

Selective Modulation of Orbitofrontal Network Activity during Negative Occasion Setting

Justin L. Shobe,¹ Konstantin I. Bakhurin,²  Leslie D. Claar,³ and Sotiris C. Masmanidis^{1,2,4,5}

¹Department of Neurobiology, ²Neuroscience Interdepartmental Program, ³Department of Bioengineering, and ⁴Integrative Center for Learning and Memory, and ⁵California Nanosystems Institute, University of California, Los Angeles, California 90095

Discrete cues can gain powerful control over behavior to help an animal anticipate and cope with upcoming events. This is important in conditions where understanding the relationship between complex stimuli provides a means to resolving situational ambiguity. However, it is unclear how cortical circuits generate and maintain these signals that conditionally regulate behavior. To address this, we established a Pavlovian serial feature-negative conditioning paradigm, where male mice are trained on a trial in which a conditioned stimulus (CS) is presented alone and followed by reward, or a feature-negative trial in which the CS is preceded by a feature cue indicating there is no reward. Mice learn to respond with anticipatory licking to a solitary CS, but significantly suppress their responding to the same cue during feature-negative trials. We show that the feature cue forms a selective association with its paired CS, because the ability of the feature to transfer its suppressive properties to a separately rewarded cue is limited. Next, to examine the underlying neural dynamics, we conduct recordings in the orbitofrontal cortex (OFC). We find that the feature cue significantly and selectively inhibits CS-evoked activity. Finally, we find that the feature triggers a distinct OFC network state during the delay period between the feature and CS, establishing a potential link between the feature and future events. Together, our findings suggest that OFC dynamics are modulated by the feature cue and its associated conditioned stimulus in a manner consistent with an occasion setting model.

Key words: negative occasion setting; orbitofrontal cortex; single-unit electrophysiology

Significance Statement

The ability of patterned cues to form an inhibitory relationship with ambiguously rewarded outcomes has been appreciated since early studies on learning and memory. However, it was often assumed that these cues, despite their hierarchical nature, still made direct associative links with neural rewarding events. This model was significantly challenged, largely by the work of Holland and colleagues, who demonstrated that under certain conditions cues can inherit occasion setting properties whereby they modulate the ability of a paired cue to elicit its conditioned response. Here we provide some of the first evidence that the activity of a cortical circuit is selectively modulated by such cues, thereby providing insight into the mechanisms of higher order learning.

Introduction

Animals routinely learn to anticipate events by extracting information from their environments. However, this can be particularly challenging when individual cues only provide partial predictive information as is often the case in naturalistic scenar-

ios. In these situations, animals will attempt to use disambiguating “features” to accurately predict outcomes (Schmajuk and Holland, 1998). A good example of this type of learning is feature-negative conditioning because behavioral success requires an animal to learn the pattern of cues that best predicts reward (Holland, 1984; Lamarre and Holland, 1987; Bueno and Holland, 2008). In the serial version of this task, animals learn that a single conditioned stimulus (CS) predicts a reward but when this same cue is preceded (with a temporal delay) by a separate feature cue, the trial goes unrewarded. Thus, the single cue elicits anticipatory behavior, but animals withhold their responses when the same cue is presented in feature-negative trials. Studies have shown that the ability to conditionally discriminate between rewarded and unrewarded trials can occur in a wide range of species, from insects to humans (Pace et al., 1980; Nallan et al., 1981; Pace and McCoy, 1981; Abramson et al., 2013), and under a variety of stimulus conditions (Holland, 1997). In the

Received March 1, 2017; revised July 18, 2017; accepted July 24, 2017.

Author contributions: J.L.S. designed research; J.L.S. and L.D.C. performed research; J.L.S., K.I.B., and S.C.M. analyzed data; J.L.S., K.I.B., L.D.C., and S.C.M. wrote the paper.

This work was supported by the Ruth Kirschstein National Research Service Award NIH T32-NS058280 to K.I.B., and a 2014 McKnight Technical Innovations in Neuroscience Award, NIH DA034178 and NIH NS100050, to S.C.M. We thank D.V. Buonomano and V. Goudar for assistance with decoding analysis.

The authors declare no competing financial interests.

Correspondence should be addressed to either Justin L. Shobe or Dr. Sotiris C. Masmanidis, University of California, 650 Charles E Young Drive South, Los Angeles, Los Angeles, CA 90095. E-mail: jshobe@gmail.com or smasmanidis@ucla.edu.

DOI:10.1523/JNEUROSCI.0572-17.2017

Copyright © 2017 the authors 0270-6474/17/379415-09\$15.00/0

mammalian brain there is evidence that these functions are mediated by specific circuits, including the retrosplenial cortex (Robinson et al., 2011), striatum, and orbitofrontal cortex (Meyer and Bucci, 2016). Despite these studies, there is still a relatively poor understanding of the relationship between feature cues and their associated conditioned stimuli that function to bias behavioral decisions.

There are two contrasting models that attempt to account for how neural circuits solve this problem. One model views the animal's ability to discriminate rewarded and unrewarded trials as a basic function of elemental conditioning, where a CS acquires a positive associative relationship to the reward to promote conditioned responding, and the feature acquires a negative relationship to suppress responding (Rescorla, 1969; Rescorla and Wagner, 1972; Rescorla and Holland, 1977). On trials in which both cues are present, the feature cue's inhibitory influence simply overrides the CS's excitatory influence, due to the feature cue's direct negative association with the reward representation (conditioned inhibition model). In the opposing model, the feature cue functions as a negative occasion setter that does not make a direct association with the reward representation (Holland, 1984, 1989, 1995a; Lamarre and Holland, 1987). Instead it modulates the ability of the CS to retrieve the reward association by acting as a kind of inhibitory gate (Holland, 1989, 1995a).

To gain mechanistic insight into these opposing models at the level of single-neuron spiking activity, we establish a Pavlovian feature-negative conditioning paradigm in head-restrained mice, which is compatible with large-scale neural recordings using silicon-based microprobes. In our task, a CS predicts the delivery of reward, but there is no reward when this CS is preceded by a feature cue (Holland, 1995a,b). We find that mice predominantly solve this task by using a strategy consistent with the second model (negative-occasion setting), because the feature acquires the ability to specifically inhibit the reward association of its paired CS (Holland, 1984, 2008). Moreover, we find that neural activity within the orbitofrontal cortex (OFC) is consistent with this model because the feature appears to selectively modulate cue-evoked firing in a manner that correlates with behavioral performance. Finally, we also observe an "activity silent" state (Stokes, 2015) in OFC network dynamics that could function to relay information during the time gap between the feature and CS cue. To our knowledge this is the first demonstration of a modulatory cortical circuit mechanism that specifically supports the occasion setting model.

Materials and Methods

Animals and surgical procedures. All procedures were approved by the University of California, Los Angeles Chancellor's Animal Research Committee. Singly housed male C57BL/6J mice ($n = 8$, 15–22 weeks old at the time of recording. The Jackson Laboratory) were used in the experiments. Animals underwent an initial head bar implantation surgery under isoflurane anesthesia in a stereotaxic apparatus to bilaterally fix, with dental cement, stainless steel head bars on the skull. After training, animals underwent a second surgery under isoflurane anesthesia on the recording day to make a single craniotomy for acute silicon microprobe recordings. An additional craniotomy was made over the posterior cerebellum for placement of an electrical reference wire. All behavioral training and recording sessions were performed in fully awake head-restrained animals.

