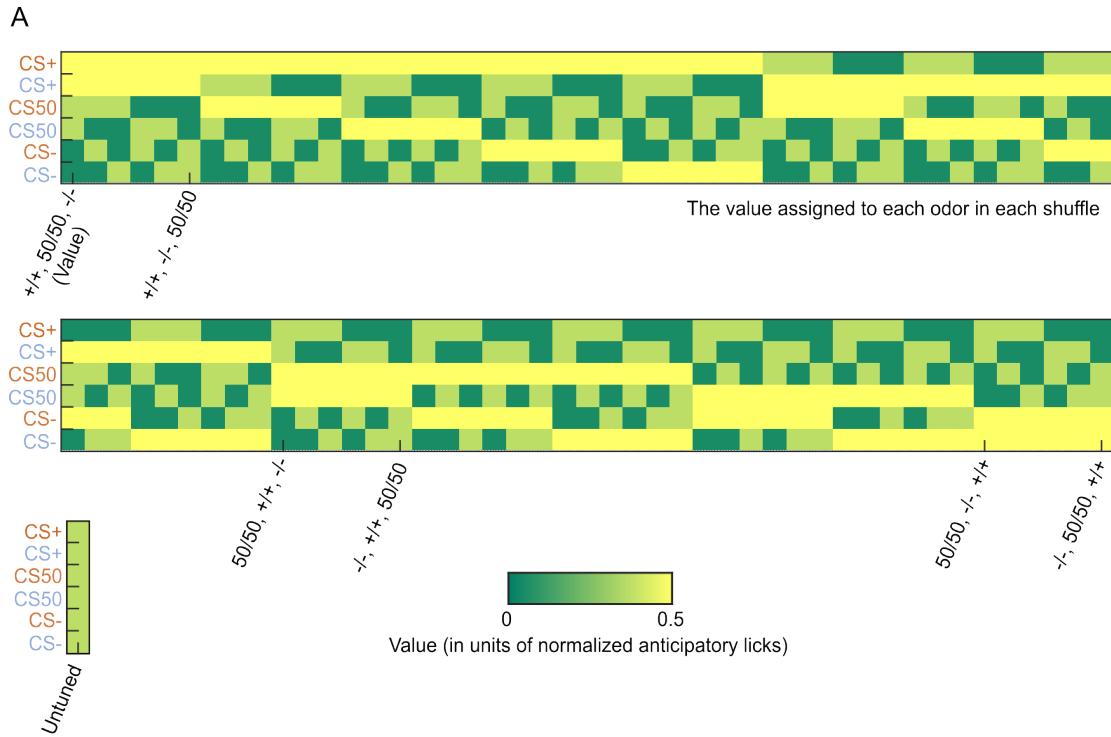


Figure S 6: Schematic of value model shuffles.

Figure S6. Schematic of value model shuffles.

(A) For each of the 90 permutations of the Value model, the value taken on by the variable cue kernel on trials corresponding to each of the 6 cue types. Value is determined in units of predicted anticipatory licks, from 0 to 1 (maximum number of anticipatory licks made). Additionally, there is an Untuned model, where all cues take on the same value.

