Prospective Representations in Rat Orbitofrontal Ensembles

Jingfeng Zhou^{1*}, Wenhui Zong¹, Chunying Jia², Matthew P.H. Gardner¹, and Geoffrey Schoenbaum^{1*}

¹ Intramural Research Program of the National Institute on Drug Abuse, Baltimore MD, USA

² Department of Computer Science and Electrical Engineering, University of Maryland, Baltimore County, Baltimore, MD, USA

*Correspondence or requests for material should be addressed to J.Z. (jingfeng.zhou@nih.gov) or G.S. (geoffrey.schoenbaum@nih.gov).

Abstract

The orbitofrontal cortex (OFC) has been proposed to encode expected outcomes, which is thought to be important for outcome-directed behavior. However, such neural encoding can also often be explained by the recall of information about the recent past. To dissociate the retrospective and prospective aspects of encoding in the OFC, we designed a non-spatial, continuous, alternating odor-sequence task that mimicked a continuous T-maze. The task consisted of two alternating sequences of four odor-guided trials (2 sequences $\times$ 4 positions). In each trial, rats were asked to make a “go” or “no-go” action based on a fixed odor-reward contingency. Odors at both the first and last positions were distinct across the two sequences, such that they resembled unique paths in the past and future, respectively; odors at positions in between were the same and thus resembled a common path. We trained classifiers using neural activity to distinguish between either sequences or positions and asked whether the neural activity patterns in the common path were more like the ones in the past or the future. We found a proximal prospective code for sequence information as well as a distal prospective code for positional information, the latter of which was closely associated with rats’ ability to predict future outcomes. This study demonstrates a prospective behaviorally-relevant predictive code in rat OFC.

26 Introduction

27 The orbitofrontal cortex (OFC) signals expected outcomes, which is believed to be fundamentally important
 28 for outcome-directed behavior (Rudebeck & Murray, 2014; Stalnaker, Cooch, & Schoenbaum, 2015; Wallis,
 29 2011). The vast majority of electrophysiological evidence supporting this assertion comes from behavioral
 30 settings where different cues predict reward outcomes with different sizes, identities, probabilities, and
 31 delays, etc (Kepecs, Uchida, Zariwala, & Mainen, 2008; Klein-Flugge, Barron, Brodersen, Dolan, & Behrens,
 32 2013; Padoa-Schioppa & Assad, 2006; Roesch, Taylor, & Schoenbaum, 2006; Tremblay & Schultz, 1999).
 33 Meanwhile, the past sensory cues, rewards, spatial directions, and behavioral choices are also reported to
 34 be reflected in the OFC neural activities (Feierstein, Quirk, Uchida, Sosulski, & Mainen, 2006; Kennerley,
 35 Behrens, & Wallis, 2011; Nogueira et al., 2017; Riceberg & Shapiro, 2017; Saez, Saez, Paton, Lau, & Salzman,
 36 2017; Young & Shapiro, 2011; Zhou, Jia, Feng, Bao, & Luo, 2015). Such findings have been taken as evidence
 37 that the OFC, together with contributions from other interconnected brain regions such as the
 38 hippocampus, might provide a neural mechanism with which animals could mentally travel through a task
 39 model in time and recall the past events and simulate future outcomes (Behrens et al., 2018; Wang,
 40 Schoenbaum, & Kahnt, 2020; Wikenheiser & Schoenbaum, 2016; R. C. Wilson, Takahashi, Schoenbaum, &
 41 Niv, 2014).

42 However, from studies where different past cues or episodes lead to different future outcomes, it is not
 43 clear whether the neural patterns observed are representing the future versus simply providing a record of
 44 past events. Additionally, in most behavioral settings, future events consist of rewards, whereas past events
 45 are normally intrinsically-neutral sensory cues. As a result, prioritized reward value processing would lead
 46 to a biased finding of stronger prospective coding (Wallis, 2007; Xie, Nie, & Yang, 2018).

47 In the present study, we resolved these confounds by recording single-unit activity in the OFC of rats
 48 performing a non-spatial, continuous odor sequence task, conceptually similar to a continuous T-maze.
 49 Combining both single-unit and neural ensemble analyses, we tested whether the OFC neural ensemble
 50 patterns during the overlapping paths in the “virtual” T-maze task resemble neural activities that occur in
 51 the past or in the future.

