nature neuroscience

Frequency-specific hippocampal-prefrontal interactions during associative learning

Scott L Brincat[1,2] & Earl K Miller[1,2]

Much of our knowledge of the world depends on learning associations (for example, face-name), for which the hippocampus (HPC) and prefrontal cortex (PFC) are critical. HPC-PFC interactions have rarely been studied in monkeys, whose cognitive and mnemonic abilities are akin to those of humans. We found functional differences and frequency-specific interactions between HPC and PFC of monkeys learning object pair associations, an animal model of human explicit memory. PFC spiking activity reflected learning in parallel with behavioral performance, whereas HPC neurons reflected feedback about whether trial-and-error guesses were correct or incorrect. Theta-band HPC-PFC synchrony was stronger after errors, was driven primarily by PFC to HPC directional influences and decreased with learning. In contrast, alpha/beta-band synchrony was stronger after correct trials, was driven more by HPC and increased with learning. Rapid object associative learning may occur in PFC, whereas HPC may guide neocortical plasticity by signaling success or failure via oscillatory synchrony in different frequency bands.

Most neurophysiological studies of PFC-HPC interactions have examined spatial memory in rodents. It seems clear, especially in primates, that the HPC and PFC have broader roles, including non-spatial explicit (declarative) memory. HPC damage causes deficits in non-spatial associative learning if implicit memory (familiarity, priming) cannot be used[1–3]. Likewise, PFC damage impairs explicit non-spatial associative memories, sparing implicit memory[4,5]. Human imaging has shown activation of both areas during associative memory[6,7]. In rodents, there is theta synchrony between the PFC and HPC during spatial memory performance[8,9] and high-frequency ripple synchrony during subsequent sleep[10], which is thought to reflect the HPC acquiring spatial information and then integrating it into cortical networks for long-term storage. A similar relationship is assumed for non-spatial memories[11], but this has not been tested in primates. We found that, in monkeys, non-spatial associations are instead acquired by the PFC. In contrast, HPC activity was consistent with the idea that it provides learning-related frequency-specific feedback to the PFC.

RESULTS

Paired associate learning task and behavioral results

We trained two adult rhesus monkeys (*Macaca mulatta*) to perform an object paired-associate learning task. The monkeys learned, through trial and error, four novel associations between arbitrary pairs of objects in each experimental session (**Fig. 1a**,**b**). Each of four cue objects (**Fig. 1a**) was randomly paired with one of two associate objects (A_1 or A_2). This 4-to-2 mapping encourages prospective recall of the associate[12] and distinguished neural activity to the cue from retrieval of its associate. Monkeys routinely learned associations in a few hundred trials (313 of 348 associations learned to criterion; **Fig. 1c**), but measures of motivation, arousal and motor function changed little with learning (**Supplementary Fig. 1**). Multiple microelectrodes were lowered daily into lateral prefrontal cortex and hippocampus (**Fig. 2**), and each recorded spiking and LFP signals while the monkeys performed the paired-associate learning task.

PFC and HPC neurons reflect learned associations and outcomes

With learning, PFC neurons increasingly showed activity after the cue that anticipated its paired associate, with an across-trial progression similar to the improvement in performance (Spearman's $\rho = 0.59$ with cue-epoch spike rate, $P = 0.04$, two-sided permutation test; **Fig. 3a** and **Supplementary Fig. 2a**,**b**). However, although HPC neuronal activity conveyed sensory signals reflecting the cue object (**Supplementary Fig. 3**), it did not reflect learning of the paired associate and showed no correlation with behavior ($\rho = -0.21$, $P = 0.73$). Instead, HPC reflected the trial outcome after the feedback (reward versus no-reward) about whether the behavioral response was correct or incorrect. This effect was stronger in the HPC than the PFC ($P = 0.049$, two-way area × learning-stage permutation ANOVA; **Fig. 3b** and **Supplementary Fig. 2c**,**d**), and was stronger in the HPC in the output subregions (CA1, subiculum) than locally projecting subregions (CA3, dentate gyrus, $P \leq 10^{-4}$; **Supplementary Fig. 4a**). With learning, HPC activity shifted from stronger activation after incorrect to correct outcomes ($P = 0.027$, two-sided permutation test on early versus late learning stages (first versus last third of learning trials), with number of incorrect versus correct trials matched across learning; **Fig. 3c** and Online Methods). This shift was present in HPC output subregions ($P = 0.005$; **Supplementary Fig. 4b**), but not local-projection subregions ($P = 0.83$; $P = 0.04$, subregion × learning-stage interaction), which may be related to a corresponding shift in communication between the HPC and PFC described below.