Behavioral task. We started food restriction 1 week after the initial head bar implantation surgery. Mice were fed daily after each training session to maintain ~90% of their baseline weight, whereas water remained freely accessible in the home cage. To begin each training session, we mounted animals on the head bar restraint bracket and placed them on a

polystyrene treadmill ball (200 mm diameter, Graham Sweet Studios) that freely rotated in a forward/backward direction. Behavioral training consisted of four successive phases: (1) habituation, (2) odor and air puff conditioning, (3) feature-negative conditioning, and (4) behavioral testing and electrophysiology. In the first phase, mice were initially habituated to the head restraint system and trained to consume a liquid reward (5 μ l, 10% sweetened condensed milk) delivered by actuation of an audible solenoid valve (Neptune Research). Licking was continuously monitored via an infrared lick meter placed in front of the reward delivery tube (Island Motion). During these sessions, animals were given rewards and exposed to a constant stream of pure air through a tube with a hole positioned in front of the nose [50 rewards per session, 13–21 s intertrial interval (ITI), 1.5 L/min air flow]. After mice learned to lick to at least 90% of the delivered rewards for 2 consecutive days, we began the second training phase. Mice received trials containing one of either two types of olfactory conditioned stimuli (CS1 or CS2, 1 s duration, 17–29 s ITI), or a mild air puff to the vibrissal pad. The air puff was odorless and thus provided a distinct (from the CS1 and CS2) but highly salient form of stimulus, which has been effectively used in head-fixed mouse behavioral paradigms (Guo et al., 2014). Aromatic compounds (isoamyl acetate in CS1, citral in CS2, Sigma-Aldrich) were diluted 1:100 in mineral oil (Sigma-Aldrich). Air (0.15 L/min) was bubbled through this liquid and combined with the 1.5 L/min stream of pure air. An additional air puff tube (which was separate from the odor delivery tubing system to prevent odors being mixed with the air puff) delivered a pulse of pure air to the vibrissal pad (0.5 s at 0.8 L/min) on the side contralateral to the recording hemisphere. This intensity level did not evoke any noticeable startle response such as blinking. CS1 and CS2 were always associated with reward, which was delivered 2.5 s after odor onset. The 1.5 s gap between the offset of the odor and the reward allows cue-evoked behavior and neural activity to be examined in the absence of potentially confounding reward stimulus signals. The air puff was not followed by any explicit outcome. Animals received 30 presentations of each trial type (CS1, CS2, air puff) in pseudorandom order during daily sessions in the second phase of training. The solenoid valves controlling the olfactory cues were sound-isolated and thus inaudible to the animal. Typically, within 2 d of training, animals began predicting the delivery of reward following CS1 or CS2 cues by exhibiting anticipatory licking during the interval between the cue and reward. After mice demonstrated anticipatory licking on at least 90% of both CS1 and CS2 trials, we began the third phase of training, in which the air puff was now set to serve as the feature cue. On unrewarded trials the air puff was presented starting 2.5 s before CS onset. The third training phase contained an equal proportion (33%) of CS1⁺, CS1⁻, and CS2⁺ trials presented in pseudorandom order (~100 trials per session; Fig. 1B, left). The superscript "+" denotes that a CS was not preceded by a feature cue and was followed by reward, whereas the superscript "-" denotes that a CS was preceded by a feature cue and was not followed by reward. The minimum reaction time for animals to initiate anticipatory licking was found to be ~0.5 s. Throughout the manuscript we define correct CS⁺ trials as those containing anticipatory licking (when licking occurred between 0.5 and 2.5 s following odor onset), correct CS⁻ trials as those in which animals withheld licking during this time period, and incorrect CS⁻ trials as those when animals licked during this time period. When mice achieved at least 90% correct CS⁺ trials and <10% incorrect CS⁻ trials, we began the last training phase, comprised of a single session which coincided with electrophysiological recordings. Here we introduced transfer trials (TT) in which the CS2 cue was preceded by an air puff feature cue (a novel pairing) and followed by reward (Fig. 1B, right). This last phase consisted of 28% CS1⁺, CS1⁻, CS2⁺ trials, and 15% transfer trials. Since the feature had never been previously associated with CS2, we used these transfer trials to determine which of two models (see Introduction) are implemented by the animals. To calculate the behavioral discrimination score, we subtracted the percentage of incorrect CS1⁻ trials from the percentage of correct CS1⁺ trials.

Electrophysiological recordings. One recording was performed per animal with a microprobe containing a total of 256 electrodes divided across four prongs that were spaced 0.2 mm apart. An array of 64 electrodes on each prong spanned 1 mm along the dorsal-ventral axis. We recorded

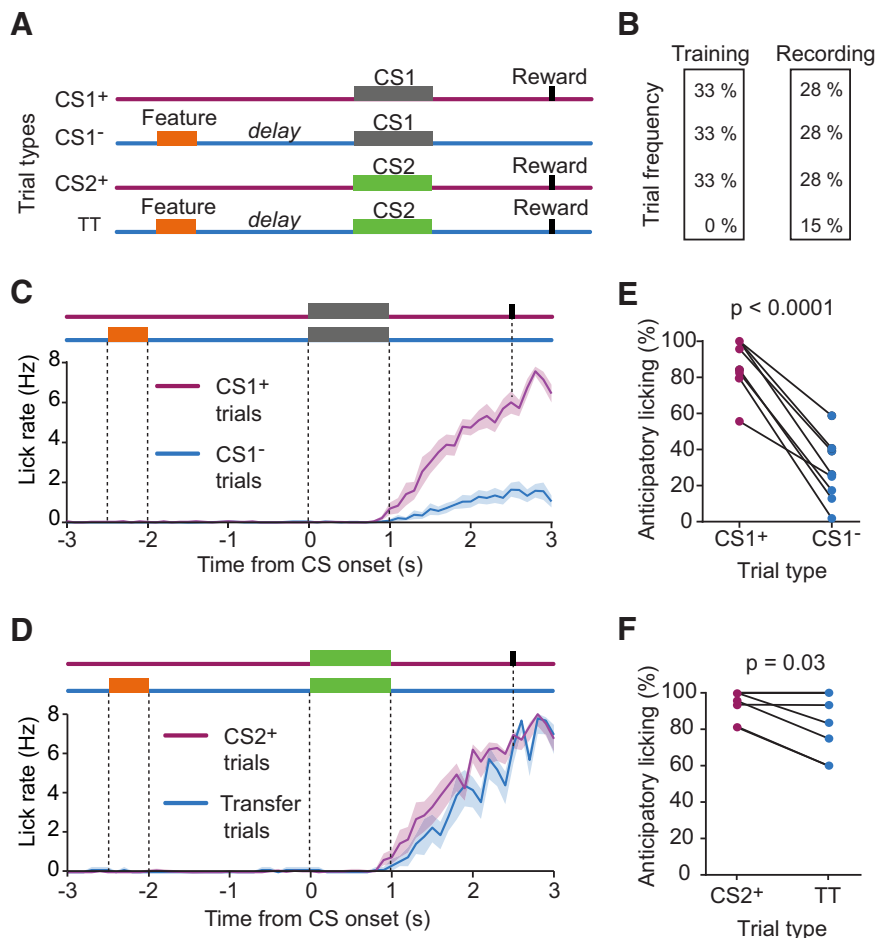


Figure 1. Distinct associations form from feature-negative conditioning. **A**, Schematic of the four distinct trial types used during training and recording sessions. In rewarded trials ($CS1^+$ and $CS2^+$), different conditioned odor stimuli (CS1 or CS2, 1 s duration) predicted the delivery of reward. In unrewarded trials ($CS1^-$), when the same odor stimuli were preceded by a feature cue (mild air puff, 0.5 s duration), there was no reward. Orange bar, Feature; gray bar, CS1; green bar, CS2; black bar, reward. **B**, Probability of presenting each trial type during initial training (left) and on the final training session corresponding to recording (right). All behavioral and electrophysiological results are from the final day. **C, D**, Average lick rate as a function of time during all rewarded and unrewarded trials. Dashed lines represent the onset and offset times of the indicated cue. Data represent mean \pm SEM ($n = 8$ mice). **E**, The feature significantly reduces the likelihood that animals express anticipatory licking ($t = 0$ –2.5 s from odor onset) in CS1 trials ($p < 0.0001$, paired t test). **F**, The feature significantly suppresses the likelihood of anticipatory licking in transfer trials ($p = 0.03$, paired t test).

from the orbitofrontal region of the prefrontal cortex (2.3–2.5 mm anterior, 0.5–1.5 mm lateral, -2.0 to -3.0 mm ventral, relative to bregma). The silicon prongs were coated with a fluorescent dye (DiD, ThermoFisher) before insertion, to facilitate *post hoc* histological reconstruction of the recording sites. Procedures for recording with silicon microprobes are described previously (Shobe et al., 2015). After the recordings, animals were overdosed with isoflurane and perfused with 10% formalin solution (Sigma-Aldrich). The brain was extracted and fixed for a minimum of 24 h at 4°C . Tissue was cut into $100\ \mu\text{m}$ sections on a vibratome and stained for DAPI ($4\ \mu\text{g}/\text{ml}$) to visualize cell nuclei. Confocal imaging of DiD and DAPI fluorescence confirmed that recordings in all mice were located in approximately the same subregions of the OFC.