Results

Continuous, alternating odor-sequence task

The behavioral task developed in this study was designed to mimic the continuous T-maze alternation task that has been commonly used to assess spatial memory (Greene & Naranjo, 1986; Verma & Moghaddam, 1996). Rather than moving through a sequence of locations in space, subjects moved through a sequence of odors. On each trial in the odor sequence (Figure 1A), the rats were presented with the appropriate odor at a central port and then had to decide whether to respond for a sucrose reward by poking into a fluid well (“Go”) or to withhold responding on non-rewarded trials (“No-Go”). The decision to respond for reward could be made correctly based simply on odor identity, or by using information available from the sequence. The task used 6 different odors, arranged in two 4-odor sequences: S1 and S2 (Figure 1B). The 4 odors in each sequence were designated as 4 positions (P1 – P4).

After training in this simple task, we recorded single-unit activity bilaterally from the lateral OFC (n = 1568 neurons; 4 rats). During recording, the rats performed the task with high accuracy as assessed by percent correct (%correct; Figure 1D) on each of the 8 trial types. Importantly, we also found the time rats spent to initiate a trial (i.e., poke latency; Figure 1E) was different depending on reward availability on current trials, suggesting the rats were using information available from the sequence structure to make predictions about future outcomes that influenced their responses.

As in a continuous T-maze alternation task, each position shares the same past, current, and future reward structure, thereby eliminating any bias that value might have on prospective versus retrospective encoding. Further, differences in sensory experience are structured, so that some positions (P1 and P4) differ locally, since the odors are unique at these positions in S1 and S2, while sharing the same recent past and future events (they come from and return to P2 and P3, each of which share a single odor cue in S1 and S2), whereas other positions (P2 and P3) are similar locally, since the local odors are the same, but differ in recent past and future events (they come from and go to P1 and P4, each of which have different odor cues in each sequence). This arrangement provides a unique opportunity to dissociate retrospective versus prospective neural representations. Specifically, activity distinguishing the sequences should weaken across P2 and P3 if it is retrospective, while it should grow stronger if it is prospective. Further retrospective activity might resemble neural activity patterns from the past (P1), while prospective activity might resemble activity in the future (P4). We analyzed activity recorded in OFC in well-trained rats performing this task to test these predictions.

Distinguishing odor sequences by OFC single-units and neural ensembles

We first examined OFC neural activity at the single-unit level. We found some neurons showing differential responses to sequences S1 and S2 at all 4 positions (Figure 2A-D). A selectivity analysis indicated that the neurons showed the most selectivity at P1 and P4 around odor sampling, with fewer neurons showing selectivity at P2 and P3 (Figure 3A). Indeed, the number of selective neurons was above the chance level (5%) only around the odor for P2 and after the outcome for P3.

Next, we examined the ability of the pattern of activity across all recorded neurons to decode sequences (S1 vs. S2) within the individual task epochs at each position. Consistent with the single-unit selectivity analysis (Figure 3A), decoding accuracy was the highest at P1 and P4, while at P2, the highest decoding happened before and after the odor delivery, and at P3, the highest decoding happened after the outcome (Figure 3B-C).

A proximal prospective code about odor sequences

To study whether the neural patterns during the delay epochs (epochs after P1 odor time and before P4 odor time) resembled the past (S1 vs. S2; odors 5+ vs. 3+ at P1) or the future (S1 vs. S2; odors 2+ vs. 4+ at P4), we trained linear support vector machine (SVM) classifiers to distinguish the sequences during the odor period at either P1 or P4 (i.e., retrospective and prospective templates; Figure 4A-B) and then used each classifier to decode the neural activity patterns in all task epochs. This analysis revealed chance decoding at most points at the delay epochs, particularly for decoding by the classifier trained with the retrospective template, which dropped to chance immediately after odor sampling in P1 (Figure 4A). Decoding by the prospective template was also at chance for most of the delay epochs, however it increased rapidly at the inter-trial interval (ITI) and initial epochs of the P4 trials before the odor was presented (Figure 4B), suggesting the emergence of a prospective representation of the impending odors.

A distal prospective code to reflect future positions

A prospective code distinguishing the odor *sequences* was apparent in the run up to the odor in P4, however the prior analysis found no evidence of a stable, sustained prospective code across the entire delay epoch. This is not completely surprising because, as we previously reported, sequences in an odor sequence task tend to be generalized in the OFC if distinguishing them is not necessary for the rats to correctly perform the task (Zhou et al., 2019). This is even true when the odors differ at a given position, thus the failure of OFC to distinguish the sequences during P2 and P3 here is consistent with that prior data.