[1]The Picower Institute for Learning and Memory, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. [2]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. Correspondence should be addressed to E.K.M. (ekmiller@mit.edu).

Received 6 October 2014; accepted 22 January 2015; published online 23 February 2015; doi:10.1038/nn.3954

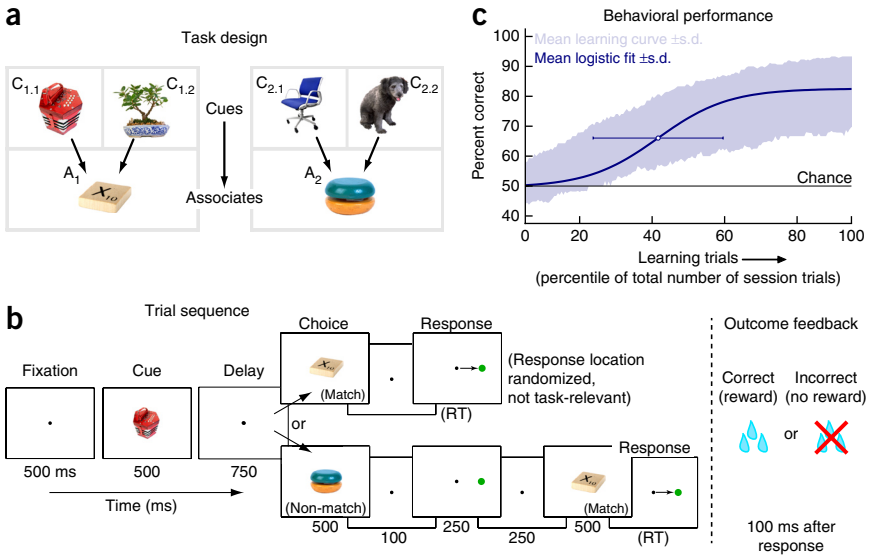

Figure 1 Paired-associate learning task. (**a**) Task design. Each session, four objects were designated as cue objects; each was arbitrarily paired with one of two associate objects. (**b**) Task trial sequence. After central fixation, a cue object was followed by a short delay and a choice object. If it was that cue's paired associate, monkeys had to saccade to a target (whose varied location was not task relevant); otherwise, they were required to withhold response through another delay until the correct associate was presented. Correct choices were rewarded with juice; incorrect choices resulted in no reward and were signaled by a red error screen and 3-s 'time-out'. Task period durations are given in milliseconds below panels (RT, reaction time). (**c**) Learning performance. Shaded area indicates mean ± s.d. of percent correct performance across all 348 associations (87 sessions), plotted as a function of the percentile of each session's trials (sessions averaged 1,117 ± 125 s.d. trials). Blue curve indicates the average sigmoidal learning curve fit to each association. White dot indicates mean ± s.d. of fit curve centers.

The learning-related change in HPC outcome bias could reflect either a sign-flip in the preference of individual neurons from incorrect to correct trials or just a relative modulation of neurons whose outcome preference is consistent throughout learning (broadly analogous to the distinction between 'global remapping' and 'rate remapping' in rodent hippocampal place cells[13]). To distinguish between these possibilities, we trained a linear (logistic regression) classifier to discriminate correct versus incorrect trials based on HPC activity early in learning and asked whether the trained classifier weights transferred to predict trial outcome late in learning. The preference-flip model predicts that early learning–derived weights and late-learning activity will have largely opposing outcome preferences, and will therefore produce prediction accuracy near (or below) chance. We instead found that an early learning–trained classifier predicted the outcome on 79% of late-learning trials, comparable to 89% for a classifier both trained and tested on (distinct) late-learning trials. This provides evidence that the observed shift with learning reflects modulation of a largely invariant hippocampal neural code for trial outcome (**Supplementary Fig. 2d**).

