

Contributions of the hippocampus to feedback learning

Kathryn C. Dickerson^{1,3} · Mauricio R. Delgado^{1,2}

Published online: 9 June 2015
© Psychonomic Society, Inc. 2015

Abstract Humans learn about the world in a variety of manners, including by observation, by associating cues in the environment, and via feedback. Across species, two brain structures have been predominantly involved in these learning processes: the hippocampus—supporting learning via observation and paired association—and the striatum—critical for feedback learning. This simple dichotomy, however, has recently been challenged by reports of hippocampal engagement in feedback learning, although the role of the hippocampus is not fully understood. The purpose of this experiment was to characterize the hippocampal response during feedback learning by manipulating varying levels of memory interference. Consistent with prior reports, feedback learning recruited the striatum and midbrain. Notably, feedback learning also engaged the hippocampus. The level of activity in these regions was modulated by the degree of memory interference, such that the greatest activation occurred during the highest level of memory interference. Importantly, the accuracy of information learned via feedback correlated with hippocampal activation and was reduced by the presence of high memory interference. Taken together, these findings provide evidence of hippocampal involvement in feedback learning by

demonstrating both its relevance for the accuracy of information learned via feedback and its susceptibility to interference.

Keywords fMRI · Probabilistic learning · Memory · Dual-task · Basal ganglia

The ability of an individual to learn and make adaptive choices takes advantage of two learning and memory systems: the medial temporal lobe (MTL) and the basal ganglia (BG). Across species, several studies have substantiated the idea that the MTL system is implicated in declarative learning, whereas the BG system, in particular the striatum, is involved in learning via feedback (for a review, see White & McDonald, 2002). Traditionally, research has focused on the dichotomy between these two learning systems, emphasizing their individual contributions to distinct learning types and the nature of their interactions (Packard & Knowlton, 2002; Poldrack & Packard, 2003; Squire, 2004; White & McDonald, 2002; White, Packard, & McDonald, 2013). However, recent studies have revealed that this dichotomy is not always observed. In fact, the extant literature has highlighted hippocampal (i.e., MTL) involvement during feedback learning (Dickerson, Li, & Delgado, 2011; Foerde et al. 2013a, b; Foerde & Shohamy, 2011a) and BG recruitment during declarative memory (for a review, see Scimeca & Badre, 2012).

One early demonstration of hippocampal engagement during feedback learning in humans was observed in a feedback-based probabilistic classification task (Poldrack et al. 2001). The authors found that the MTL was engaged early in learning but that this diminished over time, whereas the BG was initially less active and increased activation across learning. Foerde and colleagues built on this study by employing a dual-task design to investigate the roles of the MTL and BG during feedback learning concurrent with an auditory working

Electronic supplementary material The online version of this article (doi:10.3758/s13415-015-0364-5) contains supplementary material, which is available to authorized users.

✉ Kathryn C. Dickerson
kathryn.dickerson@duke.edu

¹ Department of Behavioral and Neural Sciences, Rutgers University, Newark, NJ, USA

² Department of Psychology, Rutgers University, Newark, NJ, USA

³ Center for Cognitive Neuroscience, Duke University, Durham, NC 27708, USA

memory task (Foerde, Knowlton, & Poldrack, 2006). The MTL was observed under “normal” feedback-learning conditions, and the BG was recruited during the dual-task manipulation, designed to monopolize the MTL via working memory. Activity in the MTL also correlated with performance in the “normal” feedback-learning condition, but not in the dual-task condition. Taken together, these findings documented a role for the MTL (hippocampus) during feedback learning.

However, MTL activity is not always observed during feedback learning (Delgado, Miller, Inati, & Phelps, 2005; Delgado, Nystrom, Fissell, Noll, & Fiez, 2000; Dobryakova & Tricomi, 2013; Shohamy et al. 2004a; Tricomi, Delgado, & Fiez, 2004; Wilkinson et al., 2014)—a process that seems to be more reliant on the integrity of the BG as evidenced by difficulties in learning via feedback expressed by patients with Parkinson’s disease (Foerde, Braun, & Shohamy, 2013a; Foerde, Race, et al., 2013b; Foerde & Shohamy, 2011a, b; Jahanshahi, Wilkinson, Gahir, Dharminda, & Lagnado, 2010; Knowlton, Mangels, & Squire, 1996; Myers et al. 2003; Shohamy et al. 2004a, b; Wilkinson, Lagnado, Quallo, & Jahanshahi, 2008), but not with patients with amnesia (Foerde, Race, et al., 2013b; Knowlton et al. 1996; Knowlton, Squire, & Gluck, 1994; Myers et al. 2003). Given the conflicting results and differences across paradigms, it is unclear what aspects of feedback learning engage the MTL, particularly the hippocampus. For instance, one plausible context in which the hippocampus becomes important for feedback learning is when the feedback is delayed. Indeed, patients with amnesia are impaired in probabilistic learning when the feedback is delayed by 7 s (Foerde, Race, et al., 2013b), and activation of the hippocampus in an fMRI study is observed with a similar delay (Foerde & Shohamy, 2011a). In addition, the length of the feedback delay may negatively impact recruitment of the hippocampus (Dobryakova & Tricomi, 2013), highlighting how sensitive the involvement of the hippocampus is during feedback learning.

Another important aspect of feedback learning is processing of the information conveyed by the feedback itself (i.e., positive or negative). Hippocampal signals have been shown to differentiate between correct and incorrect feedback during some forms of probabilistic, trial-and-error learning (Dobryakova & Tricomi, 2013; Li, Delgado, & Phelps, 2011)—a pattern of activation similar to responses observed in the striatum (Delgado, 2007; Delgado et al. 2000; Tricomi et al. 2004). This is consistent with processing a prediction error learning signal—a signal that indicates whether the outcome was better or worse than expected—which typically is associated with blood-oxygen-level-dependent (BOLD) responses in the striatum in humans (for reviews, see Daw & Doya, 2006; O’Doherty, 2004) and with dopamine midbrain neuronal activity in nonhuman primates (Schultz, Dayan, & Montague, 1997). A few examples have also shown a correlation between such signals and the hippocampus during

probabilistic feedback learning (Dickerson et al. 2011; Foerde & Shohamy, 2011a), although the majority of fMRI studies that have measured brain correlations with prediction error in humans have observed striatum, rather than hippocampal, involvement. Thus, although evidence that the hippocampus may contribute to feedback learning is mounting (Cincotta & Seger, 2007; Dickerson et al. 2011; Foerde, Race, et al., 2013b; Foerde & Shohamy, 2011a; Li et al. 2011; Mattfeld & Stark, 2011; Okatan, 2009; Wimmer & Shohamy, 2012), the circumstances of hippocampal recruitment and the contributions of the hippocampus to feedback learning remain uncertain and topics of great interest.

To investigate the role of the hippocampus during feedback learning, we asked the following questions: (1) Is the hippocampus recruited during a simple probabilistic feedback-learning task? (2) Is hippocampal activation correlated with the feedback cue accuracy during learning? And (3) Would the presence of a secondary task that recruits the hippocampus negatively impact feedback learning? We reasoned that one way to address these questions would be to utilize a dual-task manipulation to examine the engagement of the hippocampus under conditions of (a) baseline feedback learning and (b) feedback learning concurrent with a task competing for hippocampal engagement—namely memory interference. Our dual-task design therefore kept feedback learning constant across all sessions while varying the amount of interference (no interference, low interference, high interference). The use of a dual-task paradigm to examine related questions has been validated by others (Foerde, Braun, & Shohamy, 2013a; Foerde et al. 2006; Wimmer, Braun, Daw, & Shohamy, 2014), but it is still a novel aspect of this design, and builds on former dual-task paradigms by utilizing two dual-task conditions: one designed to explicitly interfere with the hippocampus, and the other to recruit the hippocampus at a minimal to low level. We hypothesized that the hippocampus contributes to feedback learning, and therefore that (1) the hippocampus would be engaged during feedback learning, (2) hippocampal activation would positively correlate with feedback learning, and (3) feedback learning would be adversely affected by a memory task competing for hippocampal activation (high-memory-interference condition). More specifically, we hypothesized that if the hippocampus is important for feedback learning, the accuracy of information learned via feedback should be diminished in the high-memory-interference condition relative to the other conditions. However, if hippocampal involvement is not important for feedback learning, then we should find either (a) no decrease in feedback learning in the high-memory-interference condition or (b) an improvement in feedback learning, consistent with the animal literature (Lee, Duman, & Pittenger, 2008; Packard, Hirsh, & White, 1989).

Method

Overview

We employed a dual-task design in order to examine the role of the hippocampus in feedback learning. Specifically, we probed activity in the hippocampus and in regions more typically associated with feedback learning (i.e., the striatum and midbrain) during probabilistic feedback learning concurrent with varying levels of memory interference. Participants performed a probabilistic feedback-learning task concurrent with no memory interference (control condition), low memory interference (keeping a scene in mind for a few seconds), or high memory interference (determining whether a scene was old or new). The high-memory-interference condition was designed to recruit the hippocampus for scene recognition, thereby reducing this region's ability to engage in feedback learning.