However the *positional* information was task-relevant since it was used by rats to calculate their current distance to future reward, evidenced by their different poke latencies on trial types with different

probabilities of reward (Figure 1E). Thus we next asked whether we could find a prospective code for this *positional* information (i.e., a mental simulation of future epochs collapsing across the *sequence* information at each position) during the delay epochs.

To do this, we lumped S1 and S2 together at each task epoch and built a binary SVM classifier using neural activity from odor sampling during P1 and P4 (P1 vs. P4; odor time); then we used this classifier to reexamine how the neural activity patterns evolved during the delay epochs (Figure 5). Interestingly this analysis did not reveal a clear pattern of representation when averaged across rats, however when each rat was analyzed separately it revealed significant prospective activity in 3 subjects that was masked by retrospective activity in one subject (Figure 6A-D). Importantly, such a prospective code appeared early at P2 and P3 (Rat #2 and #3), and even at P1 (Rat #4), in a phasic but not tonic manner, which is dramatically different from the proximal prospective code about sequences as shown in Figure 4. Moreover, the emergence of a distal prospective ensemble code during the delay epochs was closely associated with rats' poke latency; the stronger the prospective activity, the stronger the poke latency differed prior to rewarded versus non-rewarded trials for a subject (Figure 6E-H).

To further confirm this finding, we plotted the individual sessions based on the poke latency difference between rewarded and non-rewarded trial types (Figure 7). The bimodal distribution in Figure 7A shows that the sessions with low differences (i.e. little influence of future outcome) all came from Rat #1. When these sessions were excluded, the analysis revealed a clear prospective code similar to that was seen with individual rats (Figures 7B, 6B-D).

Together, these results suggest that both proximal and distal prospective codes exist in the OFC neural ensemble activities.

Discussion

Through analyzing neural activities recorded in the OFCs of rats performing an odor sequence alternation task, we found two types of prospective neural ensemble codes for future events. OFC ensembles used a proximal code for immediate sequence information, while used a distal neural code for future positional information. The finding provides direct electrophysiological evidence that OFC activity anticipates the future even when confounds related to differences in past events or the prioritization of future reward information are controlled.

The stringent control on possible confounds addresses gaps in our understanding that are typically ignored. For instance, differential activity after a cue but before a reward is often taken as evidence of encoding of expected outcomes and yet it could equally well be a retrospective activity, a reflection of trace memory. Further, even if multiple cues are used to predict the same reward, the value asymmetry could lead to a bias toward apparently prospective representation even where a functional bias does not exist. The task used here, though simple to perform and perhaps a bit boring even, resolves these confounds, allowing us to discern activity that is clearly prospective.

The distal prospective ensemble code is particularly intriguing because it seems to fit well with a role in prospective memory (“remembering to remember”), a higher-order brain function found in humans, non-human primates, and rodents (Beran, Evans, Klein, & Einstein, 2012; Beran, Perdue, Bramlett, Menzel, & Evans, 2012; Evans & Beran, 2012; McDaniel & Einstein, 2007; A. G. Wilson & Crystal, 2012; A. G. Wilson, Pizzo, & Crystal, 2013). There are proposed to be three phases of prospective memory: activation or initial encoding -> inactivation when subjects are engaged in other irrelevant activities -> and reactivation when relevant information (event-based or time-based) is encountered (McDaniel & Einstein, 2007; A. G. Wilson & Crystal, 2012; A. G. Wilson et al., 2013). The pattern of prospective encoding identified here in OFC matches this evolution, mostly disappearing in the ITI periods and during the delay epochs except for brief activations when relevant sensory information was delivered around the odor period. This pattern illustrates that OFC is important for anticipating future events but also illustrates that active spiking in OFC is not sufficient for the memory to be maintained. OFC is likely supported in this by other brain regions, which may hold sustained or dynamic neural activation during the long delay time. However it is likely that maintaining such information across the relatively long delay period used here requires additional mechanisms such as short-term synaptic plasticity, either in OFC or elsewhere, which is not necessarily reflected in increased or decreased neural firing rates (Barbosa et al., 2020; Mongillo, Barak, & Tsodyks, 2008; Stokes, 2015). More investigations are needed in the future to shed light on this important question.

ACKNOWLEDGMENTS

The authors thank the NIDA IRP histology core for technical assistance with histology. This work was supported by the Intramural Research Program at the National Institute on Drug Abuse (ZIA-DA000587). The opinions expressed in this article are the authors' own and do not reflect the view of the NIH/DHHS. This work used the computational resources of the NIH HPC Biowulf cluster (<http://hpc.nih.gov>).

