HHS Public Access

Author manuscript

Neuron. Author manuscript; available in PMC 2018 August 30.

Published in final edited form as:

Neuron. 2017 August 30; 95(5): 1197–1207.e3. doi:10.1016/j.neuron.2017.08.003.

Suppression of ventral hippocampal output impairs integrated orbitofrontal encoding of task structure

Andrew M. Wikenheiser[1], **Yasmin Marrero-Garcia**[1], and **Geoffrey Schoenbaum**[1,2,3,4]

[1]NIDA Intramural Research Program, Cellular Neurobiology Research Branch, Behavioral Neurophysiology Research Section, Baltimore, Maryland, 21224, USA

[2]Department of Anatomy and Neurobiology, University of Maryland, School of Medicine, Baltimore, Maryland, 21201, USA

[3]Solomon H. Snyder Department of Neuroscience, The Johns Hopkins University, Baltimore, Maryland, 21287, USA

Summary

The hippocampus and orbitofrontal cortex (OFC) both make important contributions to decision making and other cognitive processes. However, despite anatomical links between the two, few studies have tested the importance of hippocampal–OFC interactions. Here, we recorded OFC neurons in rats performing a decision making task while suppressing activity in a key hippocampal output region, the ventral subiculum. OFC neurons encoded information about expected outcomes and rats' responses. With hippocampal output suppressed, rats were slower to adapt to changes in reward contingency, and OFC encoding of response information was strongly attenuated. In addition, ventral subiculum inactivation prevented OFC neurons from integrating information about features of outcomes to form holistic representations of the outcomes available in specific trial blocks. These data suggest that the hippocampus contributes to OFC encoding of both concrete, low-level features of expected outcomes, and abstract, inferred properties of the structure of the world, such as task state.

eTOC

In the current study, Wikenheiser et al. show that the hippocampus supports the ability of orbitofrontal cortex to represent both concrete, low-level features of expected outcomes, and abstract, inferred properties of the structure of the world, such as task state.

Corresponding author: Geoffrey Schoenbaum (geoffrey.schoenbaum@nih.gov) or Andrew Wikenheiser (andrew.wikenheiser@nih.gov).

[4]Lead contact: Geoffrey Schoenbaum

Author Contributions

A.M.W. and G.S. designed the study. A.M.W. and Y.M.G. collected the data. A.M.W. and G.S. analyzed the data. A.M.W. and G.S. wrote the paper.

Introduction

Parallel research threads have linked the orbitofrontal cortex (OFC) and the hippocampus with deliberative decision making (Buzsáki and Moser, 2013; Redish, 2016; Rudebeck and Murray, 2014; Stalnaker et al., 2015; Wallis, 2007; Wikenheiser and Redish, 2015; Wilson et al., 2014). Moreover, theories of hippocampal and OFC function have converged to a surprising degree, implicating the pair in broadly similar classes of cognitive processes (Ramus et al., 2007; Shapiro et al., 2014). While hippocampus is best known as a cognitive map of space (O'Keefe and Nadel, 1978; Redish, 1999), the OFC has recently been cast in similar terms, as a map of abstract concepts such as task state (Schuck et al., 2016; Wilson et al., 2014). These ideas share an emphasis on the flexible use of previous learning to guide behavior when the reward structure of the world is unknown, ambiguous, or changing (Wikenheiser and Schoenbaum, 2016). Despite this convergence, few studies have examined how perturbing interactions between OFC and hippocampus affects behavior or neural activity in either region.

Here, we tested how suppressing neurons in one intermediary that links hippocampus to OFC—the ventral subiculum—affected decision making behavior and associated neural representations in the OFC. Neurons were recorded as rats chose between outcomes that differed in size and flavor and whose location changed or reversed across blocks of trials. Reversal tasks have been used extensively to probe OFC function (Cooch et al., 2015; Izquierdo et al., 2004; Jones and Mishkin, 1972b; Riceberg and Shapiro, 2012; Roesch et al., 2006; Stalnaker et al., 2014; Walton et al., 2010), and the hippocampus is well suited to tracking location–outcome associations. Furthermore, the arrangement of trials into blocks creates periods during which there are unique and temporarily stable rules governing how to obtain the rewarding outcomes. Such trial blocks have been argued to constitute "abstract contexts" or states (Saez et al., 2015), the use of which has been proposed to depend on OFC (Wilson et al., 2014). Here we test whether such outcome-relevant block or state representations depend on output from hippocampus (Wikenheiser and Schoenbaum, 2016).

Results

Prior to recording, rats (n = 4) were trained on an odor-guided decision making task (fig. 1a). In each trial, rats sampled one of three odor cues presented in a central odor port. The odor port was flanked by two fluid-delivery wells, where liquid rewards were dispensed. Two odors instructed rats that a response to only one fluid well, either the left or right would be rewarded on that trial ("forced-choice" trials). A third cue indicated that responses to either well would be rewarded ("free-choice" trials). The rewards delivered at each fluid well differed in flavor (chocolate or vanilla milk) and size (3 drops or 1 drop), and were presented in complementary pairs, such that the flavor and size of the reward at one fluid well unambiguously determined the flavor and size of reward available at the other well. The pairing of outcomes and fluid wells was stable for blocks of approximately 60 trials, after which the location of the large and small rewards reversed positions. Block transitions were unsignaled.

After rats reached proficiency on the task, they were injected with an AAV construct targeting the ventral subiculum bilaterally, resulting in expression of the inhibitory opsin molecule eNpHR3.0-eYFP under the control of the CamKII promoter (fig. 1b–c). At the same time, bundles of electrodes were implanted, targeting the OFC. After recovery from surgery and time to allow viral expression, the rats received several sessions of reminder training during which they were acclimated to the recording/fiber optic cable before neural recording sessions were conducted. During all neural recording sessions, light was delivered to the ventral subiculum for the full duration of each trial. In eNpHR sessions (n = 30), 620 nm light was delivered through the implanted optical fibers to activate eNpHR and suppress the activity of expressing neurons. In control sessions (n = 26), 450 nm light (which falls largely outside the range of the eNpHR3.0 frequency sensitivity range; Zhang et al., 2007) was delivered instead to control for the effects of light delivery to the tissue.

Ventral subiculum inactivation reduced preference for large rewards

In both control and eNpHR sessions, rats exhibited a preference for the large reward. This was evident on both free- (fig. 2a) and forced-choice (fig. 2b–c) trials. On forced-choice trials, the rats were significantly more accurate (eNpHR: $P = 1.20 \times 10^{-3}$; W = 1065; rank sum test; control: $P = 6.90 \times 10^{-3}$; W = 1625.5; rank sum test) and faster (eNpHR: P = 0.02; W = 713; rank sum test; control: $P = 5.40 \times 10^{-3}$; W = 1005; rank sum test) when the large reward was at stake (fig. 2b–c). Notably these effects did not differ between session types, indicating that suppressing ventral subiculum neurons did not interfere with cue-guided responses or with differences in behavior that did not require comparison of the different rewards.

By contrast, while rats also showed a strong preference for the large reward on free-choice trials in both eNpHR and control sessions, there was an effect of ventral subiculum suppression (fig. 2a). Large choice rate dropped transiently following block switches in both types of session and then rose again as rats updated their behavior to reflect the changed reward contingencies. However it was significantly higher in control than eNpHR sessions ($P = 5.00 \times 10^{-3}$; W = 734; rank sum test; fig. 2a, inset), an effect which was present even when large choice rate was computed separately for the first 20 trials and last 20 trials of blocks ($P_{early} = 0.02$; $W_{early} = 727.5$; $P_{late} = 0.02$; $W_{late} = 773$; rank sum tests). In shuffled surrogate data sets in which the identity of eNpHR and control sessions were randomly reassigned, only rarely (27 of 1000 instances; 2.7%) did the magnitude of the difference in large choice rate between the two groups equal or exceed the difference observed in the un-shuffled data set, suggesting that session-to-session variability could not account for the difference in behavior between control and eNpHR sessions.

Ventral subiculum inactivation altered the time course of OFC neural responses

We recorded activity from 229 OFC units (124 neurons in control sessions; 105 neurons in eNpHR sessions; table S1). Firing rates were significantly lower in OFC neurons during ventral subiculum inactivation (P = 0.03; $W = 1.10 \times 10^4$; rank sum test), but not during the intertrial interval when no light stimulation was active (P = 0.33; $W = 1.15 \times 10^4$; rank sum test, fig. S1). Baseline firing rates of OFC neurons were stable over time in both control and eNpHR sessions (fig. S1). To examine units' general response properties, we divided

correctly-performed forced-choice trials (which were free from the confound of behavioral differences between control and eNpHR sessions) into seven 0.5 s epochs (fig. 3a), computed firing rates in each epoch, and sorted the neurons according to the time of their peak firing rate (fig. 3b). This revealed a significant effect of subiculum inactivation on the time course of neural responses in OFC. Specifically, during control sessions, many neurons were most active during the reward anticipation epoch, the period after rats had entered one of the fluid wells, but before the outcome was delivered, whereas during eNpHR sessions, significantly fewer neurons were maximally active during this epoch (fig. 3c; $P = 7.14 \times 10^{-3}$; $z_{228} = 2.16$; z-test for population proportions, corrected for multiple comparisons). The population's average firing rate followed a similar pattern (fig. 3d): cells recorded during eNpHR sessions showed a lower mean firing rate as rats moved from the odor port to the reward well, and this difference persisted through the reward anticipation period. These data suggest that hippocampal outflow via ventral subiculum exerted its greatest influence over neural activity in OFC during reward anticipation, the interval bounded by action initiation and outcome receipt.

Ventral subiculum inactivation reduced OFC selectivity for response direction

The previous analyses showed that inhibiting ventral subiculum altered the response dynamics of OFC neurons, but did not address whether the information they encoded was affected. Consistent with previous work (Feierstein et al., 2006; Roesch et al., 2006; Stalnaker et al., 2014; Tsujimoto et al., 2009), we observed that during the reward anticipation period many OFC neurons represented features of the task related to the anticipated outcomes. Indeed, all three variables characterizing the outcomes—response direction (whether the animal chose the left or right fluid well), outcome flavor (chocolate or vanilla), and outcome size (large or small)—were prominently encoded by individual OFC neurons (Fig. 4).