Experimental design

Participants Twenty-eight right-handed adults (15 female, 13 male; mean age = 22 years, $SD = 5.0$) participated in the study and provided informed consent prior to beginning the experiment. The Institutional Review Boards of Rutgers University and the University of Medicine and Dentistry of New Jersey (UMDNJ) approved this study. Participants were screened for head injury, history of psychiatric or neurological impairments, and contraindications to MRI (e.g., metal implants, claustrophobia). All females were administered a pregnancy test, which was verified to be negative, prior to beginning the MRI session. The final analysis consisted of 24 participants (12 female, 12 male), because four participants failed to complete the fMRI session on the second day, due to scanner malfunction (two), failure to attend to stimuli in the first session (one self-reported using his/her cellphone during the session), or voluntarily withdrawing (one).

fMRI acquisition A 3-T Siemens Allegra MRI scanner at the UMDNJ Advanced Imaging Center was used to collect the structural (T1-weighted MPRAGE: 256×256 matrix, FOV = 256 mm, 176 1-mm sagittal slices) and functional images (single-shot echo EPI sequence; TR = 2,000 ms, TE = 25 ms, FOV = 192 cm, flip angle = 80° , matrix = 64×64 , slice thickness = 3 mm). Forty contiguous oblique-axial slices ($3 \times 3 \times 3$ mm voxels) were acquired parallel to the anterior commissure–posterior commissure line. Functional data preprocessing and analysis was completed using the Brain Voyager QX software (Brain Innovation, Maastricht, The Netherlands). The preprocessing consisted of slice-timing correction, three-dimensional motion correction (six parameters), spatial smoothing (4 mm, full width at half maximum), high-pass filtering of frequencies (three cycles per time course), and normalizing the data to the Talairach stereotaxic space

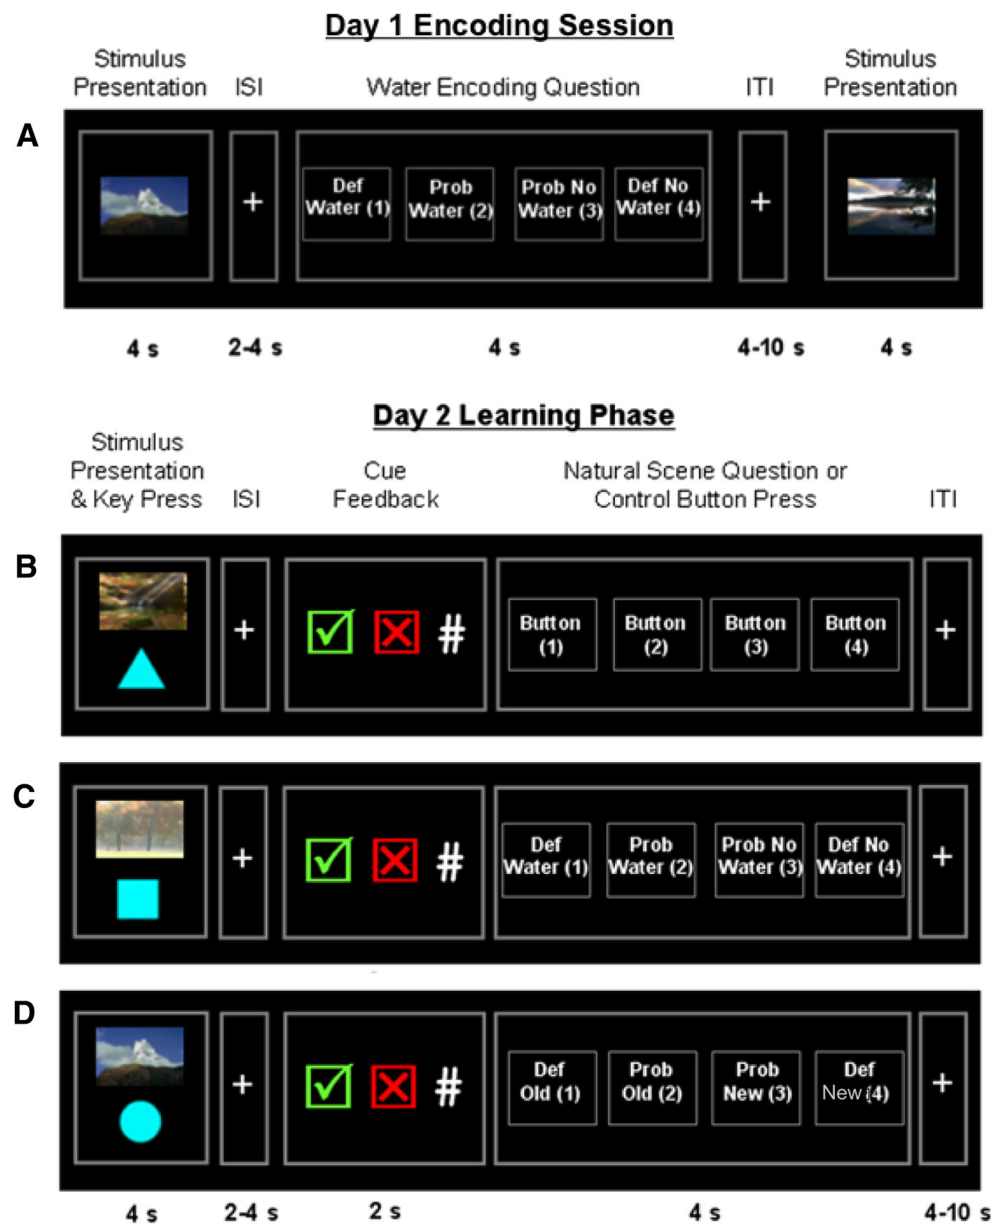
(Talairach & Tournoux, 1988). A canonical two-gamma hemodynamic response function was used to convolve the events of interest.

Experimental task E-Prime Version 2.0 (Psychology Software Tools, Pittsburgh, PA) was used to program the task. Participants viewed the task inside the MRI via the use of a back-projection system using a mirror affixed to the head coil. An MRI-compatible button box was used to record behavioral responses. Monetary compensation consisted of \$10 per hour for behavioral testing and screening, and \$25 per hour for the neuroimaging session. All participants were paid a total of \$60 and debriefed following task completion.

The experiment involved two days of participation—a behavioral and a neuroimaging session, respectively. On Day 1, participants completed a behavioral declarative-memory encoding session in the laboratory. Participants viewed 40 unique scenes (courtesy of R. Alison Adcock) on a computer and made a simple decision about each scene to ensure their attention during encoding (i.e., determining whether the scene contained water). Each scene was presented once (Fig. 1A). Participants were informed that their scene memory would be tested the following day and that their payment would be determined in part by how well they correctly discriminated between old and new scenes. Encoding trials consisted of the scene presentation (4 s), a jittered interstimulus interval (ISI; 2–4 s), the decision period (4 s), and a jittered intertrial interval (ITI; 4–10 s). During the decision period, participants indicated the presence/absence of water in the scene by pressing: 1 *definitely water in the scene*, 2 *probably water in the scene*, 3 *probably no water in the scene*, or 4 *definitely no water in the scene*. They were instructed to press the “2” or the “3” key if they had trouble remembering the scene or if it was ambiguous whether the scene contained water. A blue box highlighted the participant's choice. They were paid \$10 for participating in the encoding session.

On Day 2, participants completed the neuroimaging portion of the study. This consisted of a probabilistic feedback-learning task that varied with respect to (a) an additional behavioral requirement (session type) and (b) difficulty (adapted from Delgado et al. 2005; Delgado et al. 2000; Dickerson et al. 2011). During the experiment, participants were presented with six unique cues (shapes). The participants' goal was to learn the numerical value of these cues; they were instructed that each cue had been randomly assigned a value on the scale of 1–9. They were informed that they did not need to learn the exact value of each cue, but to categorize each one as being higher or lower than 5. Participants were told to guess and learn from feedback whether their response was correct or incorrect, and they were also instructed about the probabilistic nature of the task. Specifically, they were informed that there was no perfect answer and that they should make

Fig. 1 Experimental task design. (A) In the Day 1 encoding session, participants viewed 40 scene images in the laboratory and made a response indicating whether or not each scene contained water. In the Day 2 learning session, participants learned the probabilistic values of six cues (e.g., a triangle) via feedback while undergoing functional neuroimaging. (B) In the feedback control (FC) session, they were instructed to ignore the novel natural scenes. There was no memory interference in this session. (C) In the low-memory-interference (LMI) session, they made a judgment regarding the presence or absence of water in each new scene. (D) In the high-memory-interference (HMI) session, they judged whether each scene was old (from Day 1 encoding session) or new. ISI, interstimulus interval; ITI, intertrial interval



the response that provided them with predominantly correct feedback.

There were three distinct learning session types: no memory interference/feedback control (Fig. 1B), low memory interference (Fig. 1C), and high memory interference (Fig. 1D). The task also varied with respect to difficulty, assessed by the probabilistic values of the cues: Three of the cues were 90% predictive of their value (hereafter referred to as the easy cues), and three cues were 70% predictive of their value (hereafter referred to as the hard cues). Each session type contained one easy and one hard cue. Participants were instructed about all session types and performed a short practice version of the task (using different stimuli) prior to entering the scanner. The images for the practice session were drawn in part from the

McGill University Calibrated Colour Image Database (Olmos & Kingdom, 2004).

The experimental session on Day 2 consisted of a learning phase and a test phase. In the learning phase, participants' primary goal was to learn the numerical value of each cue. The learning phase consisted of six runs (two presentations of each session type). Each run contained 20 trials (ten easy-cue trials and ten hard-cue trials), for a total of 120 trials. The trial format in all sessions consisted of a natural scene/cue presentation (4 s); a jittered ISI (2–4 s); feedback presentation contingent on the participant's response for the cue value—√ for correct, X for incorrect, and # for missed trials (2 s); the decision period (varied by session type; 4 s); and a jittered ITI (4–10 s) before onset of the next stimulus. In all learning sessions,

the stimulus presentation period contained a natural scene and a probabilistic cue. Depending on the session type, the natural scene was either new or old. Each natural scene was presented only one time (scenes did not overlap across sessions). Participants were instructed to make a buttonpress in all sessions during the stimulus presentation period regarding the value of the probabilistic cue. Additionally, participants were instructed to optimize their responses, and their accuracy was scored according to this optimization strategy. Instructions for the natural scenes varied by session type, as described below.

In the *high-memory-interference (HMI)* session, participants were required to determine whether the natural scene presented in each trial was old or new while simultaneously learning the value of the probabilistic cues (Fig. 1D). As in all sessions, they made a buttonpress indicating the value of the probabilistic cue in the stimulus presentation period. Following the ISI and the feedback presentation period, the response screen appeared. In this session, participants were told to press 1 if they thought the scene was *definitely old*, 2 if it was *probably old*, 3 if it was *probably new*, and 4 if it was *definitely new*. Additionally, participants were instructed to press 1 only if they had a strong memory of seeing the scene on the day prior (*I remember*) and 2 if they recognized the scene to be old, but did not have a strong memory of seeing it the day before (*I know*). They were explicitly informed that half of the scenes would be old and half would be new and that this was the only session in which they would see natural scenes that had been shown on the day prior. The purpose of this session was to create a high level of memory interference by engaging the hippocampus in the scene recognition task.