AUTHOR CONTRIBUTIONS

J.Z. and G.S. designed the experiments; J.Z. collected the data; J.Z. analyzed the data with advice and technical assistance from W.Z., C.J., M.P.H.G.; and J.Z. and G.S. wrote the manuscript with input from the other authors.

COMPETING INTERESTS

The authors declare no competing interests.

MATERIALS & CORRESPONDENCE

Correspondence or requests for material should be addressed to J.Z. (jingfeng.zhou@nih.gov) or G.S. (geoffrey.schoenbaum@nih.gov).

METHODS & MATERIALS

Subjects

Subjects were 4 male Long-Evans rats (Charles River, 175 – 200 g, ~3-month-old) individually housed on a 12-h light/dark cycle and given ad libitum access to food in an animal facility at the AAALAC-accredited animal care facility at the National Institute on Drug Abuse Intramural Research Program (NIDA-IRP). Rats were water-deprived the day before any testing and received free access to water for 10 min in their home cages each afternoon after testing. If there was no testing the next day, the rats were given free access to water. All behavioral testing was carried out at the NIDA-IRP. Animal care and experimental procedures complied with the U.S. National Institutes of Health (NIH) guidelines and were approved by the Animal Care and Use Committee (ACUC) at the NIDA-IRP.

Behavioral task

The behavioral training was conducted in aluminum boxes (~18" on a side) equipped with a port for odor delivery and a well for delivery of sucrose solution. Task events were controlled by a custom-written C++ program and a system of relays and solenoid valves; entries into the odor port and the fluid well were detected by infrared beam sensors. The availability of each trial was signaled by the illumination of two house-lights located on the wall above the odor port. Nosepoke into the odor port within 5 seconds after light onset initiated the trial, leading to odor delivery after a 500-ms delay. Rats were required to remain in the port for an additional 500-ms; the trial was aborted, and the lights extinguished if the rat left the odor port before this time had elapsed. After 500-ms, the rats were free to leave the port, which terminated odor delivery. After port exit, rats had 2-s to respond at the fluid well. On rewarded trials, a response led to the delivery of a sucrose solution (10% w/v; 50 µL) after a random delay ranging from 400 to 1500-ms. Upon exit from the well, non-responding during the 2-s period, or responding on non-rewarded trials, the house lights were extinguished, indicating the end of the trial and the beginning of the inter-trial interval (ITI). A 4-s ITI followed correct trials, and an 8-s ITI followed trials on which the rat made an error.

On each trial, one of 6 odors was delivered to the odor port. The 6 odors were organized into two sequences (S1 and S2) that occurred in turn repeatedly (S1→S2→S1→S2→...→S1→S2; 40 repeats of each sequence), described as below. The odor identity is indicated by a number, and reward and non-reward is indicated by the positive (+) and negative (-) symbols, respectively:

S1: 5+ 0- 1- 2+

S2: 3+ 0- 1- 4+

Rats were trained on the full set of sequences since Day 1 until they were able to perform accurately (> 75% correct) on every trial type in a session, then electrode arrays were implanted bilaterally in OFC.

Surgical procedures

Rats were implanted with two drivable bundles of 16 electrodes (32 electrodes in total), made from nickel-chromium wires (25 μ m in bare diameter; AM Systems, WA) in bilateral OFCs (AP: 3 mm, ML: 3.2 mm). Each wire bundle was housed in a 27-gauge stainless-steel tubing and cut with a pair of fine spring scissors to extend 1.5 – 2 mm beyond the end of the tubing. The tips of wires were initially placed at 4.2 mm ventral from the brain surface. After surgery, rats were given Cephalexin (15 mg/kg) orally twice a day for two weeks to prevent any infection.

Single-unit recording

Electrophysiological signals and behavioral event timestamps were recorded with the Plexon OmniPlex System (Plexon, Dallas, TX). The initial wideband signals collected by the electrodes were amplified and digitalized at 40 kHz through a digital headstage (Digital Headstage Processor; DHP) and filtered in the control software (PlexControl) to isolate spike-band frequency (250 – 8,000 Hz) signals. Before the start of each recording session, a common median reference (CMR) for each electrode bundle (16 electrodes in each bundle; two bundles for each rat) was used to remove online noise and artifact. A threshold for each channel was set manually for each active channel to capture unsorted spikes. Spikes were sorted later offline to remove noise and isolate single units using Offline Sorter (Plexon, Dallas, TX) with a built-in template matching algorithm. Sorted files were saved as NeuroExplorer (Nex Technologies, Colorado Springs, CO) formatted files, which were exported to MATLAB (MathWorks, Natick, MA) to extract unit and behavioral event timestamps and for further analyses. Immediately after each session, the electrodes were moved 40~80 μ m ventrally in order to change the neural population being sampled.