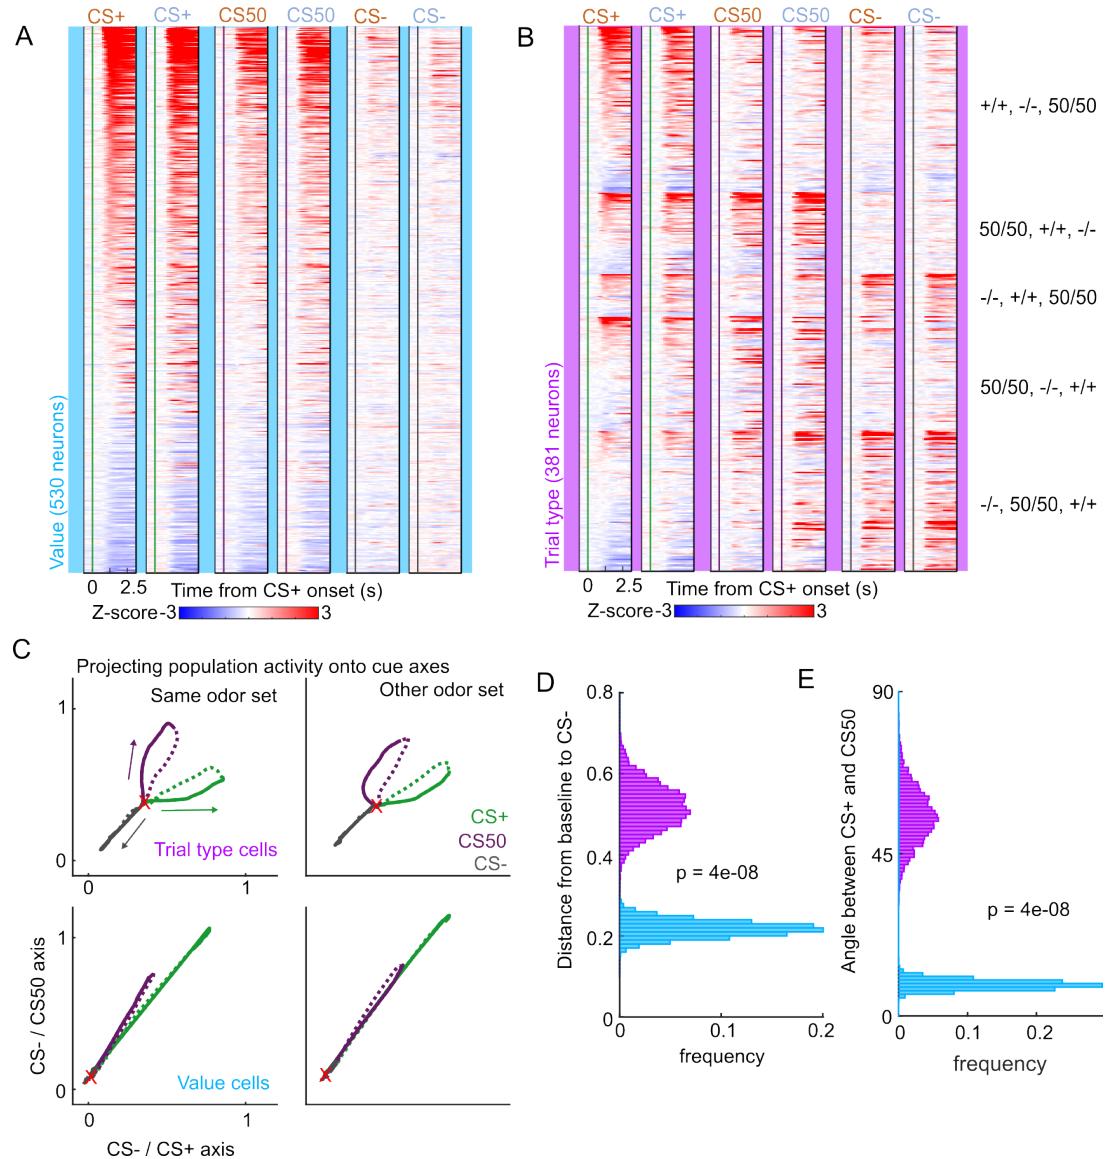


Figure S7: Additional analysis of odor coding schemes.

Figure S7. Additional analysis of odor coding schemes.

- Normalized activity of every value neuron, sorted by mean firing 0 - 1.5 s following odor set 1 CS+ onset.
- Normalized activity of every trial type neuron, sorted by model and then by mean firing 0 - 1.5 s following odor set 1 CS+ onset.
- Projecting the activity of all trial type and value cells onto the coding dimensions maximally separating CS- and CS+ (x-axis) and CS- and CS50 (y-axis). Solid line is

activity during cue, dashed line is activity following reward delivery. X marks baseline activity.