To quantify selectivity for these variables across all OFC neurons recorded during task performance, we fit generalized linear models to each units' firing rate at every point throughout correct, forced-choice trials and plotted the proportion of neurons selective for each variable (after correction for multiple comparisons) at each moment. This revealed that selectivity for each of the individual variables increased during the odor sampling epoch, the earliest point in a trial that rats had information about the available actions and outcomes, and remained elevated during action execution and reward anticipation. However, while the proportion of neurons selective for outcome size and flavor was generally similar between eNpHR and control sessions (fig. 5a), the proportion selective for response direction was significantly lower in eNpHR sessions at nearly every point after rats sampled the odor cue.

Collapsing across task epochs, the fraction of neurons that were selective for at least one variable during at least one epoch between odor sampling and reward anticipation was significantly lower in eNpHR sessions (66% control sessions; 52% eNpHR sessions; $P = 0.02$; $z_{228} = 2.38$; z-test for population proportions; fig. 5b). This difference was driven by a decrease in the proportion of cells selective for response direction (fig. 5c; $P = 4.04 \times 10^{-3}$; $z_{228} = 2.87$; z-test for population proportions). There was no significant change in the proportion of cells with significant responses to either outcome size (despite a numerical

decline; P = 0.13; z_{228} = 1.48; z-test for population proportions) or outcome flavor (P = 0.91; z_{228} = 0.12; z-test for population proportions). This pattern of selectivity was generally consistent across rats (fig. S2).

Ventral subiculum inactivation impaired OFC population encoding of integrated outcome representations

The preceding analyses showed that individual neurons in OFC represented the individual features characterizing the anticipated outcomes—their size, flavor, and location—and that ventral subiculum inactivation weakened representations of at least one of these variables, response direction. These analyses did not, however, speak to how information about outcomes was integrated at the population level. To address this question, we tested whether the chosen outcome (its direction, flavor, and size) could be decoded from OFC firing patterns during the outcome anticipation period, when many OFC neurons were active and the largest difference in neural activity between control and eNpHR sessions occurred.

The structure of the task resulted in 8 unique direction × flavor × size outcomes. We used linear discriminant analysis to measure how well the outcome chosen on a given trial could be decoded from the pattern of firing across neurons. Synthetic OFC pseudoensembles were generated by randomly selecting subsets of units from the population of recorded cells, using data from half of the correctly-executed, forced-choice trials to fit the classifier and the remainder to test classification performance. The process of drawing and testing pseudoensembles of different numbers of neurons was repeated 250 times for both eNpHR and control neurons. Classification accuracy exceeded the 1 in 8 chance level for both eNpHR and control pseudoensembles (fig. 6a). However, control ensembles outperformed ensembles composed of neurons drawn from eNpHR sessions—a two-way ANOVA with factors session type (control or eNpHR) and pseudoensemble size revealed significant main effects of session type ($P = 1.00 \times 10^{-273}$; $F_{1,4482}$ = 1441.37; Cohen's D = 0.92) and pseudoensemble size ($P < 2.23 \times 10^{-308}$; $F_{8,4482}$ = 248.30), along with a significant interaction ($P = 6.45 \times 10^{-75}$; $F_{8,4482}$ = 48.10). Post-hoc tests (paired t-tests, corrected for multiple comparisons) revealed that decoding of chosen outcomes was more accurate in control session data for pseudoensembles larger than two neurons.

Examining how errors in classification were distributed provided some additional insight into the underlying differences in information processing between control and eNpHR sessions. Misclassified control trials were typically correctly identified as left or right choices, as evidenced by the greater proportion of erroneous classifications confined within the upper-left and lower-right quadrants of the confusion matrix (fig. 6b). Consistent with the reduction in directional information observed at the level of single units (fig. 5), this pattern was attenuated in eNpHR pseudoensembles. However, if ventral subiculum suppression had only affected this direction information, then errors in eNpHR session classification should have distributed to outcomes of the same flavor and size as the rat's choice, only in the opposite location. The actual pattern of misclassification appears more complicated than this. Similarly, when we used the classifier to decode pairs of variables (fig. 6c), the disparity between eNpHR and control pseudoensembles was greatest for pairs that included response direction, for example direction and flavor ($P_{\text{session type}} = 6.45 \times$

10^{-90}; $F_{1,4482}$ = 423.12; Cohen's D = 0.56; $P_{pseudoensemble\ size}$= 3.94×10^{-158}; $F_{8,4482}$ = 103.00; $P_{interaction}$= 5.77×10^{-11}; $F_{8,4482}$ = 8.12; two-way ANOVA with factors session type and pseudoensemble size; fig. 6c, left panel) or direction and size ($P_{session\ type}$ = 3.04×10^{-111}; $F_{1,4482}$ = 531.56; Cohen's D = 0.62; $P_{pseudoensemble\ size}$= 2.46×10^{-195}; $F_{8,4482}$ = 129.02; $P_{interaction}$= 1.13×10^{-9}; $F_{8,4482}$ = 7.30; two-way ANOVA with factors session type and pseudoensemble size; fig. 6c, middle panel). Yet deficits in decoding flavor and size (ignoring direction) were also evident ($P_{session\ type}$ = 1.76×10^{-18}; $F_{1,4482}$ = 77.61; Cohen's D = 0.26; $P_{pseudoensemble\ size}$= 3.81×10^{-49}; $F_{8,4482}$ = 31.90; $P_{interaction}$= 0.01; $F_{8,4482}$ = 2.38; two-way ANOVA with factors session type and pseudoensemble size; fig. 6c, middle panel). These results hint at a more complicated effect of ventral subiculum suppression on OFC encoding of anticipated outcomes.

To explore how ventral subiculum suppression affected integrated OFC outcome encoding with more precision, beyond simply cataloging its effect on the representation of individual outcome features, we fit each OFC unit's responses during reward anticipation with a generalized linear model as described previously (fig. 5), using response direction, outcome flavor, and outcome size as predictors of firing rate. Such models use a weighted combination of predictor variables to explain firing rate across trials. The difference between model-estimated and observed firing rates—the residuals of the model—represent the portion of the neuron's firing rate response that is not explained by the predictor variables. Testing these residuals in the discriminant classification analysis described above (fig. 6) allowed us to ask whether information about chosen outcome might exist in OFC neural responses independent of the influence of size, flavor, and direction and further, whether any such integrated representations are also affected by ventral subiculum inactivation.

Surprisingly, this analysis revealed that there was significant information about the chosen outcome remaining in the residual activity of control ensembles. Despite the removal of variance in firing that encoded direction, flavor, and size, the classification accuracy remained significantly above chance for control data, increasing with pseudoensemble size (fig. 7a). A two-way ANOVA with factors session type (control or eNpHR) and pseudoensemble size revealed that classification accuracy was significantly greater for control pseudoensembles ($P_{session\ type}$ = 1.03×10^{-52}; $F_{1,4482}$ = 239.70; Cohen's D = 0.45; $P_{pseudoensemble\ size}$= 4.75×10^{-22}; $F_{8,4482}$ = 15.10; $P_{interaction}$= 3.46×10^{-27}; $F_{8,4482}$ = 18.28). Further, the residual information about chosen outcomes generated a different pattern of misclassification than was observed when all information was present. Specifically, errors in classification were frequently assigned to the *unchosen* action-flavor-size combination that accompanied the rat's actual choice within the block of trials. This tendency resulted in the X-shaped pattern present in the control (but not eNpHR) confusion matrix, with trials correctly classified falling along the main diagonal (fig. 7b; from upper left to lower right), and trials erroneously classified as the unchosen option of the same trial block falling along the minor diagonal (fig. 7b; from lower left to upper right). This suggests that the residual information represented trial block, an abstract state defined by the reward structure of the task but not otherwise overtly signaled to rats, and that this information was absent from residuals from eNpHR pseudoensembles analyzed in this manner. Thus, neurons recorded during control sessions encoded information about the task above and beyond what

they encoded about the particular features of chosen outcomes, and this additional information was lost with ventral subiculum inactivation.

But is this information equivalent to a representation of trial block or state? To confirm this directly, we attempted to decode trial block from OFC neural activity. Critically, we again used OFC responses from which information about response direction, outcome size, and outcome flavor had been removed (fig. 8a). A two-way ANOVA with factors session type (control or eNpHR) and pseudoensemble size showed that classification accuracy was significantly greater for control than for eNpHR pseudoensembles ($P_{session\ type} = 1.10 \times 10^{-154}$; $F_{1,4482} = 760.04$; Cohen's D = 0.77; $P_{pseudoensemble\ size} = 6.74 \times 10^{-93}$; $F_{8,4482} = 59.61$; $P_{interaction} = 1.22 \times 10^{-37}$; $F_{8,4482} = 24.75$). Indeed, when we fit OFC responses with regression models that included trial block as a predictor along with action and outcome information (fig. 8b), the difference in ensemble performance between the two session types was eliminated ($P_{session\ type} = 0.56$; $F_{1,4482} = 0.34$; Cohen's D = 0.01; $P_{pseudoensemble\ size} = 0.43$; $F_{8,4482} = 1.01$; two-way ANOVA with factors session type and pseudoensemble size). Together, these data show that OFC neurons recorded during control sessions contained information about trial block or state, and that this information was strongly reduced with inactivation of ventral subiculum.

Discussion

OFC neurons represent task state

Computational models of learning and decision making emphasize the importance of storing learned associations in a way that is specific to the circumstances under which they were acquired (Gershman and Niv, 2010; Redish et al., 2007; Sutton and Barto, 1998; Wilson et al., 2014). This is achieved by assigning learning to states of the world. States can be signaled overtly by observable sensory features, or implicitly by covert, unobservable properties. In either case, so long as an animal can recognize or infer the current state, it can draw on a set of previously-learned associative rules appropriate to the situation.

The hippocampus has long been associated with representing cognitive maps or state spaces (Chalmers et al., 2016; Erdem and Hasselmo, 2014; Johnson and Crowe, 2009; Johnson et al., 2007; Koene et al., 2003; O'Keefe and Nadel, 1978; Redish, 1999; Stachenfeld et al., 2016; Zilli and Hasselmo, 2008). And more recent work has suggested the OFC might also utilize or perhaps even encode these relationships (Bradfield et al., 2015; Saez et al., 2015; Schuck et al., 2016; Stalnaker et al., 2016; Wilson et al., 2014). For example, associative event-driven firing in the OFC is often specific to the cue-reward combinations appropriate for a given context or set of trials (Saez et al., 2015; Schoenbaum et al., 1999; Thorpe et al., 1983), and in one particularly nice set of parallel studies, ensembles of neurons in the OFC (Farovik et al., 2015) and the hippocampus (Komorowski et al., 2013; McKenzie et al., 2014; McKenzie et al., 2016) were found to encode a hierarchy of feature representations in rats performing a decision making task, with neurons in each area encoding abstract aspects of the task, such as the context-dependent linkages between odor cues and reward. These data are consistent with the proposal that OFC incorporates information about the state space to influence lower level associative representations, but stop short of showing representation of the states themselves.