After the recording experiments, rats were euthanized by an overdose of isoflurane and perfused with phosphate-buffered saline (PBS) followed by 4% paraformaldehyde. A small constant current was passed through each of the electrode wires to mark the final locations of electrodes. Brains were cut in 40 μ m for standard histological examination.

Quantification and statistical analyses

The number of rats and neurons were not predetermined by any statistical methods but are comparable to those reported in previous publications from our lab. All data were analyzed using MATLAB (MathWorks, Natick, MA).

Peri-event spike dynamics

Each trial was separated into 9 epochs associated with different task events: “ITI-a”, “Light”, “Poke”, “Odor”, “Unpoke”, “Choice”, and “Outcome”, “postOutcome”, “ITI-b”. “ITI-a” marked the time point 0.7 s before the house-light on. On reward trials, the time of well-entry was labeled as “Choice”. “Outcome” was at the time of reward delivery. On non-reward trials, the end of the 2-second window for responding was labeled as “Choice” and a time point 0.7 s after the “Choice” as “Outcome”. On both reward and non-reward trials, 0.7 s after the outcome was labeled as “postOutcome”, and 0.7 s after that was labeled as “ITI-b”. Behavioral performance was quantified by the percent of trials on which the rats responded correctly and the latency with which they initiated a trial after light onset. The spike train for each isolated single unit was aligned to the onset of each task event for the calculation of a peri-event time histogram (PETH). Pre-event time was set to be 200 ms, and post-event time 600 ms. Spike number was counted with a bin = 100 ms. A Gaussian kernel ($\sigma = 50$ ms) was used to smooth the PETH on each trial.

For further analyses, only 30 correct trials were randomly selected from each trial type (30 trials $\times$ 8 trial types = 240 trials in total); and the post-event firing rates (100 – 600 ms) were averaged to obtain a single measurement of neural activity for each neurons on each trial at each task epoch.

Classification analyses

The neural data at each task epoch was organized as a 2-dimensional matrix (trials $\times$ neurons) in a way that each row represents one trial and each column represents the firing rates of one neuron in all the trials. In other words, each trial is a vector in which each dimension is the firing rate of one neuron. Neurons recorded from different sessions were concatenated with alignment to the trials to form pseudo-ensembles. We shuffled trial orders within each trial type to generate a different pseudo-ensemble as well as to remove the temporal correlation between neurons. The trial-order shuffling was repeated for 10, 000 times such that 10, 000 pseudo-ensembles were generated.

We used the linear support vector machine (SVM) for classification analyses (Chang & Lin, 2011). The classification accuracy was assessed by a leave-one-out cross-validation procedure. Specifically, one trial from each trial type was left out for future testing, and all the other trials were used to train the classifier. For each pseudo-ensemble, the leave-one-out cross-validation was repeated 200 times to estimate a mean decoding accuracy. The decoding analyses were carried out on the 10, 000 pseudo-ensembles to obtain an overall mean decoding accuracy. The statistical significance of the overall mean decoding accuracy was determined by the 95% confidence interval estimated by running the same decoding process with label-

272 shuffled pseudo-ensembles. For cross-epoch classification analyses, we followed the same procedure but
273 with training and test sets from different task epochs.