Band-specific synchrony reflects trial outcome and learning

We examined outcome-related neural communication using synchrony (phase-locking) between local field potentials (LFPs; **Fig. 3a**) recorded after the behavioral response and feedback. This revealed PFC-HPC synchrony in two frequency bands: a shorter latency theta-band (~2–6 Hz) synchrony and longer latency alpha/low-beta band (~9–16 Hz) synchrony. Alpha/beta synchrony was stronger after correct trials; theta synchrony was stronger after incorrect trials (**Fig. 4a–c**). Although PFC-HPC synchrony was present before the behavioral

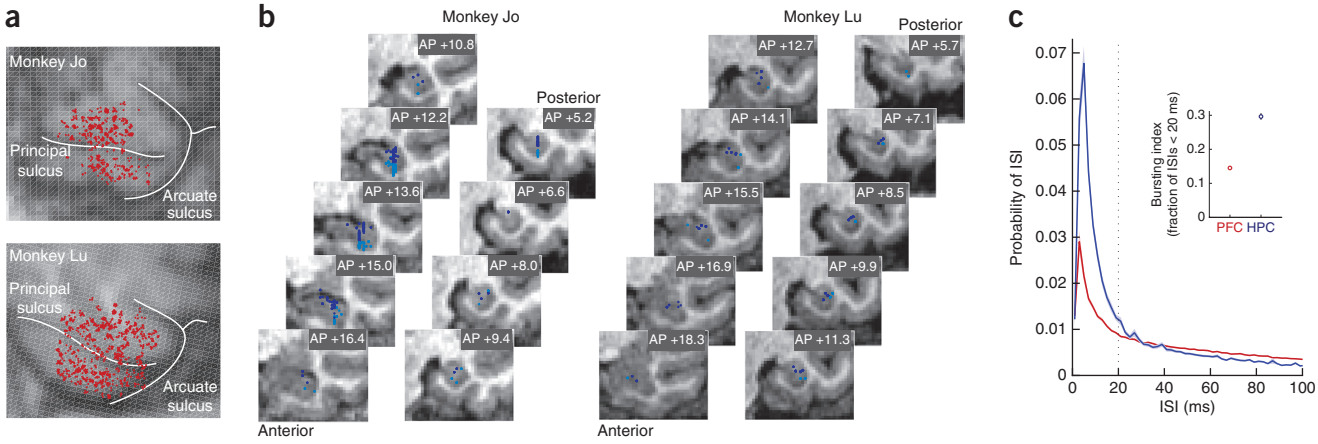

Figure 2 Prefrontal and hippocampal recording locations. (**a**) Prefrontal recording sites in two monkeys, Jo and Lu, projected onto a tangential magnetic resonance imaging (MRI) slice. Each dot represents a location at which neurons were recorded. Recordings spanned a broad region of dorsolateral and ventrolateral PFC (including parts of areas 46, 45 and 8). (**b**) Hippocampal recording sites, plotted on an anterior-to-posterior series of coronal MRI slices (labeled with millimeters relative to interaural line along the anterior-posterior axis). Dark blue indicates locally projecting subregions (dentate gyrus, CA3, CA2) and light blue represents output subregions (CA1, subiculum). Recordings spanned all subregions of the anterior ~3/4 of the hippocampal formation. (**c**) Hippocampal neurons exhibited characteristic bursty firing. Population mean (±s.e.m.) interspike interval (ISI) histograms showing the probability of each ISI averaged across all PFC (red) and HPC (blue) putative principal neurons (putative fast-spiking interneurons were discriminated on the basis of spike waveform and firing rate and excluded from this plot). Inset, population mean (±s.e.m.) spike bursting index[39]—the fraction of each neuron's ISIs < 20 ms. Both plots illustrate the distinctive burstiness of hippocampal neurons in comparison to neocortical neurons.