Here, we show that ensembles of OFC neurons directly encode information about task state, at least as defined by blocks of trials in the decision making task that we employed. Such trial blocks, defined by the momentarily stable, actionable rules for obtaining the outcomes, are a reasonable proxy for task state. These rules are not signaled by any simple, easily observable cue and thus the gestalt of the block itself becomes their defining feature. OFC neurons recorded in control sessions represented this gestalt even when information pertaining to the independent outcome features had been removed from their firing. As such, this report joins recent imaging work in showing that neural activity in OFC directly represents task state (Schuck et al., 2016). It is worth noting that while states in our task were defined by the features of outcomes—their flavor, size, and location—in principle this does not have to be the case. State information could be conveyed by other aspects of the task, either explicit and observable or implicit and inferred. Whether similar neural substrates are important for tracking and using these different types of state information is an interesting question.

OFC state representations require intact hippocampal output

Suppression of hippocampal outputs abolished the representation of task state in OFC ensembles while also causing a modest impairment in the efficiency of the rats' choice behavior in the task. The behavioral impairment on free-choice trials is consistent with classic theories of frontal cortex function, involving response inhibition or switching (Jones and Mishkin, 1972a; Miller, 2000), as well as more recent ideas that the OFC is critical to the representation of state spaces (Wilson et al., 2014). Further our data suggest this function and the associated neural integration depends substantially on hippocampal processing. Of course, the modesty of the perseverative behavioral deficit indicates that the hippocampal–OFC circuit is not the sole arbiter of either the complex or simple associative information used to perform this task in well-trained rats. This is not surprising, and in fact increases the interpretability of the neural results, since they are largely unconfounded by changes in behavior. It is also worth noting that other projections link outputs from the hippocampus to the OFC, most notably a direct projection from CA1. Simultaneous inhibition of both hippocampal and subicular projections may have resulted in an even stronger behavioral deficit.

The importance of the hippocampus in representing trial blocks or abstract states is consistent with a growing body of work showing a role for the hippocampal-frontal interactions in decision making, both spatial and not (Barron et al., 2013; Boorman et al., 2016; Ginther et al., 2011; Mack et al., 2016; Miller et al., 2016; Place et al., 2016; Spellman et al., 2015; Tavares et al., 2015). Our results also show causally how commonalities between OFC and hippocampal representations identified in previous work (Constantinescu et al., 2016; Farovik et al., 2015; Ferbinteanu et al., 2011; Johnson and Redish, 2007; McKenzie et al., 2014; McKenzie et al., 2016; Otto and Eichenbaum, 1992; Riceberg and Shapiro, 2012; Steiner and Redish, 2012; Young and Shapiro, 2011), may arise due to real-time interactions between the two regions, via either the direct projection from ventral subiculum to OFC or through indirect, multi-synaptic pathways. That optogenetic inhibition affected representations so strongly in well-trained rats on forced-choice trials indicates that hippocampal involvement is not restricted to the initial phases of training or solely to free-

choice trials on which task state information was most critical for behavioral performance, and instead reveals a pervasive and ongoing interaction between these two important regions even in extremely well established decision making settings. This is reminiscent of recent evidence for an ongoing role for hippocampus in maintaining existing informational stores that were previously thought to be hippocampal-independent once established (Tanaka et al., 2014). These data suggest that while OFC and hippocampal functions are dissociable in certain settings (Abela and Chudasama, 2013), in at least some cases interactions between these regions are synergistic and promote adaptive behavioral responses.

Hippocampal output facilitates the integration of changeable external features into stable internal representations in OFC

Hippocampal suppression also had interesting effects on the ability of ensembles in OFC to represent the properties of outcomes themselves, disrupting encoding of the location or direction of the response required to obtain the reward, and also degrading the integration of this information with information about size and flavor. While concrete features of the outcomes, such as flavor and size, were least affected by hippocampal suppression, the integration of the remaining information by OFC ensembles was inconsistent with the removal of any one component when they were deprived of hippocampal support.

In this regard, these data suggest an important role for the hippocampus in combining dissociable, oft-changing features into a single common, internal representation. This is necessarily the case for representations of trial block—block structure on this task was not defined by any one directly observable cue but was instead determined by the co-occurrence of particular responses, sizes, and flavors. Similarly, hippocampal place cells integrate a diverse range of external sensory information, internal state information, and cognitive factors to compute coherent, unified, representations of discrete locations (Jeffery et al., 2006; Kennedy and Shapiro, 2009; Kentros et al., 2004; Markus et al., 1995; Moita et al., 2003; Young et al., 1994; Zhang and Manahan-Vaughan, 2013). Though we did not manipulate such factors in a controlled way here, they likely influence directional representations in OFC and account for part of what was lost with hippocampal suppression.

However this loss of integrative coding was also revealed in how hippocampal suppression affected the representation of the outcomes themselves. Although directional information was the largest casualty of hippocampal suppression, there were modest effects on representation of flavor and size. This is evident in how outcomes were miscoded when hippocampal outputs were suppressed—ensembles did not simply misidentify the correct outcome in the wrong location.

The role of hippocampal output in supporting integration in OFC recalls its proposed contribution to learning configural associations. In configural tasks, subjects are presented with simultaneous compounds of individual cues. These cue compounds are constructed in such a way that assigning learning to the individual elements of compounds is not sufficient to solve the task. Instead, subjects must parse compounds as unique, holistic entities. Interestingly, lesion and inactivation studies, in concert with computational modeling work, have suggested that the hippocampus is an important component of the brain’s configural learning system (Gluck and Myers, 1993; Honey et al., 2014; Iordanova et al., 2009;

Iordanova et al., 2011a; Iordanova et al., 2011b; Rudy and Sutherland, 1995; Schmajuk and Blair, 1993; Schmajuk and DiCarlo, 1992; Sutherland and Rudy, 1989). Our data align well with this framework, and are consistent with the idea that higher-order integration of individual elemental representations in the OFC depends on this important hippocampal function.

STAR Methods Text

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Geoffrey Schoenbaum (geoffrey.schoenbaum@nih.gov).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Four male, experimentally-naïve, Long-Evans rats (Charles River) with normal immune function, aged approximately 3 months (175–200 g) were subjects for this experiment. Rats were singly housed in a facility accredited by the Association for Assessment and Accreditation of Laboratory Animal Care (AAALAC), and maintained on a 12-h light–dark cycle. During behavioral testing, rats were lightly water restricted; in addition to liquid reward earned on the task, rats were allowed 10 minutes of water access daily, in their home cage. All experimental and animal care procedures complied with US National Institutes of Health guidelines and were approved by that National Institutes on Drug Abuse Intramural Research Program Animal Care and Use Committee.

METHOD DETAILS

Behavior—Training and testing procedures closely followed those used in previous odor-guided choice task experiments (Cooch et al., 2015; Stalnaker et al., 2014). Briefly, experimental sessions were conducted in aluminum chambers outfitted with an odor port flanked by two fluid-delivery wells, allowing for rapid delivery of olfactory cues and liquid outcomes. Experimental sessions were conducted at the same time daily, during the light phase. The behavioral task was controlled by a custom C++ program. Task trials began with the illumination of the house light, which indicated to subjects that a new trial was available. To initiate a trial, rats nose poked in the odor port. Following a 500 ms fixation period, one of three odors was delivered to the port for 500 ms, after which animals were free to withdraw from the odor port and indicate their decision by entering the left or right fluid well. If animals withdrew from the odor port before the 1-s combined fixation and odor sampling periods had passed, the trial was aborted. The odor cues instructed animals that reward would be delivered at either the left or right fluid well, or indicated a free choice. The identity and meaning of odor cues was fixed for the duration of the experiment. Odor cues were presented in pseudo-random order (such that the same cue could not appear on more than three consecutive trials) and in equal proportions, ±1 over 250 trials.

The liquid outcomes were large (3 drops of 50 μL) or small (1 drop of 50 μL) deliveries of chocolate or vanilla milk (Nesquik; diluted 1:1 with water). Previous studies from the lab using the same strain of rats have shown that individual rats do not strongly prefer one flavor or milk over the other (Cooch et al., 2015; Stalnaker et al., 2014). The flavor and size of

outcomes delivered at each fluid well were consistent for blocks of trials, but switched across blocks. Large/small outcomes and chocolate/vanilla outcomes were always delivered from opposite side fluid wells. Thus, beginning from a random starting block, block switches alternated between size switches (where the location of the large and small reward reversed, but the location of each reward flavor remained constant) and size/flavor switches (where both the location of large and small rewards, and the location of chocolate and vanilla outcomes reversed). Each session began with a short initial block, which was followed by four blocks of approximately 65 trials each. The precise timing of block switches varied randomly to prevent animals from anticipating when they would occur, and no external signal was provided to indicate the beginning of a new block. Animals were trained on the task prior to electrode and optical fiber implantation until they performed forced-choice trials accurately and consistently completed all 5 session blocks (range: 20–33 training sessions).

Surgery—Rats were implanted with driveable bundles of 16 nickel-chromium wires (25-μm diameter; AM Systems) targeting the left lateral OFC (3 mm anterior and 3.2 mm lateral of bregma). Bundles were initially positioned 4 mm ventral from brain surface, and then lowered in small increments after each recording session to sample new neurons. During the same surgery, AAV5/CamKIIa-eNpHR3.0-eYFP (UNC virus core) was infused bilaterally in the ventral subiculum (6.5 mm posterior and ±4.5 mm lateral of bregma), and optical fibers (Plexon Inc.) were positioned over each injection site. After approximately three weeks to allow for recovery from surgery and viral expression, rats were returned to the task for recording sessions.