REFERENCES

- Barbosa, J., Stein, H., Martinez, R. L., Galan-Gadea, A., Li, S., Dalmau, J., . . . Compte, A. (2020). Interplay between persistent activity and activity-silent dynamics in the prefrontal cortex underlies serial biases in working memory. *Nature Neuroscience*, 23, 1016–1024. <https://doi.org/10.1038/s41593-020-0644-4>
- Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, 100(2), 490-509. <http://dx.dio.org/10.1016/j.neuron.2018.10.002>
- Beran, M. J., Evans, T. A., Klein, E. D., & Einstein, G. O. (2012). Rhesus monkeys (*Macaca mulatta*) and capuchin monkeys (*Cebus apella*) remember future responses in a computerized task. *The Journal of Experimental Psychology: Animal Behavior Processes*, 38(3), 233-243. <http://dx.dio.org/10.1037/a0027796>
- Beran, M. J., Perdue, B. M., Bramlett, J. L., Menzel, C. R., & Evans, T. A. (2012). Prospective memory in a language-trained chimpanzee (*Pan troglodytes*). *Learning and Motivation*, 43(4), 192-199. <http://dx.dio.org/10.1016/j.lmot.2012.05.002>
- Chang, C. C., & Lin, C. J. (2011). LIBSVM: A Library for Support Vector Machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3), 1-27. <http://dx.dio.org/10.1145/1961189.1961199>
- Evans, T. A., & Beran, M. J. (2012). Monkeys exhibit prospective memory in a computerized task. *Cognition*, 125(2), 131-140. <http://dx.dio.org/10.1016/j.cognition.2012.07.012>
- Feierstein, C. E., Quirk, M. C., Uchida, N., Sosulski, D. L., & Mainen, Z. F. (2006). Representation of spatial goals in rat orbitofrontal cortex. *Neuron*, 51(4), 495-507. <http://dx.dio.org/10.1016/j.neuron.2006.06.032>
- Greene, E., & Naranjo, J. N. (1986). Thalamic role in spatial memory. *Behavioural Brain Research*, 19(2), 123-131. [http://dx.dio.org/10.1016/0166-4328\(86\)90010-0](http://dx.dio.org/10.1016/0166-4328(86)90010-0)
- Kennerley, S. W., Behrens, T. E., & Wallis, J. D. (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nature Neuroscience*, 14(12), 1581-1589. <http://dx.dio.org/10.1038/nn.2961>
- Kepecs, A., Uchida, N., Zariwala, H. A., & Mainen, Z. F. (2008). Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455(7210), 227-231. <http://dx.dio.org/10.1038/nature07200>
- Klein-Flugge, M. C., Barron, H. C., Brodersen, K. H., Dolan, R. J., & Behrens, T. E. (2013). Segregated encoding of reward-identity and stimulus-reward associations in human orbitofrontal cortex. *The Journal of Neuroscience*, 33(7), 3202-3211. <http://dx.dio.org/10.1523/JNEUROSCI.2532-12.2013>
- McDaniel, Mark A., & Einstein, Gilles O. (2007). *Prospective memory : an overview and synthesis of an emerging field*. Thousand Oaks, Calif.: SAGE Publications.
- Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic theory of working memory. *Science*, 319(5869), 1543-1546. <http://dx.dio.org/10.1126/science.1150769>
- Nogueira, R., Abolafia, J. M., Drugowitsch, J., Balaguer-Ballester, E., Sanchez-Vives, M. V., & Moreno-Bote, R. (2017). Lateral orbitofrontal cortex anticipates choices and integrates prior with current information. *Nature Communications*, 8, 14823. <http://dx.dio.org/10.1038/ncomms14823>
- Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, 441(7090), 223-226. <http://dx.dio.org/10.1038/nature04676>
- Riceberg, J. S., & Shapiro, M. L. (2017). Orbitofrontal cortex signals expected outcomes with predictive codes when stable contingencies promote the integration of reward history. *The Journal of Neuroscience*, 37(8), 2010-2021. <http://dx.dio.org/10.1523/JNEUROSCI.2951-16.2016>
- Roesch, M. R., Taylor, A. R., & Schoenbaum, G. (2006). Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron*, 51(4), 509-520. <http://dx.dio.org/10.1016/j.neuron.2006.06.027>

