Online Research @ Cardiff

information services
gwasanaethau gwybodaeth

Post-learning hippocampal dynamics promote preferential retention of rewarding events

Matthias J. Gruber[1], Maureen Ritchey[1], Shao-Fang Wang[1], Manoj K. Doss[1], and Charan Ranganath[1,2]

[1] Center for Neuroscience, UC Davis, USA

[2] Department of Psychology, UC Davis, USA

* Correspondence to: Matthias J. Gruber

Email: mjgruber@ucdavis.edu

Phone: 530-757-8865

Fax: 530-757-8640

Address: Center for Neuroscience, University of California, 1544 Newton Court, Davis, CA 95618, USA

Abstract

Reward motivation is known to modulate memory encoding, and this effect depends on interactions between the substantia nigra/ ventral tegmental area complex (SN/VTA) and the hippocampus. It is unknown, however, whether these interactions influence offline neural activity in the human brain that is thought to promote memory consolidation. Here, we used functional magnetic resonance imaging (fMRI) to test the effect of reward motivation on post-learning neural dynamics and subsequent memory for objects that were learned in high- or low-reward motivation contexts. We found that post-learning increases in resting-state functional connectivity between the SN/VTA and hippocampus predicted preferential retention of objects that were learned in high-reward contexts. In addition, multivariate pattern classification revealed that hippocampal representations of high-reward contexts were preferentially reactivated during post-learning rest, and the number of hippocampal reactivations was predictive of preferential retention of items learned in high-reward contexts. These findings indicate that reward motivation alters offline post-learning dynamics between the SN/VTA and hippocampus, providing novel evidence for a potential mechanism by which reward could influence memory consolidation.

Introduction

Typically, a person will only retain memories for a small fraction of the events that s/he experiences on any given day. Although the selectivity of memory can be frustrating, it might be advantageous to prioritize retention of events that are salient or motivationally significant. Several theories suggest that dopaminergic input from the substantia nigra/ ventral tegmental area complex (SN/VTA) to the hippocampus is enhanced during learning of motivationally significant events (Düzel et al., 2009; Kahn and Shohamy, 2013; Lisman and Grace, 2005; Lisman et al., 2011; Shohamy and Adcock, 2010). Consistent with these models, functional magnetic resonance imaging (fMRI) studies in humans have shown that reward motivation enhances memory encoding and is associated with increased activity in, and interactions between, the SN/VTA and the hippocampus (Adcock et al., 2006; Murty and Adcock, 2014; Wittmann et al., 2005; Wolosin et al., 2012).

Most research on motivated memory effects have focused on memory encoding, whereas little is known about the effects of motivation on post-learning neural processes. According to some models of memory consolidation, neural processes that occur after learning can also influence retention (Carr et al. 2011; Káli and Dayan 2004; Sutherland and McNaughton, 2000). These models have received support from rodent single-unit recording studies showing that hippocampus-dependent memories may be subsequently reactivated after learning (Foster and Wilson 2006; Singer and Frank 2009).

In humans, recent fMRI studies have also reported evidence that changes in post-learning dynamics can predict subsequent memory performance. For example, studies on resting-state functional connectivity (RSFC) have shown that post-learning RSFC between the hippocampus and category-selective cortical regions predicted subsequent memory for stimuli from those categories (Schlichting and Preston, 2014; Tambini et al., 2010; Tompary et al., 2015). Recent studies have also used fMRI data to show that neocortical representations of specific items or item categories are reactivated during post-learning rest (Deuker et al., 2013;

Schlichting and Preston, 2014; Staresina et al., 2013). Although these studies suggest that neural processes that occur after learning can contribute to human memory, it is not clear how motivation influences post-learning activity in the hippocampus.

If motivation plays a role in prioritizing the retention of episodic memories, then it is possible that exposure to rewards would affect post-learning neural dynamics thought to play a role in memory consolidation. This idea is consistent with findings showing that reward can increase offline reactivation of hippocampal activity in rats (Lansink et al., 2009; Singer and Frank, 2009). However, it is not clear whether exposure to rewarding contexts biases post-learning interactions between the SN/VTA and hippocampus, nor whether it leads to enhanced hippocampal reactivation of rewarding events in humans. Furthermore, even if such post-learning changes occur, it is not clear whether they enhance retention of all recent events or whether they preferentially enhance retention of information learned in rewarding contexts.

Here, we used a combination of resting-state and task-based fMRI to investigate how motivational salience affects memory and post-learning dynamics in the SN/VTA and hippocampus in support of prioritizing memories of salient information. Participants were scanned during a pre-learning rest phase, a reward-motivated learning phase, and a post-learning rest phase (**Figure 1A**). During learning, participants incidentally encoded objects in one of four contexts, with contexts operationalized as 30 s mini-blocks associated with a specific incidental encoding task and a semantically matching background scene. Each context was consistently associated with either high or low immediate rewards, and the reward value was presented at the beginning of each mini-block (**Figure 1B**). After scanning, memory for objects and associated task contexts was tested (**Figure 1C**). This design enabled us to (i) relate offline post-learning dynamics to memory for motivationally significant contexts and (ii) to test whether potential post-learning dynamics contribute to later memory in a similar or different way as processes during learning. Given the established role of the hippocampus in binding item and context information (e.g. Diana et al. 2007; Eichenbaum et

4

al. 2007), and of the SN/VTA in signaling reward and salience (Haber and Knutson 2010; Shohamy and Adcock 2010), we focused our analyses on these two main regions of interest (ROIs) (for details, see Supplemental Methods and **Figure S1**).

Our experiment addressed two central questions: First, we tested whether reward-motivated learning enhances SN/VTA-hippocampus communication during post-learning rest. We hypothesized that functional connectivity between these regions would be increased following reward-motivated learning, and that the magnitude of this effect might be predictive of later memory benefits. Second, we tested whether motivational salience influences reactivation of hippocampal context representations during post-learning rest (cf. Lansink et al., 2009; Singer and Frank, 2009). We hypothesized that the hippocampus would show preferential reactivation of high-reward contexts, relative to low-reward contexts during post-learning rest, and that the number of reactivations might predict subsequent memory performance.

Results

Rewarding contexts benefit memory for object-context associations

During the study phase, correct encoding judgments were rewarded. Accuracy on these semantic encoding decisions was very high and did not differ significantly between high- and low-reward conditions (91.4% SE=±1.2 vs. 89.7% SE=±1.5, respectively; $t_{(18)}$=1.50, p=0.150), but reaction times for correct decisions were faster in high- than in low-reward contexts (910ms SE=±33 vs. 945ms SE=±33, respectively; $t_{(18)}$=-3.40, p=0.003).

Importantly, performance on the subsequent memory tests was sensitive to reward contexts learned during encoding. Memory for object-context associations (i.e. memory for which of the four contexts had been previously associated with each object) was enhanced for objects studied in high-reward contexts compared to objects studied in low-reward contexts

(i.e. HR>LR object-context memory advantage) (26.2% SE=±3.9 vs. 22.6% SE=±3.9; $t_{(18)}$=2.16, p=0.022; **Figure 1D**). Consistent with these results, estimates of recollection-based recognition were higher for objects studied in high-reward contexts compared to objects studied in low-reward contexts (i.e. HR>LR recollection advantage) (39.9% SE=±4.2 vs. 32.8% SE=±4.1, respectively; $t_{(18)}$=5.06, p<0.001). In contrast, familiarity estimates were not significantly improved for objects in high- compared to low-reward contexts (54.8% SE=±3.8 vs. 52.9% SE=±3.7, respectively; $t_{(18)}$=1.06, p=0.151). The behavioral findings demonstrate a reliable memory advantage for items learned in rewarding contexts.

