Article

Current Biology

Complementary Task Structure Representations in Hippocampus and Orbitofrontal Cortex during an Odor Sequence Task

Highlights

- Neural encoding of an odor sequence in both HPC and OFC is shaped by task demands
- HPC emphasizes early states where explicit information must be encoded in memory
- OFC emphasizes later states where hidden information must be recalled for behavior
- Thus, HPC and OFC can play complementary roles in representing task states

Authors

Jingfeng Zhou, Marlian Montesinos-Cartagena, Andrew M. Wikenheiser, Matthew P.H. Gardner, Yael Niv, Geoffrey Schoenbaum

Correspondence

jingfeng.zhou@nih.gov (J.Z.), geoffrey.schoenbaum@nih.gov (G.S.)

In Brief

Analyzing neural representations of an odor sequence, Zhou et al. show that HPC emphasizes early positions where demands to encode explicit information in memory are high, while OFC emphasizes later positions where hidden information must be recalled. These findings show that HPC and OFC can play complementary roles in representing task states.

Zhou et al., 2019, Current Biology 29, 3402–3409
October 21, 2019 Published by Elsevier Ltd.
https://doi.org/10.1016/j.cub.2019.08.040

CellPress

CellPress

Current Biology
Article

Complementary Task Structure Representations in Hippocampus and Orbitofrontal Cortex during an Odor Sequence Task

Jingfeng Zhou,[1,*] Marlian Montesinos-Cartagena,[1] Andrew M. Wikenheiser,[1] Matthew P.H. Gardner,[1] Yael Niv,[2] and Geoffrey Schoenbaum[1,3,4,5,*]
[1]Intramural Research Program of the National Institute on Drug Abuse, Baltimore, MD 21224, USA
[2]Princeton Neuroscience Institute and Department of Psychology, Princeton University, Princeton, NJ 08544, USA
[3]Department of Anatomy and Neurobiology, Maryland School of Medicine, Baltimore, MD 21201, USA
[4]Department of Neuroscience, Johns Hopkins School of Medicine, Baltimore, MD 21287, USA
[5]Lead Contact
*Correspondence: jingfeng.zhou@nih.gov (J.Z.), geoffrey.schoenbaum@nih.gov (G.S.)
https://doi.org/10.1016/j.cub.2019.08.040

SUMMARY

Both hippocampus (HPC) and orbitofrontal cortex (OFC) have been shown to be critical for behavioral tasks that require use of an internal model or cognitive map, composed of the states and the relationships between them, which define the current environment or task at hand. One general idea is that the HPC provides the cognitive map, which is then transformed by OFC to emphasize information of relevance to current goals. Our previous analysis of ensemble activity in OFC in rats performing an odor sequence task revealed a rich representation of behaviorally relevant task structure, consistent with this proposal. Here, we compared those data to recordings from single units in area CA1 of the HPC of rats performing the same task. Contrary to expectations that HPC ensembles would represent detailed, even incidental, information defining the full task space, we found that HPC ensembles—like those in OFC—failed to distinguish states when it was not behaviorally necessary. However, hippocampal ensembles were better than those in OFC at distinguishing task states in which prospective memory was necessary for future performance. These results suggest that, in familiar environments, the HPC and OFC may play complementary roles, with the OFC maintaining the subjects' current position on the cognitive map or state space, supported by HPC when memory demands are high.

INTRODUCTION

Both the hippocampus (HPC) and the orbitofrontal cortex (OFC) have been proposed to encode a cognitive map of task space, a mental model of events, and their relationships in behavioral tasks [1–9]. But their unique roles and possible interplay in this function is largely a matter of speculation since few truly parallel studies have been conducted to compare encoding in these two regions on appropriate tasks [10, 11]. Since the HPC has been shown to represent information, such as space, time, and even incidental associations [12–19], while the activity of the OFC is largely recruited to encode expected outcomes [20–22], a popular idea has arisen that the HPC contributes concrete and detailed environmental information, which is received and tailored by the OFC to signal information that is more relevant to the subject's current subjective needs [4, 23–25]. On the other hand, there is ample evidence that neural activity in HPC is also strongly shaped by the environment and the subject's current goals [26–29]. This suggests an alternative model in which the two regions cooperate in organizing information to drive behavior.

Our previously developed odor sequence task, with 24 different task states defined by both external sensory stimuli and internal memory of the sequences, provides an appropriate tool with which to examine this question. It includes explicitly and implicitly cued states, some of which are reward relevant and some of which are not, all embedded in a non-spatial map defined by clear relationships [8]. A full representation of the task structure would reflect the knowledge of both reward-relevant as well as reward-irrelevant information, while a tailored task-oriented representation would highlight information related to the behavioral goals. Indeed, our previous study found a reduced representation of the task structure in OFC that largely corresponded to the information the rats needed to correctly perform the task. Here, we set out to test whether the HPC might encode a more complete task space, less biased by current goals and therefore including more information about behaviorally irrelevant relationships.

RESULTS

We recorded from the HPC in 9 rats trained to perform the odor sequence task previously used to record in the OFC (Figures 1A and 1B) [8]. Briefly, the task consisted of a sequence of "go/no-go" odor discrimination steps. In each step in the sequence (i.e., trial), one of 16 different odor cues was presented, and the rat had to decide whether to "go" to a nearby fluid well within 2 s to obtain a sucrose reward or withhold responding to avoid a

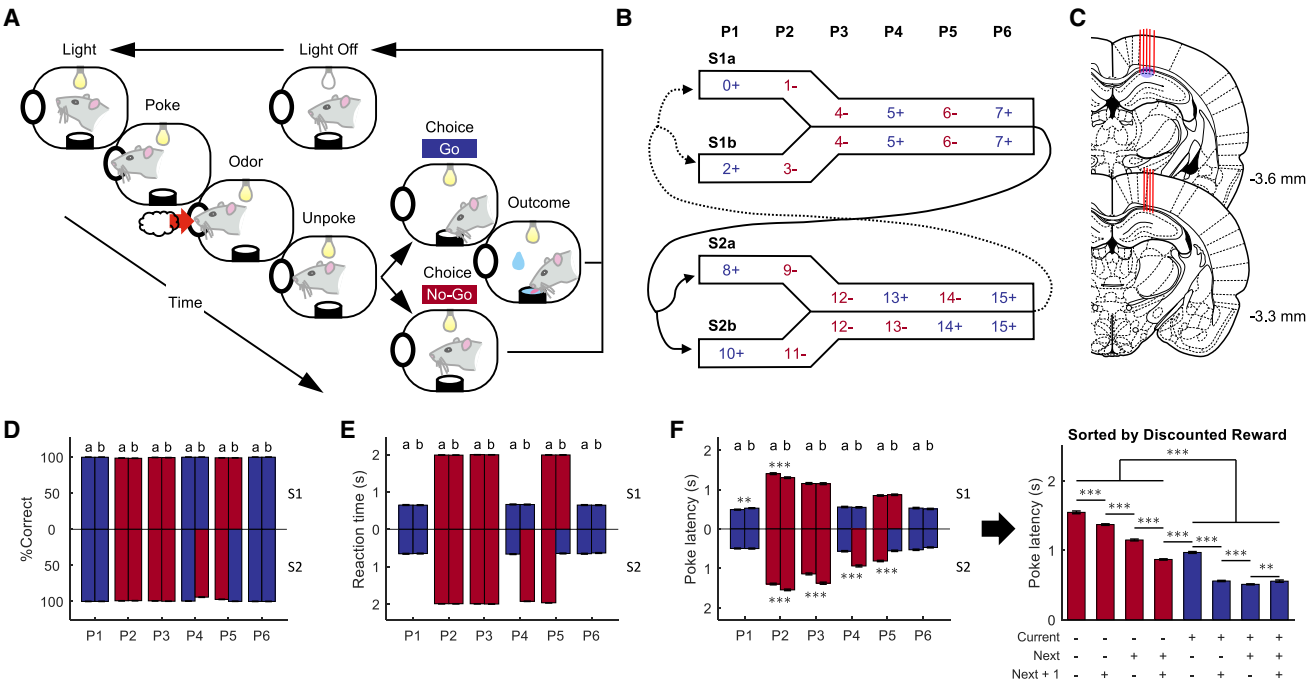

Figure 1. Odor Sequence Task Design and Behavioral Performance

(A) Depiction of a single trial in the odor sequence task. Upon the onset of two houselights on the front wall, rats initiated a trial by nose poking into an odor port and sampled one of 16 odors. Within a 2 s time window, rats decided to enter the water well or not depending on the odor identities on current and sometimes also past trials.