In the *low-memory-interference (LMI)* session, participants were instructed to attend to the natural scenes in order to perform the water detection task while simultaneously learning the value of the probabilistic cues (Fig. 1C). They were told they would not be tested on the scenes. In the LMI session, participants viewed the natural scene/cue and made a buttonpress indicating the value of the probabilistic cue during the stimulus presentation period. Following the ISI and the feedback presentation, the response screen appeared, and participants were instructed to make a response regarding the presence or absence of water in the natural scene. This response screen was identical to that presented during the Day 1 encoding session. The purpose of the LMI condition was to serve as a control for the effects of multitasking in the HMI session. Our goal was to create two dual-task conditions—one that recruited the MTL (HMI) and one that engaged the MTL at a minimal level, if at all (LMI). The minimal level of MTL recruitment could be attributed to incidental encoding that is possible during the LMI session (see the Discussion section for implications).

In the session with no memory interference, the feedback control (FC) session, participants were instructed to focus on

learning the value of the probabilistic cues and to ignore the natural scenes, which were all new (Fig. 1B). Participants were informed that they would not be tested on the scenes. They viewed the natural scene/cue and made a buttonpress indicating the value of the probabilistic cue during the stimulus presentation/response period. After the ISI, participants received feedback contingent on their response. Following the feedback, participants made a buttonpress (buttons 1–4) that had no meaningful value and served as a motor control for the other sessions. The purpose of the FC session was to assess a baseline level of participants' feedback learning in the absence of memory interference.

Equal numbers of participants completed the six sessions, in one of two orders: (1)FC, HMI, LMI, FC, HMI, LMI; or (2) FC, LMI, HMI, FC, LMI, HMI. The trial order in each block was pseudorandomized to ensure the trial types were balanced; all participants completed the identical trial order. Participants were informed that the total amount of money they would receive would be determined by how well they correctly identified old versus new natural scenes in the HMI session, and by how well they learned the value of the feedback cues in all three learning sessions. Importantly, we instructed participants in this manner so that they would try to learn both the scenes and the values of the cues equally well. However, all participants were paid the same amount of money and debriefed.

Following the learning session, participants completed a test session that contained the probabilistic cues presented in all three sessions, as well as two novel cues. Critically, no feedback or natural scenes were presented. Participants were required to indicate the value of the probabilistic cues that they had learned in the learning phase to test for accuracy. The test session contained 80 trials total, ten presentations of each cue, presented in random order. The trial structure and timing consisted of a self-timed stimulus–response period during which participants indicated the value of the cue, followed by a 6- to 14-s jittered ITI before onset of the next cue.

Data analyses

Behavioral data analysis D prime (d') analyses were performed to examine the participants' ability to correctly identify scenes that contained water versus those that did not in the Day 1 and Day 2 sessions, and to identify old versus new scenes in the Day 2 task (Macmillan & Creelman, 1991). Such analyses are informative because they account for a participant's response bias (e.g., being more likely to choose old vs. new) when calculating a participant's ability to correctly identify a stimulus (sensitivity). We also report hit rates, false alarm rates, misses, and correct rejections, and focus our analyses and interpretations on hit rates. Finally, consistent with other probabilistic-learning tasks (Dickerson et al. 2011; Foerde et al. 2006; Foerde, Race, et al., 2013b; Foerde &

Shohamy, 2011a), when analyzing the feedback cue learning data, we calculated accuracy scores based on optimizing the cue probabilities. For example, for a cue that was 70% higher than 5, participants received a score of 100% correct if they responded “higher than 5” on every trial.

Neuroimaging data analyses: General linear model The primary analysis of interest was to examine activation within the brain across the three distinct sessions (FC, LMI, and HMI) to investigate the role of the hippocampus in feedback learning under different levels of memory interference. For the learning phase, we conducted a random-effects general linear model (GLM) containing three predictors, one representing each session type: FC, LMI, and HMI. Missed trials and six motion parameters were also included as regressors of no interest. The BOLD signal was examined during the stimulus presentation period (when the probabilistic cue and natural scene were presented on the screen). During this time period, participants made a choice regarding the value of the probabilistic cue and presumably either ignored or attended to the natural scene, depending on the session (ignore the scene in the FC session; look for water in the LMI session; or try to remember the scene in the HMI session). This time period was examined in order to probe for differences in learning and memory processing within the brain, with the primary focus on the hippocampus, but also investigating the striatum and midbrain as regions particularly involved in feedback-based learning and reward-related processing (for reviews, see Daniel & Pollmann, 2014; Delgado, 2007; Schultz, 1998).

First, we performed a contrast of FC versus baseline to probe for hippocampal activation during feedback learning with no memory interference. Second, we performed a repeated measures analysis of variance (ANOVA) using Session Type (FC, LMI, HMI) as a within-subjects factor. We examined the main effect of session type to determine how the hippocampus, ventral striatum, and midbrain varied according to the level of memory interference. Third, we performed an exploratory ANOVA during the feedback period, again using Session Type (FC, LMI, HMI) as a within-subjects factor to determine how the hippocampus, ventral striatum, and midbrain were engaged during the processing of cue feedback. The results from this analysis should be interpreted with caution, however, because there were insufficient incorrect trials to model the data appropriately. The resulting statistical maps in all analyses were false-discovery rate (FDR) corrected.

Post-hoc and correlation analyses conducted on the behavioral and neuroimaging data consisting of more than two tests within a family of comparisons were corrected for multiple comparisons with the sequential Bonferroni technique (Holm, 1979; Rice, 1989). All significant

statistical tests survived correction for multiple comparisons, unless otherwise noted.

Results

Behavioral results

Water detection task: Participants correctly identified scenes Participants were able to successfully discriminate between scenes that contained water versus those that did not. On both days, participants’ hit rates for identifying scenes with water were greater than the false alarm rates, with a bias toward labeling scenes as containing water (for statistics, see Table 1).

Scene recognition: Participants successfully remembered the scenes In the HMI session, participants’ hit rate was significantly greater than the false alarm rate, with a bias toward labeling scenes as new, suggesting that participants were able to successfully discriminate between old and new scenes. Furthermore, participants’ hit rate was greater than chance (for statistics, see Table 1).

Dual-task comparison: Participants performed better on the LMI session To assess the difficulty differences across our memory interference manipulations, we compared participants’ hit and false alarm rates across the LMI and HMI sessions. We found a significantly higher hit rate in the LMI session than in the HMI session [$t(23) = 9.21, p < .0001, d = 2.159$] and a lower false alarm rate in the LMI session than in the HMI session [$t(23) = 3.12, p < .005, d = 0.640$]. These results suggest that the LMI condition (identifying and keeping in mind whether or not scenes contained water) was an easier task than the HMI recognition task (discriminating between old and new scenes).

Interactions between session type and feedback learning: Memory interference effects on feedback learning We examined feedback learning as it varied across levels of memory interference (session type); for instance, did participants learn the value of the cues better in the FC than in the HMI session? To examine feedback learning as it varied by memory interference, we performed a 3 (Session Type: FC vs. LMI vs. HMI) \times 2 (Cue Difficulty: easy vs. hard) within-subjects, repeated measures ANOVA (Fig. 2). The results revealed a main effect of session type [$F(1.65, 37.87) = 16.58, p < .001$], a main effect of cue difficulty [$F(1, 23) = 9.57, p < .01$], and no significant interaction [$F(1.56, 35.82) = 0.44, p = .60$]. Participants performed better on cues learned in the FC and LMI sessions than on cues in the HMI session [$t(23) = 5.75, p < .0001, d = -1.204$, and $t(23) = 3.54, p < .005, d = -0.730$, respectively]. No accuracy differences were observed between

Table 1 Low- and high-memory-interference behavioral results

	Day 1	Day 2
LMI: Water Detection Task		
Mean hit rate (<i>SEM</i>)	.97 (.00)	.93 (.01)
Mean false alarm rate (<i>SEM</i>)	.16 (.02)	.11 (.02)
Hit vs. false alarm rate	$t(23)= 48.14, p < .0001, d= 12.535$	$t(23)= 32.39, p < .0001, d= 6.907$
d' prime	$d'= 2.95, C= -.41$	$d'= 2.95, C= -.12$
HMI: Scene Retrieval		
Mean hit rate (<i>SEM</i>)		.61 (.04)
Mean false alarm rate (<i>SEM</i>)		.24 (.03)
Hit vs. false alarm rate	N.A.	$t(23)= 9.58, p < .0001, d= 1.973$
d' prime		$d'= 1.11, C= .25$
Hit rate vs. chance		$t(23)= 3.10, p < .01$

Abbreviations indicate low-memory-interference (LMI) and high-memory-interference (HMI) sessions

the FC and LMI sessions [$t(23)= 1.60, p= .12, d= 0.396$], suggesting that the LMI condition did not negatively affect feedback learning. Participants performed better in the FC and LMI sessions than in the HMI session, suggesting that the scene memory manipulation in the HMI session negatively impacted feedback learning. As expected, overall performance was significantly greater for easy than for hard cues [$t(23)= 3.13, p < .005, d= 0.639$]. Importantly, we observed no significant interaction with cue difficulty, indicating that difficulty had similar effects across session types. See Supplementary Fig. 1 for mean accuracy plotted across learning blocks.

Feedback-learning effects on memory performance We also examined memory performance as it varied by feedback cue accuracy (correct vs. incorrect trials). To examine how participants' scene memory was influenced by feedback cue accuracy, we compared participants' hit, false alarm, miss, and correct rejection rates, as well as d' and C for correct versus incorrect trials for 22 participants (two participants were excluded from this analysis due to a hit rate of 0 on incorrect trials). There was a marginally higher false alarm rate on incorrect trials [$t(21)= 1.76, p= .09, d= 0.376$]; no other tests were significant [hit rate, $t(21)= 0.07, p= .95, d= 0.014$; miss rate, $t(21)= 0.07, p= .95, d= 0.014$; correct rejection rate, $t(21)= 1.48, p= .15, d= 0.316$; d' , $t(21)= 1.39, p= .18, d= 0.296$; C , $t(21)= 1.45, p= .16, d= 0.309$]. Thus, memory performance did not differ significantly with respect to feedback-learning accuracy.