Hippocampal encoding-related activation is independent of reward context

Because object-context memory has been shown to strongly depend on hippocampal encoding-related processes (e.g. Diana et al. 2007; Eichenbaum et al. 2007), our first fMRI analyses investigated whether reward value affected the relationship between hippocampal activity during object encoding and later performance on the object-context memory task. To have a sensitive hippocampal ROI, we used the NeuroSynth tool (Yarkoni et al., 2011) to identify hippocampal voxels that are sensitive to reward motivation (see Supplemental Methods and **Figure S1**). Activity in this independently created bilateral hippocampus ROI was enhanced during encoding of objects that were later associated with correct object-context memory, relative to objects that were not associated with accurate object-context memory (main effect SME: $F_{(1,18)}$=6.02; p=0.025), but this effect did not significantly vary as a function of reward context (SME x Reward interaction: $F_{(1,18)}$=2.46; p=0.134). Additional whole-brain, voxel-based SME analyses did also not reveal any significant SME x Reward interactions, but only showed main effects of SME (independent of reward) (**Table S1**). In addition, results on cue-elicited activity across the whole brain and in our ROIs were in line with previous findings on reward anticipation (**Table S1**).

Hippocampal activity patterns during encoding reflect reward context

We next tested whether activity patterns in the hippocampal ROI differentiated between the different contexts during encoding. To address this question, we used multi-voxel pattern analyses (MVPA) to determine whether hippocampal activity patterns carried information about unique scene contexts or more generally about contexts associated with high or low reward (i.e. independent of the precise scene context).

First, we trained a 4-way classifier to distinguish between the four specific scene contexts. We found that overall classifier performance was not significantly above chance (25.4% SE=±0.3; $t_{(18)}$=1.60, *p*=0.063). Closer inspection of classifier accuracy for each of the four contexts, however, revealed that hippocampal ROI activity patterns successfully distinguished between contexts associated with high and low rewards, but they were insensitive to differences between specific scene contexts. That is, the 4-way classifier made significantly more errors within same-reward contexts (i.e. contexts that shared the same reward but differed in scene context) than for other-reward contexts (i.e. contexts that differed in both reward and scene context) (25.8% SE=±0.3 vs. 24.4% SE=±0.2, respectively; $t_{(18)}$=3.20, *p*=0.002). This result suggests that activity patterns in the hippocampal ROI were sensitive to the overall reward value associated with the contexts, and insensitive to the differences between contexts that shared the same reward value. Consistent with this impression, a binary classification analysis revealed above chance (51.2% SE=±0.5; $t_{(18)}$=2.58, *p*=0.010) accuracy in classification of high- vs. low-reward contexts (for additional analyses based on voxel patterns in other ROIs and across the whole brain, see **Table S2**).

Post-learning SN/VTA-hippocampus functional connectivity changes predict later preferential memories for highly motivational information

In the next analyses, we investigated whether reward-motivated learning enhanced interactions between SN/VTA and hippocampus during post-learning rest. To address this

question, we quantified resting-state functional connectivity (RSFC) between the bilateral SN/VTA and hippocampus ROIs, separately for the pre- and the post-learning rest periods (**Figure 2A**). RSFC between SN/VTA and hippocampus numerically increased from pre- to post-learning rest (pre-learning rest: *r*=0.208 SE=±0.020; post-learning rest: *r*=0.230 SE=±0.037), but the increase was not statistically significant ($t_{(18)}$=0.60, *p*=0.280), potentially due to high inter-subject variability. In line with recent studies that have shown that individual differences in post-learning RSFC are predictive of memory performance (Tambini et al., 2010; Tompary et al., 2015), we hypothesized that inter-subject variability in RSFC changes reflect individual differences in the extent to which reward value influenced memory performance. We therefore computed a *Pearson's* correlation indexing the relationship between SN/VTA-hippocampal RSFC changes (i.e. post-learning > pre-learning RSFC) and HR>LR object-context memory advantages. The analysis revealed that SN/VTA-hippocampal RSFC changes were significantly associated with HR>LR object-context memory advantages (*r*=0.465, *p*=0.023, **Figure 2B**) (for other ROI analyses and exploratory whole-brain, voxel-based analyses, see **Table S3**).

Increased hippocampal reactivation of high-reward contexts during post-learning rest predict later preferential memory for high-reward information

Our next analyses focused on whether reward value influences the likelihood of offline reactivation of learning contexts. If so, we might expect to see preferential reactivation of hippocampal context representations that were associated with high-reward value, relative to those associated with low-reward value. To test this prediction, we trained a pattern classifier to differentiate between the hippocampal ROI activity patterns associated with high- and low-reward contexts during the encoding phase, and we "tested" the classifier on scans acquired during the post-learning rest period. Across all time points during the rest period, we counted how often the classifier associated the resting-state activity pattern with the high reward

context pattern (**Figure 3A**). If motivational salience influenced reactivation of recent learning events, we would expect hippocampal activity patterns to be associated with high-reward contexts more often than what would be expected by chance (i.e., 50%, or 244 time points). Consistent with our prediction, the number of time points classified as high-reward contexts was significantly higher than what would be expected by chance during the post-learning rest period (248.2 SE=±1.37, $t_{(18)}$=3.07, p=0.003; **Figure 3B**). Importantly, during the pre-learning rest period, the number of instances in which the classifier output predicted high-reward contexts was not greater than chance (241.9 SE=±1.70, $t_{(18)}$=-1.20, p>0.05; **Figure 3B**). Further analyses showed that the number of classifications of high reward contexts was significantly higher during post- compared to pre-learning rest ($t_{(18)}$=3.15, p=0.003). Because enhanced classification of high-reward contexts was specific to the post-learning rest period, the results are consistent with the idea that high reward contexts were preferentially reactivated following reward-motivated learning.

We next explored whether preferential reactivation of high-reward contexts could account for preferential memory for objects that were learned in these contexts. To test the relationship between hippocampal reactivation and memory, we computed the change between the numbers of classifications of high-reward contexts from pre- to post-learning rest. This "high-reward reactivation index" captures the effect of high-reward context reactivation, while controlling for overall classifier prediction biases. The magnitude of the hippocampal high-reward reactivation index was significantly predictive of the HR>LR object-context memory advantage (*Pearson's r*=0.495, *p*=0.016; **Figure 3C**), such that participants who showed more reactivation of high-reward contexts showed a larger memory bias for items learned in these contexts.

As reported in the Supplemental Results, the findings remained unchanged when classifier output was thresholded to identify high- and low-reward reactivation (Supplemental

Results and **Figure S2**). Reactivation analyses on voxel patterns in other ROIs and across the whole brain are summarized in **Table S2**.

SN/VTA-hippocampal RSFC and hippocampal reactivations explain reward-related memory advantage independently of encoding activity and overall memory

The analyses described above demonstrate that both SN/VTA-hippocampal RSFC changes and hippocampal reactivation events were associated with the later reward-related memory advantages for object-context associations. We next tested whether these post-learning dynamics could account for subsequent memory performance over and above what could be accounted for by encoding activity. To address this question, we performed hierarchical multiple regression analysis in which individual differences in the HR>LR object-context memory effect were modeled in the following steps. In the first step, we tested the extent to which these individual differences could be accounted for by hippocampal encoding-related activity alone (i.e. the Reward x SME interaction; SME: correct object-context memory vs. incorrect object-context memory). This *encoding-only model* only explained 8.9% of the variance in contributing to the later HR>LR object-context memory advantage ($F_{(1,17)}$=1.67; p=0.214). In the next step, we added individual pre-to-post learning changes in SN/VTA-hippocampal RSFC to the model. This *encoding + RSFC model* explained 34.8% of the variance in contributing to the HR>LR object-context memory advantage, which is significantly higher than what was explained by the *encoding-only model* ($F_{(1,16)}$=6.35; p=0.023). Finally, we added individual differences in the high-reward reactivation index to the model. This *encoding + RSFC + reactivations model* explained 52.0% of the variance, which is significantly higher than what was explained by the *encoding + RSFC model* ($F_{(1,15)}$=5.37; p=0.035). Examination of this three-predictor model showed that SN/VTA-hippocampal RSFC changes and hippocampal reactivations significantly predicted the HR>LR object-context

memory advantage ($t_{(15)}$=2.81 and $t_{(15)}$=2.32, respectively; p's≤0.018), but the encoding data did not significantly contribute to the overall model ($t_{(15)}$=1.19, p=0.127).