Firing rate analysis, and identification of significantly discriminating or modulated cells. Spike sorting was performed using custom, semi-automated scripts written in MATLAB (MathWorks) for the identification of putative single units. The analysis combined all types of units (putative pyramidal cells and interneurons). The mean firing rate per unit was calculated by binning spike count data into 5 ms time steps, convolving with a Gaussian kernel ($SD = 25\ \text{ms}$), and averaging across trials of the same stimulus type (either $CS1^+$, $CS1^-$, $CS2^+$, transfer). To determine whether a unit's activity significantly discriminated between $CS1^+$ and

$CS1^-$ trials, we used a permutation test to detect significant differences in observed firing rate for each time step between these trials (Bakhrin et al., 2016). The firing rate was sampled from $t = 0$ to 1 s post- $CS1$ onset in time steps of 5 ms. For each time step, the data from $CS1^+$ and $CS1^-$ trials were shuffled, and a new absolute difference in firing rate was calculated. This was repeated 10,000 times to obtain a distribution of permuted differences in firing rates. A unit was defined as being discriminating if the absolute value of the observed rate difference was higher than the 99th percentile of the permuted distribution ($p = 0.01$). To calculate whether a unit's activity was significantly modulated we applied the same permutation analysis to compare cue-related firing with baseline activity. In each case, we used a 1 s period, corresponding to the duration of the cue, to determine cue-related firing, and compared this to a 4 s within-trial baseline period (-7 to -3 s, 4 s duration chosen to provide a smooth baseline average).

Onset/offset cell and population overlap analysis. Latency to peak firing during the period between the feature cue and CS ($t = -2$ to 0 s from CS onset) was estimated from the maximum average firing rate using 5 ms time bins and a Gaussian kernel convolution. Firing rate was calculated from the average of both $CS1^-$ and transfer trials (i.e., all trials containing a feature cue). The observed latency distribution across all recorded cells (see Fig. 4C) showed a good fit to the sum of two Lorentzian distributions. We defined the cutoff between onset and offset cells at the local minimum in the latency distribution, which occurred at $t = -1.9$ s from CS onset. The range of latency values was bounded from -2.5 to -1 s. To determine the overlapping population size predicted by chance between the feature, $CS1^-$ and $CS1^+$ cues, we first calculated the percentage of neurons per animal ($n = 8$) that was significantly modulated in response to these three individual cues. We then multiplied these three percentage values together to determine each animal's percentage of overlapping cells predicted by chance. This, in turn, was statistically compared with the observed overlap value of the corresponding animal using a paired t test.

Network state prediction analysis. Analysis of cortical network state (see Fig. 5B, C) was performed separately for each animal, using both $CS1^-$ and transfer trials (i.e., all trials containing a feature cue). For the network state analysis, these two trial types were behaviorally indistinguishable because during the delay period, the animal had no prior knowledge of which CS it would subsequently receive. For each trial, the spike count for each unit was calculated for the 1 s period before the feature presentation [defined as the baseline (BL)], and for the 1 s period occurring before the odor stimulus presentation [defined as the delay (DL)]. This resulted in two paired population rate vectors for each trial to be used in the classification algorithm. We used a binary support vector machine (SVM) classifier with a linear kernel, implemented in the LIBSVM library v3.21 (Chang and Lin, 2011). The classifier was trained to distinguish between population rate vectors on BL and DL periods (Fig. 5B). We used a repeated fivefold cross-validation strategy, so that each training set contained four folds of trials, leaving the remaining fold for testing. Each fold of the data was used once for testing, ensuring that each trial was tested exactly once. During testing, each population rate vector in the tested fold was classified as belonging to either BL or DL

periods. The classifier's performance was defined as the percentage of correctly classified BL and DL periods across all tested folds. We repeated this procedure 500 times, each time shuffling the order of trials allocated to the folds, to account for potential variability across trials in the population and to ensure the most accurate estimate of classifier performance. The average of all 500 accuracy scores was defined as the decoder accuracy score for each dataset. To maximize decoder performance, we determined the optimal SVM misclassification cost parameter, C , via an iterative search across a range of parameters (also using five-fold cross-validation). The final value of C ranged from 0.002 to 0.0625. To determine the chance level of performance for each population, we shuffled the BL and DL labels on the data. We then applied the binary classifiers that were trained on observed data to the randomized datasets in a parallel cross-validation procedure. The mean decoder accuracy score on the randomized data ($\sim 50\%$) was used as chance level for each dataset.

We used a similar approach to classify whether delay period activity before incorrect CS1⁻ trials, was more similar to the baseline period before correct CS1⁺ trials, or the delay period before correct CS1⁻ trials (see Fig. 5D,E). For each trial, the spike count for each unit was calculated for the 1 s baseline period before the odor presentation during correct CS1⁺ trials [defined as the baseline before licking (BLL)], the 1 s delay period occurring before the odor stimulus presentation during correct CS1⁻ trials [defined as the delay before lick withholding (DLW)], and the 1 s delay period occurring before the odor stimulus presentation during incorrect CS1⁻ trials [defined as the delay before errant licking (DLL)]. This resulted in a population rate vector for each trial of each class to be used in the classification algorithm. Because there were an uneven number of correct trial observations (unlike in the paired situation described for the BL vs DL activity classification) we equalized the numbers of correct trials by randomly subsampling the larger population down to the size of the smaller population. This ensured that classification would not be biased toward the type of trial that contained a greater number of observations. After training the classifier on balanced data from the two correctly performed trial types, we then tested all of the DLL observations on the model and asked whether the classifier was more likely to identify activity in the DLL period as a BLL or DLW period. We repeated this procedure 500 times, each time shuffling the order of trials before subsampling, thus creating a new classifier on new combinations of training trials. To maximize decoder performance, we determined the optimal SVM misclassification cost parameter, C . The optimal parameter for each dataset was determined by first subsampling from BLL and DLW trials, and performing fivefold cross validation decoding while systematically varying C . This procedure was performed 100 times, with each iteration containing a new combination of subsampled trials. Thus, we chose the C parameter that resulted in the highest BLL and DLW separation. The final values of C ranged from 0.001 to 0.125.

Experimental design and statistical analyses. All statistical tests were performed in MATLAB or Prism (GraphPad) software. The sample size, type of test used, and probability value is reported in the text and figure legends. All p values < 0.0001 are reported as $p < 0.0001$. One subject (animal 1) was excluded from the analysis of Figure 5E for having only one DLL trial, which prevented a statistically sound analysis.