Relationship between feedback cue learning and scene memory We examined whether there was a significant relationship between participants' performance in feedback learning and scene memory by performing across-subjects correlations of the cue accuracy in each session with the hit rate in the HMI session. No significant relationships were observed in

any session [HMI, $r(22)= .24, p= .26$; LMI, $r(22)= .33, p= .11$; FC, $r(22)= .28, p= .19$], indicating that performance on these tasks may be orthogonal.

Test phase A test phase was included to examine accuracy differences between the probabilistic cues (easy and hard) acquired across the three session types during the learning phase. The first analysis probed for potential differences in session types, irrespective of cue difficulty, by conducting a repeated measures, within-subjects ANOVA using Session Type as a within-subjects factor (FC, LMI, HMI, and novel). A significant main effect of session type was observed [$F(2.11, 48.62)= 13.21, p < .001$; Supplementary Fig. 2]. Post-hoc t tests indicated significantly greater performance for the cues presented in all sessions than for novel cues [FC vs. novel, $t(23)= 5.18, p <$

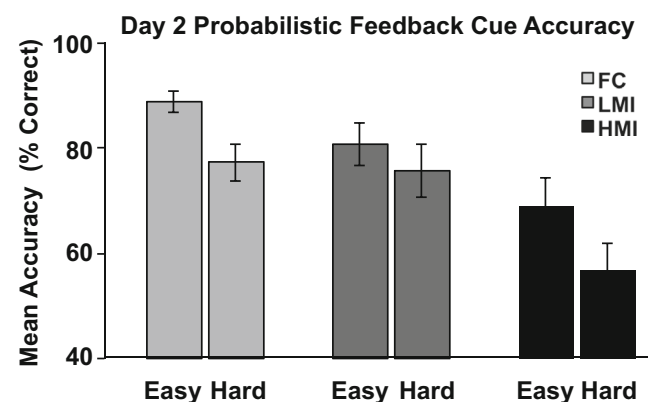


Fig. 2 Participants' mean accuracy for learning the values of the probabilistic cues in the Day 2 learning session. A 3 (Session Type: FC vs. LMI vs. HMI) \times 2 (Cue Difficulty: easy vs. hard) within-subjects, repeated measures ANOVA revealed a main effect of session type ($p < .001$), a main effect of cue difficulty ($p < .01$), and no significant interaction ($p= .60$). Error bars represent standard errors of the means. Abbreviations indicate the feedback control (FC), low-memory-interference (LMI), and high-memory-interference (HMI) sessions

.0005, $d = 1.061$; LMI vs. novel, $t(23) = 4.54$, $p < .0005$, $d = 0.932$; HMI vs. novel, $t(23) = 3.19$, $p < .005$, $d = 0.653$]. Marginally significantly better accuracy was observed for cues previously presented in the FC session than for those in the HMI session [$t(23) = 2.19$, $p = .04$, $d = -0.447$; trend after sequential Bonferroni correction]. No other significant differences were observed between cues previously presented in the various learning sessions (all p values $\geq .11$).

In the second analysis, we examined differences across session types and levels of cue difficulty, excluding novel cues. A 3 (Session Type: FC vs. LMI vs. HMI) \times 2 (Cue Difficulty: easy vs. hard) repeated measures, within-subjects ANOVA was performed. The results revealed a significant main effect of session type [$F(1.21, 27.87) = 3.88$, $p = .05$], no significant main effect of cue difficulty [$F(1, 23) = 0.07$, $p = .80$], and no significant interaction of cue difficulty and session type [$F(1.79, 41.13) = 0.60$, $p = .54$]. As we reported in the preceding paragraph, the main effect of session type was driven by nearly better performance in the FC than in the HMI session [$t(23) = 2.19$, $p = .04$, $d = -0.447$; trend after sequential Bonferroni correction].

Neuroimaging results

Contrast of feedback learning versus baseline: Hippocampal engagement during feedback learning To first investigate whether the hippocampus was engaged during feedback learning, we performed a whole-brain contrast of feedback learning (using our FC session) versus baseline. This comparison revealed extensive engagement of the bilateral MTL, including the hippocampus (right: $x, y, z = 21, -23, -4$; left: $x, y, z = -25, -23, -5$) and parahippocampal cortex (right: $x, y, z = 23, -29, -14$; left: $x, y, z = -22, -30, -14$) [$q(\text{FDR}) < .05$; Fig. 3]. These results support the hypothesis that the hippocampus is engaged in feedback learning.

Session type GLM: Engagement of subcortical learning and memory regions Having established that the hippocampus is engaged in feedback learning, we next

examined how activity in the hippocampus and in other subcortical learning and memory regions varied in terms of the level of interference by competing tasks. To do so, we conducted a repeated measures ANOVA using Session Type (FC, LMI, HMI) as a within-subjects factor. Examining the main effect of session type revealed the engagement of subcortical regions previously reported during learning and memory processes (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Shohamy & Adcock, 2010), including the hippocampus, ventral striatum, and midbrain [$q(\text{FDR}) < .001$; Table 2]. We extracted the mean BOLD signal from these functionally defined regions of interest (ROIs) and then conducted post-hoc t tests to examine the pattern of activity within the bilateral hippocampus, bilateral ventral striatum, and midbrain.

All examined a priori ROIs (identified from the whole-brain ANOVA) displayed similar patterns of activity showing modulation by memory interference, with the greatest amount of BOLD signal being exhibited during the HMI condition (Fig. 4; see Table 3 for all t statistics). Activity within the right hippocampus ($x, y, z = 18, -12, -18$; Fig. 4B) scaled according to session type. The left hippocampus ($x, y, z = -20, -19, -13$) showed greater activation in the HMI than in the LMI and FC conditions, and marginally greater activity in the LMI than in FC conditions. A similar pattern of activity was observed in areas more typically associated with feedback learning (ventral striatum and midbrain). The right ventral striatum ($x, y, z = 10, 3, 3$; Fig. 4A) showed enhanced activity in the HMI condition relative to the LMI and FC sessions, with no significant difference between the LMI and FC sessions. The left ventral striatum ($x, y, z = -11, 4, 2$) displayed a similar pattern of results, with approaching greater activity for LMI than for FC. Lastly, activity within the midbrain ($x, y, z = 0, -15, -9$; Fig. 4C) scaled according to session type.

Session type GLM control analyses We constructed three additional models to examine whether differences in difficulty, learning, and/or processing speed across session types

Fig. 3 Whole-brain contrast of feedback control versus baseline, revealing extensive engagement of the medial temporal lobe (MTL), including the bilateral hippocampus and parahippocampal cortex

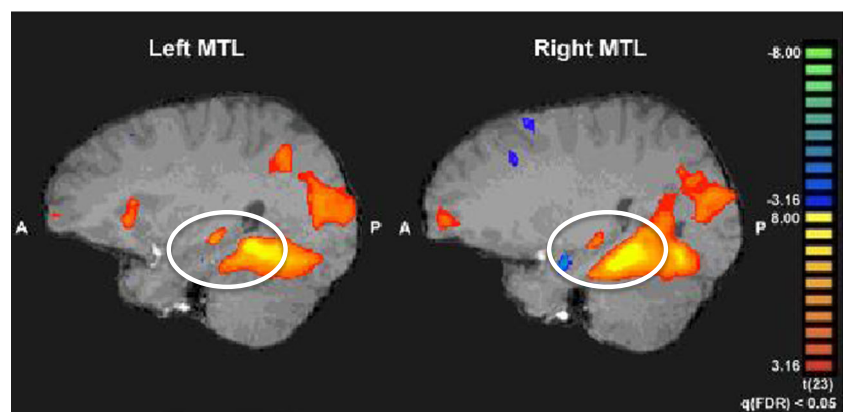


Table 2 Learning phase: Whole-brain session type model

Region of Activation	Brodmann Area (BA)	Laterality	Talairach Coordinates			# Voxels (mm ³)	<i>F</i> stat
			<i>x</i>	<i>y</i>	<i>z</i>		
Precentral gyrus	BA 6	Left	−28	−11	54	563	28.53
Middle frontal gyrus	BA 6	Right	29	−5	51	482	21.35
Middle frontal gyrus	BA 6	Left	−4	1	48	743	29.81
Precentral gyrus	BA 6	Right	41	−2	33	554	34.49
Supramarginal gyrus	BA 40	Right	50	−53	30	539	24.16
Precentral gyrus	BA 6	Left	−40	−2	30	791	29.41
Supramarginal gyrus	BA 40	Left	−55	−56	30	219	16.28
Temporal parietal junction	BA 13	Right	44	−35	24	172	20.89
Insula	BA 13	Left	−31	19	9	831	28.59
Insula		Right	32	16	3	642	31.93
Ventral striatum		Right	10	3	3	566	36.96
Ventral striatum		Left	−11	4	2	728	37.38
Midbrain			0	−15	−9	420	20.99
Hippocampus		Left	−20	−19	−13	101	19
Ventral visual processing stream		Left	−31	−47	−15	1,000	91.15
Hippocampus		Right	18	−12	−18	289	22.48

significantly contributed to the observed results. The results from all three models were consistent with our original findings, as we detail below.

Session type GLM with cue difficulty First we included cue difficulty as a regressor in the model that contained session type, allowing us to examine BOLD responses in our ROIs for both easy and hard cues. If difficulty differences were driving the results, one might expect to see a different pattern of results for easy than for hard cues. In the behavioral data, differences emerged across the HMI and LMI session for the hard cues only. Therefore, it could be argued that our results may have been caused by differential responses in our ROIs specifically to the hard cues. If this were the case, we should observe a significant interaction of session type with cue difficulty. Importantly, there was no significant interaction in our a priori ROIs. As can be seen in Supplementary Fig. 3, the BOLD responses were not significantly different for easy- and hard-cue trials in the hippocampus or ventral striatum (all *p* values > .17). The midbrain showed greater activation for the hard than for the easy cues in the HMI session only [$t(23) = 2.43$, $p < .05$, $d = 0.497$], but the overall pattern of activation remained the same in this region.