(B) The 16 different odors were presented in four sequences. Arrows indicate the transitions between sequences.

(C) Recoding sites of CA1. Red lines indicate cannula and electrode tracks from 9 rats. The blue shaded area shows one example of electrolytic lesion made during the perfusion.

(D) Percent correct on each trial type for the two S1 sequences (above the middle) and two S2 sequences (below the middle). Blue, rewarded ("go") states; red, non-rewarded ("no-go") states. Error bars within this figure are SEMs.

(E) Reaction time measured the period from the odor port exit to the water well entry. Correct no-go responses were encoded as a reaction time of 2 s (the response time window).

(F) Poke latency was the time from the onset of houselights to the entry of the odor port. The poke latencies were significantly different between S2a and S2b at P2, P3, P4, and P5 ($p = 5.9 \times 10^{-8}$, 1.5×10^{-10}, 5.2×10^{-30}, and 7.1×10^{-29}; W = 91,048, 88,717, 76,979, and 130,169, respectively; two-sided Wilcoxon rank-sum test; n = 322 sessions) and also between S1a and S1b at P1 and P2 ($p = 8 \times 10^{-3}$ and 8×10^{-4}; W = 97,624 and 111,755, respectively; two-sided Wilcoxon rank-sum test; n = 322 sessions). This latency was highly modulated by both current and future rewards, as seen in the right plot, which averages over states that have a similar three-trial future reward and is sorted by discounted future reward (from lowest to highest). Poke latency was shorter when more reward was forthcoming ($p = 0$; W = 9,050,017 for the comparison between negative and positive states; $p = 1.5 \times 10^{-13}$, 1.4×10^{-33}, 3×10^{-63}, 1×10^{-6}, 1.9×10^{-98}, 9.2×10^{-7}, and 5×10^{-3}; W = 314,497, 1,636,876, 1,345,843, 1,567,283.5, 1,774,243, 1,164,216, and 1,016,706 for consecutive pairs of bars from left to right, respectively; two-sided Wilcoxon rank-sum test). A three-way ANOVA analysis revealed that reward on current, next, and next + 1 trials significantly affected poke latency ("current": F(1, 7,721) = 7,697.73, $p = 0$; "next": F(1, 7,721) = 921.61, $p = 0$; "next + 1": F(1, 7,721) = 146.28, $p = 0$). There were also significant interactions between them ("current × next": F(1, 7,721) = 675.92, $p = 0$; "current × next + 1": F(1, 7,721) = 256.64, $p = 0$; "next × next + 1": F(1, 7,721) = 10.58, $p = 1 \times 10^{-3}$). + indicates a rewarded trial and − a non-rewarded trial; **p < 0.01; ***p < 0.001.

See also Figures S1, S2, and S3.

prolonged inter-trial interval (ITI). The 16 odor cues were presented in predictable sequences. There were two major sequences (S1 and S2), each of which nested two subsequences (a and b). Each sequence consisted of six different odors or positions (P1–P6). Odors at the first two positions (P1 and P2) were unique for each subsequence (8 odors; two for each of the 4 subsequences), while odors at each of the last four positions (P3–P6) were common across the two major sequences (8 remaining odors). Thus, S1a and S1b shared one set of odors at these positions, and S2a and S2b shared another set. All but two of the odors had a fixed association with reward or non-reward, so for most of the trial types or positions, correct behavior did not require information about the sequence. However, for the odors at P4 and P5 in S2, the significance of the odor differed for the two subsequences S2a and S2b. This required the rats to maintain information across several trials to perform correctly.

A bundle of 16 stereotrodes (32 channels) was implanted in the dorsal HPC after the rats were extensively trained. After recovery from surgery, single-unit activity was recorded from CA1 as rats performed the task (Figure 1C; n = 2,056 neurons). Rats were highly proficient on the task during recording, performing accurately on the discrimination at all positions, including P4 and P5 in S2, where performance required memory