Results

Behavioral responses reveal a negative occasion setting strategy

In the feature-negative conditioning task, mice ($n = 8$) are exposed to conditioned odor stimuli (CS1 and CS2, 1 s duration)

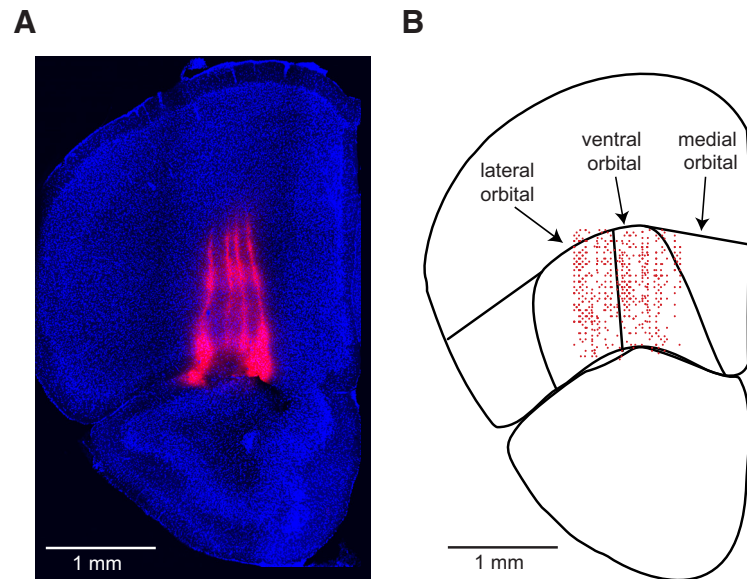


Figure 2. Silicon microprobe recordings in the OFC. **A**, Representative confocal image of a coronal section showing the recording position of the silicon microprobe containing four prongs. Before insertion, the prongs were painted with DiI (red) to facilitate visualization. The section was stained with DAPI (blue). **B**, Coronal section from the Franklin and Paxinos (1997) mouse brain atlas (2.35 mm anterior to bregma) annotated with the estimated position of each putative unit (red dot) in relation to the OFC structure.

that are either followed by reward if no feature cue (mild air puff) was present, or not followed by reward if a feature cue was present before the odor stimulus (Fig. 1A). Therefore, the presence or absence of the feature cue determines the outcome on that trial. On training sessions, we presented three trial types with equal likelihood: CS1⁺, CS1⁻, and CS2⁺ (Fig. 1B, left). Thus, during this training period, the feature cue was presented in half of the CS1 trials, but never paired with the CS2 trials. The final training session, which coincided with electrophysiological recordings, included transfer trials in the form of the same feature cue followed by the CS2 cue (Fig. 1B, right).

On the final training session, the percentage of CS1⁻ trials with licking was significantly reduced relative to CS1⁺ trials (Fig. 1C,E; $p < 0.0001$, paired t test). Thus, mice learned that the feature predicts an unrewarded outcome with respect to the CS1 cue. To determine the specificity of the feature-CS association, we introduced a small percentage (15%) of transfer trials, which animals encountered for the first time during the recording session. Animals showed a reduction in licking on transfer trials relative to CS2⁺ trials (Fig. 1D,F; $p = 0.03$, paired t test). However, the inhibitory effect of the feature on licking in CS1⁻ trials [62% median reduction, 28% interquartile range (IQR)] was significantly greater than its effect on transfer trials (9% median reduction, 21% IQR, $p < 0.0001$, paired t test). Thus, the feature cue primarily suppressed CS1 elicited anticipatory licking behavior (compared with CS2), as predicted by the negative occasion setting model. This selectivity also suggests that information about the feature cue's presence is maintained during the delay period, to guide the animal's decision about whether to lick following the CS presentation.

Feature cues selectively inhibit OFC encoding of conditioned stimuli

Previous studies suggest that the OFC regulates feature-negative behavior (Meyer and Bucci, 2016). However, the neural activity correlates of this behavior have not been studied in this brain area. We used silicon-based microprobes (4 silicon prongs with

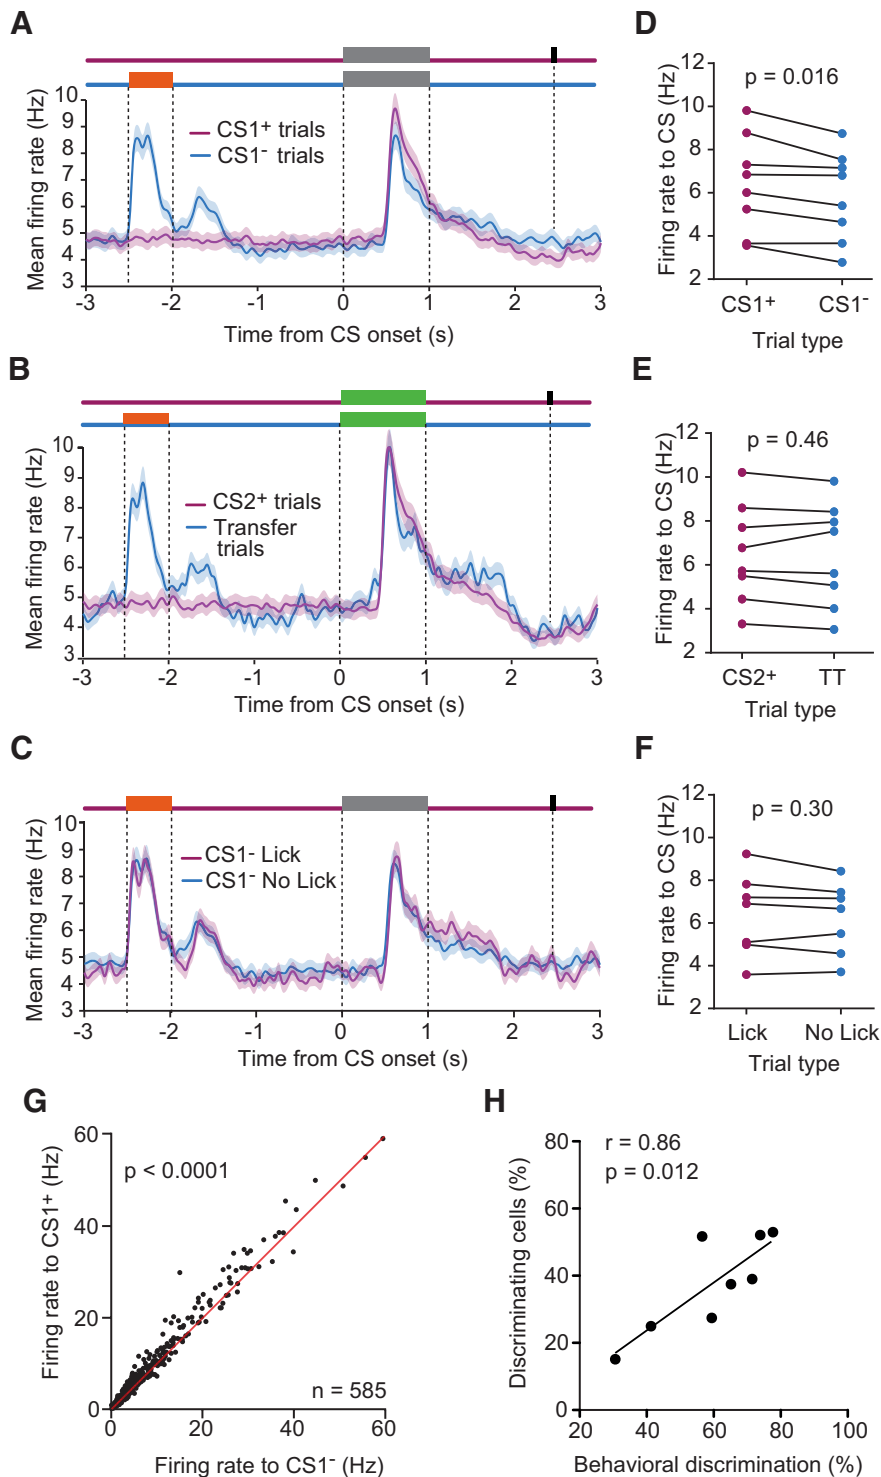


Figure 3. Cue-dependent modulation of OFC activity. **A–C**, Mean firing rate as a function of time in different trial types. Dashed lines represent the onset and offset times of the indicated cue. Data represent mean \pm SEM ($n = 585$ units). Gray bar, CS1; green bar, CS2; orange bar, feature. **A**, Comparison of CS1⁺ with CS1⁻ trials. **B**, Comparison of CS2⁺ with transfer trials. **C**, Comparison of CS1⁻ trials with licking or without anticipatory licking. **D–F**, Mean firing rate per animal during the CS presentation period ($t = 0–1$ s), in different trial conditions. Data represent individual animals ($n = 8$). **D**, CS1⁻ trials exhibit significantly lower firing than CS1⁺ trials ($p = 0.016$, paired t test). **E**, There is no significant difference in mean firing between transfer trials and CS2⁺ trials ($p = 0.46$, paired t test). **F**, There is no significant difference in mean firing between CS1⁻ trials with licking and those without licking ($p = 0.3$, paired t test). **G**, Comparison of the average firing rate per unit during the CS cue presentation period ($t = 0–1$ s) between CS1⁺ and CS1⁻ trial types. Across the population ($n = 585$) there was a significant bias toward lower firing during CS1⁻ trials ($p < 0.0001$, paired t test). **H**, Behavioral discrimination (percentage correct CS1⁺ trials minus percentage incorrect CS1⁻ trials) is significantly correlated with the percentage of OFC units per animal that discriminate between CS1⁺ and CS1⁻ trials (Pearson $r = 0.82$, $p = 0.012$).