- (D) For the odor set 1 projection, distribution of 5000 bootstrapped distances between baseline activity (prior to odor onset) and the CS- representation (0, 0). Value cells were closer to CS- at baseline than trial type cells.
- (E) For the odor set 1 projection, distribution of 5000 bootstrapped angles between CS+ and CS50 vectors (baseline to peak). Value cells had a smaller angle than trial type cells.

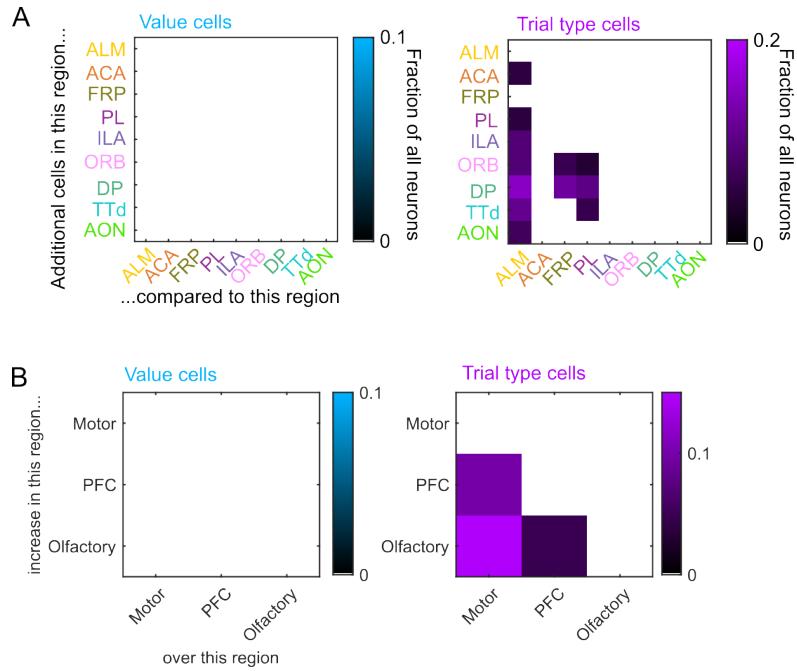


Figure S 8: Relative proportions of value and trial type cells across regions.

Figure S8. Relative proportions of value and trial type cells across regions.

(A) Additional cue value (left) or trial type (right) neurons in region on Y-axis compared to region on X-axis as a fraction of all neurons, for regions with non-overlapping 95% confidence intervals.

(B) As in (A), for region groups.