At the end of the experiment, the final electrode position was marked by passing a small current through the wires, and rats were euthanized by and overdose of isoflurane. Rats were perfused intracardially, and brains were processed for histological examination. Brains were cut at 40 μm coronal sections. Free-floating sections were incubated for 1 h in 0.1 M phosphate buffer (PB) supplemented with 4% bovine serum albumin and 0.3% Triton X-100. Sections were then incubated overnight with a mouse anti-YFP (1:1,000, 632381, Clontech Laboratories). After rinsing three times in PB, sections were incubated for 2 hours with donkey anti-mouse conjugated with Alexa Flour-488 (Jackson Immunoresearch Laboratories). After rinsing, the sections were mounted on slides and Flouromount-G with DAPI (Electron Microscopy Sciences) mounting media was used to preserve fluorescence and counterstain the sections.

Single unit recording—Neural activity was collected using Plexon Multichannel Acquisition Processor systems (Plexon Inc.). Voltage signals from electrodes were amplified and filtered following standard procedures used in previous studies. After each recording session in which at least one putative single unit was detected, electrodes were advanced at least 40 μM. If no units were present, electrodes were advanced 40–80 μM. Neural activity was manually sorted into putative single units offline, using MClust (http://redishlab.neuroscience.umn.edu/MClust/MClust.html), following standard spike sorting procedures. Table S1 reports the number of units recorded from each rat.

Optogenetic stimulation—Light stimulation was delivered using a combined optogenetic/electrical commutator interfaced with custom-made 2.5 mm FC ferrules (Plexon Inc.). Stimulation was continuous (i.e. not pulsed) for the duration of trials, beginning with the illumination of the house light, and terminating after outcome delivery. In control stimulation sessions, 450 nm light (6–10 mW power output) was delivered. This wavelength falls outside the optimal frequency sensitivity range of the eNpHR3.0 molecule, and served as a control for light delivery to brain tissue. In eNpHR stimulation sessions, 620 nm (6–10 mW power output) light was delivered to activate eNpHR3.0 and suppress neural activity. Subjects alternated pseudorandomly between eNpHR and control stimulation sessions, with no more than two of the same type of session occurring on consecutive days.

QUANTIFICATION AND STATISTICAL ANALYSIS

All data were analyzed using Matlab (Mathworks). Instances of multiple comparisons were corrected for with the Benjamini–Hochberg procedure. Error bars in figures denote the standard error of the mean. The number of subjects was chosen based on previous similar single-unit recording studies in rats.

Behavioral epochs—For many neural analyses, we segmented task trials into seven 0.5 s behavioral epochs (fig. 3a). The pre-trial epoch began 0.5 s before the house light was turned on to indicate the availability of a new trial, and ended when the house light was illuminated. The fixation epoch began when rats made a nose poke in the central odor port. No odor cues were delivered during this time; rats were simply required to keep their nose in the odor well for 0.5 s. The odor sampling epoch was contiguous to the fixation period, and comprised the 0.5 s for which an odor cue was delivered to the odor port. Following odor delivery rats were free to make a choice between the two fluid wells. The movement epoch was 0.5 s preceding rats' entry to one of the fluid wells. Note that while the movement epoch was always synchronized to rats' arrival at their chosen fluid well, because movement speed was self-paced (and therefore variable) the odor sampling and movement epochs sometimes overlapped. The reward anticipation epoch began with rats' entry into one of the fluid wells and terminated 0.5 s later with delivery of fluid outcomes on correct trials. The outcome delivery epoch began when fluid was dispensed and ended 0.5 s later. Finally, the post-trial period began when the house light was extinguished, indicating the end of the trial. As described above, light stimulation was present during all epochs, except for the pre- and post-trial epochs.

OFC neuron firing rate dynamics—To examine the general response profile of OFC neurons, firing rates for each neuron were computed in 45 ms bins, averaged across correct, forced-choice trials (as this subset of trials was counterbalanced across outcome flavor/size and response direction), and peak normalized (fig. 3b). Cells were counted as maximally active for a given epoch if, for at least one bin during that epoch, the unit's normalized firing rate exceeded 95% of its absolute maximal value (fig. 3c). Population firing rates (fig. 3d) were computed over all neurons and trials (i.e. regardless of individual units' selectivity or response properties) in 45 ms time bins.

Regression analyses—To examine the dynamics of neural selectivity for task variables, we fit linear regression models (Matlab function *fitglm*) to each neuron's firing rate throughout the course of trials (fig. 5a). These analyses included data from correct, forced-choice trials. Response direction, outcome size, and outcome flavor were treated as categorical predictors. Firing rates were computed in 45 ms bins within each task epoch, and for each bin a separate regression model was calculated. P-values were corrected for multiple comparisons within neuron, and the fraction of the total population of recorded neurons significant for each behavioral variable (corrected p-value < 0.05) was computed for each time bin. This approach allowed us to visualize how population selectivity fluctuated over the course of trials with good time resolution. To examine selectivity more generally, we used the same approach, but with a single 0.5 s time bin for each epoch. For the purposes of the Venn diagrams (fig. 5a) and bar charts (fig. 5b), neurons were considered selective for a behavioral variable if corrected regression model p-values for that predictor were less than 0.05 for at least one epoch between odor sampling and reward anticipation, inclusive.

Classification analyses—We used linear discriminant analysis (Matlab function: *classify*) to decode information about response direction and outcome size/flavor from pseudoensembles of OFC neurons. These analyses focused on the reward anticipation epoch, as single unit firing rates and selectivity differed the most between eNpHR and control sessions during this task epoch. The observation data used to fit and test the discriminant classifier was trial-by-trial spike counts from each neuron during the 0.5 s reward anticipation epoch. As in previous analyses, data was restricted to the counterbalanced set of correct, forced-choice trials. Only sessions with a minimum of 20 completed trials in each trial block were included for analysis. Spike count patterns across neurons in the pseudoensemble were classified into 8 unique *action* × *flavor* × *size* categories. Pseudoensembles were created by randomly selecting a subset of units for inclusion from population of cells recorded on the task. The process of generating pseudoensembles and testing their classification performance was conducted separately for eNpHR and control session neurons. Half of the trials were used to fit the classifier (the training set), and the other half (the test set) were withheld to evaluate classification accuracy. The process of drawing and testing pseudoensembles in this way was repeated 250 times. Because trial block lengths were not deterministic, the number of trials for each outcome triplet could vary slightly across sessions (and therefore across neurons in each pseudoensemble, which were drawn across sessions). To decide how many trials to include in the training set, we identified the smallest, even number of trials across all units in each pseudoensemble over all outcome triplets, and set the trial number to that value. This ensured that no outcome triplet was overrepresented in training or testing data, and that training and testing sets were equal in size. For each pseudoensemble, we also trained a classifier with trial type labels randomly re-assigned to training set firing rate observations (but maintaining the same division of training and testing data used to fit and test the unshuffled classifier). This shuffling procedure ensured that classification of randomized data was at the theoretical chance level (1/8 correct). To remove selectivity for either trial type (fig. 7) or trial block (fig. 8) from OFC responses, we fit firing rate data with regression models as in previous analyses (fig. 5), including as predictors any variables whose impact on firing rate we wished to remove. The residuals of these models represent the portion of OFC neuron firing

rate variability that was not be accounted for by linear combinations of the predictor variables.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank members of the Schoenbaum lab for helpful comments on this work, and Carlos Mejias-Aponte and the NIDA IRP histology core for technical support. This work was supported by funding from the US National Institute on Drug Abuse at the Intramural Research Program. The opinions expressed in this article are the authors' own and do not reflect the view of the US National Institutes of Health, the US Department of Health and Human Services or the US government.

References

Abela AR, Chudasama Y. Dissociable contributions of the ventral hippocampus and orbitofrontal cortex to decision-making with a delayed or uncertain outcome. European Journal of Neuroscience. 2013; 37:640–647. [PubMed: 23190048]

Barron HC, Dolan RJ, Behrens TEJ. Online evaluation of novel choices by simultaneous representation of multiple memories. Nature neuroscience. 2013; 16:1492–1498. [PubMed: 24013592]

Boorman, Erie D., Rajendran, Vani G., O'Reilly, Jill X., Behrens, Tim E. Two Anatomically and Computationally Distinct Learning Signals Predict Changes to Stimulus-Outcome Associations in Hippocampus. Neuron. 2016; 89:1343–1354. [PubMed: 26948895]

Bradfield LA, Dezfouli A, van Holstein M, Chieng B, Balleine BW. Medial Orbitofrontal Cortex Mediates Outcome Retrieval in Partially Observable Task Situations. Neuron. 2015; 88:1268–1280. [PubMed: 26627312]

Buzsáki G, Moser EI. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. Nature Neuroscience. 2013; 16:130–138. [PubMed: 23354386]

Chalmers E, Luczak A, Gruber AJ. Computational Properties of the Hippocampus Increase the Efficiency of Goal-Directed Foraging through Hierarchical Reinforcement Learning. Frontiers in Computational Neuroscience. 2016; 10

Constantinescu AO, O'Reilly JX, Behrens TEJ. Organizing conceptual knowledge in humans with a gridlike code. Science. 2016; 352:1464–1468. [PubMed: 27313047]

Cooch NK, Stalnaker TA, Wied HM, Bali-Chaudhary S, McDannald MA, Liu TL, Schoenbaum G. Orbitofrontal lesions eliminate signalling of biological significance in cue-responsive ventral striatal neurons. Nat Commun. 2015; 6

Erdem UM, Hasselmo ME. A biologically inspired hierarchical goal directed navigation model. Journal of Physiology-Paris. 2014; 108:28–37.

Farovik A, Place RJ, McKenzie S, Porter B, Munro CE, Eichenbaum H. Orbitofrontal Cortex Encodes Memories within Value-Based Schemas and Represents Contexts That Guide Memory Retrieval. The Journal of Neuroscience. 2015; 35:8333–8344. [PubMed: 26019346]

Feierstein CE, Quirk MC, Uchida N, Sosulski DL, Mainen ZF. Representation of Spatial Goals in Rat Orbitofrontal Cortex. Neuron. 2006; 60:495–507.