Furthermore, we asked whether the relationship between post-learning effects and later HR>LR memory advantage could be explained by overall memory performance. In a further model, we controlled for overall object-context memory by adding a fourth *overall memory* predictor to the previous three-predictor model. The overall fit of this *encoding + RSFC + reactivations + overall memory model* did not significantly differ from the *encoding + RSFC + reactivations model* ($F_{(1,14)}$=0.25; p=0.624). Most importantly, both the SN/VTA-hippocampal RSFC changes and hippocampal reactivations still remained significant predictors of the HR>LR object-context memory advantage ($t_{(14)}$=2.60 and $t_{(14)}$=2.27, respectively; p's≤0.020). These findings demonstrate that post-learning RSFC changes and reactivation of high-reward contexts contributed to the HR>LR memory advantage even after controlling for encoding activity and overall memory performance.

Relationship between post-learning dynamics and reward-related memory advantage was limited to the hippocampus

Given that the hippocampal ROI exhibited preferential reactivation of high-reward contexts, we conducted exploratory analyses to see if the same effect could be observed in other brain regions. For example, one might expect to see preferential reactivation in the nucleus accumbens (NAcc), a region that shows high intrinsic functional connectivity with the SN/VTA and the hippocampus (Kahn and Shohamy, 2013). In addition, it is conceivable that other reward-sensitive areas such as the ventromedial prefrontal cortex (vmPFC) or even visual areas such as V1 (which should show strong dissociations between encoding contexts) could show enhanced reactivations for high-reward contexts. However, exploratory analyses in these ROIs did not reveal any significant effects of reward context on reactivation (encoding classification and reactivation analyses on voxel patterns in these ROIs are summarized in

Table S2). In addition, using an approach from a previous reactivation fMRI study by Deuker et al. (2013) on whole-brain voxel patterns did also not reveal significant high-reward context reactivations (**Table S2**). Importantly, individual differences in high-reward reactivations in these regions did not show any significant correlations with the later HR>LR object-context memory advantage. Thus preferential reactivation of high-reward contexts that predicted the later memory advantage was not seen in widespread areas across the brain. These findings demonstrate the selectivity of high-reward context reactivation predicting later memory advantage in the hippocampus.

We next investigated whether these additional ROIs (i.e. NAcc, vmPFC, and V1) showed increased RSFC with the hippocampus or SN/VTA ROI after reward-motivated learning. However, we did not find any significant RSFC changes in these additional ROIs and individual differences in RSFC changes did also not predict the later HR>LR object-context memory advantage **(Table S3)**. In addition, exploratory whole-brain, voxel-based analyses comparing RSFC (with the hippocampal ROI as a seed) between pre- and post-learning rest periods did also not show any significant RSFC changes surviving our cluster-corrected threshold (for exploratory analyses using a liberal statistical threshold, see **Table S3**).

Discussion

Results from this study demonstrate that: (1) interactions between the SN/VTA and the hippocampus after learning and (2) hippocampal reactivation of high-reward contexts are associated with preferential memory for items learned in these contexts. Importantly, the relationship between hippocampal *post*-learning dynamics and later preferential memory was limited to the hippocampus and it could not be explained by overall memory performance or encoding-related activity. These findings provide strong support for the idea that post-learning neural dynamics in the hippocampus can prioritize the retention of salient events.

Hippocampal reactivation of salient contexts during post-learning rest

Single unit recording studies in rodents have provided ample evidence for post-learning reactivation of hippocampal spatial context representations. Specifically, sequences of place cells that are activated during spatial navigation and learning are spontaneously replayed during sleep (e.g., Ji and Wilson, 2007) and wakeful rest (e.g., Foster and Wilson 2006; Singer and Frank 2009). It is notable that place cell sequences associated with acquisition of food rewards (e.g. the sequence representing the arm of a maze directly leading to a food reward) are preferentially reactivated, and researchers have speculated that this could be a mechanism for enhanced consolidation of salient events (Lansink et al., 2009; Singer and Frank, 2009).

Previous studies in humans have demonstrated that post-learning reactivation events can be observed via multivariate pattern analyses on fMRI data (Deuker et al., 2013; Schlichting and Preston, 2014; Staresina et al., 2013). These studies have reported reactivation of specific *items or item categories* in entorhinal, retrosplenial, (Staresina et al. (2013), and occipitotemporal cortex (Deuker et al., 2013; Schlichting and Preston, 2014), but reactivation has not been observed in the hippocampus. In contrast, single-unit recording studies have only demonstrated hippocampal replay of spatial *context* information, with the strongest effects for place cell sequences associated with rewards. Motivated by these findings, we therefore designed this experiment such that reward values were associated with specific contexts in which objects were previously encountered, so that we could examine reactivation of these contexts during post-learning rest. We found that hippocampal representations of high-reward contexts were preferentially reactivated, and that the magnitude of this effect strongly predicted subsequent memory for objects learned in high-reward contexts. These findings closely parallel results showing that preferential reactivation

of place cell sequences associated with rewards is related to preferential memory for reactivated contexts (Lansink et al., 2009; McNamara et al., 2014; Singer and Frank, 2009).

In rodents, replay events are compressed in time and take place during hippocampal sharp-wave ripples (SWRs) that last approximately 50 ms (for a review, see Girardeau and Zugaro, 2011). Given the sluggish time course of the blood oxygenation level dependent (BOLD) response measured with fMRI, it is reasonable to question whether BOLD activity patterns could reflect such brief neural events. One important factor to consider is that SWRs are sometimes isolated events, but they can also occur in clusters (Papatheodoropoulos, 2010; for a review, see Buzsáki, 2015). In direct intracranial recordings from depth electrodes implanted in the human hippocampus, researchers have found evidence for SWR events during sleep (Staresina et al., 2015) and awake rest (Axmacher et al., 2008). Consistent with the evidence in rodents, SWR events in humans during awake rest occurred in clusters with gaps of 100-200 ms between SWRs (Axmacher et al., 2008). In light of this evidence, it is conceivable that a train of hippocampal SWR events could drive hemodynamic activity patterns that could be detected with BOLD fMRI. Indeed, hippocampal SWR events in monkeys have been shown to trigger BOLD responses in the hippocampus and neocortex (Logothetis et al., 2012; Ramirez-Villegas et al., 2015). Taken together, the evidence for clustered SWRs during awake rest and the preferential replay of reward-associated sequences provide reason to believe that clustered SWRs could underlie the observed hippocampal reactivation events observed in the present study.

Hippocampal representations during learning

During learning, we found that hippocampal activity patterns could not be used to classify the four specific encoding contexts, but they reliably distinguished between low- and high-reward contexts. The finding is surprising, because it has been shown that hippocampal patterns encode specific context information (Chadwick et al., 2011; Libby et al., 2014;

14

Ritchey et al., 2015). It is possible that the predominance of reward value in the classification analysis was a consequence of the fact that the hippocampal ROI included only reward-selective hippocampal voxels (independently derived in a meta-analysis; cf. Gruber et al., 2014). That said, it is notable that studies in rats and humans have indicated that rewards play a significant role in hippocampal context representations (McKenzie et al., 2014; Wolosin et al., 2013). For instance, using recordings from the rat hippocampus, McKenzie et al. (2014) showed that population-level response patterns were more similar during events that shared the same reward value (even if they were associated with different spatial locations) than they were between events that did not share the same reward value (even if they appeared in the same location). Consistent with these results, human hippocampal activity patterns have been shown to be more similar for reward cues with the same reward-value compared to cues with a different reward-value (Wolosin et al., 2013). Building on these findings, future studies could test whether hippocampal representations of spatial or non-spatial contexts are altered after these contexts are associated with different motivational factors (e.g. reward or shock).