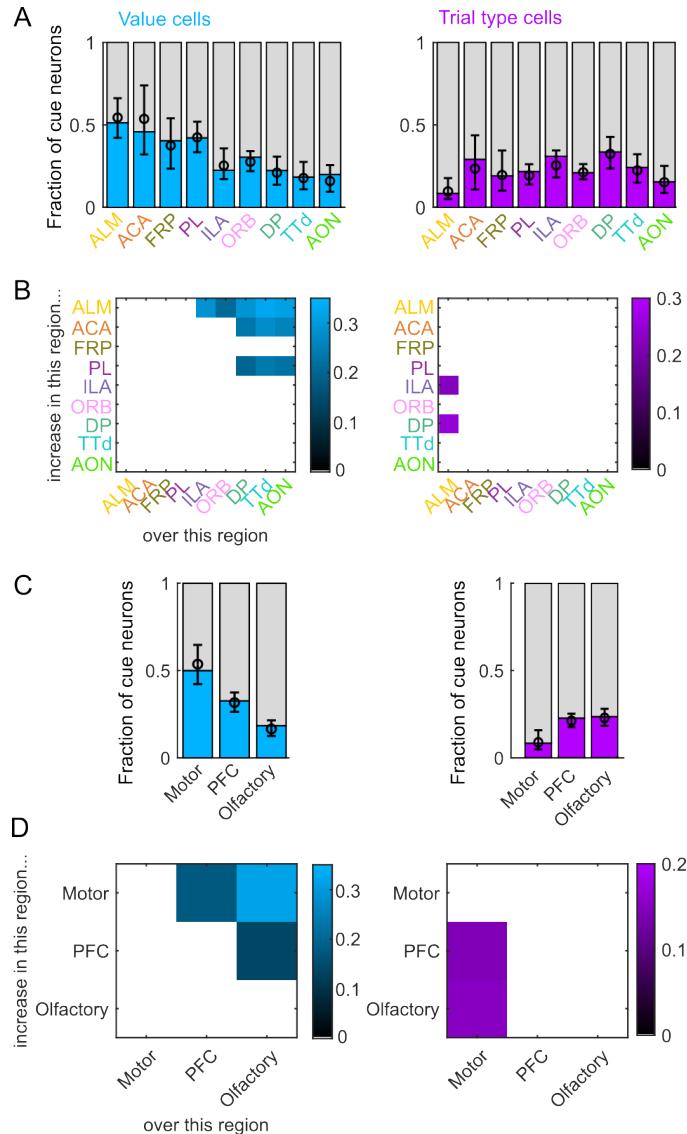


Figure S 9: Value coding as a proportion of cue cells.

Figure S9. Value coding as a proportion of cue cells.

- (A) Fraction of cue neurons in each region classified as coding value (left) or trial type (right), as well as estimated fraction(+/-95% CI) with random effect of session (see Methods).
- (B) Additional value/trial type cue neurons in region on Y-axis compared to region on X-axis as a fraction of all cue neurons, for regions with non-overlapping 95% confidence intervals.
- (C) As in (A), for region groups.

(D) As in (B), for region groups.

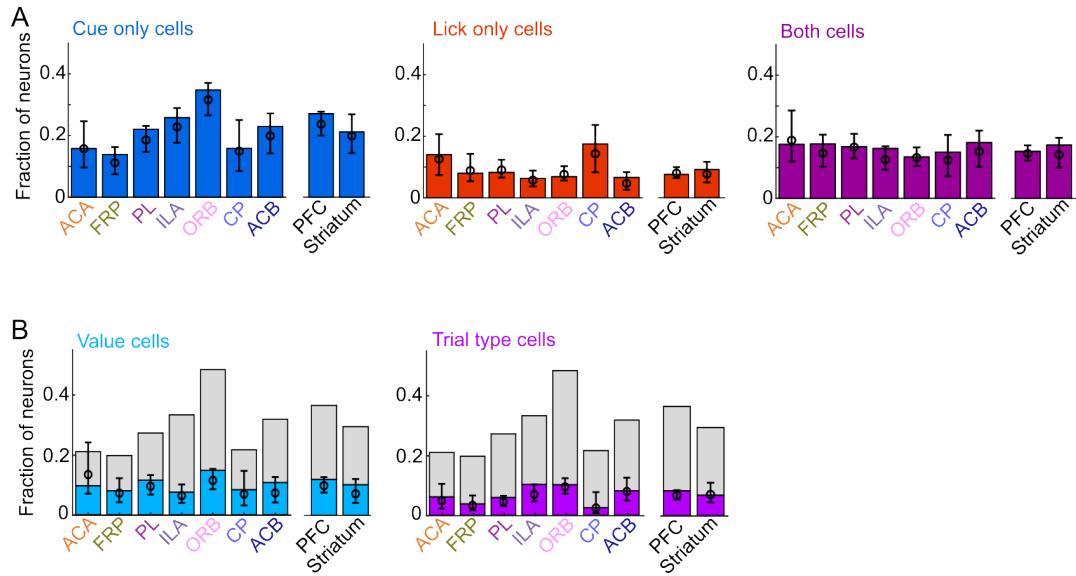


Figure S 10: Comparing PFC and striatum.

Figure S10. Comparing PFC and striatum.

- (A) Fraction of neurons in each region and region group classified as coding cues (left), licks (middle), or both (right), as well as estimated fraction(+/-95% CI) with random effect of session (see Methods).
- (B) Fraction of neurons in each region and region group classified as coding value (left) or trial type (right), as well as estimated fraction(+/-95% CI) with random effect of session. Light gray bars are remaining cue neurons not in that category.