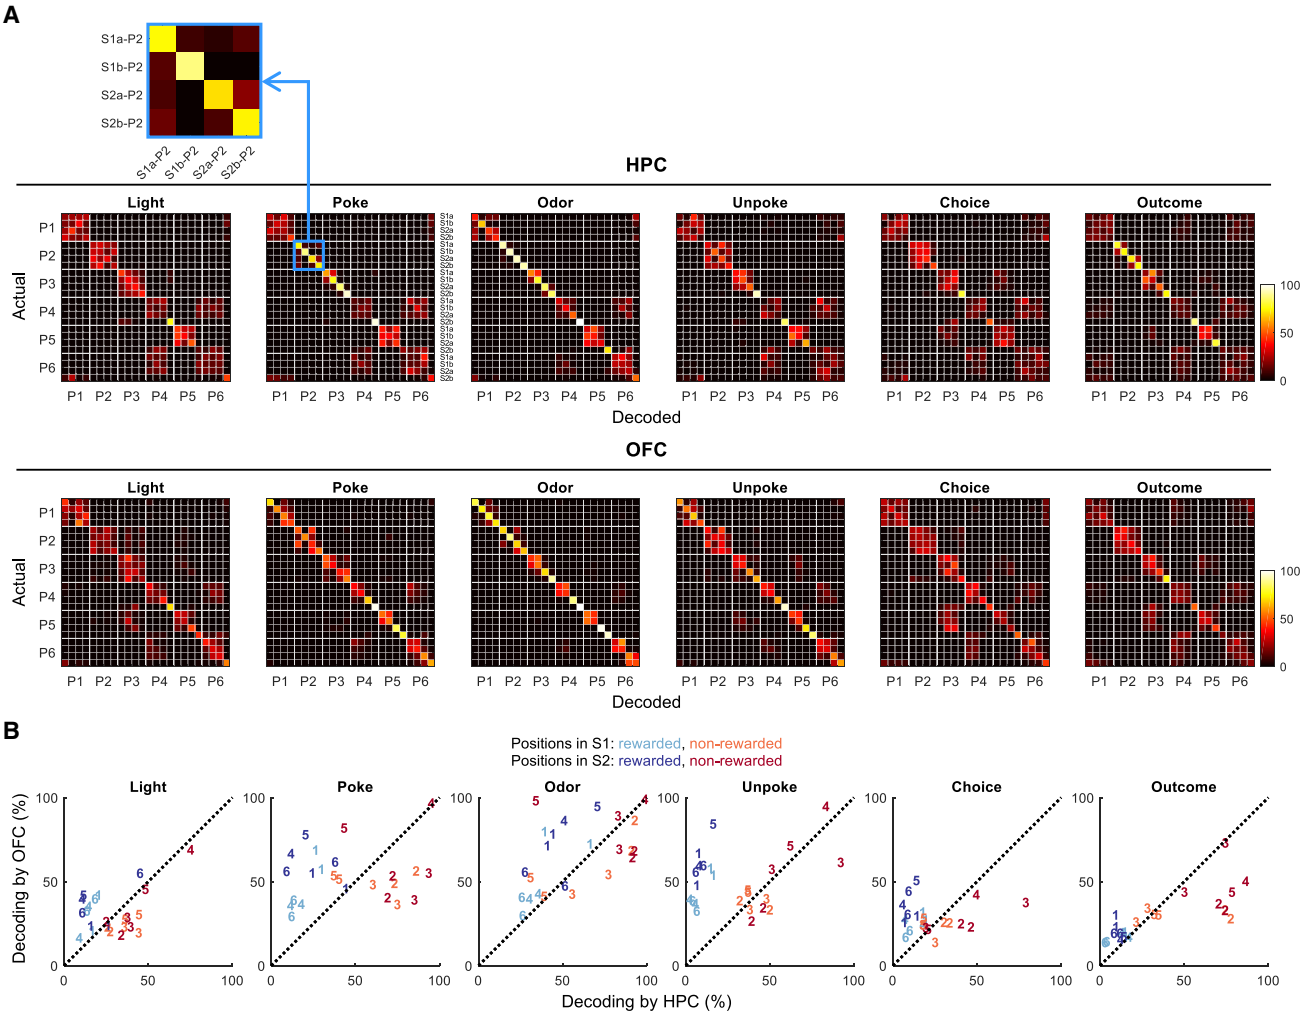

Figure 2. Comparisons of Decoding Accuracy of 24 Task States between the HPC and OFC

(A) Confusion matrices obtained from the decoding of 24 states from HPC ensemble activity at six different task positions. The y axis for each panel indicates the rats' actual state, and the x axis shows the state decoded from the neural ensemble activities. The brighter color indicates a higher probability of decoding the state. These plots show what states are confused with what states; for example, at the time of the light (left matrix), each of the four states in P2 (one in each sequence) is confused with all other states in P2, but not with states in other positions. In contrast, at the time of the poke, these states are well distinguished. (B) Scatterplots show the accuracy in decoding each of the 24 states by HPC and OFC ensembles. The colored numbers (1–6) indicate positions (P1–P6). Light blue and light red, rewarded and non-rewarded positions in S1, respectively; dark blue and dark red, rewarded and non-rewarded positions in S2, respectively. See also Figure S4.

for the events of prior trials (Figures 1D and 1E). The influence of the current position in the sequence was apparent in rats' latency to initiate trials following light onset. This latency was also tightly modulated by future rewards, reflecting knowledge of the sequence structure (Figure 1F). These behavioral results were all similar to those obtained under the same task design and conditions during recording from OFC [8]. At the single-unit level, both areas showed substantial neuronal selectivity to 24 task states as well as reward values of not only the current but also the past and future trials and sequence pairs at each position (Figures S1, S2, and S3). Therefore, we combined the newly recorded HPC data and existing OFC data (n = 1,078 neurons) to directly compare task structure representations between the two areas.

The Precision of Hippocampal Encoding of Task States Was Dampened by Proximal Reward and Task Irrelevance

To compare task structure representations, we constructed pseudo-ensembles of CA1 single units and tested for representation of the 24 positions in the 4 sequences, using data from epochs corresponding to each of the 6 different events in a trial ("light," "poke," "odor," "unpoke," "choice," and "outcome"). Like ensembles constructive of OFC neurons, HPC ensembles contained significant information about sequence positions shown in confusion matrices (Figure 2A), particularly during the periods surrounding odor sampling that were key to correct responding. In both areas, the patterns evident in the confusion matrices developed gradually as ensemble size increased and

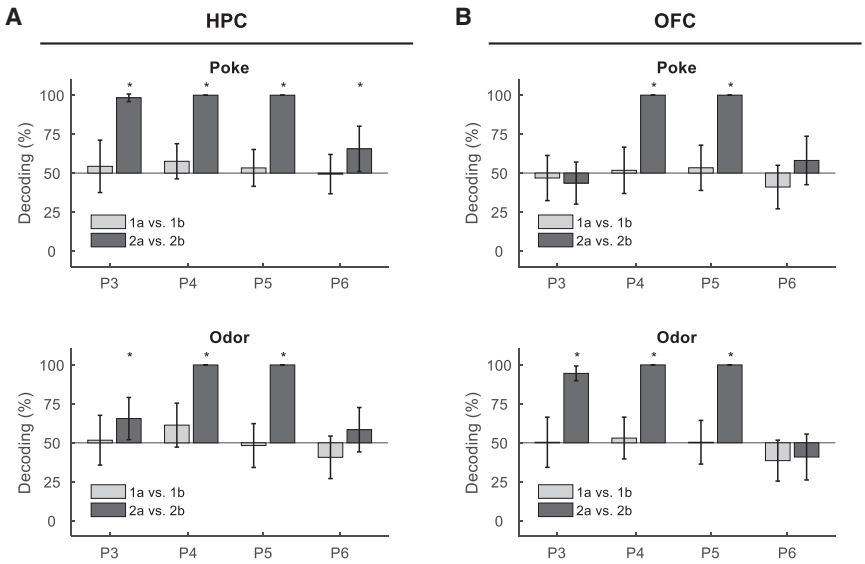

Figure 3. Comparisons of Decoding Accuracy of Sequence Pairs at Odor-Overlapping Positions between the HPC and OFC

(A and B) Binary decoding of subsequences S1a vs. S1b and S2a vs. S2b during "poke" (upper panel) and "odor" (lower panel) time by HPC (A) and OFC (B). Asterisks indicate the mean decoding accuracy exceeds the 95% confidence interval constructed with label-shuffled data. Error bars indicate standard deviations (SDs).

did not depend on recordings from any particular subject (Figure S4), suggesting that the observed effects reflect information distributed across relatively large portions of the population. This is consistent with the single-unit analysis presented above and also with the varied patterns evident in individual single units (Figures S1 and S2; see also Figure S3 in [8] for OFC single units).

However, contrary to the hypothesis that HPC would represent sequence information with higher fidelity than OFC—and with less bias for behavioral relevance, internal goals, or reward—HPC ensemble activity at different positions (P1–P6) performed generally no better and often worse than OFC ensembles at correctly identifying the rat's specific state (i.e., one of 24; Figure 2B). Indeed, rather than being less sensitive to behavioral relevance, HPC ensembles seemed, if anything, more sensitive to this factor in what they represented (Figures 3 and 4).