Ferbinteanu J, Shirvalkar P, Shapiro ML. Memory Modulates Journey-Dependent Coding in the Rat Hippocampus. Journal of Neuroscience. 2011; 31:9135–9146. [PubMed: 21697365]

Gershman SJ, Niv Y. Learning latent structure: carving nature at its joints. Current Opinion in Neurobiology. 2010; 20:251–256. [PubMed: 20227271]

Ginther MR, Walsh DF, Ramus SJ. Hippocampal Neurons Encode Different Episodes in an Overlapping Sequence of Odors Task. Journal of Neuroscience. 2011; 31:2706–2711. [PubMed: 21325539]

Gluck MA, Myers CE. Hippocampal Mediation of Stimulus Representation: A computational Theory. Hippocampus. 1993; 3:491–516. [PubMed: 8269040]

Honey RC, Iordanova MD, Good M. Associative structures in animal learning: Dissociating elemental and configural processes. Neurobiology of Learning and Memory. 2014; 108:96–103. [PubMed: 23769767]

Iordanova MD, Burnett DJ, Aggleton JP, Good M, Honey RC. The role of the hippocampus in mnemonic integration and retrieval: complementary evidence from lesion and inactivation studies. European Journal of Neuroscience. 2009; 30:2177–2189. [PubMed: 20128853]

Iordanova MD, Burnett DJ, Good M, Honey RC. Pattern memory involves both elemental and configural processes: Evidence from the effects of hippocampal lesions. Behavioral Neuroscience. 2011a; 125:567–577. [PubMed: 21574677]

Iordanova MD, Good M, Honey RC. Retrieval-Mediated Learning Involving Episodes Requires Synaptic Plasticity in the Hippocampus. The Journal of Neuroscience. 2011b; 31:7156. [PubMed: 21562278]

Izquierdo A, Suda RK, Murray EA. Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. Journal of Neuroscience. 2004; 24:7540–7548. [PubMed: 15329401]

Jeffery KJ, Anand RL, Anderson MI. A role for terrain slope in orienting hippocampal place fields. Experimental brain research. 2006; 169:218–225. [PubMed: 16273400]

Johnson A, Crowe D. Revisiting Tolman: Theories and cognitive maps. Cognitive Critique. 2009; 1:43–72.

Johnson A, vd Meer MAA, Redish AD. Integrating hippocampus and striatum in decision-making. Current Opinion in Neurobiology. 2007; 17:692–697. [PubMed: 18313289]

Johnson A, Redish AD. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. Journal of Neuroscience. 2007; 27:12176–12189. [PubMed: 17989284]

Jones B, Mishkin M. Limbic lesions and the problem of stimulus-reinforcement associations. Experimental Neurology. 1972a; 36:362–377. [PubMed: 4626489]

Jones B, Mishkin M. Limbic lesions and the problem of stimulus—Reinforcement associations. Experimental Neurology. 1972b; 36:362–377. [PubMed: 4626489]

Kennedy PJ, Shapiro ML. Motivational states activate distinct hippocampal representations to guide goal-directed behaviors. Proceedings of the National Academy of Sciences. 2009; 106:10805–10810.

Kentros CG, Agnihotri NT, Streater S, Hawkins RD, Kandel ER. Increased Attention to Spatial Context Increases Both Place Field Stability and Spatial Memory. Neuron. 2004; 42:283–295. [PubMed: 15091343]

Koene RA, Gorchetchnikov A, Cannon RC, Hasselmo ME. Modeling goal-directed spatial navigation in the rat based on physiological data from the hippocampal formation. Neural Networks. 2003; 16:577–584. [PubMed: 12850010]

Komorowski RW, Garcia CG, Wilson A, Hattori S, Howard MW, Eichenbaum H. Ventral Hippocampal Neurons Are Shaped by Experience to Represent Behaviorally Relevant Contexts. The Journal of Neuroscience. 2013; 33:8079. [PubMed: 23637197]

Mack ML, Love BC, Preston AR. Dynamic updating of hippocampal object representations reflects new conceptual knowledge. Proceedings of the National Academy of Sciences. 2016; 113:13203–13208.

Markus EJ, Qin Y, Leonard B, Skaggs WE, McNaughton BL, Barnes CA. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. Journal of Neuroscience. 1995; 15:7079–7094. [PubMed: 7472463]

McKenzie S, Frank AJ, Kinsky NR, Porter B, Rivière PD, Eichenbaum H. Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. Neuron. 2014; 83:202–215. [PubMed: 24910078]

McKenzie S, Keene CS, Farovik A, Bladon J, Place R, Komorowski R, Eichenbaum H. Representation of memories in the cortical–hippocampal system: results from the application of population similarity analyses. Neurobiology of learning and memory. 2016; 134:178–191. [PubMed: 26748022]

Miller EK. The prefrontal cortex and cognitive control. Nature Neuroscience Reviews. 2000; 1:59–65.

Miller KJ, Botvinick MM, Brody CD. Dorsal hippocampus plays a causal role in model-based planning. bioRxiv. 2016

Moita MA, Rosis S, Zhou Y, LeDoux JE, Blair HT. Hippocampal place cells acquire location-specific responses to the conditioned stimulus during auditory fear conditioning. Neuron. 2003; 37:485–497. [PubMed: 12575955]

O'Keefe, J., Nadel, L. The Hippocampus as a Cognitive Map. Clarendon Press; 1978.

Otto T, Eichenbaum H. Complementary roles of the orbital prefrontal cortex and the perirhinal-entorhinal cortices in an odor-guided delayed nonmatching-to-sample task. Behavioral Neuroscience. 1992; 106:762–775. [PubMed: 1445656]

Place R, Farovik A, Brockmann M, Eichenbaum H. Bidirectional prefrontal-hippocampal interactions support context-guided memory. Nature Neuroscience. 2016

Ramus SJ, Davis JB, Donahue RJ, Discenza CB, Waite AA. Interactions Between the Orbitofrontal Cortex and Hippocampal Memory System During the Storage of Long-Term Memory. Annals of the New York Academy of Sciences. 2007; 1121:216–231. [PubMed: 17872388]

Redish, AD. Beyond the Cognitive Map: From Place Cells to Episodic Memory. MIT Press; 1999.

Redish AD. Vicarious trial and error. Nat Rev Neurosci. 2016; 17:147–159. [PubMed: 26891625]

Redish AD, Jensen S, Johnson A, Kurth-Nelson Z. Reconciling reinforcement learning models with behavioral extinction and renewal: Implications for addiction, relapse, and problem gambling. Psychological Review. 2007; 114:784–805. [PubMed: 17638506]

Riceberg JS, Shapiro ML. Reward stability determines the contribution of orbitofrontal cortex to adaptive behavior. The Journal of Neuroscience. 2012; 32:16402–16409. [PubMed: 23152622]

Roesch MR, Taylor AR, Schoenbaum G. Encoding of Time-Discounted Rewards in Orbitofrontal Cortex Is Independent of Value Representation. Neuron. 2006; 51:509–520. [PubMed: 16908415]

Rudebeck, Peter H., Murray, Elisabeth A. The Orbitofrontal Oracle: Cortical Mechanisms for the Prediction and Evaluation of Specific Behavioral Outcomes. Neuron. 2014; 84:1143–1156. [PubMed: 25521376]

Rudy JW, Sutherland RJ. Configural association theory and the hippocampal formation: An appraisal and reconfiguration. Hippocampus. 1995; 5:375–389. [PubMed: 8773252]

Saez A, Rigotti M, Ostojic S, Fusi S, Salzman CD. Abstract context representations in primate amygdala and prefrontal cortex. Neuron. 2015; 87:869–881. [PubMed: 26291167]

Schmajuk NA, Blair HT. Stimulus configuration, spatial learning, and hippocampal function. Behavioural brain research. 1993; 59:103–117. [PubMed: 8155276]

Schmajuk NA, DiCarlo JJ. Stimulus configuration, classical conditioning, and hippocampal function. Psychological review. 1992; 99:268. [PubMed: 1594726]

Schoenbaum G, Chiba AA, Gallagher M. Neural Encoding in Orbitofrontal Cortex and Basolateral Amygdala during Olfactory Discrimination Learning. J Neurosci. 1999; 19:1876–1884. [PubMed: 10024371]

Schuck NW, Cai MB, Wilson RC, Niv Y. Human orbitofrontal cortex represents a cognitive map of state space. Neuron. 2016; 91:1402–1412. [PubMed: 27657452]

Shapiro, ML., Riceberg, JS., Seip-Cammack, K., Guise, KG. Space, Time and Memory in the Hippocampal Formation. Springer; 2014. Functional interactions of prefrontal cortex and the hippocampus in learning and memory; p. 517-560.

Spellman T, Rigotti M, Ahmari SE, Fusi S, Gogos JA, Gordon JA. Hippocampal-prefrontal input supports spatial encoding in working memory. Nature. 2015; 522:309–314. [PubMed: 26053122]

Stachenfeld KL, Botvinick MM, Gershman SJ. The hippocampus as a predictive map. bioRxiv. 2016

Stalnaker TA, Berg BA, Aujla N, Schoenbaum G. Cholinergic interneurons use orbitofrontal input to track beliefs about current state. Journal of Neuroscience. 2016; 36:6242–6257. [PubMed: 27277802]

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

Stalnaker TA, Cooch NK, Schoenbaum G. What the orbitofrontal cortex does not do. Nature Neuroscience. 2015; 18:620–627. [PubMed: 25919962]

Steiner AP, Redish AD. The road not taken: neural correlates of decision making in orbitofrontal cortex. Frontiers in neuroscience. 2012; 6

Sutherland RJ, Rudy JW. Configural association theory: The role of the hippocampal formation in learning, memory, and amnesia. Psychobiology. 1989; 17:129–144.

Sutton, RS., Barto, AG. Reinforcement Learning: An introduction. MIT Press; 1998.

Tanaka KZ, Pevzner A, Hamidi AB, Nakazawa Y, Graham J, Wiltgen BJ. Cortical representations are reinstated by the hippocampus during memory retrieval. Neuron. 2014; 84:347–354. [PubMed: 25308331]

Tavares RM, Mendelsohn A, Grossman Y, Williams CH, Shapiro M, Trope Y, Schiller D. A map for social navigation in the human brain. Neuron. 2015; 87:231–243. [PubMed: 26139376]

Thorpe S, Rolls E, Maddison S. The orbitofrontal cortex: Neuronal activity in the behaving monkey. Experimental Brain Research. 1983; 49:93–115. [PubMed: 6861938]

Tsujimoto S, Genovesio A, Wise SP. Monkey orbitofrontal cortex encodes response choices near feedback time. The Journal of Neuroscience. 2009; 29:2569–2574. [PubMed: 19244532]

Wallis JD. Orbitofrontal Cortex and Its Contribution to Decision-Making. Annual Review of Neuroscience. 2007; 30:31–56.