SN/VTA and hippocampus interactions during post-learning rest

The second major finding from this study was that SN/VTA-hippocampal resting-state functional connectivity (RSFC) increases from pre- to post-learning rest periods were predictive of preferential memory for objects from high-reward contexts. This finding builds on results from previous studies, which found that post-learning RSFC between the hippocampus and category-selective cortical regions (i.e., fusiform face area, parahippocampal place area, and lateral occipital complex) predicted subsequent memory for stimuli from those categories (Schlichting and Preston, 2014; Tambini et al., 2010). In addition, the mere continuation of task-related hippocampal activity into a post-learning period has also been associated with later overall memory (Peigneux et al., 2006; Tambini and Davachi, 2013). Finally, Tompary et al. (2015) showed that the increase of SN/VTA-hippocampal RSFC from pre- to post-learning

periods predicted later overall memory for object-object associations. The findings of Tompary and colleagues are particularly relevant because they demonstrate that, in the absence of an explicit reward manipulation, changes in RSFC between SN/VTA and hippocampus are related to overall memory. Critically, our findings show that when events differ in reward value, these RSFC changes are not associated with overall memory performance, and instead they are associated with preferential memory for items learned in high-reward contexts. Accordingly, the present results are consistent with the idea that post-learning interactions between the SN/VTA and hippocampus prioritize retention of motivationally salient information over other less salient information.

Our findings align well with single-unit studies in rodents that show a relationship between VTA and hippocampal neurons during post-learning rest periods (Gomperts et al., 2015; McNamara et al., 2014; Valdés et al., 2015). For example, Gomperts et al. (2015) showed that spiking activity of reward-responsive VTA neurons coincided with hippocampal SWRs. Importantly, this VTA-hippocampal relationship was specific to awake rest periods, and not evident during sleep suggesting again a predominant role of SWRs during awake rest periods. In addition, a recent study by McNamara et al. (2014) used optogenetics to manipulate dopaminergic fibers from VTA neurons leading into the hippocampal CA1 region and showed that optogenetic manipulation *during* learning enhanced *post*-learning reactivation of novel environments. In a separate set of experiments, the authors showed that such optogenetic manipulation during learning stabilized performance during a memory probe test. Taken together, these findings provide a foundation for understanding at the neural level how post-learning interactions between the SN/VTA and hippocampus (indirectly measured here through BOLD fMRI) could prioritize the consolidation of memories for rewarding events.

Selectivity of post-learning dynamics in the hippocampus

We ran several exploratory analyses to determine whether reward-motivated learning affected post-learning dynamics in regions other than the hippocampus. We explored reactivation in other dopaminergic ROIs such as the nucleus accumbens or ventromedial prefrontal cortex and in a non-dopaminergic, visual cortex ROI (i.e. V1), but these analyses did not reveal evidence for preferential reactivation of high-reward contexts. Thus, our results do not reflect an artifact or a global tendency of the brain to reactivate representations of rewarding events. It is possible, however, that other reward-motivated learning paradigms could reveal evidence of post-learning reactivation biases in other regions. For example, recordings from the ventral striatum in rodents indicate that reactivation events can be observed in the nucleus accumbens during the performance of a spatial delayed alternation task (Pennartz et al., 2004) and during active foraging (Lansink et al., 2008; 2009). Pennartz et al. (2004) speculated that the accumbens may serve as a "limbic-motor interface", and if this view is correct, then one might expect the accumbens to exhibit preferential reactivation of action sequences that lead to reward. Further research is needed to investigate how other cortical and subcortical areas contribute to post-learning dynamics and how they complement hippocampal reactivations.

Future directions

Post-learning hippocampal reactivations and post-learning SN/VTA-hippocampal RSFC changes independently contributed to preferential memory for motivationally salient information. Importantly, the relationship between post-learning dynamics and later preferential object-context memory could not be explained by learning-related hippocampal activity. The findings highlight the critical role of post-learning dynamics for later retention and provide strong evidence against the idea that post-learning might reflect a mere by-product of encoding-related processes. Post-learning contributions to later memory that are independent

of encoding processes have also been shown in recent experiments investigating post-learning dynamics (e.g., Staresina et al., 2013; Tompary et al, 2015). In line with the idea that dopamine affects late LTP (Lisman et al., 2011), McNamara et al. (2014) demonstrated that dopaminergic stimulation via optogenetics during hippocampus-dependent learning led specifically to increased reactivation of hippocampal firing patterns during a post-learning period, but dopaminergic stimulation did not affect learning itself. Together, these findings suggest a critical role of post-learning dynamics for later preferential retention. Recent accounts propose that salience (e.g. via reward or emotion) might be the key determinant that leads to the prioritization of some information over other less relevant information in support of enhanced memory consolidation (Atherton et al., 2015; Mather et al., 2015). In line with this, our findings suggest hippocampal reactivations do not necessarily boost memories in general but bias memory retention towards salient information. Future studies would need to address how such prioritization during post-learning periods differs from other forms of prioritization or strengthening that could occur during learning itself.

The present study adds to a growing body of evidence to suggest that spontaneous brain activity is influenced by prior learning experiences (Deuker et al., 2013; Peigneux et al., 2006; Schlichting and Preston, 2014; Staresina et al., 2013; Tabmbini and Davachi, 2013; Tambini et al., 2010; Tompary et al., 2015). Crucially, our study extends those findings showing that a manipulation such as reward, which is thought to enhance consolidation processes (Murayama and Kitagami, 2014; Murayama and Kuhbandner, 2011), affects post-learning hippocampal activity, and that those hippocampal changes, in turn, preferentially support retention of salient information. Future studies could build on the present findings by investigating how direct manipulations of brain activity (via brain stimulation or pharmacological interventions) during post-learning rest affect later memory. Initial evidence suggests that manipulations of post-learning activity can "rescue" memories for non-salient

events (Feld et al., 2014; Oudiette et al, 2013). For instance, Feld et al. (2014) showed that administration of a dopamine agonist during sleep boosted later memory for *low*-reward information up to the level of high-reward information. In addition, Oudiette et al. (2013) showed that covert reactivation of information also "rescued" memory specifically for *low*-reward information. Another fruitful approach for future studies might be to relate physiological markers of sleep or consolidation to post-learning RSFC and reactivations. Recently, Igloi et al. (2015) showed that *learning*-related functional connectivity between the hippocampus and the caudate correlated with the number of cortical spindles (recorded via scalp EEG) during a post-learning nap.

Conclusion

In summary, the results from this study suggest that post-learning neural dynamics might be a mechanism by which the brain prioritizes retention of events that lead to rewarding outcomes. Given that dopaminergic functions change over the course of healthy aging and that these functions can be impaired by neurological (e.g. Parkinson's disease) and psychiatric (e.g. schizophrenia) disorders, it would be fruitful to investigate post-learning SN/VTA-hippocampal interactions in these populations. The findings reported here provide a potentially promising starting point for clarifying how changes in dopaminergic function affect learning and consolidation processes that determine whether a past event will ultimately be retained.

Experimental Procedures

The details about the participants, stimulus material, behavioral analyses, fMRI acquisition and preprocessing, regions-of-interest approach, and whole-brain MVPA analyses are presented in the Supplemental Experimental Procedures.