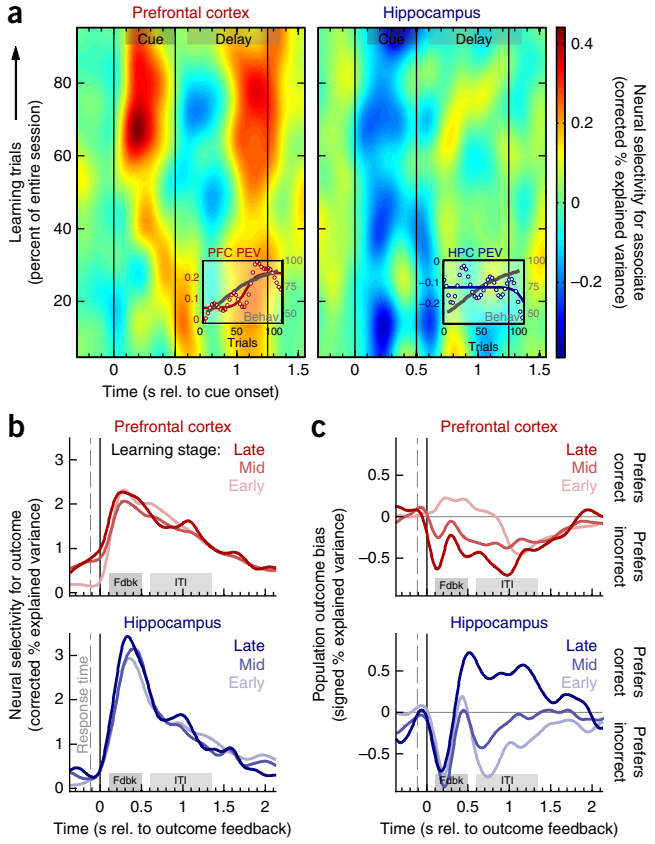

Figure 3 Prefrontal neurons reflect learned associations and hippocampal neurons reflect trial outcome. (**a**) Mean percent of variance explained (PEV) by learned associate objects in PFC (left, $n = 319$ neurons) and HPC (right, $n = 199$) spiking activity, plotted across time after cue onset and learning trials. Bias correction resulted in negative values for some trials and times at which values were less than expected based on selectivity for random combinations of cue objects (Online Methods). Gray bars represent analytical epochs focusing on cue and delay periods. Insets show behavioral (gray) and neural "learning curves"—mean cue-epoch PEV across trials. Only PFC showed learning of associates in parallel with behavior. (**b**) Mean percent of variance in PFC (top) and HPC (bottom) neurons explained by trial outcome (correct versus incorrect), plotted across time after outcome feedback (reward versus no-reward) for early, middle and late learning stages (light-to-dark colors). Gray bars represent analytical epochs focusing on transient responses to outcome feedback and sustained activity during the inter-trial interval. Outcome was represented more strongly in HPC in the outcome feedback epoch. (**c**) Mean bias (signed PEV, Online Methods) in PFC (top) and HPC (bottom) neurons for correct (positive values) versus incorrect (negative) outcomes as a function of time and learning stages. HPC shifted from incorrect to correct outcomes with learning. Although there was a significant area × learning-stage interaction ($P = 0.03$), PFC showed no significant change with learning ($P = 0.3$).

response, it did not robustly reflect trial outcome (**Supplementary Fig. 5**). Theta synchrony following incorrect outcomes decreased with learning ($P \leq 10^{-4}$; **Fig. 4d,e**), whereas alpha/beta synchrony following correct outcomes increased with learning ($P = 5 \times 10^{-4}$, two-sided permutation test on early versus late learning). Although the theta effect was similar across HPC subregions, the alpha/beta increase with learning only occurred for synchrony between PFC and HPC output subregions (**Supplementary Fig. 6**). Thus, with learning, there was a shift in PFC-HPC synchrony from theta toward higher frequencies, paralleling the shift in HPC spiking activity from incorrect to correct trials (see above).