The influence of behavioral relevance on the encoding of positions by HPC ensembles was particularly evident in two aspects of the results. One was in how the HPC ensembles represented P3–P6 in S1, where the rats did not need to distinguish subsequences (Figure 3). OFC performed at chance at distinguishing these positions (Figure 3B), providing an effective floor against which to see a better representation of incidental sensory associations in HPC. However, HPC ensembles also performed at chance at decoding these states; like OFC, they were much better at distinguishing these positions in S2, where

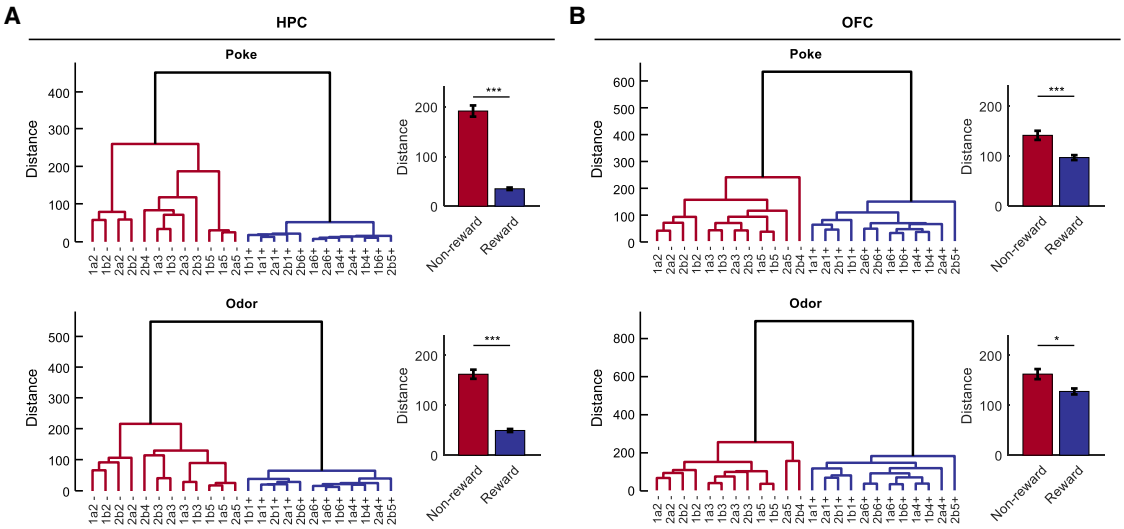

Figure 4. Hierarchical Clustering of 24 Task States by the HPC and OFC

(A and B) Hierarchical clustering of 24 states based on population neural activities at the time of the trial-initiating nose poke (top) and odor sampling (bottom) in HPC (A) and OFC (B). The Mahalanobis distance was used to create a dissimilarity matrix that measured the dissimilarity or distance between each pair of task states in the population activity space. The dissimilarity matrix was used to construct a hierarchical clustering tree by an unweighted average linkage method. The clustering results were shown in dendrograms. In the HPC, the average pairwise distance between the different non-rewarded states was larger than the distance between the different rewarded states (middle), indicating better encoding of non-rewarded states ("poke time": mean difference = 156.4, $p = 6 \times 10^{-20}$, W = 2,381; "odor time": mean difference = 112.7, $p = 2.4 \times 10^{-17}$, W = 2,527; two-sided Wilcoxon rank-sum test; n = 66 pairs). The same effect also appeared in the OFC (right) but was much less prominent (poke time: mean difference = 44.6, $p = 6 \times 10^{-5}$, W = 3,514; odor time: mean difference = 34.8, p = 0.03, W = 3,908; two-sided Wilcoxon rank-sum test; n = 66 pairs). *p < 0.05; ***p < 0.001. Error bars are SEMs.

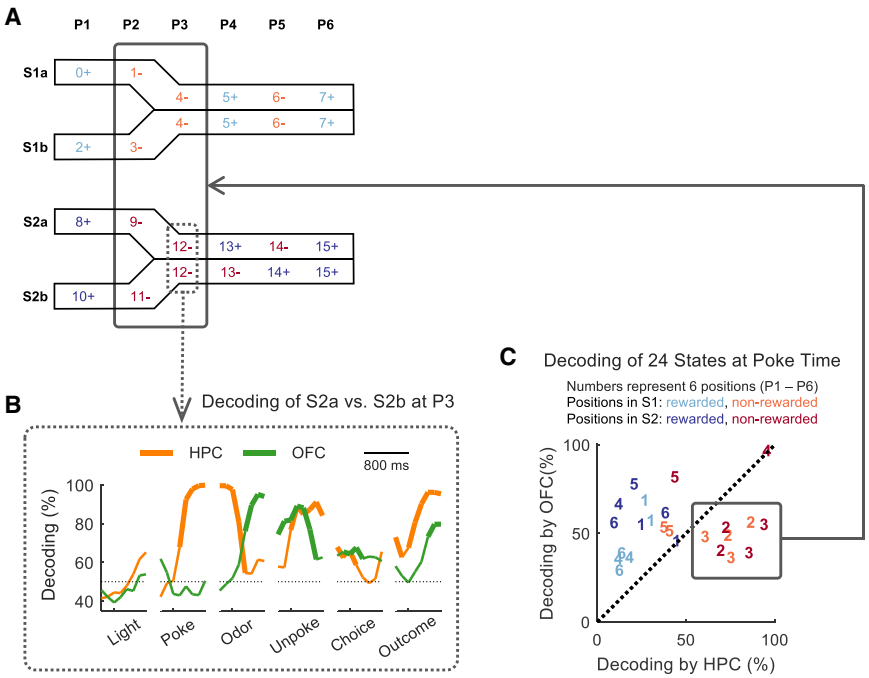

Figure 5. Decoding of States at P2 and P3 from HPC Ensemble Activity at Poke Time Was Better than from OFC

(A) Positions (P1–P6) in S1 were colored with light blue (rewarded) and light red (non-rewarded). Positions (P1–P6) in S2 were colored with dark blue (rewarded) and dark red (non-rewarded). Dotted square highlights two task states 2a3 and 2b3. Solid square highlights 8 task states at P2 and P3.

(B) Binary decoding of sequences 2a versus 2b at P3 across time by HPC (orange) and OFC (green) ensembles. Bold lines indicate the mean decoding accuracy was above the 95% confidence interval estimated from the decoding with label-shuffled data. Time bin = 100 ms.

(C) Comparison of decoding accuracies of 24 task states by HPC and OFC identified 8 states (highlighted by a gray square; $p < 0.001$; two-sided Wilcoxon rank-sum test) where HPC showed better decoding than OFC. In addition, OFC showed significantly better decoding than HPC for the rest task states ($p < 0.001$; two-sided Wilcoxon rank-sum test) except S2b1 ($p = 0.35$; two-sided Wilcoxon rank-sum test). Task states were color coded the same way as in (A).

doing so was relevant to the rats' behavior (Figure 3A). The second area where the impact of behavioral relevance was evident was in the strong effect of current reward on encoding by the HPC ensembles (Figure 4). HPC ensembles showed much better state discrimination between non-rewarded positions than between rewarded positions during both poke and odor time (Figure 4A). OFC ensembles also showed similar trends but with smaller effect sizes (Figure 4B).

HPC Activity Distinguishes Different Sequences at P2 and P3 Better Than OFC

To further investigate the differences between HPC and OFC, we next looked at the representation of P3 in sequence 2a versus 2b, where the rats were required to maintain the memory of past odors without discriminative sensory cues in order to know what to do in the following position (P4). Distinguishing P3 in sequence 2a versus 2b was significantly better for HPC ensembles as compared to OFC ensembles during the poke time (Figures 5A and 5B). Ensemble performance at poke time was particularly intriguing because, during this period, there was no confound from bottom-up sensory-driven neural responses as may occur during the odor period. Thus, any information encoded during the poke time must be supported purely by the sequence information at all positions. We therefore asked whether encoding by HPC ensembles was better than by OFC ensembles at all 24 positions at poke time. This analysis identified 8 positions or states that were significantly better represented by ensembles from HPC than from OFC (Figure 5C). These 8 states all occurred at P2 and P3—critical positions in each sequence where the rats transitioned from the odor-unique positions, where external information was provided regarding sequence, to the odor-common arms, where tracking sequence required memory for prior events.

Complementary Emphases on State Representations by HPC and OFC

To further test whether this transition from the unique to the common arms captured a complementary pattern of encoding in HPC and OFC, we sorted the 24 states based on how well each was represented by the two areas (Figures 6A and 6B) at the poke time. Task states best represented by each area were largely non-overlapping, with HPC better representing states around the transition and OFC better representing states in the common arms after the transition.