Session type GLM with correct trials only The second model contained correct trials only as regressors of interest (including incorrect trials as predictors of no interest). This independent analysis again allowed us to examine

whether differences in accuracy or learning were driving our effects. Specifically, if incorrect trials, the proportion of incorrect to correct trials, or another metric of difficulty that existed across session types contributed to our originally observed results, one might expect a different pattern to emerge in an analysis that modeled correct trials only. Importantly, the results from this model recapitulated our original findings, again suggesting that differences in accuracy/learning were not a significant contributor to the original results (see Supplementary Fig. 4).

Session type GLM with reaction time The third analysis contained trial-by-trial reaction times included as a confound regressor in the model. This independent analysis again allowed us to examine whether differences in performance, as assessed by reaction times, were driving our effects. This model was particularly important, because reaction time differences were observed across session types (see Supplementary Fig. 5). Therefore, if reaction time differences across session types contributed to our originally observed results, one might expect a different pattern to emerge in this model. Critically, the results from this model recapitulated our original findings, providing converging evidence that differences in processing speed/accuracy/learning were not significant contributors to the original results (see Supplementary Fig. 6). Together, these data support the argument that differences in difficulty across sessions did not significantly contribute to the observed patterns of BOLD activation.

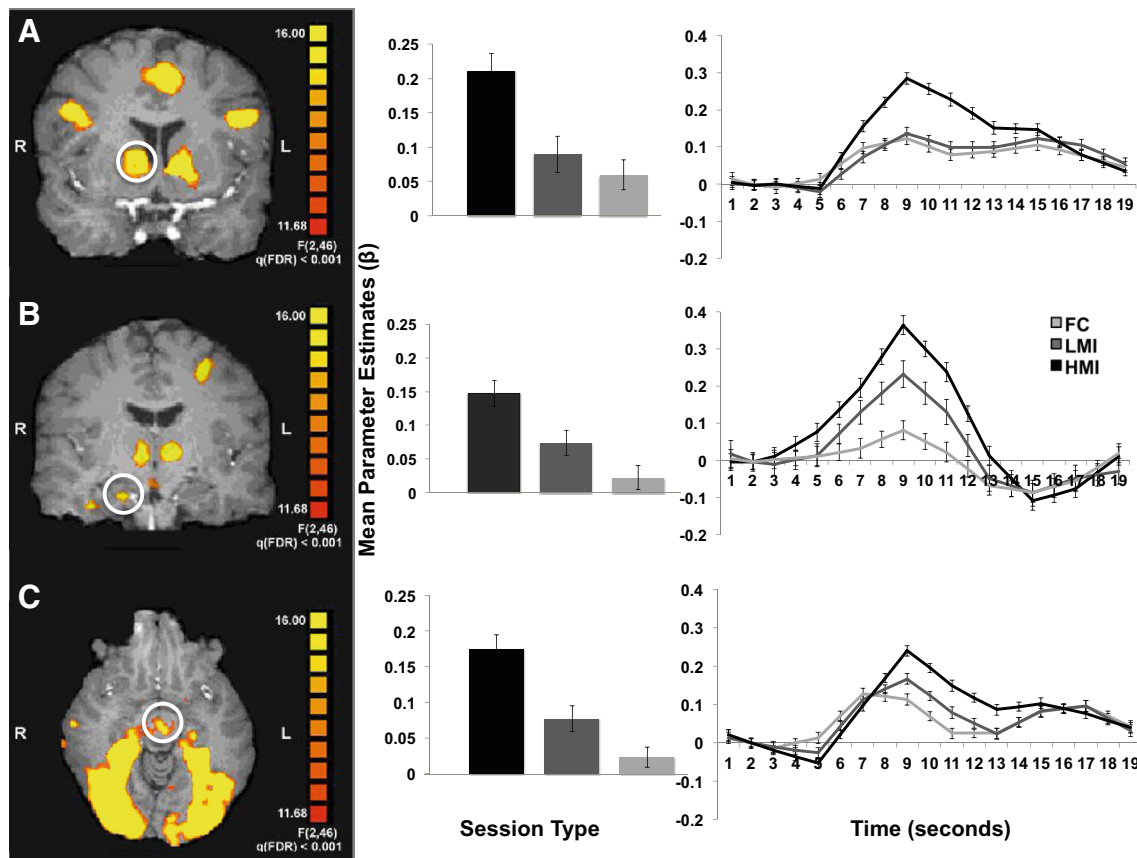


Fig. 4 Whole-brain results from general linear model of the main effect of session type: Mean BOLD signal in bilateral ventral striatum(A), right hippocampus(B), and midbrain(C). All regions exhibited the greatest amount of activation in the HMI session, as can be seen in the bar plots (middle panel) and the event-related averages (ERA; right panel). Time

Point 3 denotes trial onset in the ERA plots. For simplicity, activation from the right ventral striatum and right hippocampus only is displayed here. Error bars represent standard errors of the means. Abbreviations in the legend indicate the feedback control (FC), low-memory-interference (LMI), and high-memory-interference (HMI) sessions

Brain–behavior correlations: Hippocampus significantly correlated with feedback learning in the FC session and with memory performance in the HMI session In order to probe the behavioral relevance of the observed activation, Pearson’s correlations were conducted between the mean parameter estimates from our a priori ROIs and the behavioral accuracy measures. To test our hypothesis that the hippocampus was engaged in feedback learning, we examined whether the hippocampal BOLD signal correlated with feedback cues learning during the learning phase, specifically for the FC

session. According to our hypothesis, the mean BOLD signal and feedback cue accuracy were positively correlated in the right hippocampus [$r(22) = .44, p < .05$] and marginally positively correlated in the left hippocampus [$r(22) = .37, p = .08$] in the FC session (Fig. 5). These results demonstrate the functional significance of the engagement of the hippocampus during probabilistic feedback learning by relating the BOLD signal with behavioral accuracy.


To test our hypothesis that hippocampal activation was related to scene memory in the HMI condition, we performed

Table 3 Post-hoc *t* tests in a priori ROIs: Comparison across session types

Regions	HMI > LMI	HMI > FC	LMI > FC
R hippocampus	$t(23) = 3.80, p < .001, d = 0.775$	$t(23) = 6.94, p < .0001, d = -1.419$	$t(23) = 2.92, p < .01, d = -0.597$
L hippocampus	$t(23) = 4.26, p < .0005, d = 0.870$	$t(23) = 5.50, p < .0001, d = 1.134$	$t(23) = 1.84, p = .08, d = 0.381$
R ventral striatum	$t(23) = 4.33, p < .0005, d = 0.884$	$t(23) = 7.05, p < .0001, d = 1.460$	$t(23) = 1.09, p = .29, d = 0.225$
L ventral striatum	$t(23) = 6.20, p < .0001, d = 1.281$	$t(23) = 8.21, p < .0001, d = 1.804$	$t(23) = 2.01, p = .06, d = 0.419$
Midbrain	$t(23) = 3.90, p < .001, d = 0.798$	$t(23) = 7.58, p < .0001, d = 1.606$	$t(23) = 2.47, p < .05, d = 0.510$

The table shows planned post-hoc comparisons following significant ANOVA results; all significant tests survived correction for multiple comparisons. Abbreviations indicate the feedback control (FC), low-memory-interference (LMI), and high-memory-interference (HMI) sessions

Fig. 5 The mean BOLD signal was significantly positively correlated with feedback cue accuracy during the feedback control (FC) session in the right hippocampus [$r(22) = .44, p < .05$], and marginally positively correlated in the left hippocampus [$r(22) = .37, p = .08$]



an across-subjects analysis correlating participants' hit rates with the BOLD signal in the HMI condition from the bilateral hippocampal ROIs. Consistent with our hypothesis, the right hippocampus positively correlated with participants' hit rates [$r(22) = .54, p < .01$], suggesting that the HMI manipulation was successful [note that there was no significant relationship in the left hippocampus: $r(22) = .14, p = .52$; see Supplementary Fig. 7]. Notably, no significant relationships emerged between memory performance and BOLD signal in the ventral striatum [right, $r(22) = .22, p = .30$; left, $r(22) = .14, p = .50$] or the midbrain [$r(22) = .12, p = .58$]. To test whether the relationship between the hippocampus and scene memory was different from the relationship of the striatum or midbrain with scene memory, we performed Steiger's Z test (Meng, Rosenthal, & Rubin, 1992). Specifically, we performed a one-tailed test, since we hypothesized that the hippocampus would be significantly more positively correlated with scene memory than would the other two regions. We observed that the relationship between hippocampal activation and memory performance was distinct from relationships in the other regions: hippocampus versus left ventral striatum, $Z = 1.85, p <$

$.05$; versus right ventral striatum, $Z = 1.65, p = .05$; versus midbrain, $Z = 1.82, p < .05$.

Feedback period analysis: Session type GLM with correct trials only—Engagement of hippocampus and striatum Lastly, we performed an exploratory analysis of the feedback period, when participants received feedback regarding their choices about the cue values. This analysis was deemed exploratory because it included only correct trials, given the scarcity of incorrect trials. This model revealed the engagement of both the hippocampus and striatum, although the loci of activation and the patterns of results differed from the results observed during the stimulus presentation period. The hippocampus bilaterally (right: $x, y, z = 19, -16, -14$; left: $x, y, z = -22, -15, -17$) was deactivated and showed the following pattern: FC and LMI > HMI. The right putamen ($x, y, z = 26, 2, 4$) showed a pattern similar to the one observed in the stimulus presentation period (HMI = LMI > FC), whereas the left ventral striatum ($x, y, z = -15, 3, -3$) showed a blunted response to the HMI session (FC and LMI > HMI) (see Supplementary Fig. 8).