- 318 Rudebeck, P. H., & Murray, E. A. (2014). The orbitofrontal oracle: cortical mechanisms for the prediction and
319 evaluation of specific behavioral outcomes. *Neuron*, 84(6), 1143-1156.
320 <http://dx.dio.org/10.1016/j.neuron.2014.10.049>
- 321 Saez, R. A., Saez, A., Paton, J. J., Lau, B., & Salzman, C. D. (2017). Distinct roles for the amygdala and orbitofrontal
322 cortex in representing the relative amount of expected reward. *Neuron*, 95(1), 70-77 e73.
323 <http://dx.dio.org/10.1016/j.neuron.2017.06.012>
- 324 Stalnaker, T. A., Cooch, N. K., & Schoenbaum, G. (2015). What the orbitofrontal cortex does not do. *Nat*
325 *Neurosci*, 18(5), 620-627. <http://dx.dio.org/10.1038/nn.3982>
- 326 Stokes, M. G. (2015). 'Activity-silent' working memory in prefrontal cortex: a dynamic coding framework. *Trends*
327 *in Cognitive Sciences*, 19(7), 394-405. <http://dx.dio.org/10.1016/j.tics.2015.05.004>
- 328 Tremblay, L., & Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature*,
329 398(6729), 704-708. <http://dx.dio.org/10.1038/19525>
- 330 Verma, A., & Moghaddam, B. (1996). NMDA receptor antagonists impair prefrontal cortex function as assessed
331 via spatial delayed alternation performance in rats: modulation by dopamine. *The Journal of*
332 *Neuroscience*, 16(1), 373-379. <https://dx.doi.org/10.1523/JNEUROSCI.16-01-00373.1996>
- 333 Wallis, J. D. (2007). Orbitofrontal cortex and its contribution to decision-making. *Annual Review of Neuroscience*,
334 30, 31-56. <http://dx.dio.org/10.1146/annurev.neuro.30.051606.094334>
- 335 Wallis, J. D. (2011). Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nat Neurosci*,
336 15(1), 13-19. <http://dx.dio.org/10.1038/nn.2956>
- 337 Wang, F., Schoenbaum, G., & Kahnt, T. (2020). Interactions between human orbitofrontal cortex and
338 hippocampus support model-based inference. *PLoS Biololgy*, 18(1), e3000578.
339 <http://dx.dio.org/10.1371/journal.pbio.3000578>
- 340 Wikenheiser, A. M., & Schoenbaum, G. (2016). Over the river, through the woods: cognitive maps in the
341 hippocampus and orbitofrontal cortex. *Nature Review Neuroscience*, 17(8), 513-523.
342 <http://dx.dio.org/10.1038/nrn.2016.56>
- 343 Wilson, A. G., & Crystal, J. D. (2012). Prospective memory in the rat. *Animal Cognition*, 15(3), 349-358.
344 <http://dx.dio.org/10.1007/s10071-011-0459-5>
- 345 Wilson, A. G., Pizzo, M. J., & Crystal, J. D. (2013). Event-based prospective memory in the rat. *Current Biology*,
346 23(12), 1089-1093. <http://dx.dio.org/10.1016/j.cub.2013.04.067>
- 347 Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task
348 space. *Neuron*, 81(2), 267-279. <http://dx.dio.org/10.1016/j.neuron.2013.11.005>
- 349 Xie, Y., Nie, C., & Yang, T. (2018). Covert shift of attention modulates the value encoding in the orbitofrontal
350 cortex. *Elife*, 7. <http://dx.dio.org/10.7554/eLife.31507>
- 351 Young, J. J., & Shapiro, M. L. (2011). Dynamic coding of goal-directed paths by orbital prefrontal cortex. *The*
352 *Journal of Neuroscience*, 31(16), 5989-6000. <http://dx.dio.org/10.1523/JNEUROSCI.5436-10.2011>
- 353 Zhou, J., Gardner, M. P. H., Stalnaker, T. A., Ramus, S. J., Wikenheiser, A. M., Niv, Y., & Schoenbaum, G. (2019).
354 Rat orbitofrontal ensemble activity contains multiplexed but dissociable representations of value and
355 task structure in an odor sequence task. *Current Biology*, 29(6), 897-907.
356 <http://dx.dio.org/10.1016/j.cub.2019.01.048>
- 357 Zhou, J., Jia, C., Feng, Q., Bao, J., & Luo, M. (2015). Prospective coding of dorsal raphe reward signals by the
358 orbitofrontal cortex. *The Journal of Neuroscience*, 35(6), 2717-2730.
359 <http://dx.dio.org/10.1523/JNEUROSCI.4017-14.2015>

Figure 1. Task design, histology, and behavioral performance. (A) Single trial of the behavioral task. Rats sampled one of 6 odors from an odor port on each trial and made a "Go" choice by poking into a nearby fluid well or a "No-Go" choice by withholding their responses. (B) The 6 odors were organized into two 4-odor sequences, named S1 and S2. The four odors in each sequence represent four positions (P1 – P4). S1 and S2 alternated like a "figure eight". (C) Reconstruction of recording sites. Red squares indicate locations of electrodes. (D) Percent correct (%correct) on each trial type during single-unit recording sessions. Blue indicates trial types with reward, while red indicates trial types without reward. Error bars are standard errors of the mean (SEMs). A two-way ANOVA ($n = 64$ sessions) with factors sequence ($F_{1,504} = 1.03$; $p = 0.31$; $\eta^2 = 0.0014$) and position ($F_{3,504} = 84.0$; $p = 4.3 \times 10^{-44}$; $\eta^2 = 0.33$) was performed. No significant interaction was observed ($F_{3,504} = 0.66$; $p = 0.57$; $\eta^2 = 0.0026$). (E) Poke latency measures the time from light onset to odor port entry. Error bars are SEMs. A two-way ANOVA ($n = 64$ sessions) with factors sequence ($F_{1,504} = 0.10$; $p = 0.75$; $\eta^2 = 4.0 \times 10^{-5}$) and position ($F_{3,504} = 681.0$; $p = 8.0 \times 10^{-177}$; $\eta^2 = 0.8$) was performed. No significant interaction was observed ($F_{3,504} = 0.79$; $p = 0.5$; $\eta^2 = 9.3 \times 10^{-4}$).