We next asked whether this pattern was also evident in the ability of the ensembles to distinguish positions (P2–P5) across each sequence pair (S1a versus S1b and S2a versus S2b; Figure 6C). We defined P2 and P3 as the "transition" phase and positions P4 and P5 as the "after-transition" phase of the sequence. Consistent with data shown in Figure 3, both HPC and OFC did poorly at distinguishing positions in S1a versus S1b once the odors and outcomes were the same (P3–P5), confirming that neither region maintained information that was not required for the behavioral performance (Figure 6C, left panel). By contrast, in each of the other comparisons, the HPC ensembles performed significantly better at discriminating positions at the transition than did ensembles composed of OFC neurons. After the transition, this pattern was largely reversed, with OFC ensembles providing better decoding than HPC ensembles at P4 and P5 as long as the value was the same (Figure 6D; S2a versus either S1a or S1b). When the value of these positions differed in the comparison (Figure 6C, right panel; S2a versus S2b), both regions performed well after the transition. However, for HPC, this performance is deceptive, since it is based not on true representation of the two states at each position but rather on the relatively uniform representation of all the rewarded positions in both sequences (see dendrograms; Figure 4). In the comparisons, when the reward availability was

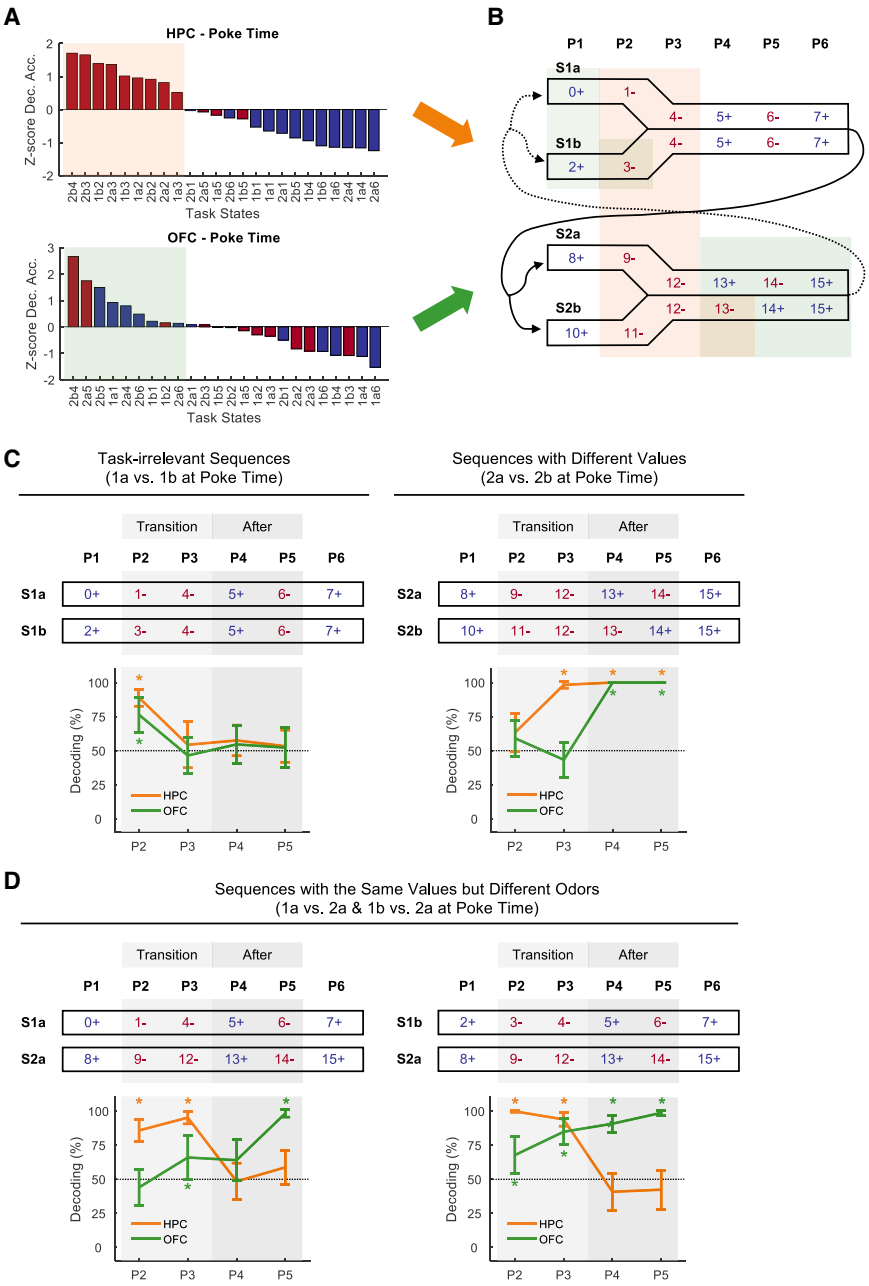

Figure 6. States Best Encoded in HPC Were Predominantly at P2 and P3, and Those Better Encoded in OFC Were Later in the Sequence

(A) Twenty-four task states sorted based on the *Z* scores of their decoding accuracies at the poke time from HPC activity (top panel) or OFC activity (bottom panel).

(B) The sorting in (A) identified the first 9 states where HPC (orange background) or OFC (green background) showed the best decoding accuracy.

(C) Binary decoding of 1a versus 1b (left) and 2a versus 2b (right) at four positions (poke time; P2–P5; 200 repeats of decoding processes). Orange and green asterisks (meaning HPC and OFC, respectively) indicate the mean decoding accuracy is above the 95% confidence interval from the decoding with label-shuffled data. Error bars are SDs.

(D) Binary decoding of 1a versus 2a (left) and 1b versus 2a (right) at four positions (poke time; P2–P5; 200 repeats of decoding). The colors have the same meanings as in (C). Error bars are SDs. In addition, for 1a versus 2a, a two-way ANOVA analysis ("area" as factor 1: HPC or OFC; "phase" as factor 2: "transition" or "after transition") revealed main effects by both brain areas ($F(1, 1{,}596) = 24.9$; $p = 6.8 \times 10^{-7}$) and task phases ($F(1, 1{,}596) = 45.5$; $p = 2.2 \times 10^{-11}$). An interaction was also observed ($F(1, 1{,}596) = 1{,}575.9$; $p = 2.6 \times 10^{-240}$). For the HPC, the decoding accuracy of subsequences significantly decreased from the "transition" phase to "after transition" phase ($p = 3.8 \times 10^{-9}$; post hoc Tukey's HSD (honestly significant difference) test); while for the OFC, it significantly increased ($p = 3.8 \times 10^{-9}$; post hoc Tukey's HSD test). Similar results were also seen for 1b versus 2a by a two-way ANOVA analysis. Specifically, both brain areas ($F(1, 1{,}596) = 913.5$; $p = 4.7 \times 10^{-159}$) and task phases ($F(1, 1{,}596) = 1{,}166.5$; $p = 2.2 \times 10^{-192}$) showed main effects. An interaction ($F(1, 1{,}596) = 4{,}678.8$; $p = 0$) exists between the two factors. The decoding of subsequences in the HPC decreased from the transition phase to after transition phase ($p = 3.8 \times 10^{-9}$; post hoc Tukey's HSD test), while that in the OFC increased ($p = 3.8 \times 10^{-9}$; post hoc Tukey's HSD test).

the same, HPC showed the best decoding during the transition but decreased its ability in distinguishing sequences after the transition, while OFC showed the opposite pattern (Figure 6D).