We also examined within-area LFP phase-locking and power. Although within-PFC synchrony followed a similar pattern as between-area synchrony (**Supplementary Fig. 7**), intra-hippocampal

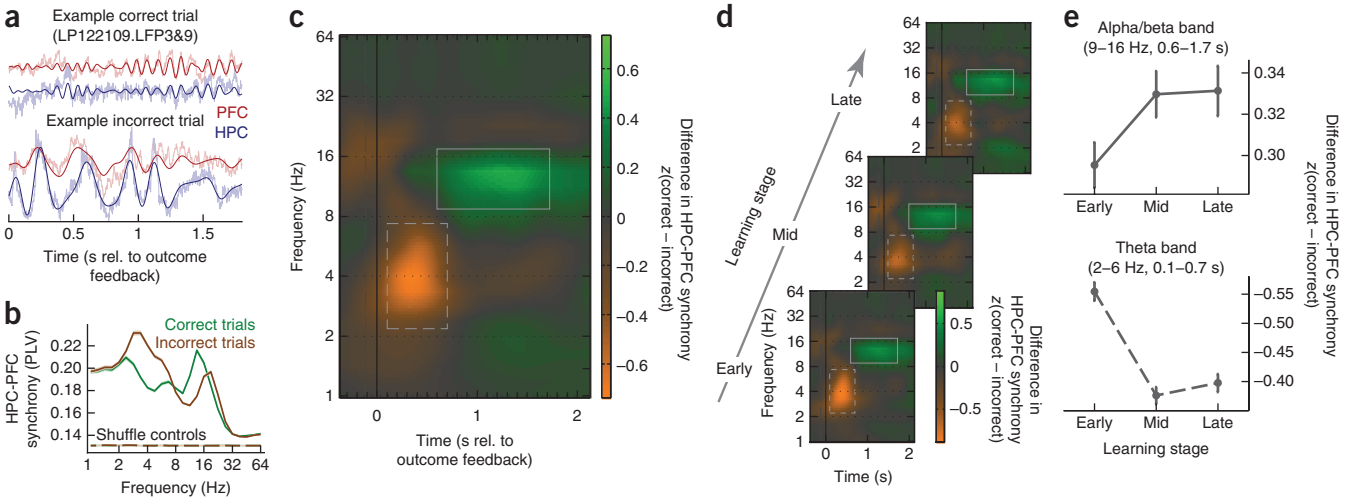

Figure 4 Hippocampal-prefrontal oscillatory synchrony carries learning-related information about trial outcome. (**a**) Example LFPs from a pair of sites in HPC (blue) and PFC (red) following a correct (top) and an incorrect (bottom) trial. Lighter colors represent raw LFPs and darker colors represent LFPs filtered in the 9–16-Hz (top) and 2–6-Hz (bottom) bands. (**b**) Mean synchrony (±s.e.m.) between HPC and PFC LFPs, plotted as a function of frequency, following correct (green) and incorrect (brown) outcomes. Synchrony was computed as the across-trial phase-locking value (PLV), calculated in the feedback and ITI epochs (100–1,725 ms after outcome feedback), and was averaged across all 970 electrode pairs and sessions. Dashed curves represent mean synchrony (±s.e.m.) expected by chance (based on shuffling HPC and PFC signals across trials), which was nearly identical across trial outcome and frequency. (**c**) Mean *z*-scored difference in HPC-PFC synchrony (dPLV) between correct and incorrect trials, plotted as spectrograms across time and frequency. Theta-band (dashed rectangle) and alpha/beta-band (solid rectangle) synchrony were stronger for incorrect and correct outcomes, respectively. (**d**) Mean dPLV as a function of learning (bottom to top, early, middle and late learning stages). (**e**) Summary of synchrony learning effects: mean dPLV (±s.e.m.) pooled in the alpha/beta-band (top) and theta-band (bottom, note that higher values in this plot reflect stronger negative PLV differences) regions of interest as a function of learning stage. Theta (incorrect) synchrony decreased with learning, whereas alpha/beta (correct) increased.