64 electrodes each) to simultaneously record from dozens of orbitofrontal units during the final training session ($n = 8$ mice, 48–119 single units per animal). After each recording, we verified the silicon prong locations using confocal microscopy (Fig. 2A), and used these images to estimate the recording site and corresponding unit positions. We found that the measurements were primarily located in the ventral and lateral subregions of the OFC (Fig. 2B).

Based on the finding that the feature cue predominantly diminished levels of anticipatory licking in response to the CS1, we hypothesized that the feature cue would modulate odor stimulus-evoked cortical activity. Consistent with this prediction, we observed that the presence of the feature, on CS1⁻ trials, suppressed the OFC population's mean firing rate relative to CS1⁺ trials during the CS presentation period ($n = 585$ units pooled across 8 mice; Fig. 3A). We then separately examined the mean firing rate in each animal and found that the feature caused a significant reduction in firing rate during the 1 s CS1 presentation period (Fig. 3D; $p = 0.016$, paired t test). In contrast, we did not see any feature effect on mean CS2-evoked firing rate during transfer trials (Fig. 3B,E; $p = 0.46$, paired t test). Furthermore, we found a small but statistically significant difference ($p = 0.045$, paired t test) between the feature-induced reduction in firing rate on CS1 compared with CS2 cues, demonstrating that the feature selectivity inhibits the encoding of the CS1 representation. We also found the OFC does not appear to encode choice, because we did not observe any difference in mean firing rate between CS1⁻ trials with anticipatory licking and CS1⁻ trials without licking (Fig. 3C,F; $p = 0.30$, paired t test).

To further examine the feature cue's effect on OFC neuronal responses to CS1 cues, we compared the firing rates between the CS1⁺ and CS1⁻ trials for each individual neuron during the 1 s cue presentation period ($n = 585$ units pooled across 8 mice). We found that a significant fraction of neurons had a lower firing rate in the CS1⁻ trials (Fig. 3G; $p < 0.0001$, paired t test), suggesting that the feature suppressed the response of a large proportion of OFC neurons. We also found that the percentage of cells per animal that could discriminate between CS1⁺ and CS1⁻ trials during the 1 s CS presentation period was significantly correlated with behavioral discrimination (Fig. 3H; $n = 8$ mice, Pearson $r = 0.86$, $p = 0.012$). Thus,

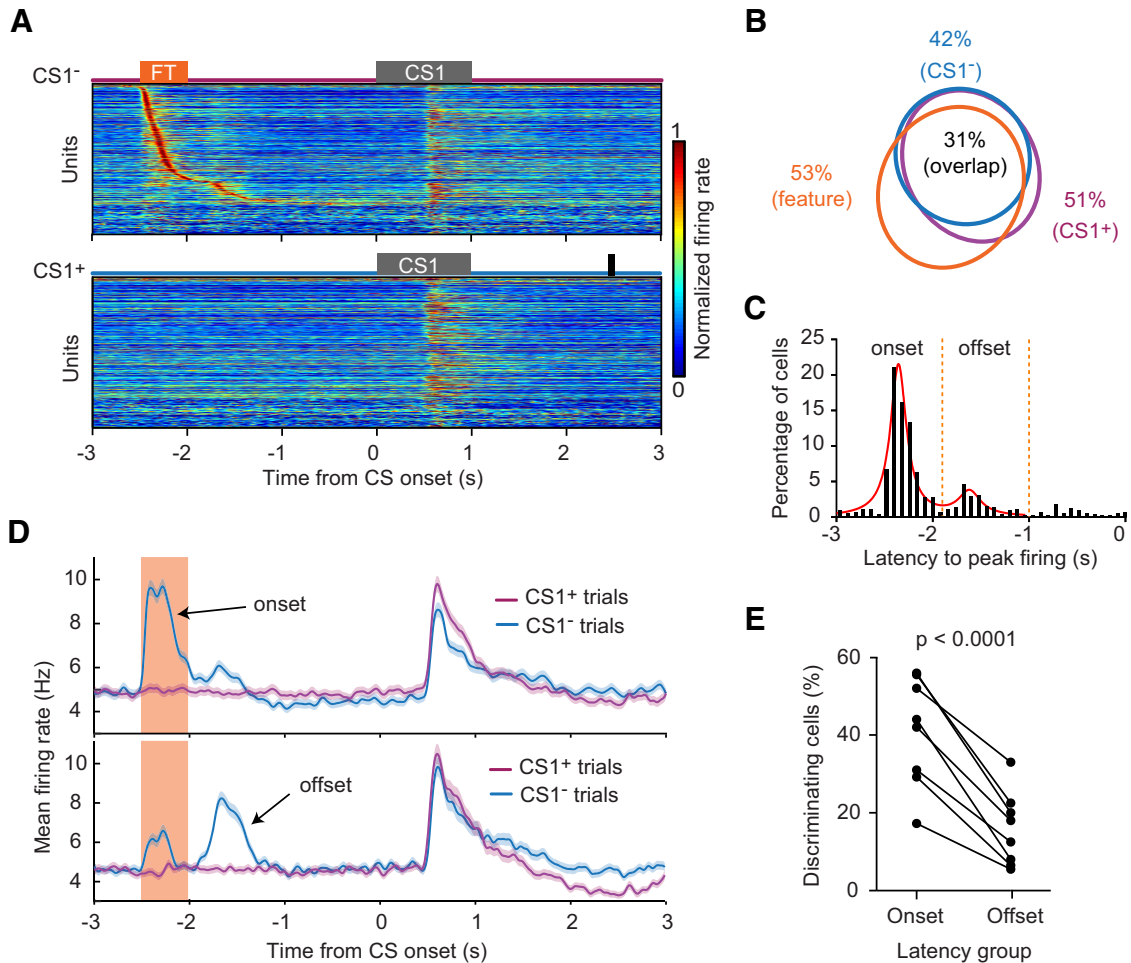


Figure 4. Identification of temporally distinct feature encoding populations. **A**, Mean normalized firing rate as a function of time of the recorded population ($n = 585$ cells). Each cell's firing rate is normalized to its peak firing rate on CS^- trials (top) and CS^+ trials (bottom). Units are ordered by latency to peak firing relative to onset of the feature cue (FT). Units are plotted in the same order in the top and bottom panels (red indicates high firing rate). **B**, Venn diagram showing the overlapping relationship between units that were significantly modulated by the feature cue (orange), the CS1 cue during CS^+ trials (magenta), and the CS1 cue during CS^- trials (blue). Values represent the median percentage of modulated cells across $n = 8$ animals. **C**, Distribution of the latency to peak firing for the recorded population ($n = 585$ units). Two major peaks were resolved using Lorentzian curve fits (red line). Dashed orange lines demarcate the onset and offset cell populations. **D**, Mean firing rate as a function of time during CS^+ and CS^- trials. The top and bottom panels are comprised of onset and offset cells, respectively. Data represent mean \pm SEM ($n = 585$ units). The orange shaded area represents the time during the feature cue presentation. **E**, The percentage of cells that discriminated between CS^+ and CS^- trials was significantly higher in the onset cell population ($n = 8$ mice, $p < 0.0001$, paired t test).

the greater the proportion of OFC units that distinguished between non-feature and feature trial types, the better the animal was at correctly licking to CS^+ trials and correctly withholding to CS^- trials. Therefore, these electrophysiological measurements, together with the corresponding behavioral tests, support the negative occasion setting model by showing that the feature cue selectively suppresses OFC activity and anticipatory behavior following CS1 cues, but not CS2 cues.