DISCUSSION

The current study sought to test the hypothesis that the HPC represents detailed, sometimes incidental, information about the cognitive map of a task, while such information is tailored in the OFC to highlight distinctions more directly relevant to behavioral goals [4, 5, 24, 30]. In the odor sequence task used here, this hypothesis predicts that the identity of the 24 task states should be better encoded in activity recorded in the HPC as compared to the OFC. In particular, we expected that HPC ensembles would perform as well as those from OFC at distinguishing behaviorally relevant states and much better at distinguishing states—such as those in the common arm of S1—whose discrimination was unimportant for performing the task correctly. However, our data showed that the task representations in both OFC and HPC were closely tied to behavioral demands; neither maintained sequence information that was not required for behavior (Figure 3), and the HPC ensembles actually performed more poorly than those in OFC at distinguishing different rewarded states (Figure 4).

In evaluating the significance of these findings, it is important to note that the task environment was familiar to the rats in both

studies. This is by necessity to some extent, since it would be difficult to analyze neural representations during poor and highly variable performance earlier in learning, when such representations also may be changing as the task becomes familiar. This leaves open the possibility that, earlier in learning, the HPC-OFC system might operate differently. However, the current results suggest that, once an environment is familiar, neither region maintains much information about the structure of the task that is incidental to task performance.

What then distinguishes the cognitive maps in OFC and HPC in familiar environments? One striking difference between the HPC and OFC was that the HPC encoding was highly biased toward the non-rewarded task states [31]. This did not seem to simply reflect non-reward, however, since the non-rewarded positions at P5 (1a5−, 1b5−, and 2a5−) were poorly represented by the HPC (Figure 6D). Instead, this finding seems to be better explained by the proximity of these states to entry into the common arms of the sequences. This transition point is where information must be encoded in episodic memory, so that later states in the common arm can be properly distinguished. The bias in HPC to strongly represent these positions, and the relatively poor encoding of rewarded positions, is reminiscent of evidence that the HPC represents the start of journeys through a mapped environment [32–34] and is also similar to distinctions between HPC and prefrontal representations seen in prior work on context [35–40]. For example, theta oscillations in the HPC have been shown to precede those in the medial prefrontal cortex (mPFC) when subjects first enter a unique context, while mPFC theta activity precedes that in the HPC later in the trial when decisions are made within this context [41]. In our task, the unique arms may be thought of as setting the context for determining the rules or map to apply in the common arms.

For the most part, discriminating between sequences on the common arms had to be based on latent or hidden information, available only in memory. In this regard, it is interesting that, once in the common arms, the HPC generalized across states with similar external information, whereas OFC became better at discriminating these hidden task states. This suggests that one distinction between the cognitive maps represented in the two areas may be the degree to which inference or encoding of latent information underlies what is represented. While this may be interpreted as a bias in OFC for task-relevant information, the seeming task relevance of OFC representations may simply reflect the importance of latent states in determining goals and behavior in such tasks. By providing a task in which sections were identical except for the necessity for inferring latent causes, our results highlight this difference.

Whether the superior performance of the OFC at representing the hidden or latent information represents a true difference between this area and HPC is less clear from our data, and we would view that as unlikely. Both areas have been clearly and repeatedly implicated in both representing and using such information [26–29, 42–44], so the important question to us seems to be not which one does it better or more strongly but rather how they cooperate in this general representational function and under what circumstances one or the other is dominant. In this regard, the current results provide a rare chance to directly compare when and how strongly hidden information is represented in the two areas under the same specific conditions. Additional work will be required to suss out what about the current task drives the observed dichotomy, but at least two possibilities spring to mind. One, discussed earlier, is that it may reflect the relatively large amount of training given to these rats. The HPC may play a different role in supporting the representation of latent information earlier in learning, which then is minimized or unnecessary in advanced stages of training. A second possibility is that the representation of latent information in the HPC was rendered less prominent by the position of this information near the end of the sequences. Perhaps if there was a need to represent this information at the start of the sequences, then it would be better represented in HPC.

STAR★METHODS

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

- KEY RESOURCES TABLE
- LEAD CONTACT AND MATERIALS AVAILABILITY
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Behavioral Task
  - Surgical Procedures
  - Single-Unit Recording
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Task Events and Peri-event Spike Train Analysis
  - Classification Analyses
  - Hierarchical Clustering Analyses
- DATA AND CODE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j.cub.2019.08.040.

ACKNOWLEDGMENTS

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.

AUTHOR CONTRIBUTIONS

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

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: May 18, 2019
Revised: July 31, 2019
Accepted: August 16, 2019
Published: October 3, 2019