Discussion

Learning via feedback is well known to engage BG circuits, as has been evidenced by an array of research, from animal studies to impairments observed in patients afflicted with Parkinson's disease (e.g., Dickerson et al. 2011; Foerde, Braun, & Shohamy, 2013a; Foerde et al. 2006; Foerde, Race, et al., 2013b; Foerde & Shohamy, 2011a; Jahanshahi et al. 2010; Knowlton et al. 1996; Packard & Knowlton, 2002; Poldrack et al. 2001; Shohamy, Myers, Grossman, et al., 2004a; Shohamy, Myers, Onlaor, & Gluck, 2004b; Wilkinson et al. 2008). Less is known about the role of the MTL, and particularly the hippocampus, in feedback learning. In this experiment, we utilized a dual-task manipulation to engage the hippocampus with varying levels of memory interference that were concurrent with feedback learning, in order to characterize the contribution of the hippocampus to this type of learning. We hypothesized that the hippocampus contributes to feedback learning, and therefore that (1) the hippocampus would be engaged during feedback learning, (2) hippocampal activation would positively correlate with feedback learning, and (3) feedback learning would be adversely affected by a memory task competing for hippocampal activation. In accordance with our hypotheses, we observed (1) engagement of the hippocampus in all feedback-learning sessions, (2) that feedback-learning accuracy correlated with hippocampal function in the control condition (FC), and (3) that feedback learning was compromised in a condition of memory interference, where hippocampal activation significantly correlated with scene memory. These results suggest that hippocampal engagement in feedback learning may be an important contributor to the accuracy of learned information.

The present study corroborates recent evidence in the literature that the hippocampus is involved in feedback learning (Cincotta & Seger, 2007; Dickerson et al. 2011; Foerde, Race, et al., 2013b; Foerde & Shohamy, 2011a; Li et al., 2011; Mattfeld & Stark, 2011; Okatan, 2009; Wimmer & Shohamy, 2012), and critically offers the demonstration that hippocampal engagement is important for the accuracy of the information acquired during feedback learning. Timing is also an important aspect of feedback learning, since prior reports have documented the involvement of the hippocampus when a delay separates the presentation of a stimulus to be learned from its related feedback (Foerde, Race, et al., 2013b; Foerde & Shohamy, 2011a) and during subsequent testing post-learning (which may index memory, rather than learning, of cue values; Foerde et al. 2006). In the present study, we showed that hippocampal activity also correlates with feedback learning after a relatively short delay. We hypothesized that the hippocampus would correlate with feedback learning under conditions of no memory interference (FC session). Consistent with our hypothesis, correlations with behavioral measures indicated that the right hippocampus correlated with

feedback learning under normal conditions (FC session). Furthermore, when feedback learning was diminished in a memory interference session (HMI), hippocampal activation positively correlated with scene memory (hit rate). This result suggests that our manipulation was, at least in part, successful, and highlights the utility of using dual-task manipulations to engage a region of interest.

We additionally performed an exploratory analysis of the feedback period, when participants received feedback about their choices regarding the values of probabilistic cues (higher or lower than 5). A whole-brain ANOVA revealed the engagement of the hippocampus and of regions of the striatum, at different loci than in the stimulus presentation period. Interestingly, the hippocampus was deactivated during the feedback period. We reason that these data provide one possible explanation of why feedback cue accuracy was worse in the HMI session, and support the interpretation that the hippocampus plays a meaningful role in feedback learning. Note that the results from this time period should be interpreted with a degree of caution, however, since we were only able to appropriately model correct trials.

Role of the hippocampus in feedback learning

Converging evidence from prior studies (Dickerson et al. 2011; Foerde et al. 2006; Foerde, Race, et al., 2013b; Foerde & Shohamy, 2011a; Li et al. 2011; Poldrack et al. 2001) suggested that the hippocampus contributes to feedback-based learning of cue–reward associations. The hippocampus is believed to play a role in encoding episodes (Gluck, Meeter, & Myers, 2003; Schacter & Wagner, 1999), stimulus–stimulus representations (Bunsey & Eichenbaum, 1995; Eichenbaum & Bunsey, 1995; Gluck & Myers, 1993), as well as flexibly adapting information to use in novel contexts (Myers et al. 2003; Shohamy & Adcock, 2010; Shohamy & Wagner, 2008; Wimmer & Shohamy, 2012). In category learning, it has been suggested that the hippocampus functions both to compress overlapping stimulus information and to differentiate between distinct stimulus information (Gluck et al. 2003; Gluck & Myers, 1993). It is thought these hippocampal representations are formed over time through exposure to many trials during training, and are subsequently communicated to cortical regions. It is possible, therefore, that the hippocampus was compressing and updating feedback-related information acquired throughout learning in the present task. This hypothesis fits with prior literature that has shown hippocampal activation to feedback-related information, including information differentiating between correct and incorrect trials (Dobryakova & Tricomi, 2013; Li et al. 2011), encoding prediction errors (Dickerson et al. 2011; Foerde & Shohamy, 2011a), and prediction-error-like signals (Chen, Olsen, Preston, Glover, & Wagner, 2011). Furthermore, the hippocampus may express information about correct associations even

before this is detected in the subject (Okatan, 2009). As such, the hippocampus may be among the first brain areas to express learning and to drive behavioral change, a hypothesis consistent with evidence that the hippocampus exhibits enhanced activation prior to information that is later remembered (Adcock et al. 2006; Shohamy & Adcock, 2010).

Evidence by Foerde and colleagues (Foerde, Race, et al., 2013b; Foerde & Shohamy, 2011a), combined with our present findings, suggests that timing may play a critical role in moderating hippocampal engagement in feedback learning. In their prior work, Foerde and Shohamy (2011a) demonstrated that introducing a 6-s delay between the cue and the outcome period was enough to recruit the hippocampal system as the predominant learning system, rather than the BG. Here we showed that a delay of as little as 2 s was enough to recruit hippocampal activation. In our paradigm, the time window between the cue and the outcome period varied between 2 and 4 s. In addition, there was a delay between participants' viewing and making memory judgments about the scenes (old/new). Therefore, one potentially important shared feature across these distinct tasks was timing—a short delay period prior to feedback and prior to recording memory choice. Interestingly, a recent study by Dobryakova and Tricomi (2013) used a very long delay (~25 min) in a feedback-learning task, and they did not observe hippocampal activation, suggesting that performance after very long delays may not rely on this system. Thus, an interesting future direction would be to better understand the relationship between timing and hippocampal engagement during feedback learning, which may delineate exactly what circumstances are required in order to recruit this system.

Relation to prior dual-task studies

The results from the present study complement the rich existing literature on dual-task paradigms (Craik, Govoni, Naveh-Benjamin, & Anderson, 1996; Foerde, Poldrack, & Knowlton, 2007; Frensch, Lin, & Buchner, 1998; Jiménez & Vázquez, 2005; Klingberg, 2000; Pashler, 1994; Poldrack et al. 2005; Shanks & Channon, 2002; Waldron & Ashby, 2001), as well as prior work examining the effects of working memory and cognitive load on behavioral accuracy performance (Kahneman, Treisman, & Burkell, 1983; Longstaffe, Hood, & Gilchrist, 2014; Treisman, Kahneman, & Burkell, 1983). In particular, our study is similar to the experiment of Foerde and colleagues (2006), who utilized a feedback-based version of the weather prediction task, consisting of both a single-task control condition and a dual-task condition. In the probe session (no feedback), the MTL BOLD signal correlated with accuracy and declarative knowledge for single-task cues, whereas the BG BOLD signal correlated with accuracy for dual-task cues. The authors concluded that this task

could be learned via either the MTL or BG, with engagement of these regions depending on the task demands.

Although it is similar to that of Foerde and colleagues (2006), our experiment differs from theirs in a few key ways: First, we included an important control condition—the LMI condition. Inclusion of this second dual-task condition allowed us to compare hippocampal activation and behavioral accuracy across dual-task conditions that manipulated hippocampal engagement to greater (HMI) and lesser (LMI) degrees, while controlling for multitasking. Second, we utilized a less difficult probabilistic-learning task, one that does not require the integration of information across multiple cues. Third, we employed a different type of memory interference task, to allow us to explore memory interference in a distinct domain—namely visual rather than auditory.

Importantly, we replicated Foerde and colleagues' (2006) findings of the MTL's involvement in feedback learning under single-task conditions and the BG's engagement in dual-task conditions, as well as their observations of correlated activity between the MTL and feedback learning under “normal” conditions (FC condition). Unlike Foerde and colleagues (2006), however, the hippocampus and striatum were not exclusively engaged during the single- and dual-task conditions, respectively. In our paradigm, the ventral striatum, hippocampus, and midbrain exhibited increased activation in the memory interference conditions. We replicated these results in three models—one that contained correct trials only, one that modeled easy- and hard-probabilistic-cue trials, and one that included trial-by-trial reaction time measures—thereby reducing the concern that these data were confounded by differences in accuracy or learning across session types. Foerde et al. (2006) observed greater BG activation for single-task than for dual-task correct trials, which differs from our observation of greater BG activation in the dual-task conditions than in the control condition.

We attribute these divergent findings to differences in (a) experimental paradigms, (b) dual-task conditions, and (c) reward incentives across conditions. As we previously mentioned, unlike the weather prediction task, our feedback task did not require participants to integrate associations across multiple stimuli. This multicue integration may more heavily recruit the MTL during learning. Alternatively, it is possible that the use of our relatively simple probabilistic-learning task resulted in recruitment of the hippocampus, as compared with the more complex probabilistic-learning tasks previously used (e.g., that of Knowlton et al. 1996). It is conceivable that participants in our study may have used a more declarative approach of memorizing the cue–outcome associations explicitly, thereby recruiting the hippocampus. However, given that the behavioral performance was on average around 80% for the FC and LMI sessions, it seems unlikely that they used a completely declarative approach. If participants were explicitly memorizing the correct answers and optimizing their

responses (as instructed), we would expect near 100% accuracy. Future work directly comparing probabilistic-learning tasks that vary in difficulty and the amount of necessary cue integration may help clarify this question.

Our study also shares features with two recent studies examining the contributions of incidental episodic memory encoding to reward learning (Foerde, Braun, & Shohamy, 2013a; Wimmer et al., 2014). In these paradigms, better memory for incidentally encoded items co-occurred with worse performance from immediate feedback in Parkinson's patients (Foerde, Braun, & Shohamy, 2013a), and correlated with reduced prediction error signaling in the striatum in healthy individuals (Wimmer et al. 2014). Our findings complement these studies by showing impaired feedback learning concurrent with a memory task. Importantly, our experiment differed from these paradigms in the following key ways: (1) We utilized intentional episodic retrieval as a dual-task condition, rather than incidental episodic encoding, and (2) we used two dual-task conditions—one serving as a control for cognitive demand and the level of hippocampal interference. Nonetheless, these paradigms highlight the usefulness of dual-task designs to begin answering important open questions in the field regarding the interplay between the hippocampus and striatum in multiple learning types.