Task Procedures

There were four critical stages in the paradigm: (1) a pre-learning rest phase, (2) a reward-motivated incidental learning phase, (3) a post-learning rest phase, and (4) a surprise recognition memory test (**Figure 1A**). The delay between the study and memory test phase was on average 36 minutes (range: 29-43 minutes) so that the time between the initial encoding of a trial and its test was on average 1h 23min.

Pre- and post-learning rest phases. During the rest phases (10 min each), participants lay in the MRI scanner and the room was kept as dark as possible. Participants were instructed to stay relaxed and to keep their eyes open and to look straight ahead (which was monitored with a camera).

Reward-motivated incidental learning phase. In the incidental learning phase, participants made semantic judgments during highly or lowly rewarded mini-blocks (**Figure 1B**). During the whole mini-block (30 s), a scene was presented as background that semantically matched the encoding judgment within a mini-block. Each mini-block started with high or low reward cue ($2.00 or $0.02) and a particular encoding judgment. Participants then performed the particular encoding judgment on four consecutively presented objects. If participants made a correct judgment, a reward was given with a probability of 80% for each correct encoding judgment. This ensured a level of uncertainty about the judgment accuracy in an attempt to keep attention levels high throughout the encoding session. At the end of a mini-block, the accumulated reward within a mini-block was presented.

We used the following four encoding judgments: "Does this item weigh more than a basketball?", "Would this item float?", "Is this item bigger than a laptop screen?", "Can this item be juggled?". The background that was presented throughout a whole mini-block (i.e. 30 s) was one of the following scenes, semantically matched to the four encoding judgments, respectively: basketball court, swimming pool, open space office, and circus tent (for exemplary trials, see **Figure 1B**). In half of the participants, the first two judgments were always associated with high reward and the latter two judgments with low reward, and vice versa for the other half of participants. Participants were instructed in more detail about the criteria for a yes/no response for each question. Before entering the MRI scanner, participants practiced the encoding phase. They were given feedback about their practice responses in order to ensure that all participants had the chance from the start of the encoding phase to get the maximum reward. Importantly, the practice phase mimicked the encoding phase during scanning with the critical exception that participants were not told until the beginning of the scanned encoding phase which judgments would be associated with high and low reward. Participants were also told about the 80% probability to obtain a reward during the scanned encoding phase. Therefore, during the pre-encoding rest phase, participants knew about the nature of the encoding phase but were unaware of the associated reward values for the particular judgments. 'Yes' responses were given with the right index finger and 'no' responses with the right middle finger on an MRI compatible response box. The study phase was divided into ten scanning runs (4.5 min each). In each run, two mini-blocks for each of the four encoding tasks were presented, mini-blocks alternated between high and low reward conditions and the order of the four encoding judgments was fixed within a run.

Reward cues were presented in the middle of the screen on a white square ("$ 2.00" written in green, "$ 0.02" written in black). Underneath the reward cue, the encoding judgment for the particular mini-block was presented. Both cue and encoding judgment were presented

for 2 s. Objects and feedback (e.g. "You just earned $6!") were presented within a white square in the middle of the screen for a duration of 2 s. The ITIs within a mini-block (following the cue, the four objects, and the feedback) were jittered (mean: 3 s; range: 1-9 s) and each mini-block was fixed at 30 s. The interval between mini-blocks was also fixed to 2 s displaying a cross hair in the middle of the screen with a grey background.

Memory test. On average 36 min after the end of the encoding phase, memory was tested for all encoded objects and the context in which they were presented during encoding. The memory test was self-paced and the encoded objects were intermixed with 120 new objects in a random order. A trial started with the presentation of an object (within a white rectangle) on a grey background (**Figure 1C**). Underneath the object, the possible responses to the first judgment were shown: R, 5, 4, 3, 2, and 1 (representing the following responses: R = "confidently remembered with specific details", 5 = "confidently familiar without any details", 4 = "unconfident familiar", 3 = "guessing", 2 = "unconfident new", 1 = "confident new"). After a response was given, the possible responses to a second memory judgment were shown. Underneath the object, the scenes (i.e. background scenes during encoding) representative of the four encoding judgments were shown. Participants were asked to indicate in which encoding context the object was studied. In addition to responding to one of the four encoding contexts, participants could also indicate with a fifth response key if they were "not sure" about the object-context association. In addition, participants were instructed to also press this fifth response key for the second memory judgment, if they gave a "new" response for the first memory judgment.

FMRI resting-state functional connectivity (RSFC) analyses

To address how communication between the SN/VTA and hippocampus changed from pre- to post-learning rest period, we investigated RSFC between our two ROIs. We used in-house scripts in MATLAB (The Math Works, Inc., USA) to compute correlations between the time courses of our ROIs. We did this separately for the pre- and the post-learning rest run. For all runs in each participant, we extracted voxels from the ROI masks, a white matter mask, and a cerebrospinal fluid mask from the normalized functional images and averaged the time course within masks. First, these mean functional time courses were corrected for linear trends. Scans with excessive motion artifacts (identified by ART) were replaced with an average of the neighboring scans. Then, the time courses were band-pass filtered for frequencies of 0.01 Hz to 0.1 Hz, mean-centered, and the previously interpolated scans were scrubbed from the time courses. The first three scans in each phase were discarded to allow for T1 equilibration effects. Note that we found the same pattern of results when we ran a whole-brain voxel-by-voxel analysis and then averaged the time course across voxels within the ROI masks as a final preprocessing step. Pair-wise correlations (*Pearson's r*) of ROI time courses were computed controlling for white matter and cerebrospinal fluid time courses and six motion parameters obtained by rigid body correction in SPM8. This resulted in a RSFC estimate for pre- and post-learning rest separately. One-tailed paired-sample *t*-tests were performed to test whether RSFC between our ROIs increased from pre- to post-learning rest. All statistical analyses were performed on Fisher *z*-transformed correlation values (i.e. *Pearson's r*). We focused the analyses on RSFC between the bilateral SN/VTA and hippocampus ROIs. Furthermore, one-tailed *Pearson's* correlations tested whether a RSFC increase from pre- post-learning rest showed a relationship with later memory advantages (memory accuracy: high-reward – low-reward). In addition, we performed exploratory whole-

23

brain, voxel-based analyses with the bilateral hippocampal ROI as seed (for details, see **Table S4**).

FMRI multivariate pattern analyses (MVPA)

Preprocessing. Because our primary interest was the investigation of activity patterns during offline rest periods (i.e. in the absence of any stimuli), we used the preprocessed functional images (unsmoothed). We further used SPM8 to correct for motion artifacts via regressing out motion-related activity (spikes and six motion parameters similar as in the RSFC analyses approach) and used a high-pass filter set at 0.008 Hz (128 s) along with an AR(1) model to correct for auto-correlations. We then selected the data from all voxels within our hippocampal ROI and loaded them into the Princeton MVPA toolbox (Polyn et al., 2005) (http://www.csbmb.princeton.edu/mvpa) that was used for all further MVPA pre-processing and classification. Because the hippocampal ROI only included voxels that are maximal sensitive to reward, no additional feature selection was performed and all ROI voxels were included in the analysis (the same approach was used for further exploratory ROIs). Depending on the classifier, each scan was associated with either a unique scene context regressor (i.e. four different regressors) or a reward regressor (i.e. high or low reward regressor). The regressors were assigned to all time points during the entire scene context (i.e. cue period, objects, and feedback) in order to train the classifier on multiple aspects of the encoding context. All regressors were shifted by four time points to account for the hemodynamic lag. For each run, the data were de-trended and z-scored. Time between mini-blocks was not fed into the classifier.