Temporally specific feature encoders have unique discriminatory properties

To further understand the encoding properties of feature and CS1 cues, we examined the firing patterns of individual neurons under different stimulus conditions. Across the recorded population, we found that a large proportion of cells appeared to respond to individual cues (feature, CS^+ , or CS^-), or a combination of these cues (Fig. 4A). To quantify this relationship, we calculated the proportion of units that were significantly modulated by single cues or different cue combinations. We found that across $n = 8$ mice, 53% (median, 20% IQR) of the neurons re-

sponded to the feature, whereas 51% (median, 15% IQR) and 42% (median, 28% IQR) of neurons responded to the CS1 in the CS^+ and CS^- trials, respectively (Fig. 4B). Notably, 31% (median, 19% IQR) of the neurons responded to all three cues. This overlap is significantly higher than chance levels (10%, 11% IQR), based on the total number of identified units in the OFC (paired t test, $p < 0.0001$), suggesting a common representation of the cells that encode these stimuli. These results suggest that not only is OFC encoding of a reward-associated stimulus (CS1) modulated by the feature cue, but that this circuit is strongly tuned by stimuli that activate overlapping neuronal subpopulations.

In the population of feature responsive cells, we found evidence for heterogeneous response properties, with some cells responding early, and others later to the feature cue (Fig. 4A). We calculated each unit's latency to peak firing during the feature period, and found that the latency values appeared to cluster into two distinct firing groups (Fig. 4C). One group of neurons fired maximally around the feature onset time (onset cells), whereas another group preferentially fired around the feature offset time

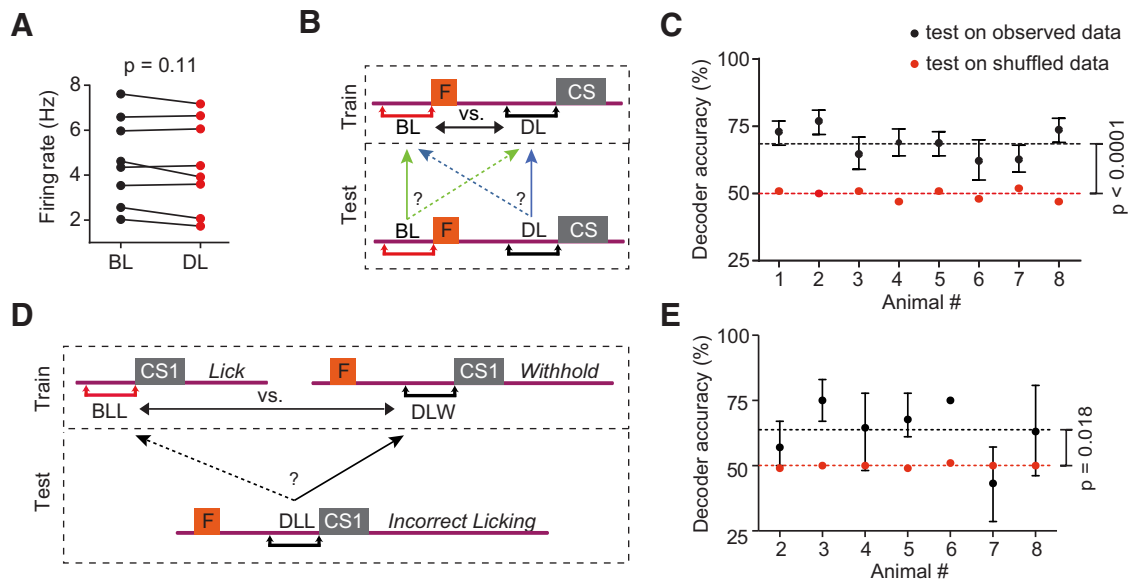


Figure 5. A distinct network state initiated by the feature cue. **A**, There is no significant difference in mean OFC firing rate during the final 1 s of the DL period and a 1 s BL period before feature cue presentation ($p = 0.11$, paired t test). **B**, Strategy used to determine whether the network state in the BL period is distinct from that of the DL period. This two-step process required training (top dashed box) and testing (bottom dashed box) a binary classifier. During testing, each period (BL, green arrows; DL, blue arrows) was classified as either a correct match (e.g., BL classified as BL, solid arrow) or an incorrect match (e.g., BL classified as DL, dashed arrow). **C**, Mean classifier accuracy per animal of the classifier in **B** (accuracy defined as the percentage of correctly classified BL and DL periods across all tested folds, black) was significantly above chance levels shown in red ($n = 8$, $p < 0.0001$, paired t test). The average accuracy across the experimental group was $69 \pm 2\%$ (mean \pm SEM, dashed black line). **D**, Strategy used to classify whether delay period activity before incorrect CS1⁻ trials (DLL), was more similar to the baseline period before correct CS1⁺ trials (BLL), or the delay period before correct CS1⁻ trials (DLW). The classifier was trained (top dashed box) to distinguish population activity during BLL periods from DLW periods. During testing (bottom dashed box) DLL activity was compared with BLL and DLW activity and classified as more similar to either BLL (dashed line) or DLW (solid line). **E**, Mean classifier accuracy per animal of the classifier in **D** (accuracy defined as the percentage of DLL periods that were labeled as DLW, black) was significantly above chance levels shown in red ($n = 7$, $p = 0.018$, paired t test). The average accuracy across the experimental group was $64 \pm 4\%$ (mean \pm SEM, dashed black line). Note that animal 1 only had one DLL trial and was excluded from the analysis in **E**. Error bars in **C** and **E** represent 95% confidence intervals across all iterations and dashed lines represent the average values across all animals.

(offset cells). We separately examined the mean CS1-triggered firing rate of the onset and offset cells, and found that they appeared to show different responses during the CS1 presentation period (Fig. 4D). Specifically, the mean firing rate of onset cells appeared markedly reduced in CS1⁻ relative to CS1⁺ trials (Fig. 4D, top). This suggests that the CS1 representation associated with the onset population is highly susceptible to suppressive properties of the feature. In contrast, the response of offset cells to CS1 was less perturbed by the feature (Fig. 4D, bottom). To quantify these differences, we compared the number of cells within each group that significantly discriminated between the CS1⁺ and CS1⁻ trial types during the 1 s CS1 presentation period. We found that the onset group contained a significantly larger proportion of discriminating cells relative to the offset group (Fig. 4E; $p < 0.0001$, paired t test).

Feature cues trigger a distinct network state in the delay period

If the feature cue influences subsequent OFC encoding of reward-conditioned stimuli, we hypothesized that information about whether the feature cue was present is maintained in the OFC throughout the delay period. Previous work suggests that delay periods in working memory tasks often coincide with persistent firing patterns in the prefrontal cortex (Fuster and Alexander, 1971; Goldman-Rakic, 1995; Fuster, 2005; Riley and Constantinidis, 2015). However, our data revealed that the average firing rate in the OFC returns to baseline levels before the CS1 onset (Figs. 3A, 4D), suggesting that the feature cue does not trigger persistent changes in mean spiking activity. In support of this observation, there was no significant difference in mean firing rate between the final 1 s of the DL period, and a 1 s BL period before feature cue

onset (Fig. 5A; $n = 8$ mice, $p = 0.11$, paired t test). We therefore wondered whether the OFC could still maintain the information about the feature cue's presence during the delay period without any significant persistent activity signal. We speculated that if the OFC is maintaining this information, it does so through an activity silent but distinct network state (Stokes et al., 2013; Stokes, 2015), that does not give rise to an overt change in mean firing rate. An alternative possibility is that another region outside the OFC is exclusively responsible for maintaining the feature cue information. To determine whether OFC networks exhibit dynamics during the delay period that are distinct from the BL period, we used a decoder to distinguish between population activity in the DL and BL periods from the same trial containing a feature cue (Fig. 5B). The decoder was applied to simultaneously recorded populations of cells from individual animals. Our results reveal that for all animals tested, the decoder performed significantly above chance levels in discriminating between activity in the BL and DL periods (Fig. 5C; $n = 8$ mice, $p < 0.0001$, paired t test). The average accuracy was $69 \pm 2\%$ (mean \pm SEM, dashed black line). To rule out any differential interaction between the paired BL and DL periods and the previous trial, we also compared the BL and DL periods from separate trials: CS1⁺ and CS1⁻, respectively. In this case, the decoder also performed significantly above chance levels in discriminating between activity in the BL and DL periods ($n = 8$ mice, $p < 0.0001$, paired t test, data not shown). The average accuracy was $69 \pm 3\%$ (mean \pm SEM), which is very close to our value using the paired period method. A direct comparison revealed no significant differences ($n = 8$ mice, $p = 0.96$, paired t test), indicating that both approaches produce the same result. Together, these findings suggest that, despite the absence of an overt change in mean

population firing rate, the feature induces a distinct network state in the OFC during the delay period.