Limitations and alternative interpretations

A design limitation of the present study was that the FC session always occurred as the first and fourth runs. Therefore, the mean position of the FC session was earlier than those of the other conditions. Two task constraints limited us to this design: First, it was important to acquire a baseline measure of feedback learning—so the FC session needed to be first. Second, we did not want participants to complete the same condition type consecutively or to complete more than two dual-task conditions in a row, in order to limit fatigue. However, it is possible that proactive interference and/or fatigue could have contributed to our findings. Future iterations of this paradigm counterbalancing the condition order should address this concern.

Another design choice to consider for future iterations is that we did not test participants' memory for the novel scenes presented in the LMI session; participants were accurately informed that their memory for the novel scenes would not be tested. However, it may have been helpful to perform a surprise memory test to examine whether participants had incidentally encoded any of the scenes. This would lend another interpretation of the data: Hippocampal activation in the LMI session could have been due to incidental encoding of the novel images. This interpretation is consistent with prior work in the literature supporting hippocampal activation during incidental encoding (see Schacter & Wagner, 1999, for a review). However, it has also been shown that intentional

encoding elicits greater hippocampal activation than incidental encoding (for a discussion, see Martin, 1999). Even if our LMI session recruited some amount of hippocampal activation due to incidental encoding, this would not invalidate our argument that the hippocampus contributes to feedback learning. In our task, intentional encoding (the HMI session) recruited more hippocampal activation and disrupted feedback learning to a greater extent than incidental encoding (the LMI session). However, future applications of this paradigm or of other related paradigms should consider directly testing for incidental encoding.

An alternative interpretation of the activity observed in our a priori ROIs is that these effects were driven by task difficulty. The two memory interference sessions, HMI and LMI, were not equated in difficulty, with behavioral results indicating that the HMI session was more difficult for participants. Therefore, one interpretation is that these regions were engaged according to the level of task difficulty, with the increased BOLD signal being observed in the most difficult condition (HMI). However, the results from three additional analyses—including cue difficulty, correct trials only, and trial-wise reaction time measures—replicated our original finding, thereby suggesting that the observed pattern of activation was not significantly driven by differences in difficulty, learning/accuracy, or processing speed. It could be, however, that some other construct—such as cognitive load, attention, effort, and so on—significantly contributed to the results, which will be an important topic for future studies. In addition, future experiments should examine within-subjects measures relating hippocampal BOLD activation with feedback learning to be sure that this effect was not driven by a subset of participants. A caveat of this experiment is that it examined only across-subjects relationships between BOLD activation and our learning measures (feedback cue accuracy and scene memory); however, examination of the correlation plots revealed that the effects were not driven by one or a few individuals.

Another important note is that the pattern of activation that we observed was not specific to our a priori regions. We performed a whole-brain analysis that revealed the engagement of multiple regions, noted in Table 2. Although some regions displayed a pattern of activation similar to those in our areas of interest (e.g., the cingulate cortex), it is important to note that the pattern of results that we observed was not ubiquitous across task-engaged areas. For example, task-evoked activity in bilateral premotor cortex ($x, y, z = 28, -5, 47$; $x, y, z = -27, -5, 50$) did not show differential responses to the HMI and LMI conditions, and the opposite pattern of activation ($HMI < LMI < FC$) was observed in the dorsomedial prefrontal cortex ($x, y, z = 18, 52, 32$). Future studies focused on whole-brain network analyses may explore the contributions of multiple regions to feedback learning.

An additional possibility is that the pattern of BOLD activity that we observed in the hippocampus, striatum, and midbrain is a learning signal associated with acquiring the probabilistic cue values. Participants' accuracy was better in the FC and LMI sessions than in the HMI session. Therefore, the enhanced activity to the HMI session could have been due to the fact that participants were still acquiring the cue value associations, whereas in the other sessions they acquired them more quickly. This interpretation is consistent with literature suggesting that activity within the striatum, hippocampus, and midbrain corresponds to a learning signal (Delgado et al. 2005; Dickerson et al. 2011; Mattfeld & Stark, 2011; Schultz, 1998; Tricomi et al. 2004).

A last point to consider when interpreting the data is the possibility that participants may have engaged in serial processing in the present task. Given the relatively long allotted response time in our paradigm (4 s), it is possible that participants engaged in serial processing of the probabilistic cues and scenes. One hypothesis is that in the dual-task conditions, particularly in the HMI session, participants may have spent less time deciding the value of the probabilistic cues. This hypothesis offers one explanation for the observed reduced accuracy in this session relative to the others. Contrary to this hypothesis, however, the actual cue reaction times were significantly greater in the HMI session (see Supplementary Fig. 5). This result suggests that if participants were engaged in serial processing, they were making a decision about the scenes first. If so, this provides an explanation for why cue accuracy was negatively impacted by the presence of the HMI task, but memory performance was unaffected by cue accuracy. In support of this idea, memory performance did not change across runs, as assessed by a *t* test on the hit rate values calculated for Run 1 and Run 2 [$t(23) = 1.44, p = .16, d = 0.303$].

Conclusions

The primary finding from this experiment suggests that the hippocampus is important for learning via feedback. The hippocampus was engaged in all feedback-learning sessions, both with and without additional memory interference. Importantly, hippocampal BOLD signal positively correlated with feedback-learning accuracy. Furthermore, feedback accuracy was reduced in the HMI session in which the hippocampal BOLD signal correlated with declarative recognition, suggesting both that our manipulation was successful and that the hippocampus may contribute to the accuracy of feedback information. This result complements existing literature showing that the hippocampus tracks feedback prediction errors (Dickerson et al. 2011; Foerde & Shohamy, 2011a) and discriminates between correct and incorrect trials (Dobryakova & Tricomi, 2013; Li et al. 2011). These and related emerging

data (Cincotta & Seger, 2007; Dickerson et al. 2011; Foerde, Race, et al., 2013b; Foerde & Shohamy, 2011a; Li et al. 2011; Mattfeld & Stark, 2011; Okatan, 2009; Poldrack et al. 2001; Wimmer & Shohamy, 2012) strongly support a role for the hippocampus in feedback learning.

Author Notes K.C.D. is now at the Center for Cognitive Neuroscience, Duke University. The authors thank Drs. Elizabeth Tricomi, R. Alison Adcock, Dominic Fareri, Vishnu Murty, Christopher Coutlee, Lila Davachi, and Bart Krekelberg, as well as Jane Chen for analysis assistance and helpful discussion, and the staff at the University Heights Imaging Center. The research reported in this article was supported by the National Institute of Mental Health of the National Institutes of Health, under Award Numbers MH08408 and F32MH100764. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The authors declare no conflicts of interest.

References

- Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron*, *50*, 507–517. doi:10.1016/j.neuron.2006.03.036
- Bunsey, M., & Eichenbaum, H. (1995). Selective damage to the hippocampal region blocks long-term retention of a natural and nonspatial stimulus–stimulus association. *Hippocampus*, *5*, 546–556.
- Chen, J., Olsen, R. K., Preston, A. R., Glover, G. H., & Wagner, A. D. (2011). Associative retrieval processes in the human medial temporal lobe: Hippocampal retrieval success and CA1 mismatch detection. *Learning and Memory*, *18*, 523–528. doi:10.1101/lm.2135211
- Cincotta, C. M., & Seger, C. A. (2007). Dissociation between striatal regions while learning to categorize via feedback and via observation. *Journal of Cognitive Neuroscience*, *19*, 249–265. doi:10.1162/jocn.2007.19.2.249
- Craik, F. I. M., Govoni, R., Naveh-Benjamin, M., & Anderson, N. D. (1996). The effects of divided attention on encoding and retrieval processes in human memory. *Journal of Experimental Psychology: General*, *125*, 159–180. doi:10.1037/0096-3445.125.2.159
- Daniel, R., & Pollmann, S. (2014). A universal role of the ventral striatum in reward-based learning: Evidence from human studies. *Neurobiology of Learning and Memory*, *114*, 90–100. doi:10.1016/j.nlm.2014.05.002
- Daw, N. D., & Doya, K. (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, *16*, 199–204. doi:10.1016/j.conb.2006.03.006
- Delgado, M. R. (2007). Reward-related responses in the human striatum. *Annals of the New York Academy of Sciences*, *1104*, 70–88. doi:10.1196/annals.1390.002
- Delgado, M. R., Miller, M. M., Inati, S., & Phelps, E. A. (2005). An fMRI study of reward-related probability learning. *NeuroImage*, *24*, 862–873. doi:10.1016/j.neuroimage.2004.10.002
- Delgado, M. R., Nystrom, L. E., Fissell, C., Noll, D. C., & Fiez, J. A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of Neurophysiology*, *84*, 3072–3077.
- Dickerson, K. C., Li, J., & Delgado, M. R. (2011). Parallel contributions of distinct human memory systems during probabilistic learning. *NeuroImage*, *55*, 266–276. doi:10.1016/j.neuroimage.2010.10.080
- Dobryakova, E., & Tricomi, E. (2013). Basal ganglia engagement during feedback processing after a substantial delay. *Cognitive, Affective, & Behavioral Neuroscience*, *13*, 725–736. doi:10.3758/s13415-013-0182-6