Walton ME, Behrens TEJ, Buckley MJ, Rudebeck PH, Rushworth MFS. Separable Learning Systems in the Macaque Brain and the Role of Orbitofrontal Cortex in Contingent Learning. Neuron. 2010; 65:927–939. [PubMed: 20346766]

Wikenheiser AM, Redish AD. Decoding the cognitive map: ensemble hippocampal sequences and decision making. Current Opinion in Neurobiology. 2015; 32:8–15. [PubMed: 25463559]

Wikenheiser AM, Schoenbaum G. Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex. Nat Rev Neurosci. 2016; 17:513–523. [PubMed: 27256552]

Wilson RC, Takahashi YK, Schoenbaum G, Niv Y. Orbitofrontal cortex as a cognitive map of task space. Neuron. 2014; 81:267–279. [PubMed: 24462094]

Young BJ, Fox GD, Eichenbaum H. Correlates of hippocampal complex-spike cell activity in rats performing a nonspatial radial maze task. Journal of Neuroscience. 1994; 14:6553–6563. [PubMed: 7965059]

Young JJ, Shapiro ML. Dynamic Coding of Goal-Directed Paths by Orbital Prefrontal Cortex. The Journal of Neuroscience. 2011; 31:5989–6000. [PubMed: 21508224]

Zhang F, Wang LP, Brauner M, Liewald JF, Kay K, Watzke N, Wood PG, Bamberg E, Nagel G, Gottschalk A, Deisseroth K. Multimodal fast optical interrogation of neural circuitry. Nature. 2007; 446:633–639. [PubMed: 17410168]

Zhang S, Manahan-Vaughan D. Spatial Olfactory Learning Contributes to Place Field Formation in the Hippocampus. Cerebral Cortex. 2013:bht239.

Zilli EA, Hasselmo ME. Modeling the role of working memory and episodic memory in behavioral tasks. Hippocampus. 2008; 18:193–209. [PubMed: 17979198]

- Hippocampus and OFC both encode task structure to anticipate events.
- OFC neurons normally integrate information about expected outcomes.
- Without hippocampal output, these representations are degraded and less integrated.

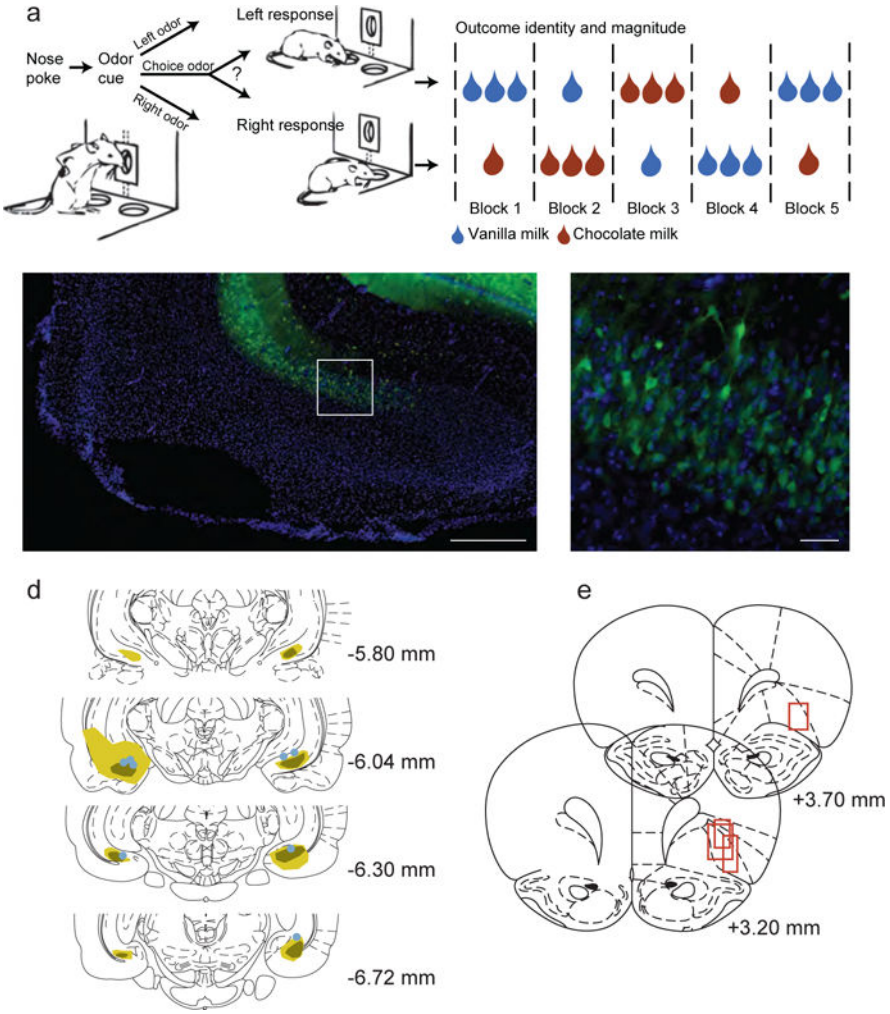

Figure 1. Task and histology

a) Rats performed an odor-guided decision making task. Odor cues delivered to the central port instructed rats on which action (go left, go right, go either direction) would be rewarded on that trial. Rats responded to either the left or right fluid well, where reward was delivered if they chose correctly on forced-choice trials, or for a response to either well on free-choice trials. The outcomes delivered at each fluid well differed in size and flavor, and changed across blocks of trials as illustrated in the example block sequence. b) The ventral subiculum was injected bilaterally with an AAV virus carrying the eNpHR3.0-eYFP construct under the control of the CamKII promoter. Immunohistochemistry was used to identify eNpHR expressing neurons. Green staining indicates eNpHR expression, while the blue DAPI counterstain labels nuclei. Ventral subiculum (vSUB) and dentate gyrus (DG) cell layers are labelled. Scale bar indicates 1 mm. c) A magnified view of the area marked by the white box in panel b, showing individual neurons expressing eNpHR. Scale bar indicates 100 microns. d) Expression was confirmed in the ventral subiculum for all rats. Green shading indicates the maximal (light) and minimal (dark) extent of expression. Dots indicate optical fiber placements. e) Approximate neural recording locations in the OFC are indicated with red boxes.

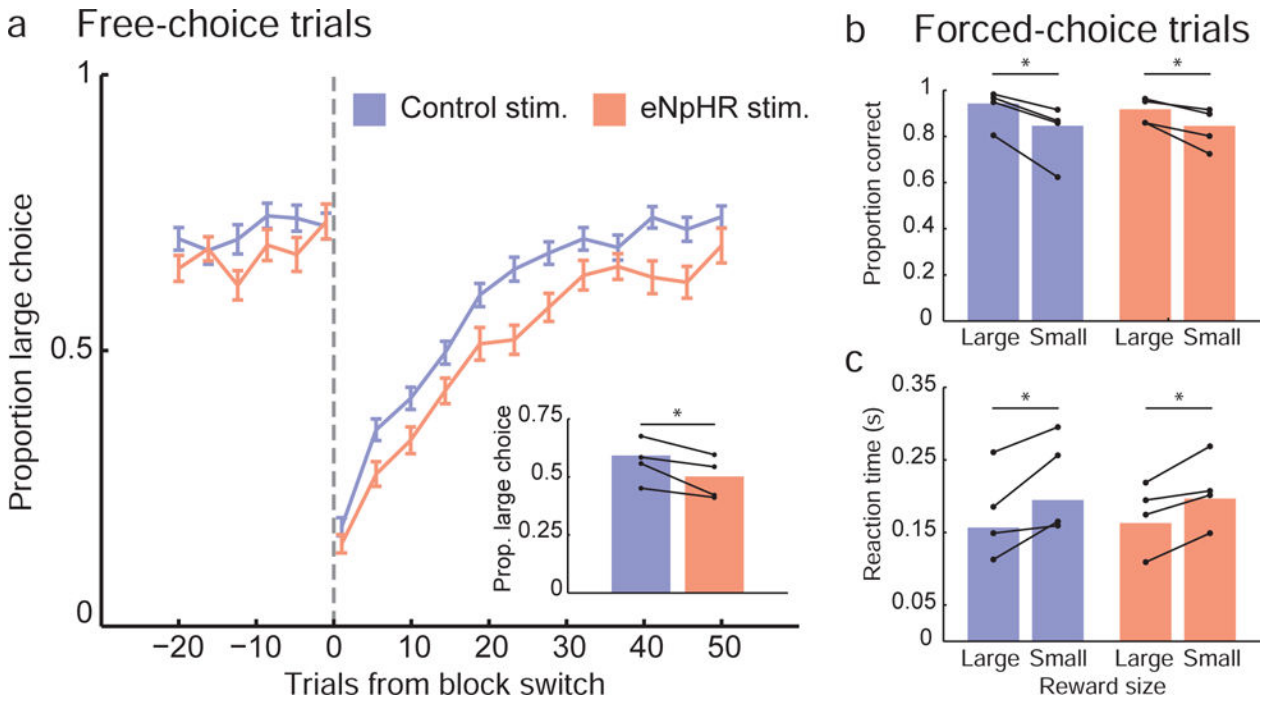

Figure 2. Ventral subiculum inactivation impaired free-choice but not forced-choice behavior

a) The fraction of large choices rats made on free-choice trials was computed and aligned to block switches (when the location of large and small reward outcomes reversed positions; dashed vertical line). With eNpHR stimulation, rats adjusted their behavior in response to block switches more slowly than in control stimulation sessions. Overall, the large outcome choice rate was lower in eNpHR sessions, compared to control sessions (panel a inset). b) Rats performed correctly on a high fraction of forced-choice trials, when the odor cue instructed them which fluid well to select. There was no difference in accuracy between eNpHR and control sessions, although in both session types rats were significantly more likely to select correctly when cued to select the large outcome. Similarly, during both control and eNpHR sessions, reaction times were faster on forced-choice trials directing rats to the large outcome. For all bar graph plots, bars indicate mean values computed across all rats and sessions, and connected dots show mean values computed separately for each rat.