REFERENCES

1. Tolman, E.C. (1948). Cognitive maps in rats and men. Psychol. Rev. *55*, 189–208.

2. O'Keefe, J., and Nadel, L. (1978). The Hippocampus as a Cognitive Map (Oxford University Press).
3. Wilson, R.C., Takahashi, Y.K., Schoenbaum, G., and Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. Neuron *81*, 267–279.
4. Wikenheiser, A.M., and Schoenbaum, G. (2016). Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex. Nat. Rev. Neurosci. *17*, 513–523.
5. Behrens, T.E.J., Muller, T.H., Whittington, J.C.R., Mark, S., Baram, A.B., Stachenfeld, K.L., and Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. Neuron *100*, 490–509.
6. Constantinescu, A.O., O'Reilly, J.X., and Behrens, T.E.J. (2016). Organizing conceptual knowledge in humans with a gridlike code. Science *352*, 1464–1468.
7. Schuck, N.W., Cai, M.B., Wilson, R.C., and Niv, Y. (2016). Human orbitofrontal cortex represents a cognitive map of state space. Neuron *91*, 1402–1412.
8. Zhou, J., Gardner, M.P.H., Stalnaker, T.A., Ramus, S.J., Wikenheiser, A.M., Niv, Y., and Schoenbaum, G. (2019). Rat orbitofrontal ensemble activity contains multiplexed but dissociable representations of value and task structure in an odor sequence task. Curr. Biol. *29*, 897–907.e3.
9. Theves, S., Fernandez, G., and Doeller, C.F. (2019). The hippocampus encodes distances in multidimensional feature space. Curr. Biol. *29*, 1226–1231.e3.
10. McKenzie, S., Frank, A.J., Kinsky, N.R., Porter, B., Rivière, P.D., and Eichenbaum, H. (2014). Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. Neuron *83*, 202–215.
11. Farovik, A., Place, R.J., McKenzie, S., Porter, B., Munro, C.E., and Eichenbaum, H. (2015). Orbitofrontal cortex encodes memories within value-based schemas and represents contexts that guide memory retrieval. J. Neurosci. *35*, 8333–8344.
12. O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res. *34*, 171–175.
13. Stark, C.E.L., and Okado, Y. (2003). Making memories without trying: medial temporal lobe activity associated with incidental memory formation during recognition. J. Neurosci. *23*, 6748–6753.
14. Kumaran, D., and Maguire, E.A. (2006). The dynamics of hippocampal activation during encoding of overlapping sequences. Neuron *49*, 617–629.
15. Moser, E.I., Kropff, E., and Moser, M.B. (2008). Place cells, grid cells, and the brain's spatial representation system. Annu. Rev. Neurosci. *31*, 69–89.
16. Zhou, W., Hohmann, A.G., and Crystal, J.D. (2012). Rats answer an unexpected question after incidental encoding. Curr. Biol. *22*, 1149–1153.
17. Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. Nat. Rev. Neurosci. *15*, 732–744.
18. Hsieh, L.T., Gruber, M.J., Jenkins, L.J., and Ranganath, C. (2014). Hippocampal activity patterns carry information about objects in temporal context. Neuron *81*, 1165–1178.
19. Eichenbaum, H. (2017). On the integration of space, time, and memory. Neuron *95*, 1007–1018.
20. Schoenbaum, G., Chiba, A.A., and Gallagher, M. (1998). Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. Nat. Neurosci. *1*, 155–159.
21. Rudebeck, P.H., and Murray, E.A. (2014). The orbitofrontal oracle: cortical mechanisms for the prediction and evaluation of specific behavioral outcomes. Neuron *84*, 1143–1156.
22. Mainen, Z.F., and Kepecs, A. (2009). Neural representation of behavioral outcomes in the orbitofrontal cortex. Curr. Opin. Neurobiol. *19*, 84–91.
23. Stalnaker, T.A., Cooch, N.K., and Schoenbaum, G. (2015). What the orbitofrontal cortex does not do. Nat. Neurosci. *18*, 620–627.
24. Lisman, J., Buzsáki, G., Eichenbaum, H., Nadel, L., Ranganath, C., and Redish, A.D. (2017). Viewpoints: how the hippocampus contributes to memory, navigation and cognition. Nat. Neurosci. *20*, 1434–1447.
25. Rudebeck, P.H., and Rich, E.L. (2018). Orbitofrontal cortex. Curr. Biol. *28*, R1083–R1088.
26. Battaglia, F.P., Sutherland, G.R., and McNaughton, B.L. (2004). Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. J. Neurosci. *24*, 4541–4550.
27. Ferbinteanu, J., and Shapiro, M.L. (2003). Prospective and retrospective memory coding in the hippocampus. Neuron *40*, 1227–1239.
28. Frank, L.M., Brown, E.N., and Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. Neuron *27*, 169–178.
29. Kennedy, P.J., and Shapiro, M.L. (2004). Retrieving memories via internal context requires the hippocampus. J. Neurosci. *24*, 6979–6985.
30. Wiltgen, B.J., Zhou, M., Cai, Y., Balaji, J., Karlsson, M.G., Parivash, S.N., Li, W., and Silva, A.J. (2010). The hippocampus plays a selective role in the retrieval of detailed contextual memories. Curr. Biol. *20*, 1336–1344.
31. Lee, I., Griffin, A.L., Zilli, E.A., Eichenbaum, H., and Hasselmo, M.E. (2006). Gradual translocation of spatial correlates of neuronal firing in the hippocampus toward prospective reward locations. Neuron *51*, 639–650.
32. Hölscher, C., Jacob, W., and Mallot, H.A. (2003). Reward modulates neuronal activity in the hippocampus of the rat. Behav. Brain Res. *142*, 181–191.
33. Gauthier, J.L., and Tank, D.W. (2018). A dedicated population for reward coding in the hippocampus. Neuron *99*, 179–193.e7.
34. Baraduc, P., Duhamel, J.R., and Wirth, S. (2019). Schema cells in the macaque hippocampus. Science *363*, 635–639.
35. Siapas, A.G., Lubenov, E.V., and Wilson, M.A. (2005). Prefrontal phase locking to hippocampal theta oscillations. Neuron *46*, 141–151.
36. Sigurdsson, T., Stark, K.L., Karayiorgou, M., Gogos, J.A., and Gordon, J.A. (2010). Impaired hippocampal-prefrontal synchrony in a genetic mouse model of schizophrenia. Nature *464*, 763–767.
37. Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P.L., Gioanni, Y., Battaglia, F.P., and Wiener, S.I. (2010). Coherent theta oscillations and reorganization of spike timing in the hippocampal- prefrontal network upon learning. Neuron *66*, 921–936.
38. Backus, A.R., Schoffelen, J.M., Szebényi, S., Hanslmayr, S., and Doeller, C.F. (2016). Hippocampal-prefrontal theta oscillations support memory integration. Curr. Biol. *26*, 450–457.
39. Eichenbaum, H. (2017). Prefrontal-hippocampal interactions in episodic memory. Nat. Rev. Neurosci. *18*, 547–558.
40. Elliott Wimmer, G., and Büchel, C. (2019). Learning of distant state predictions by the orbitofrontal cortex in humans. Nat. Commun. *10*, 2554.
41. Place, R., Farovik, A., Brockmann, M., and Eichenbaum, H. (2016). Bidirectional prefrontal-hippocampal interactions support context-guided memory. Nat. Neurosci. *19*, 992–994.
42. Riceberg, J.S., and Shapiro, M.L. (2012). Reward stability determines the contribution of orbitofrontal cortex to adaptive behavior. J. Neurosci. *32*, 16402–16409.
43. Young, J.J., and Shapiro, M.L. (2009). Double dissociation and hierarchical organization of strategy switches and reversals in the rat PFC. Behav. Neurosci. *123*, 1028–1035.
44. Young, J.J., and Shapiro, M.L. (2011). Dynamic coding of goal-directed paths by orbital prefrontal cortex. J. Neurosci. *31*, 5989–6000.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Long-Evans rat	Charles River	RRID: RGD_2308852
Software and Algorithms		
MATLAB	MathWorks	RRID: SCR_001622
Offline Sorter	Plexon	RRID: SCR_000012
Neuroexplorer	Nex Technologies	RRID: SCR_001818
Other		
Nickel-chromium wire	AM Systems	Cat No. 761000
Plexon standard commutator	Plexon	http://plexon.com/products/plexonstandard-commutator
Plexon headstage cable	Plexon	http://plexon.com/products/headstage-cables
Plexon headstage	Plexon	http://plexon.com/products/headstages

LEAD CONTACT AND MATERIALS AVAILABILITY

This study did not generate any unique reagents, however 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

Male Long-Evans rats (Charles River, 175 – 200 g, ~3-month-old) were individually housed on a 12-h light/dark cycle with *ad libitum* access to food in an animal facility, which was accredited by the Association for Assessment and Accreditation of Laboratory Animal Care (AAALAC). Rats were deprived of water before they performed the task then had free access to water for 10 min in their home cages after each session. All behavioral testing was conducted at the NIDA-IRP. Animal care and experimental procedures complied with US National Institutes of Health (NIH) guidelines and were approved by National Institutes on Drug Abuse Intramural Research Program (NIDA-IRP) Animal Care and Use Committee (ACUC).

METHOD DETAILS

Behavioral Task

Rats were trained in aluminum boxes (~18" on a side) with an odor port and a fluid well for odor and sucrose solution delivery, respectively. Odor and fluid delivery were enabled by solenoids that were controlled by a custom-written C++ program. Odor port and fluid well entries were monitored through infrared beam sensors. The availability of each trial was signaled by the illumination of two house-lights in the front wall above the odor panel. Rats initiated the trial by poking into the odor port within 5 s after light onset. If a trial was initiated successfully, an odor would be delivered to the odor port after a 500-ms delay. Odor delivery continued until the rats withdrew their noses from the odor port. The trial was aborted (indicated by light off) if the rat left the odor port in less than 500 ms. Otherwise, at the time of leaving the odor port, rats had a 2 s time window to respond at the fluid well. On trials with rewards, responding at the fluid well in time led to the delivery of a sucrose solution (10% w/v; 50 μL) after a delay of 400 – 1500 ms. After the rats consumed the sucrose solution and left the well, the house lights were turned off, indicating the end of the trial and the beginning of the ITI. If the rat did not respond within the 2 s window, the house lights were turned off at the end of the 2 s period. If the rat responded at the fluid well within the 2 s time window on a non-reward trial, an exceptionally rare event in recording, the house lights were extinguished, an indication of the end of a trial, and no reward was delivered. A 4 s ITI followed correct "Go" or "No-Go" trials, and an 8 s ITI followed trials on which the rat made an error.

On each trial, one of 16 odors were delivered. The 16 odors were organized into two pairs of sequences (S1a, S1b, S2a, and S2b) as below. Reward and non-reward trials were labeled as positive (+) and negative (–), respectively, following the odor numbers.