Figure 2. Example of single-units. (A-D) Four example neurons that exhibited differential firing to sequences S1 versus S2 at different positions (P1 – P4, indicated by grey arrows). Blue and red colors mean reward and non-reward trial types, respectively. Shaded areas indicate SEMs.

Figure 3. Discriminating sequences S1 versus S2 at both the single-unit and neural ensemble levels. (A) The neuronal selectivity to sequences S1 versus S2 of each neuron was calculated at different epochs [“ITI-a”, “Light”, “Poke”, “Odor”, “Unpoke”, “Choice”, “Outcome”, “postOut. (postOutcome)”, “ITI-b”] for all the four positions ($p < 0.05$; One-way ANOVA). The p values were not corrected; the dotted lines indicate the chance or baseline level of selectivity given this criterion. Note that only at two task epochs [“Odor” at P3 and P4; black bars] could the rats discriminate the two sequences (S1 versus S2) only based on current sensory information. At all other task epochs, shown as grey bars, an internal memory of the sequences had to be used. (B) Accuracy of decoding S1 versus S2 in each of the 9 task epochs within each position (P1 – P4). Error bars are standard deviations (SDs) and each asterisk indicates that the mean decoding accuracy exceeds a 95% confidence interval (CI) calculated using the same decoding process but with label-shuffled data. The meaning of black bars is the same as in panel A. (C) Accuracy of decoding S1 versus S2 at each position with different ensemble sizes. Task epochs used for P1, P2 and P4 are “Poke” and “Odor”, while task epochs used for P3 are “Post-Out.” and “ITI-b”. Error bars show SDs and the asterisks indicate that the mean decoding accuracy exceeds 95% CIs as results of decoding with shuffled trial labels.

Figure 4. Cross-epoch decoding of sequences (A) Each dot represents a single trial in the high-dimensional ensemble activity space ($n = 30$ trials for each trial type). A binary linear SVM classifier was trained to discriminate sequences S1 vs. S2 by using neural activities during odor sampling during P1 (30 trials for each trial type). The trained classifier was used to test how well S1 vs. S2 could be decoded by using the neural ensemble activities at all task epochs (9 epochs within each positions; P1 – P4). Error bars show SDs and asterisks over bar plots indicates that the mean decoding accuracy exceeds 95% CIs calculated using the same decoding process but with label-shuffled data. (B) The same as in (A) except that the SVM classifier was trained by using neural activities at P4 odor time.

Figure 5. Cross-epoch decoding of positions. Trials in S1 and S2 were combined in each task epoch at 4 positions. Each dot represents a single trial in the high-dimensional ensemble activity space ($n = 60$ trials for each trial type). A binary linear SVM classifier was trained to discriminate positions P1 vs. P4 by using neural activities at odor sampling. The trained classifier was then used to decode neural ensemble activities at all task epochs (9 epochs within each positions; P1 – P4). Error bars show SDs and asterisks over bar plots indicates that the mean decoding accuracy exceeds 95% CIs calculated using the same decoding process but with label-shuffled data. Positive values for decoding accuracy indicate that activity at a particular position/epoch was more often classified as P4 than P1, while negative values means the opposite.

Figure 6. Cross-epoch decoding of positions for each rat. (A-D) Cross-epoch decoding of P1 versus P4 was performed on each rat, using the same approach and conventions as in Figure 5. (E-H) Poke latency as in Figure 1E, calculated separately for each rat. Two-way ANOVAs were performed with two factors (sequence and position). Statistical results were shown on the right side of bar plots.

Figure 7. Cross-epoch decoding of positions based on performance. (A) The histogram of poke latency differences between rewarded and non-rewarded trial types. Rats #2-4 showed substantially larger differences in poke latencies than Rat #1 (B) Cross-epoch decoding of P1 versus P4 was performed on sessions from Rats #2-4, which showed large differences in their poke latencies prior to rewarded vs. non-rewarded trials, consistent with use of the sequences to predict upcoming reward. This analysis used the same approach and conventions as in Figure 5.