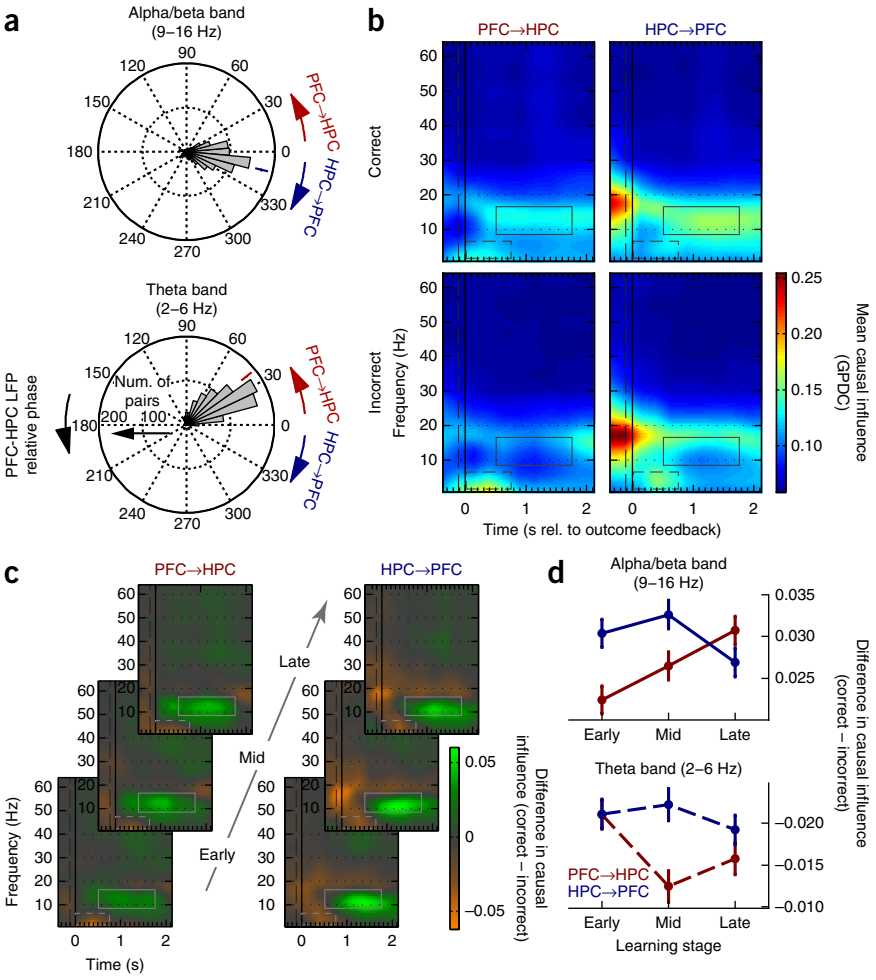

Figure 5 Stronger theta PFC→HPC directional influence; stronger HPC→PFC alpha/beta influence. (**a**) Angular histograms of mean PFC-HPC LFP phase lag (n = 970 electrode pairs), pooled in alpha/beta-band (top) and theta-band (bottom) regions of interest. Colored tick marks indicate mean across all pairs. HPC led for alpha/beta frequencies, whereas PFC led for theta. (**b**) Frequency-domain directional influences (GPDC) from PFC to HPC (left, n = 970 electrode pairs) and from HPC to PFC (right, n = 970), following correct (top) and incorrect (bottom) trials. GPDC was plotted as spectrograms across time and frequency. Overall, alpha/beta-band influences (solid rectangles) were stronger from HPC to PFC and for correct trials, whereas theta-band influences (dashed rectangles) were stronger from PFC to HPC and for incorrect trials. (**c**) Cross-area directional influences across learning. Selectivity for trial outcome was quantified by the difference in directional strength (GPDC) between correct and incorrect outcomes, and is plotted separately for each direction (PFC→HPC, left; HPC→PFC, right) and learning stage (bottom to top). (**d**) Summary of learning effects: mean difference (correct – incorrect) in causal influence (±s.e.m.) pooled in the alpha/beta-band (top) and theta-band (bottom) regions of interest, as a function of learning stage. With learning, theta interactions showed a decreasing trend, whereas alpha/beta interactions shifted from a HPC→PFC to PFC→HPC directionality.

synchrony exhibited a distinct pattern in which theta synchrony increased, rather than decreased, with learning. This indicates the observed learning-related synchrony changes do not simply reflect state changes with global effects. LFP power, reflecting local synchrony, exhibited a pattern broadly similar to cross-area synchrony (**Supplementary Fig. 8**), as would be expected from an interacting system with causal links between local and long-range synchrony. Outcome selectivity in cross-area synchrony could not, however, be fully attributed to local power differences, as it remained significant even when band-specific power was balanced across correct and incorrect trials ($P \leq 10^{-4}$, for both frequency bands, one-sample bootstrap test; **Supplementary Fig. 9a**).