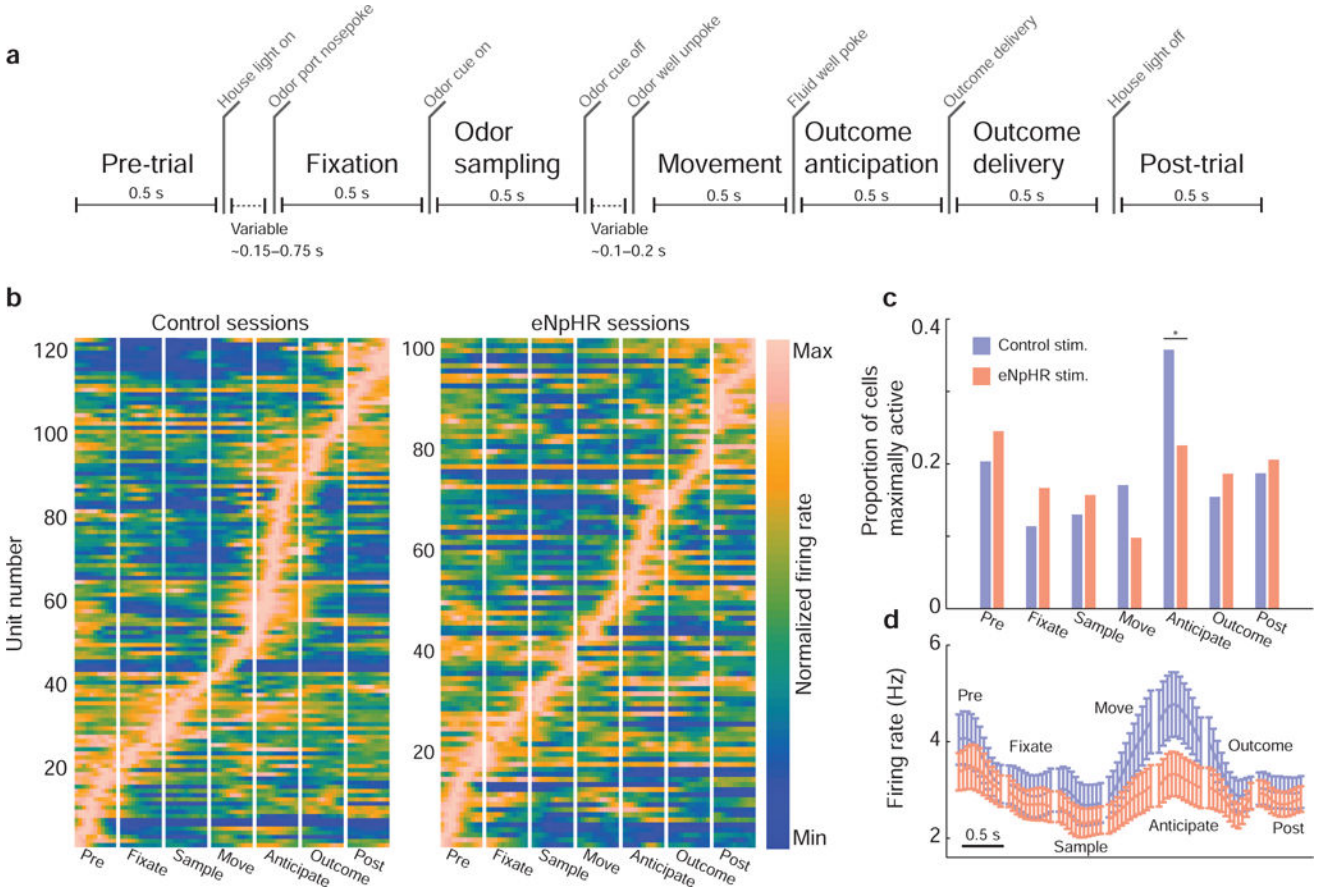

Figure 3. OFC firing rate dynamics were altered by ventral subiculum inactivation

a) Trials were divided into seven 0.5 s epochs bounded by important task events. b) Average, peak-normalized firing rates were computed for cells within task epochs. Each row represents the firing rate of a single unit (brighter colors = stronger activation), with time on the x-axis (each panel = 0.5 s; bin size was 45 ms). c) In control sessions, most OFC neurons reached their maximal firing rate during the outcome anticipation epoch—after rats indicated their decision, but before outcomes were delivered. With subiculum inactivated, significantly fewer OFC neurons reached their peak firing rate during reward anticipation. d) Raw, un-normalized population firing rates showed a similar pattern—in eNpHR sessions the OFC population firing rate was lower than in control sessions. This difference began during the movement epoch (as animals left the odor port and moved towards one of the fluid wells), and was sustained through the reward anticipation period. Bin size was 45 ms; error bars indicated the SEM. See also Figure S1.

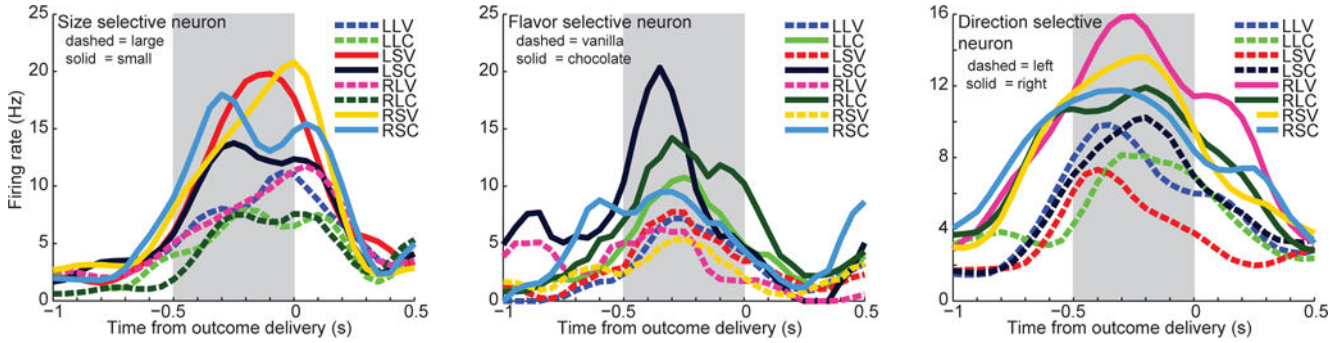

Figure 4. Single unit encoding of action and outcome information

The average firing rates for three example units (recorded during control sessions) are plotted, with activity aligned to outcome delivery. The plots show the unit's average firing rate for each of eight possible outcomes animals selected on forced-choice trials. The shaded region indicates the reward anticipation task epoch. The unit on the left discriminated the size of outcomes; dashed lines group all trials types in which the animal selected the large reward outcome, and solid lines indicate trial types when the rat selected the small outcome. Note that neurons can encode outcome size information with either increases or decreases in firing rate. The unit in the middle discriminated outcome flavor, firing more for trials in which chocolate was selected (solid lines) than trials in which vanilla (dashed lines) was chosen. The unit on the right distinguished response direction, firing faster when the animal chose the right fluid well (solid lines) than when it selected the left fluid well (dashed lines). The bin size used to compute firing rate was 45 ms. Outcome abbreviations: LLV = left, large, vanilla; LLC = left, large, chocolate; LSV = left, small, vanilla; LSC = left, small, chocolate; RLV = right, large, vanilla; RLC = right, large, chocolate; RSV = right, small, vanilla; RSC = right, small, chocolate.

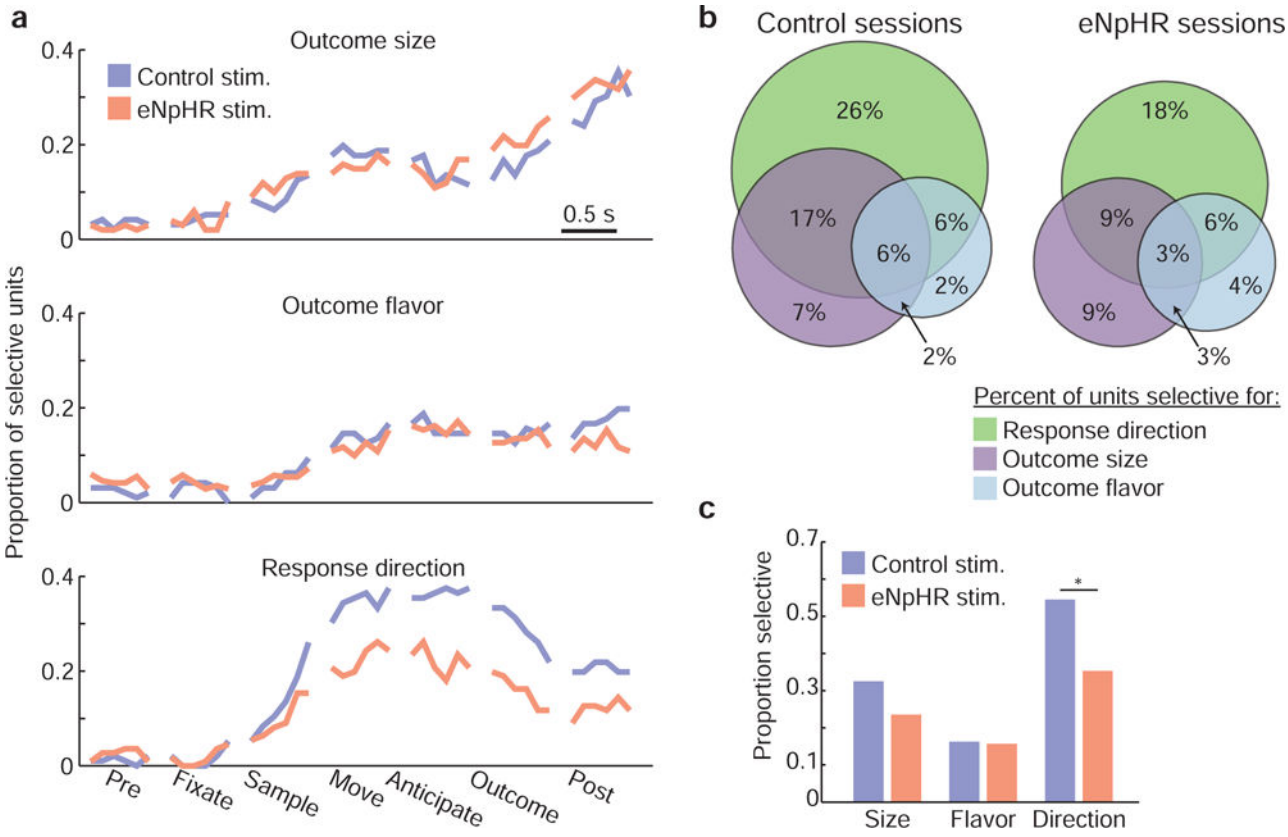

Figure 5. Ventral subiculum inactivation reduced OFC selectivity for response direction

a) We fit regression models to neural data to examine how well the firing rates of individual OFC neurons were explained by the variables outcome size, outcome flavor, and response direction throughout the course of task trials. The fraction of neurons whose responses were significantly modulated by each of these variables increased during odor sampling, and was sustained throughout the trials. While similar proportions of neurons were selective for size and flavor during control and eNpHR sessions, response direction selectivity was strongly attenuated in eNpHR sessions. The bin size for the neural activity used to fit regression models was 45 ms. b) Venn diagrams show the fraction of neurons that were selective for size, flavor, and/or direction during odor sampling, movement, or reward anticipation. c) Ventral subiculum inactivation selectively reduced the proportion of neurons modulated by response direction. See also Figure S2.