Finally, we examined whether OFC network dynamics during the delay period also provide information about the subsequent behavioral choice of the animal on that trial. In other words, is there a prospective code during the delay period that predicts whether or not the mouse will lick? To test this, we took advantage of the observation that mice sometimes licked incorrectly during CS1⁻ trials (Fig. 1C). We trained a classifier to distinguish between population activity occurring during the baseline period before correct CS1⁺ trials (BLL), and the delay period before correct CS1⁻ trials (DLW). First, using cross-validation, we found that the classifier could distinguish these periods above chance levels ($n = 8$ mice, $p < 0.0001$, paired t test, data not shown), consistent with a distinct network state during DL and BL periods shown in Figure 5C. We next examined whether OFC population activity in the delay period before incorrect CS1⁻ trials (DLL) was classified more frequently as a BLL or DLW period (Fig. 5D). There was a significant preference for the classifier to label DLL as a DLW period (Fig. 5E; $n = 7$ mice, $p = 0.018$, paired t test). The average accuracy was $64 \pm 4\%$ (mean \pm SEM, dashed black line). Thus, it appears that in the OFC, the feature rather than the behavioral outcome (i.e., licking) dictates the delay period network state. This is consistent with our earlier findings showing no significant difference in mean firing rate between correct and incorrect CS1⁻ trials (Fig. 3C,F). These findings suggest that the feature triggers a network state that is maintained throughout the delay period, which could function to downregulate the network's response to the CS1 stimulus.

Discussion

This is, to our knowledge, the first study to show the neural dynamics that may underlie an occasion setter's ability to modulate behavior. A key insight that this study reveals is the selective nature of the association between the feature cue and the conditioned odor stimulus (Holland, 1984). The feature causes animals to suppress their conditioned responding in the form of anticipatory licking to a trained stimulus (CS1). However, the ability of the feature to suppress conditioned responding does not transfer to another stimulus (CS2) that had never previously been paired with the feature. Neural recordings in the OFC complement this finding by showing that the feature-negatively modulates activity triggered by CS1, but this modulation effect does not transfer to the CS2 cue. This lack of transfer, observed in both our behavioral and neurophysiological data, rules out the simple Rescorla–Wagner model because this model posits that the feature's inhibitory properties should transfer to any CS paired with that reward. The fact that we did not observe transfer thus provides strong evidence against a direct inhibitory link between the feature cue and reward. Furthermore, our data suggest that the OFC may be involved in the task, because a measure of the level of OFC modulation by the feature (percentage of cells per animal that discriminate between CS1⁺ and CS1⁻ trials) significantly correlates with an individual animal's behavioral discrimination. Together, our findings provide strong evidence for the negative occasion setting model (Holland, 1984; Lamarre and Holland, 1987) in which a feature cue can modulate the ability of a separate cue to retrieve its reward association.

Our data also suggest a possible OFC information transfer mechanism between the feature and conditioned odor stimulus during the delay period. Many studies on working memory have found persistent changes in mean population firing activity that accompany the delay period (Fuster and Alexander,

1971; Goldman-Rakic, 1995; Miller et al., 1996; Miller and Cohen, 2001; Pasternak and Greenlee, 2005; Liu et al., 2014). Although we found that many cortical neurons were activated within ~ 1.5 s of the feature cue's presentation, this activity did not appear to persist into the final 1 s of the delay period, suggesting that the OFC subregions that were targeted here do not exhibit sustained changes in activity. Of course, this observation does not rule out the possibility that persistent activity occurs in other brain areas. On the other hand, a number of studies suggest that sustained activity is not necessary to retain task-relevant information (Jensen and Tesche, 2002; Howard et al., 2003; Riggall and Postle, 2012; Ester et al., 2015; Lundqvist et al., 2016). Intriguingly, an activity silent model of working memory raises the possibility that information is retained in the patterns of network-level activity (Stokes et al., 2013; Stokes, 2015). To examine whether such an effect could be taking place in the OFC during the final 1 s of the delay period in our task, we used a machine learning-based decoding algorithm to assess whether this time period coincides with a distinct network state. In all mice tested the decoder was able to accurately distinguish delay from baseline period activity at above chance levels, consistent with the activity silent working memory model (Stokes et al., 2013; Stokes, 2015). Thus, our data indicate that the OFC has the potential to transfer information about the feature cue across the delay period.

Our results suggest that the OFC uses the feature as a source of rule information to regulate behavioral responses. As discussed above, the degree to which the feature cue suppresses anticipatory licking correlates with its ability to modulate neural activity to the conditioned odor stimulus. In contrast, we found no change in OFC activity during trials when animals incorrectly lick during a feature-negative trial. Moreover, our classifier results suggest that the network state during the delay period before incorrect CS1⁻ trials (DLL) is significantly different from the state during the baseline period before correct CS1⁺ trials (BLL), even though both types of trials contain licking. These two pieces of evidence suggest that the OFC code is relatively insensitive to behavioral choice. Thus, our data are consistent with a number of other studies indicating the importance of rule encoding in the OFC (Buckley et al., 2009; Tsujimoto et al., 2009, 2012; Johnson et al., 2016; Sleezer et al., 2016).

The information coding properties revealed here provide insight into how the brain could quickly manipulate information at more abstract levels to regulate behavior. The feature appears to trigger a distinct network state that specifically interacts with its trained conditioned odor stimulus. This may occur by inducing a temporary functional reweighting of synaptic connections within OFC microcircuits (Fujisawa et al., 2008; Stokes, 2015). As a whole, this model fits well with the viewpoint that the OFC provides the animal with a cognitive map of task space (Roesch et al., 2006; Wilson et al., 2014; Cooch et al., 2015; Sharpe et al., 2015; Lopatina et al., 2016; Wikenheiser and Schoenbaum, 2016) because the extent to which the conditioned odor stimulus alters neural activity is mediated by the network state set by the feature. Together, our observations provide a potential mechanism that helps to explain how animals can rapidly interpret the meaning of a conditionally rewarded cue to make timely behavioral decisions.

References

- Abramson CI, Cakmak I, Duell ME, Bates-Albers LM, Zuniga EM, Pendegrift L, Barnett A, Cowo CL, Warren JJ, Albritton-Ford AC, Barthell JF, Hranitz JM, Wells H (2013) Feature-positive and feature-negative learning in honey bees. *J Exp Biol* 216:224–229. [CrossRef Medline](#)