S1a: 0+ 1- 4- 5+ 6- 7+
S1a: 2+ 3- 4- 5+ 6- 7+
S2a: 8+ 9- 12- 13+ 14- 15+
S2b: 10+ 11- 12- 13- 14+ 15+

Before training with any odors, rats were first shaped to nosepoke at the odor port and then respond at the well for a reward. Then, they were trained to discriminate an odor pair (one rewarded and one non-rewarded odor) from sequence 1a or 1b. Sessions consisted of a maximum of 480 trials. After rats reached a high criterion in performance (> 90% correct ratio), additional odor pairs were added until the rats were able to perform well in a session containing sequences 1a and 1b. After learning sequences 1a and1b, rats were trained to discriminate between odors 13/14 from sequence 2, including several reversals of the valence of the pair. After the third reversal, additional odor pairs were added from sequence 2 if the rats were able to maintain accurate performance (> 75% correct) on each trial type. Once sequence 2 had been fully introduced in this manner, the rats began sessions containing both pairs of sequences (S1a, S1b, S2a, and S2b).

In this final training phase, each sequence was repeated for 20 times to make 480 trials in total. Sequences 1a and 1b were always followed by 2a or 2b with roughly equal probability (0.55 and 0.45, respectively). Sequence 2a and 2b were always followed by 1a or 1b also with slightly more dissymmetry in probability (0.67 and 0.37, respectively). The overall sequence was repeated from start to finish in each session.

1b 2a 1a 2a 1a 2b 1b 2b 1b 2b 1b 2a 1b 2b 1a 2a 1a 2b 1a 2a 1a 2b 1b 2a 1b 2a 1b 2b 1b 2a 1b 2b 1a 2a 1a 2b 1a 2a 1a 2b 1a 2a 1b 2a 1b 2a 1a 2b 1a 2a 1a 2a 1a 2b 1b 2b 1b 2b 1b 2b 1b 2a 1a 2a 1a 2b 1b 2b 1b 2b 1b 2a 1b 2b 1a 2a 1a 2b 1a 2a

Before electrode implantation, animals were kept training on the full sequence for at least three weeks until they were able to perform well (> 75% correct) on every trial type.

Surgical Procedures

Rats were implanted with a drivable bundle of 16 stereotrodes made from nickel-chromium wires (17.8 μm in bare diameter; AM Systems, WA) that targeted the left hippocampal CA1 (AP: −3.5 mm; ML: 2 mm). Wire bundles were housed in a cannula and cut with a pair of fine surgical scissors to extend 1.5 – 2 mm beyond the end of the cannula inside the brain. The tips of wires were initially placed at 1 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 were recorded with the Plexon Multichannel Acquisition Processor (MAP) systems (Plexon, Dallas, TX). The initial signals collected by the electrodes were sequentially amplified through a headstage (20 ×), a differential preamplifier (50 ×), and a final acquisition processor (1 – 32 ×). A low pass filter (300 Hz) was used to record field potentials, while a band-pass filter (250 – 8, 000 Hz) was used to isolate spike activities. For spike recordings, a threshold was set manually for each active channel to capture unsorted spikes. Timestamps for behavioral events were sent to the Plexon system, synchronized and recorded alongside the neural activity. Spikes were sorted to identify single units offline using Offline Sorter (Plexon, Dallas, TX) with a template matching algorithm. Sorted files were exported as NeuroExplorer (Nex Technologies, Colorado Springs, CO) formatted files to extract unit and behavioral event timestamps, which were then exported as MATLAB (MathWorks, Natick, MA) files for further analysis.

After two-week recovery from the surgery, the electrodes were slowly moved down (∼10 μm each time) over ∼2 weeks to search for strong theta and sharp-wave signals in the local field potentials, indicators of the hippocampal CA1 area. The electrodes were considered to be in the hippocampal CA1 area when both the field potential signatures (theta and sharp-wave signals) and typical pyramidal neurons (wide waveforms and low baseline firing rates) appeared. Recording sessions continued daily until pyramidal activity disappeared. The electrodes were moved down by ∼40 μm every two weeks or by ∼20 μm due to the lack of pyramidal activity. At the end of testing, rats were euthanized by an overdose of isoflurane. The final positions of electrodes were marked by passing a small constant current through the wires, and the brains were processed with Nissl staining for histological examination.

QUANTIFICATION AND STATISTICAL ANALYSIS

The number of rats and the number of 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).

Task Events and Peri-event Spike Train Analysis

Each trial was divided into 6 epochs associated with different task events: "light," "poke," "odor," "unpoke," "choice," and "outcome." On rewarded trials, the time of well-entry was labeled as "choice." The "outcome" was at the time of reward delivery. On unrewarded trials, the end of the 2 s window for responding was labeled as "choice" and a time point 1.5 s after the "choice" as the "outcome." Behavioral performance was quantified as the percent of trials on which the rats responded correctly, their reaction time from the odor port to the fluid well, and the latency with which they initiated a trial after light onset. The data analysis only included sessions in which the percent correct was above 75% for every trial type. Error trials were removed for all the analyses. The spike train for each isolated single unit was aligned to the onset of each task event. 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 (σ = 50 ms) was used to smooth the spike train on each trial.

Classification Analyses

We used a linear discriminant analysis (LDA) algorithm (MATLAB function: fitcdiscr) to classify 24 trial types or states for each one of six task events. Neurons recorded from different sessions were aligned to form pseudo-ensembles. Firing rates during a period of 100 – 600 ms after each task event were used for classification of 24 states. Each trial was seen as an observation that contained firing rates from 2, 056 neurons from the HPC or 1, 078 neurons from the OFC (480 trials in total but only 360 trials were used in this study as a result of the removal of error trials; 360 trials = 24 trial types × 15 trials). The classification accuracy was assessed by leave-one-out cross-validation. Specifically, one trial from each trial type was left out for future testing, and all the other trials were used for the training. Principal component analysis (PCA) was used for feature extraction and dimension reduction in the training set. The classifier was trained on the principal components (PCs) necessary to explain 80% variance (HPC: 103.7 ± 0.5 PCs; OFC: 151.2 ± 0.4 PCs). The same PCA transformation from the training set was applied to the test set. Trial order for each neuron was shuffled within the same trial type. The trial-order shuffling was repeated for 100 times. For each time of trial-order shuffling, the leave-one-out cross-validation was repeated for 500 times to estimate the decoding accuracy. The mean decoding accuracy for each trial type as shown in the confusion matrix was the mean across 100 runs (corresponding to 100 times of trial-order shuffling). The statistical significance of the mean decoding accuracy was determined by the 95% confidence interval estimated by running the same decoding process with label-shuffled data. For the binary classification of subsequences at each time point (bin size: 100 ms) or after task event (time period: 100 - 600 ms), the processes were the same except only the first three PCs were used, and the trial-order shuffling was repeated for 200 times.

Hierarchical Clustering Analyses

The hierarchical agglomerative clustering was performed on data that was projected onto the LDA space. Each trial was organized as a vector with firing rates of neurons as dimensions. The original data was first transformed to PCA space with 80% variance retained, then transformed to LDA space supervised by 24 trial-type labels. A dissimilarity matrix was computed by measuring the Mahalanobis distance between each pair of trial-type means. Based on the dissimilarity matrix, an agglomerative hierarchical cluster tree was generated with the unweighted average distance method. The clustering results were plotted as dendrograms.

DATA AND CODE AVAILABILITY

The dataset and MATLAB scripts used in this study will be made available upon request by the lead contact, Geoffrey Schoenbaum (geoffrey.schoenbaum@nih.gov).