Band-specific directionality of PFC-HPC causal influence

Theta and alpha/beta synchrony differed in the direction of putative causal influence. For theta frequencies, the phase of HPC LFPs lagged behind PFC (mean rel. phase = 39°, $P \leq 10^{-4}$, bootstrap test versus zero phase lag; **Fig. 5a**), consistent with a PFC to HPC directionality; the reverse was true for alpha/beta frequencies (−13.5°, $P \leq 10^{-4}$). We confirmed this using generalized partial directed coherence (GPDC), a frequency-domain analog of Granger causality that measures the degree to which signals can predict each other's future values. It also revealed oscillatory interactions in the theta and alpha/beta bands (**Fig. 5b**), with stronger theta influence (dashed lines) from PFC to HPC (P = 0.01) and stronger alpha/beta influence (solid lines) from the HPC to PFC ($P \leq 10^{-4}$, direction factor in two-way causal direction × trial outcome permutation ANOVA). These differences remained significant when band-specific LFP power was balanced across correct and incorrect trials ($P \leq 10^{-4}$ for both frequency bands, direction factor in two-way causal direction × trial outcome permutation ANOVA; **Supplementary Fig. 9b**), indicating that the observed directionality cannot be explained by differences in local power. As above, theta and alpha/beta interactions were stronger for incorrect and correct trials, respectively ($P \leq 10^{-4}$ for both; **Fig. 5b**). With learning, there were significant decreases in incorrect-reflecting PFC to HPC theta influences (P = 0.021; **Fig. 5c,d**) and correct-reflecting HPC to PFC alpha/beta influences (P = 0.04, two-sided permutation test on early versus late learning), suggesting these interactions may be most important during the early stages of learning. In contrast, initially weak PFC to HPC alpha/beta influences reflecting correct outcomes increased with learning ($P \leq 10^{-4}$), eventually becoming even stronger than the HPC to PFC direction ($P = 10^{-3}$, interaction in two-way causal-direction × learning-stage permutation ANOVA).

DISCUSSION

These results suggest different roles and interactions between the PFC and HPC during object associative learning. Only PFC neurons showed neural correlates of learning the paired associates. The HPC was more engaged when feedback was given about whether the trial was correct or incorrect. Early in learning, incorrect outcomes activated HPC neurons and promoted cross-area theta synchrony with a stronger influence from the PFC to the HPC. Correct outcomes, in contrast, promoted alpha/beta-band synchrony that was initially stronger in the HPC to PFC direction. However, as learning progressed, correct outcomes increasingly evoked PFC to HPC alpha/beta-band influences and HPC neuronal spiking. This shift in

HPC outcome coding (and other properties; **Supplementary Fig. 10**) distinguishes it from unipolar positive and negative reward-prediction-error signals in the midbrain dopaminergic nuclei[14] and lateral habenula[15], respectively. It may, however, reflect a functional shift in the importance of negative and positive feedback. Early on in trial-and-error learning, errors provide critical information about which pairs of objects should not be associated. But once associations are learned, errors are more likely to simply reflect lapses in response inhibition or attention rather than true errors of associative choice. In contrast, positive feedback following correct trials[16], and the likely resulting dopamine release in hippocampus[17], have been shown to preferentially enhance long-term consolidation of new learning. Thus, the shift in HPC bias from incorrect to correct outcomes as learning progresses is consistent with a transition from neural signals that support acquisition to those that promote consolidation.