- Eichenbaum, H., & Bunsey, M. (1995). On the binding of associations in memory: Clues from studies on the role of the hippocampal region in paired-associate learning. *Current Directions in Psychological Science*, 4, 19–23. doi:10.1111/1467-8721.ep10770954
- Foerde, K., Braun, E. K., & Shohamy, D. (2013a). A trade-off between feedback-based learning and episodic memory for feedback events: Evidence from Parkinson's disease. *Neurodegenerative Diseases*, 11, 93–101. doi:10.1159/000342000
- Foerde, K., Knowlton, B. J., & Poldrack, R. A. (2006). Modulation of competing memory systems by distraction. *Proceedings of the National Academy of Sciences*, 103, 11778–11783. doi:10.1073/pnas.0602659103
- Foerde, K., Poldrack, R. A., & Knowlton, B. J. (2007). Secondary-task effects on classification learning. *Memory & Cognition*, 35, 864–874.
- Foerde, K., Race, E., Verfaellie, M., & Shohamy, D. (2013b). A role for the medial temporal lobe in feedback-driven learning: Evidence from amnesia. *Journal of Neuroscience*, 33, 5698–5704. doi:10.1523/JNEUROSCI.5217-12.2013
- Foerde, K., & Shohamy, D. (2011a). Feedback timing modulates brain systems for learning in humans. *Journal of Neuroscience*, 31, 13157–13167. doi:10.1523/JNEUROSCI.2701-11.2011
- Foerde, K., & Shohamy, D. (2011b). The role of the basal ganglia in learning and memory: Insight from Parkinson's disease. *Neurobiology of Learning and Memory*, 96, 624–636. doi:10.1016/j.nlm.2011.08.006
- Frensch, P. A., Lin, J., & Buchner, A. (1998). Learning versus behavioral expression of the learned: The effects of a secondary tone-counting task on implicit learning in the serial reaction task. *Psychological Research*, 61, 83–98. doi:10.1007/s004260050015
- Gluck, M. A., & Myers, C. E. (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus*, 3, 491–516. doi:10.1002/hipo.450030410
- Gluck, M. A., Meeter, M., & Myers, C. E. (2003). Computational models of the hippocampal region: linking incremental learning and episodic memory. *Trends in Cognitive Sciences*, 7, 269–276. doi:10.1016/S1364-6613(03)00105-0
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6, 65–70. Retrieved from www.jstor.org/stable/4615733
- Jahanshahi, M., Wilkinson, L., Gahir, H., Dharminda, A., & Lagnado, D. A. (2010). Medication impairs probabilistic classification learning in Parkinson's disease. *Neuropsychologia*, 48, 1096–1103. doi:10.1016/j.neuropsychologia.2009.12.010
- Jiménez, L., & Vázquez, G. A. (2005). Sequence learning under dual-task conditions: Alternatives to a resource-based account. *Psychological Research*, 69, 352–368. doi:10.1007/s00426-004-0210-9
- Kahneman, D., Treisman, A., & Burkell, J. (1983). The cost of visual filtering. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 510–522. doi:10.1037/0096-1523.9.4.510
- Klingberg, T. (2000). Limitations in information processing in the human brain: Neuroimaging of dual task performance and working memory tasks. In *Progress in brain research* (Vol. 126, pp.95–102). Amsterdam, The Netherlands: Elsevier. doi:10.1016/S0079-6123(00)26009-3
- Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, 273, 1399–1402.
- Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning and Memory*, 1, 106–120.
- Lee, A. S., Duman, R. S., & Pittenger, C. (2008). A double dissociation revealing bidirectional competition between striatum and hippocampus during learning. *Proceedings of the National Academy of Sciences*, 105, 17163–17168. doi:10.1073/pnas.0807749105
- Li, J., Delgado, M. R., & Phelps, E. A. (2011). How instructed knowledge modulates the neural systems of reward learning. *Proceedings of the National Academy of Sciences*, 108, 55–60. doi:10.1073/pnas.1014938108
- Longstaffe, K. A., Hood, B. M., & Gilchrist, I. D. (2014). The influence of cognitive load on spatial search performance. *Attention, Perception, & Psychophysics*, 76, 49–63. doi:10.3758/s13414-013-0575-1
- Macmillan, N., & Creelman, C. D. (1991). *Detection theory: A user's guide*. New York, NY: Cambridge University Press.
- Martin, A. (1999). Automatic activation of the medial temporal lobe during encoding: Lateralized influences of meaning and novelty. *Hippocampus*, 9, 62–70. doi:10.1002/(SICI)1098-1063(1999)9:1<62::AID-HIPO7>3.0.CO;2-K
- Mattfeld, A. T., & Stark, C. E. L. (2011). Striatal and medial temporal lobe functional interactions during visuomotor associative learning. *Cerebral Cortex*, 21, 647–658. doi:10.1093/cercor/bhq144
- Meng, X.-L., Rosenthal, R., & Rubin, D. B. (1992). Comparing correlated correlation coefficients. *Psychological Bulletin*, 111, 172–175. doi:10.1037/0033-2909.111.1.172
- Myers, C. E., Shohamy, D., Gluck, M. A., Grossman, S., Kluger, A., Ferris, S., . . . Schwartz, R. (2003). Dissociating hippocampal versus basal ganglia contributions to learning and transfer. *Journal of Cognitive Neuroscience*, 15, 185–193. doi:10.1162/089892903321208123
- O'Doherty, J. P. (2004). Reward representations and reward-related learning in the human brain: Insights from neuroimaging. *Current Opinion in Neurobiology*, 14, 769–776. doi:10.1016/j.conb.2004.10.016
- Okatan, M. (2009). Correlates of reward-predictive value in learning-related hippocampal neural activity. *Hippocampus*, 19, 487–506. doi:10.1002/hipo.20535
- Olmos, A., & Kingdom, F. A. (2004). A biologically inspired algorithm for the recovery of shading and reflectance images. *Perception*, 33, 1463–1473.
- Packard, M. G., & Knowlton, B. J. (2002). Learning and memory functions of the basal ganglia. *Annual Review of Neuroscience*, 25, 563–593. doi:10.1146/annurev.neuro.25.112701.142937
- Packard, M. G., Hirsh, R., & White, N. M. (1989). Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: Evidence for multiple memory systems. *Journal of Neuroscience*, 9, 1465–1472.
- Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin*, 116, 220–244. doi:10.1037/0033-2909.116.2.220
- Poldrack, R. A., Clark, J., Paré-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. *Nature*, 414, 546–550. doi:10.1038/35107080
- Poldrack, R. A., & Packard, M. G. (2003). Competition among multiple memory systems: Converging evidence from animal and human brain studies. *Neuropsychologia*, 41, 245–251.
- Poldrack, R. A., Sabb, F. W., Foerde, K., Tom, S. M., Asanow, R. F., Bookheimer, S. Y., & Knowlton, B. J. (2005). The neural correlates of motor skill automaticity. *Journal of Neuroscience*, 25, 5356–5364. doi:10.1523/JNEUROSCI.3880-04.2005
- Rice, W. R. (1989). Analyzing tables of statistical tests. *Evolution*, 43, 223–225.
- Schacter, D. L., & Wagner, A. D. (1999). Medial temporal lobe activations in fMRI and PET studies of episodic encoding and retrieval. *Hippocampus*, 9, 7–24. doi:10.1002/(SICI)1098-1063(1999)9:1<7::AID-HIPO2>3.0.CO;2-K
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80, 1–27.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275, 1593–1599.

- Scimeca, J. M., & Badre, D. (2012). Striatal contributions to declarative memory retrieval. *Neuron*, *75*, 380–392. doi:10.1016/j.neuron.2012.07.014
- Shanks, D. R., & Channon, S. (2002). Effects of a secondary task on “implicit” sequence learning: Learning or performance? *Psychological Research*, *66*, 99–109.
- Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Sciences*, *14*, 464–472. doi:10.1016/j.tics.2010.08.002
- Shohamy, D., Myers, C. E., Grossman, S., Sage, J., Gluck, M. A., & Poldrack, R. A. (2004a). Cortico-striatal contributions to feedback-based learning: Converging data from neuroimaging and neuropsychology. *Brain*, *127*, 851–859. doi:10.1093/brain/awh100
- Shohamy, D., Myers, C. E., Onlaor, S., & Gluck, M. A. (2004b). Role of the basal ganglia in category learning: How do patients with Parkinson’s disease learn? *Behavioral Neuroscience*, *118*, 676–686. doi:10.1037/0735-7044.118.4.676
- Shohamy, D., & Wagner, A. D. (2008). Integrating memories in the human brain: Hippocampal–midbrain encoding of overlapping events. *Neuron*, *60*, 378–389. doi:10.1016/j.neuron.2008.09.023
- Squire, L. R. (2004). Memory systems of the brain: a brief history and current perspective. *Neurobiology of Learning and Memory*, *82*, 171–177. doi:10.1016/j.nlm.2004.06.005
- Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional system. An approach to cerebral imaging*. Stuttgart, Germany: Thieme.
- Treisman, A., Kahneman, D., & Burkell, J. (1983). Perceptual objects and the cost of filtering. *Perception & Psychophysics*, *33*, 527–532. doi:10.3758/BF03202934
- Tricomi, E. M., Delgado, M. R., & Fiez, J. A. (2004). Modulation of caudate activity by action contingency. *Neuron*, *41*, 281–292. doi:10.1016/S0896-6273(03)00848-1
- Waldron, E. M., & Ashby, F. G. (2001). The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychonomic Bulletin & Review*, *8*, 168–176.
- White, N. M., & McDonald, R. J. (2002). Multiple parallel memory systems in the brain of the rat. *Neurobiology of Learning and Memory*, *77*, 125–184.
- White, N. M., Packard, M. G., & McDonald, R. J. (2013). Dissociation of memory systems: The story unfolds. *Behavioral Neuroscience*, *127*, 813–834. doi:10.1037/a0034859
- Wilkinson, L., Lagnado, D. A., Quallo, M., & Jahanshahi, M. (2008). The effect of feedback on non-motor probabilistic classification learning in Parkinson’s disease. *Neuropsychologia*, *46*, 2683–2695. doi:10.1016/j.neuropsychologia.2008.05.008
- Wilkinson, L., Tai, Y. F., Lin, C. S., Lagnado, D. A., Brooks, D. J., Piccini, P., & Jahanshahi, M. (2014). Probabilistic classification learning with corrective feedback is associated with in vivo striatal dopamine release in the ventral striatum, while learning without feedback is not. *Human Brain Mapping*, *35*, 5106–5115. doi:10.1002/hbm.22536
- Wimmer, G. E., & Shohamy, D. (2012). Preference by association: How memory mechanisms in the hippocampus bias decisions. *Science*, *338*, 270–273. doi:10.1126/science.1223252
- Wimmer, G. E., Braun, E. K., Daw, N. D., & Shohamy, D. (2014). Episodic memory encoding interferes with reward learning and decreases striatal prediction errors. *Journal of Neuroscience*, *34*, 14901–14912. doi:10.1523/JNEUROSCI.0204-14.2014