MVPA classifier. Two different MVPA were performed. First, we tested whether activity patterns were dissociable during encoding. To do this, we selected all ten encoding runs and performed a cross validation approach. That is, we tested each run separately while training

the classifier on the nine remaining runs. Classifier performance was averaged across all ten iterations. Second, in order to test for high-reward reactivation during pre- and post-learning rest, we selected again all ten encoding runs but now trained the classifier on all ten runs and then "tested" the classifier on time points from both pre- and post-learning rest. We used a ridge regression classifier that has been used successfully in the neuroimaging literature (Detre et al., 2013; Newman and Norman, 2010; Poppenk and Norman, 2012). In comparison to a standard multiple regression, ridge regression has the advantage of using an L2 regularization term to minimize the sum of the squared feature weights (penalty term was set at 2) (for details of this analysis, see Detre et al., 2013).

Analyses of classifier output. For the classifier that tested encoding accuracy, we investigated whether the overall classifier performance (averaged across all ten runs) was above chance by using a one-tailed one-sample *t*-test. For the classifier that tested for reactivation during post-learning rest periods, we first investigated the classifier's output on all time points. That is, we counted how often the classifier predicted that a time point was associated with high reward and then used one-tailed paired-sample *t*-tests to determine whether this number was greater than chance for post-learning and pre-learning rest, separately. Additionally, we tested whether there was more evidence for high-reward contexts during post-learning compared to pre-learning rest (i.e. evidence for increased reactivation compared to baseline). The final analyses tested whether increased reactivations correlate with the later HR>LR object-context memory advantage. We therefore computed across-participants one-tailed *Pearson's* correlations (one-tailed) to test whether the reward-specific increase in hippocampal reactivation during post-encoding rest relative to baseline (i.e. [high reward: post-learning rest – pre-learning rest]) with the behavioral HR>LR object-context memory advantage.
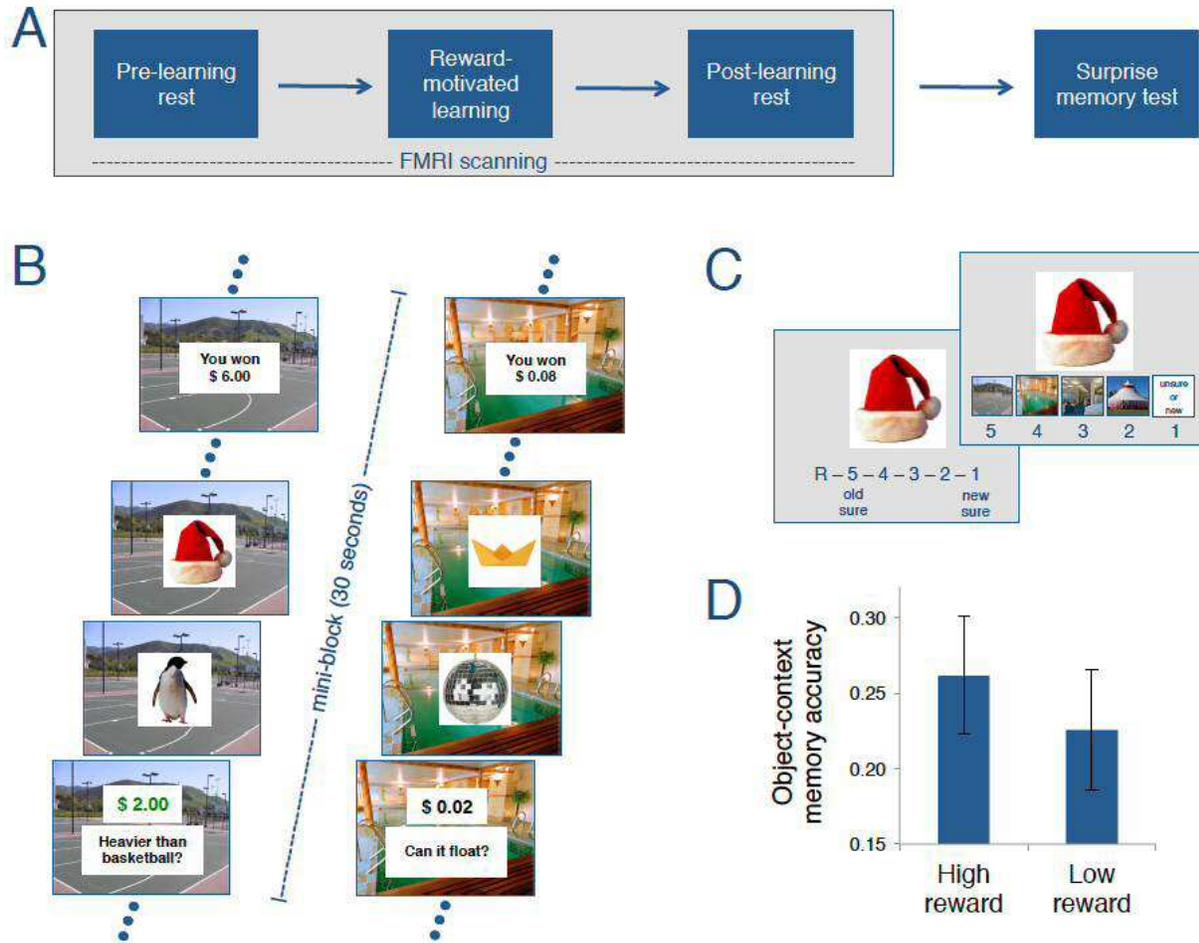
Figure 1. Overview of study design. **(A)** fMRI data were collected during a pre-learning rest phase (10 min), a reward-motivated learning task (48 min), and a post-learning rest phase (10 min). After scanning, participants performed a surprise recognition memory test (on average, the interval between initial learning of an item and presentation of that item at test was around 83 min). **(B)** During learning, participants encoded objects in 30s trial blocks, with each block associated with a particular scene background and encoding task. Each block started with a high ($2.00) or low ($0.02) reward cue and instructions about which encoding task should be performed for the next four consecutively presented objects. Reward was contingent on accurate judgment performance, and feedback about the accumulated reward was given at the end of the block. To maintain the continuity of each encoding context, the background scene semantically matched the relevant task (e.g. For the task: "Does this item