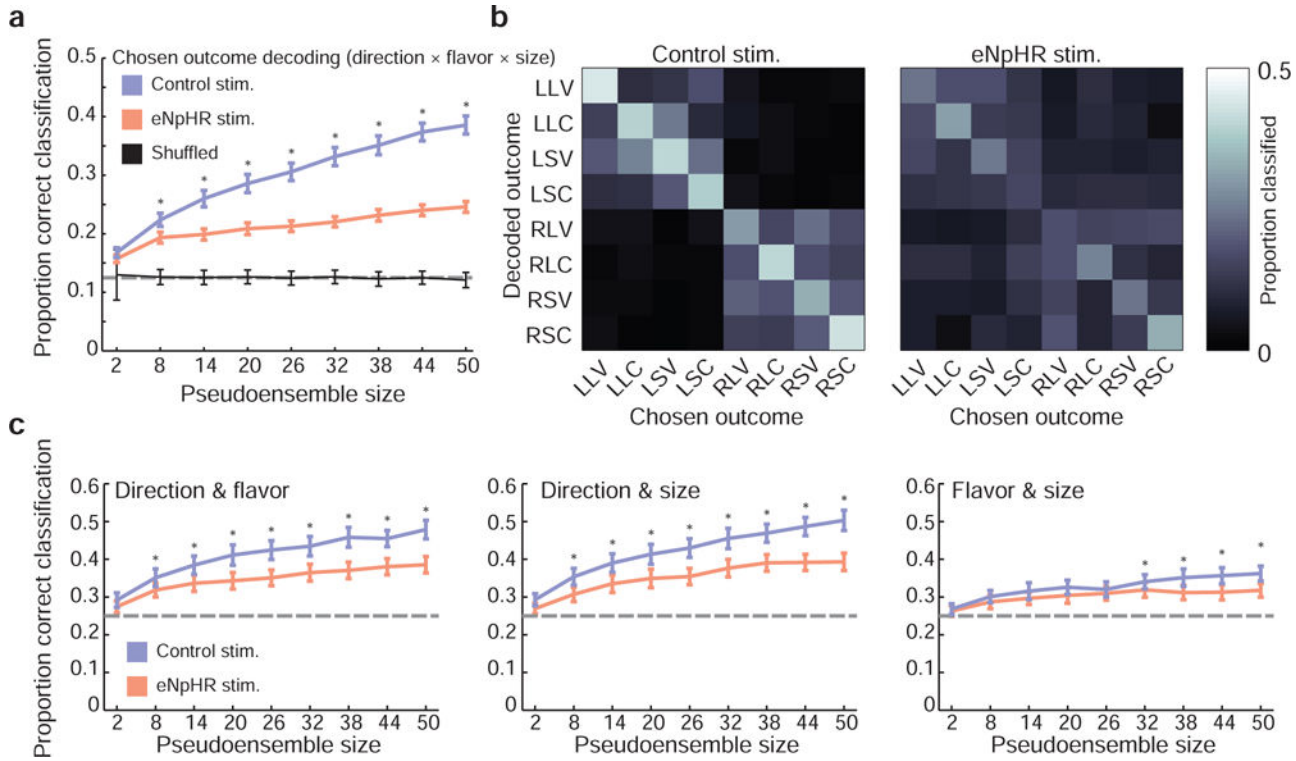

Figure 6. Ventral subiculum inactivation altered OFC outcome representations

a) Decoding chosen outcome was more accurate for control neurons than for eNpHR neurons for most pseudoensemble sizes (paired t-tests, corrected for multiple comparisons). Error bars indicate the SEM in classification accuracy across 250 random pseudoensembles; dashed horizontal line indicates chance level for classification. b) Confusion matrices show the proportion of test samples classified correctly (along the main diagonal, from the upper-left to the lower-right corner) and incorrectly (off the main diagonal) for each trial type using pseudoensembles of 50 neurons. Control session errors in classification were largely confined to the upper left and lower right quadrants, indicating that direction information was generally encoded correctly even when flavor or outcome information were not. In contrast, classification errors in eNpHR sessions were more dispersed. c) Decoding variables in pairs (rather than the full action–outcome triplet) revealed pronounced deficits in eNpHR pseudoensembles when pairs included response direction, along with a more subtle deficit in decoding flavor and size. Error bars indicate the SEM in classification accuracy across 250 random pseudoensembles; dashed horizontal line indicates chance level for classification. Asterisks indicate significant differences between control and eNpHR pseudoensemble classification performance (paired t-tests, corrected for multiple comparisons).

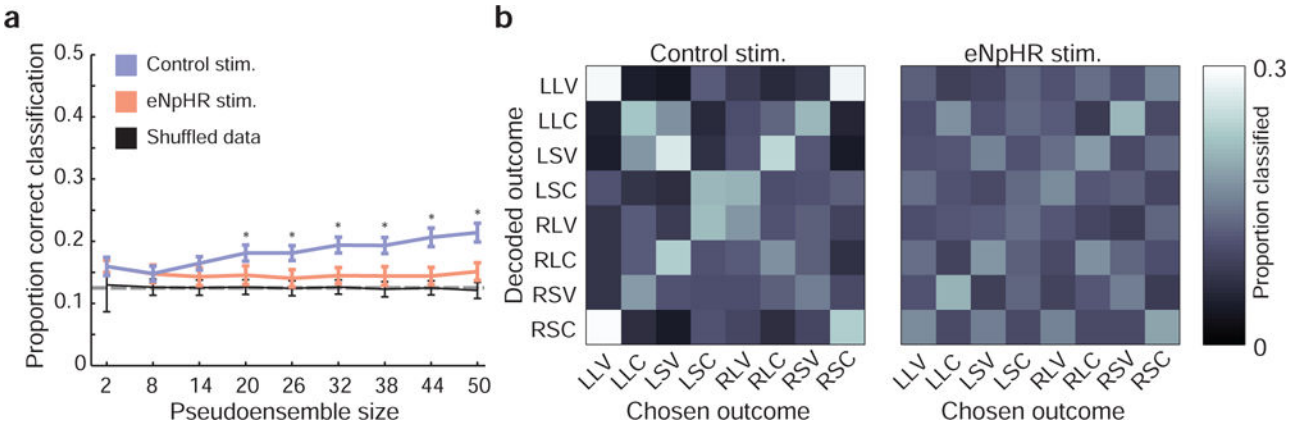

Figure 7. OFC represents task structure information beyond outcome features

a) With selectivity for action and outcome features removed, pseudoensembles drawn from control session neurons achieved significantly better classification performance than eNpHR pseudoensembles. Error bars indicate the SEM in classification accuracy across 250 random pseudoensembles; dashed horizontal line indicates chance level for classification. Asterisks indicate significant differences between control and eNpHR pseudoensemble classification performance (paired t-tests, corrected for multiple comparisons). b) For control data, errors in classification were concentrated along the minor diagonal from the lower-left to upper-right corner, corresponding the unchosen action and outcome combination that accompanied the option rats chose. This pattern was substantially weaker in eNpHR sessions.

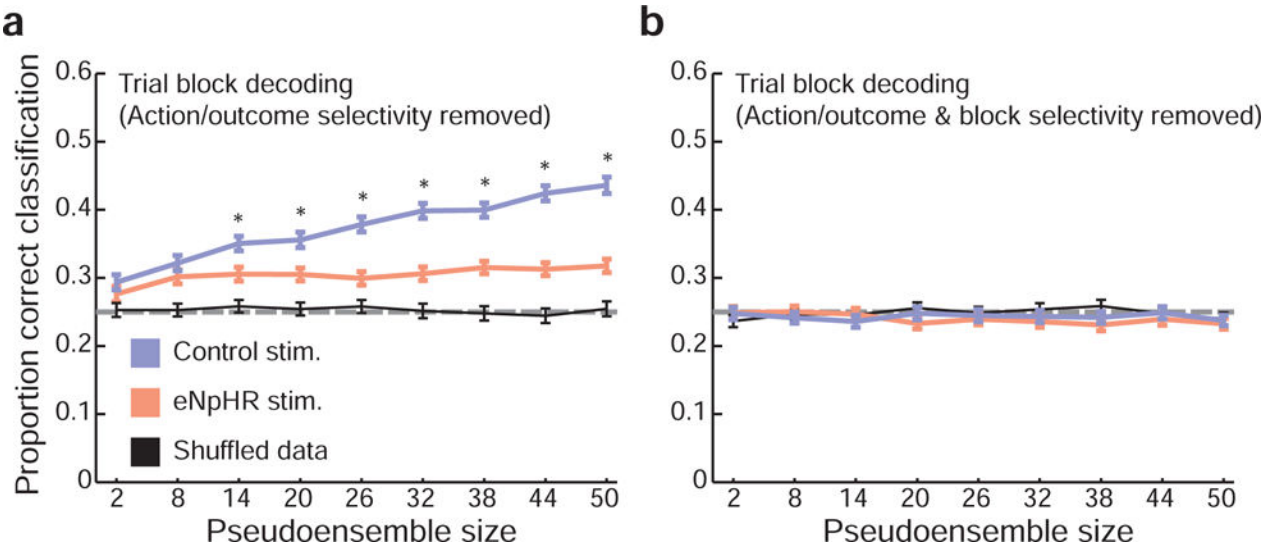

Figure 8. Ventral subiculum inactivation abolishes OFC task block representations

a) With selectivity for action and outcome features removed, control session pseudoensembles were significantly better at trial block decoding. Error bars indicate the SEM in classification accuracy across 250 random pseudoensembles; dashed horizontal line indicates chance level for classification. Asterisks indicate significant differences between control and eNpHR pseudoensemble classification performance (paired t-tests, corrected for multiple comparisons). b) However, when trial block was included as a predictor in the regression models, no differences in performance were observed between eNpHR and control data.