- Bakhurin KI, Mac V, Golshani P, Masmanidis SC (2016) Temporal correlations among functionally specialized striatal neural ensembles in reward conditioned mice. *J Neurophysiol* 115:1521–1532. [CrossRef Medline](#)
- Buckley MJ, Mansouri FA, Hoda H, Mahboubi M, Browning PG, Kwok SC, Phillips A, Tanaka K (2009) Dissociable components of rule-guided behavior depend on distinct medial and prefrontal regions. *Science* 325:52–58. [CrossRef Medline](#)
- Bueno JL, Holland PC (2008) Occasion setting in Pavlovian ambiguous target discriminations. *Behav Processes* 79:132–147. [CrossRef Medline](#)
- Chang CC, Lin CJ (2011) LIBSVM: a library for support vector machines. *ACM Trans Intell Syst Technol* 2:27. [CrossRef](#)
- Cooch NK, Stalnaker TA, Wied HM, Bali-Chaudhary S, McDannald MA, Liu TL, Schoenbaum G (2015) Orbitofrontal lesions eliminate signalling of biological significance in cue-responsive ventral striatal neurons. *Nat Commun* 6:7195. [CrossRef Medline](#)
- Ester EF, Sprague TC, Serences JT (2015) Parietal and frontal cortex encode stimulus-specific mnemonic representations during visual working memory. *Neuron* 87:893–905. [CrossRef Medline](#)
- Franklin KB, Paxinos G (1997) *The mouse brain in stereotaxic coordinates*. San Diego: Academic.
- Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11:823–833. [CrossRef Medline](#)
- Fuster JM (2005) *Memory*. In: *Cortex and mind*, pp 111–142. New York: Oxford UP.
- Fuster JM, Alexander GE (1971) Neuron activity related to short-term memory. *Science* 173:652–654. [CrossRef Medline](#)
- Goldman-Rakic PS (1995) Cellular basis of working memory. *Neuron* 14:477–485. [CrossRef Medline](#)
- Guo ZV, Hires SA, Li N, O'Connor DH, Komiyama T, Ophir E, Huber D, Bonardi C, Morandell K, Gutnisky D, Peron S, Xu NL, Cox J, Svoboda K (2014) Procedures for behavioral experiments in head-fixed mice. *PLoS One* 9:e88678. [CrossRef Medline](#)
- Holland PC (1984) Differential effects of reinforcement of an inhibitory feature after serial and simultaneous feature negative discrimination training. *J Exp Psychol Anim Behav Process* 10:461–475. [CrossRef Medline](#)
- Holland PC (1989) Transfer of negative occasion setting and conditioned inhibition across conditioned and unconditioned stimuli. *J Exp Psychol Anim Behav Process* 15:311–328. [CrossRef Medline](#)
- Holland PC (1995a) Transfer of occasion setting across stimulus and response in operant feature positive discriminations. *Learn Motiv* 26:239–263. [CrossRef](#)
- Holland PC (1995b) The effects of intertrial and feature-target intervals on operant serial feature-positive discrimination learning. *Anim Learn Behav* 23:411–428. [CrossRef](#)
- Holland PC (1997) Brain mechanisms for changes in processing of conditioned stimuli in Pavlovian conditioning: implications for behavior theory. *Anim Learn Behav* 25:373–399. [CrossRef](#)
- Holland PC (2008) Cognitive versus stimulus-response theories of learning. *Learn Behav* 36:227–241. [CrossRef Medline](#)
- Howard MW, Rizzuto DS, Caplan JB, Madsen JR, Lisman J, Aschenbrenner-Scheibe R, Schulze-Bonhage A, Kahana MJ (2003) Gamma oscillations correlate with working memory load in humans. *Cereb Cortex* 13:1369–1374. [CrossRef Medline](#)
- Jensen O, Tesche CD (2002) Frontal theta activity in humans increases with memory load in a working memory task. *Eur J Neurosci* 15:1395–1399. [CrossRef Medline](#)
- Johnson CM, Peckler H, Tai LH, Wilbrecht L (2016) Rule learning enhances structural plasticity of long-range axons in frontal cortex. *Nat Commun* 7:10785. [CrossRef Medline](#)
- Lamarre J, Holland PC (1987) Transfer of inhibition after serial feature negative discrimination training. *Learn Motiv* 18:319–342. [CrossRef](#)
- Liu D, Gu X, Zhu J, Zhang X, Han Z, Yan W, Cheng Q, Hao J, Fan H, Hou R, Chen Z, Chen Y, Li CT (2014) Medial prefrontal activity during delay period contributes to learning of a working memory task. *Science* 346:458–463. [CrossRef Medline](#)
- Lopatina N, McDannald MA, Styer CV, Peterson JF, Sadacca BF, Cheer JF, Schoenbaum G (2016) Medial orbitofrontal neurons preferentially signal cues predicting changes in reward during unblocking. *J Neurosci* 36:8416–8424. [CrossRef Medline](#)
- Lundqvist M, Rose J, Herman P, Brincat SL, Buschman TJ, Miller EK (2016) Gamma and beta bursts underlie working memory. *Neuron* 90:152–164. [CrossRef Medline](#)
- Meyer HC, Bucci DJ (2016) Imbalanced activity in the orbitofrontal cortex and nucleus accumbens impairs behavioral inhibition. *Curr Biol* 26:2834–2839. [CrossRef Medline](#)
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202. [CrossRef Medline](#)
- Miller EK, Erickson CA, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16:5154–5167. [Medline](#)
- Nallan GB, Brown MB, Edmonds C, Gillham V, Kowalewski K, Miller JS (1981) Transfer effects in feature-positive and feature-negative learning by adult humans. *Am J Psychol* 94:417–429. [CrossRef Medline](#)
- Pace GM, McCoy DF (1981) Effects of stimulus contact on the feature-positive effect. *Am J Psychol* 94:153–158. [CrossRef Medline](#)
- Pace GM, McCoy DF, Nallan GB (1980) Feature-positive and feature-negative learning in the rhesus monkey and pigeon. *Am J Psychol* 93:409–427. [CrossRef Medline](#)
- Pasternak T, Greenlee MW (2005) Working memory in primate sensory systems. *Nat Rev Neurosci* 6:97–107. [CrossRef Medline](#)
- Rescorla RA (1969) Pavlovian conditioned inhibition. *Psychol Bull* 72:77–94. [CrossRef](#)
- Rescorla RA, Holland PC (1977) Associations in Pavlovian conditioned inhibition. *Learn Motiv* 8:429–447. [CrossRef](#)
- Rescorla RA, Wagner AR (1972) A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: *Classical conditioning II: Current research and theory*, pp 64–99. New York: Appleton-Century-Crofts.
- Riggall AC, Postle BR (2012) The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. *J Neurosci* 32:12990–12998. [CrossRef Medline](#)
- Riley MR, Constantinidis C (2015) Role of prefrontal persistent activity in working memory. *Front Syst Neurosci* 9:181. [CrossRef Medline](#)
- Robinson S, Keene CS, Iaccarino HF, Duan D, Bucci DJ (2011) Involvement of retrosplenial cortex in forming associations between multiple sensory stimuli. *Behav Neurosci* 125:578–587. [CrossRef Medline](#)
- Roesch MR, Taylor AR, Schoenbaum G (2006) Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron* 51:509–520. [CrossRef Medline](#)
- Schmajuk NAE, Holland PCE (1998) Occasion setting: associative learning and cognition in animals. (Schmajuk NA, Holland PC, eds). Washington: American Psychological Association.
- Sharpe MJ, Wikenheiser AM, Niv Y, Schoenbaum G (2015) The state of the orbitofrontal cortex. *Neuron* 88:1075–1077. [CrossRef Medline](#)
- Shobe JL, Claar LD, Parhami S, Bakhurin KI, Masmanidis SC (2015) Brain activity mapping at multiple scales with silicon microprobes containing 1,024 electrodes. *J Neurophysiol* 114:2043–2052. [CrossRef Medline](#)
- Sleezer BJ, Castagno MD, Hayden BY (2016) Rule encoding in orbitofrontal cortex and striatum guides selection. *J Neurosci* 36:11223–11237. [CrossRef Medline](#)
- Stokes MG (2015) “Activity-silent” working memory in prefrontal cortex: a dynamic coding framework. *Trends Cogn Sci* 19:394–405. [CrossRef Medline](#)
- Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J (2013) Dynamic coding for cognitive control in prefrontal cortex. *Neuron* 78:364–375. [CrossRef Medline](#)
- Tsujimoto S, Genovesio A, Wise SP (2009) Monkey orbitofrontal cortex encodes response choices near feedback time. *J Neurosci* 29:2569–2574. [CrossRef Medline](#)
- Tsujimoto S, Genovesio A, Wise SP (2012) Neuronal activity during a cued strategy task: comparison of dorsolateral, orbital, and polar prefrontal cortex. *J Neurosci* 32:11017–11031. [CrossRef Medline](#)
- Wikenheiser AM, Schoenbaum G (2016) Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex. *Nat Rev Neurosci* 17:513–523. [CrossRef Medline](#)
- Wilson RC, Takahashi YK, Schoenbaum G, Niv Y (2014) Orbitofrontal cortex as a cognitive map of task space. *Neuron* 81:267–279. [CrossRef Medline](#)