weigh more than a basketball?", a basketball court remained on the screen). **(C)** During the memory test, participants indicated (i) whether they could remember an object and how confident they are about their response and (ii) whether they could remember the associated encoding context (i.e. particular task semantically matched with background scene). **(D)** Participants showed enhanced memory for object-context associations that were learned in high-reward compared to low-reward contexts (i.e. HR>LR object-context memory advantage).
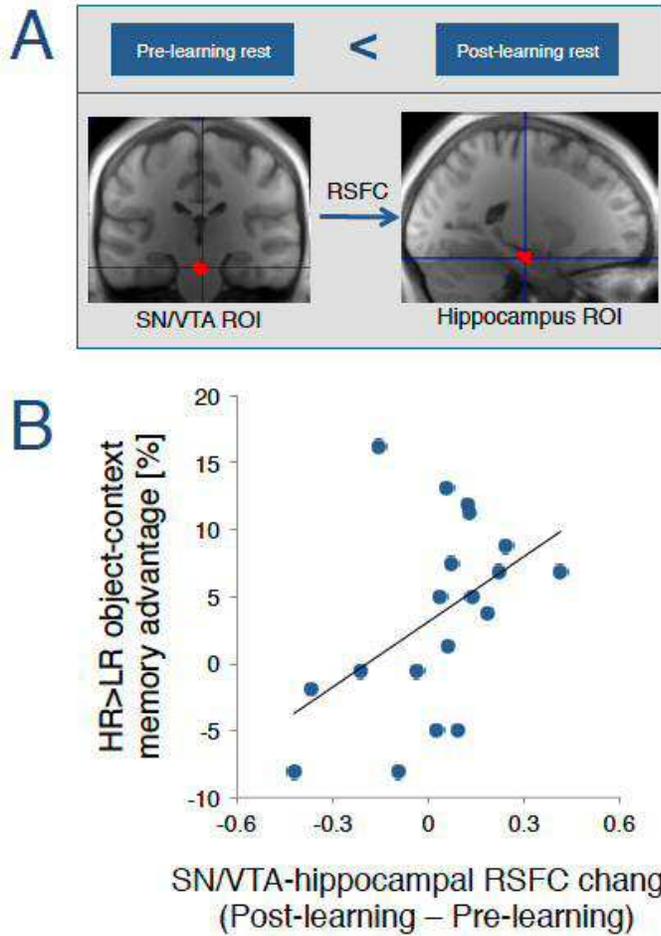
Figure 2. Increases in SN/VTA-hippocampal resting-state functional connectivity (RSFC) correlate with the HR>LR object-context memory advantage. **(A)** We asked whether RSFC changes following learning (i.e. Pre-learning rest < Post-learning rest) between the SN/VTA and hippocampus ROI predicted later the reward-related memory advantage. **(B)** Across subjects, changes in SN/VTA-hippocampal RSFC positively correlated with the HR>LR object-context memory advantage.
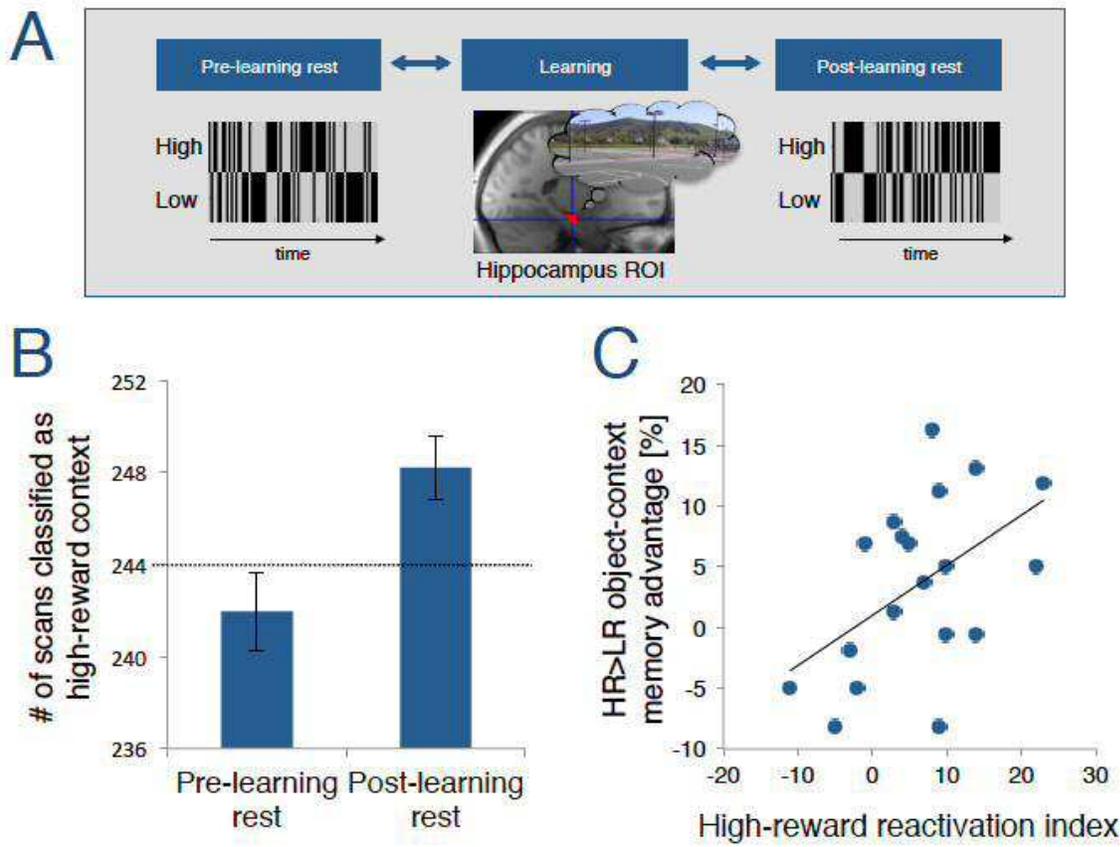
Figure 3. Hippocampal reactivation of high-reward contexts predict the HR>LR object-context memory advantage. **(A)** A 2-way classifier was trained on hippocampal activity patterns to dissociate between high- and low-reward contexts and then 'tested' on all time points during the pre- and post-learning rest period. We counted the number of time points that were classified as for high-reward context. A schematic is shown how timepoints might be labeled for one participant's pre- and post-learning rest period. **(B)** Classifier results showing that the hippocampal activity patterns were more likely than chance-level (dotted line) to be associated with a high-reward context during the post-learning rest, but not during the pre-learning rest period. **(C)** The 'high-reward reactivation index', which is the difference in reactivation of high-reward contexts in post-learning rest compared to pre-learning rest, correlated with the HR>LR object-context memory advantage.

References

Adcock, R.A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., and Gabrieli, J.D.E. (2006). Reward-motivated learning: mesolimbic activation precedes memory formation. Neuron *50*, 507–517.

Atherton, L.A., Dupret, D., and Mellor, J.R. (2015). Memory trace replay: the shaping of memory consolidation by neuromodulation. Trends Neurosci. *38*, 560–570.

Axmacher, N., Elger, C.E., and Fell, J. (2008). Ripples in the medial temporal lobe are relevant for human memory consolidation. Brain *131*, 1806–1817.

Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. Hippocampus *25*, 1073–1188.

Carr, M.F., Jadhav, S.P., and Frank, L.M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. Nat Neurosci *14*, 147–153.

Chadwick, M.J., Hassabis, D., and Maguire, E.A. (2011). Decoding overlapping memories in the medial temporal lobes using high-resolution fMRI. Learn. Mem. *18*, 742–746.

Detre, G.J., Natarajan, A., Gershman, S.J., and Norman, K.A. (2013). Moderate levels of activation lead to forgetting in the think/no-think paradigm. Neuropsychologia *51*, 2371–2388.

Deuker, L., Olligs, J., Fell, J., Kranz, T.A., Mormann, F., Montag, C., Reuter, M., Elger, C.E., and Axmacher, N. (2013). Memory consolidation by replay of stimulus-specific neural activity. J. Neurosci. *33*, 19373–19383.

Diana, R.A., Yonelinas, A.P., and Ranganath, C. (2007). Imaging recollection and familiarity in the medial temporal lobe: a three-component model. Trends Cogn. Sci. (Regul. Ed.) *11*, 379–386.

Düzel, E., Bunzeck, N., Guitart-Masip, M., Wittmann, B., Schott, B.H., and Tobler, P.N. (2009). Functional imaging of the human dopaminergic midbrain. Trends Neurosci. *32*, 321–328.

Eichenbaum, H., Yonelinas, A.P., and Ranganath, C. (2007). The Medial Temporal Lobe and Recognition Memory. Annu. Rev. Neurosci. *30*, 123–152.

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.

Feld, G.B., Besedovsky, L., Kaida, K., Münte, T.F., and Born, J. (2014). Dopamine D2-like receptor activation wipes out preferential consolidation of high over low reward memories during human sleep. J Cogn Neurosci *26*, 2310–2320.

Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature *440*, 680–683.

Girardeau, G., and Zugaro, M. (2011). Hippocampal ripples and memory consolidation. Curr. Opin. Neurobiol. *21*, 452–459.

Gomperts, S.N., Kloosterman, F., and Wilson, M.A. (2015). VTA neurons coordinate with the hippocampal reactivation of spatial experience. Elife 10.7554/eLife.05360

Gruber, M.J., Gelman, B.D., and Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. Neuron *84*, 486–496.

Haber, S.N., and Knutson, B. (2010). The reward circuit: linking primate anatomy and human imaging. Neuropsychopharmacology *35*, 4–26.

Kahn, I., and Shohamy, D. (2013). Intrinsic connectivity between the hippocampus, nucleus accumbens, and ventral tegmental area in humans. Hippocampus *23*, 187-192.

Igloi, K., Gaggioni, G., Sterpenich, V., and Schwartz, S. (2015). A nap to recap or how reward regulates hippocampal-prefrontal memory networks during daytime sleep in humans. Elife 10.7554/eLife.07903

Ji, D., and Wilson, M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat Neurosci *10*, 100–107.

Káli, S., and Dayan, P. (2004). Off-line replay maintains declarative memories in a model of hippocampal-neocortical interactions. Nat Neurosci *7*, 286–294.

Knutson, B., Adams, C.M., Fong, G.W., and Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. J. Neurosci. *21*, RC159.

Lansink, C.S., Goltstein, P.M., Lankelma, J.V., Joosten, R.N.J.M.A., McNaughton, B.L., and Pennartz, C.M.A. (2008). Preferential reactivation of motivationally relevant information in the ventral striatum. J. Neurosci. *28*, 6372–6382.

Lansink, C.S., Goltstein, P.M., Lankelma, J.V., McNaughton, B.L., and Pennartz, C.M.A. (2009). Hippocampus leads ventral striatum in replay of place-reward information. PLoS Biol. *7*, e1000173.

Libby, L.A., Hannula, D.E., and Ranganath, C. (2014). Medial temporal lobe coding of item and spatial information during relational binding in working memory. J. Neurosci. *34*, 14233–14242.

Lisman, J.E., and Grace, A.A. (2005). The Hippocampal-VTA Loop: Controlling the Entry of Information into Long-Term Memory. Neuron *46*, 703–713.

Lisman, J., Grace, A.A., and Düzel, E. (2011). A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. Trends Neurosci. *34*, 536–547.

Logothetis, N.K., Eschenko, O., Murayama, Y., Augath, M., Steudel, T., Evrard, H.C., Besserve, M., and Oeltermann, A. (2012). Hippocampal-cortical interaction during periods of subcortical silence. Nature *491*, 547–553.

Mather, M., Clewett, D., Sakaki, M., and Harley, C.W. (in press). Norepinephrine ignites local hot spots of neuronal excitation: How arousal amplifies selectivity in perception and memory. Behav Brain Sci.

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.

McNamara, C.G., Tejero-Cantero, Á., Trouche, S., Campo-Urriza, N., and Dupret, D. (2014). Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. Nat Neurosci *17*, 1658–1660.

Murayama, K., and Kitagami, S. (2014). Consolidation power of extrinsic rewards: Reward cues enhance long-term memory for irrelevant past events. J Exp Psychol Gen *143*, 15–20.

Murayama, K., and Kuhbandner, C. (2011). Money enhances memory consolidation – But only for boring material. Cognition *119*, 120–124.

Murty, V.P., and Adcock, R.A. (2014). Enriched Encoding: Reward Motivation Organizes Cortical Networks for Hippocampal Detection of Unexpected Events. Cereb. Cortex *24*, 2160-2168.

Newman, E.L., and Norman, K.A. (2010). Moderate excitation leads to weakening of perceptual representations. Cereb. Cortex *20*, 2760–2770.

Oudiette, D., Antony, J.W., Creery, J.D., and Paller, K.A. (2013). The role of memory reactivation during wakefulness and sleep in determining which memories endure. J. Neurosci. *33*, 6672–6678.

Papatheodoropoulos, C. (2010). Patterned activation of hippocampal network (approximately 10 Hz) during in vitro sharp wave-ripples. Neuroscience *168*, 429–442.

Peigneux, P., Orban, P., Balteau, E., Degueldre, C., Luxen, A., Laureys, S., and Maquet, P. (2006). Offline persistence of memory-related cerebral activity during active wakefulness. PLoS Biol. *4*, e100.

Pezzulo, G., van der Meer, M.A.A., Lansink, C.S., and Pennartz, C.M.A. (2014). Internally generated sequences in learning and executing goal-directed behavior. Trends Cogn. Sci. (Regul. Ed.) *18*, 647–657.

Polyn, S.M., Natu, V.S., Cohen, J.D., and Norman, K.A. (2005). Category-specific cortical activity precedes retrieval during memory search. Science *310*, 1963–1966.

Poppenk, J., and Norman, K.A. (2012). Mechanisms supporting superior source memory for familiar items: a multi-voxel pattern analysis study. Neuropsychologia *50*, 3015–3026.

Ramirez-Villegas, J.F., Logothetis, N.K., and Besserve, M. (2015). Diversity of sharp-wave-ripple LFP signatures reveals differentiated brain-wide dynamical events. Proceedings of the National Academy of Sciences *112*, E6379–E6387.

Ritchey, M., Montchal, M.E., Yonelinas, A.P., and Ranganath, C. (2015). Delay-dependent contributions of medial temporal lobe regions to episodic memory retrieval. Elife *4*:e05025

Schlichting, M.L., and Preston, A.R. (2014). Memory reactivation during rest supports upcoming learning of related content. Proceedings of the National Academy of Sciences *111*, 15845–15850.

Shohamy, D., and Adcock, R.A. (2010). Dopamine and adaptive memory. Trends Cogn. Sci. (Regul. Ed.) *14*, 464–472.

Singer, A.C., and Frank, L.M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. Neuron *64*, 910–921.

Staresina, B.P., Alink, A., Kriegeskorte, N., and Henson, R.N. (2013). Awake reactivation predicts memory in humans. Proceedings of the National Academy of Sciences *110*, 21159–21164.

Staresina, B.P., Bergmann, T.O., Bonnefond, M., van der Meij, R., Jensen, O., Deuker, L., Elger, C.E., Axmacher, N., and Fell, J. (2015). Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. Nat Neurosci *18*, 1679–1686.

Sutherland, G.R., and McNaughton, B. (2000). Memory trace reactivation in hippocampal and neocortical neuronal ensembles. Curr. Opin. Neurobiol. *10*, 180–186.

Tambini, A., and Davachi, L. (2013). Persistence of hippocampal multivoxel patterns into postencoding rest is related to memory. Proceedings of the National Academy of Sciences *110*, 19591–19596.

Tambini, A., Ketz, N., and Davachi, L. (2010). Enhanced brain correlations during rest are related to memory for recent experiences. Neuron *65*, 280–290.

Tompary, A., Duncan, K., and Davachi, L. (2015). Consolidation of Associative and Item Memory Is Related to Post-Encoding Functional Connectivity between the Ventral Tegmental Area and Different Medial Temporal Lobe Subregions during an Unrelated Task. J. Neurosci. *35*, 7326–7331.

Valdés, J.L., McNaughton, B.L., and Fellous, J.-M. (2015). Offline reactivation of experience-dependent neuronal firing patterns in the rat ventral tegmental area. J. Neurophysiol. *114*, 1183–1195.

Wittmann, B.C., Schott, B.H., Guderian, S., Frey, J.U., Heinze, H.-J., and Düzel, E. (2005). Reward-related FMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. Neuron *45*, 459–467.

Wolosin, S.M., Zeithamova, D., and Preston, A.R. (2012). Reward modulation of hippocampal subfield activation during successful associative encoding and retrieval. J Cogn Neurosci *24*, 1532–1547.

Wolosin, S.M., Zeithamova, D., and Preston, A.R. (2013). Distributed hippocampal patterns that discriminate reward context are associated with enhanced associative binding. J Exp Psychol Gen *142*, 1264–1276.

Yarkoni, T., Poldrack, R.A., Nichols, T.E., Van Essen, D.C., and Wager, T.D. (2011). Large-scale automated synthesis of human functional neuroimaging data. Nat. Methods *8*, 665–670.