The HPC is critical for formation of explicit memories. Rodent neurophysiological studies suggest that it acquires spatial memories and consolidates them in the neocortex, including PFC[8,10]. The primate HPC shows rapid activity changes related to spatial associative learning[18]. But the HPC is also known to be critical for non-spatial memory in rodents[11,19] and especially in primates, where it has a general role in explicit memory formation[1,3]. Lesion studies have suggested that perirhinal cortex—part of the medial temporal lobe system[3] that includes the HPC—may be more critical for object associative learning than the HPC[19,20], and neural correlates of object associations have been seen in perirhinal, prefrontal and inferotemporal cortex[12,21–23]. However, these studies examined associations that were familiar or learned gradually (over days or weeks), situations known to favor neocortical representation. Our results suggest that rapid acquisition of object associations also occurs in the neocortex, not the HPC, perhaps in the PFC in particular given its importance for behavioral flexibility. Object associations may lack the context required for explicit HPC representation[24].

Both the HPC and PFC signal trial outcome, differentiating between correct and error trials[25,26]. Our results suggest that this information is communicated between HPC and PFC via synchrony at different frequencies: theta for incorrect and alpha/beta for correct. Human and animal studies suggest that oscillatory activity is associated with memory encoding and retrieval[27–29], as well as other cognitive processes[30–32]. Higher frequency (gamma) oscillations are thought to underlie the transient formation of local neuronal ensembles, whereas lower frequencies may recruit larger networks resulting from their longer integration times[31,33,34]. Thus, the lower frequency (theta and alpha/beta) synchrony we observed may reflect formation of larger networks connecting PFC, HPC, and likely many other cortical and subcortical structures. Further experimentation will be necessary to delineate the extent of these networks and dissect out how each of their nodes function in learning. Human EEG also shows theta oscillations with a frontal source reflecting conflict or error[35]; our results suggest that these oscillations are propagated to hippocampus during learning. We did not, however, observe the sustained bouts of theta oscillations typically seen in locomoting rodents[8,9,36]; it remains an open question whether our theta-band synchrony reflects a distinct or related phenomenon.

But what computational role might these particular frequencies have? Growing evidence suggests that beta oscillations are ideal for maintaining active cell assemblies and their associated cognitive states[37]. This hypothesis is consistent with the idea that alpha/beta oscillations might have a role in maintaining neural representations active during correct associations. Studies of synaptic plasticity have also shown that low-frequency synaptic stimulation fosters long-term depression, whereas high-frequency stimulation fosters long-term potentiation[38], with the crossover point at ~8–10 Hz. PFC-HPC theta interactions may therefore have weakened synapses active during incorrect associations, whereas alpha/beta interactions strengthened those active for the correct associations.

In sum, these observations show that rapid formation of non-spatial associations may occur in the PFC, not the HPC. The main role of the HPC appears to be signaling trial outcome, signals which are communicated with PFC via band-specific oscillatory synchrony and may be involved in guiding neocortical learning. The results also provide further support for the idea that synchrony in different frequency bands may have functionally different roles in neural communication[30,34].

METHODS

Methods and any associated references are available in the online version of the paper.

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

We thank E. Antzoulatos, A. Bastos, J. Donoghue, N. Kopell, S. Kornblith, R. Loonis, M. Lundqvist, M. Moazami, V. Puig, J. Rose, J. Roy, A. Salazar-Gómez, L. Tran and M. Wilson for helpful comments and suggestions, and D. Altschul, B. Gray, M. Histed, D. Ouellette and the MIT veterinary staff for technical assistance. This work was supported by NIMH Conte Center grant P50-MH094263-03 (E.K.M.), NIMH fellowship F32-MH081507 (S.L.B.), and The Picower Foundation.

AUTHOR CONTRIBUTIONS

S.L.B. and E.K.M. designed the experiments. S.L.B. trained the monkeys, performed the experiments and analyzed the data. S.L.B. and E.K.M. wrote the paper.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

1. Scoville, W.B. & Milner, B. Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* **20**, 11–21 (1957).
2. Cohen, N.J. & Squire, L.R. Preserved learning and retention of pattern-analyzing skill in amnesia: dissociation of knowing how and knowing that. *Science* **210**, 207–210 (1980).
3. Squire, L.R., Stark, C.E.L. & Clark, R.E. The medial temporal lobe. *Annu. Rev. Neurosci.* **27**, 279–306 (2004).
4. Gutnikov, S.A., Ma, Y.-Y. & Gaffan, D. Temporo-frontal disconnection impairs visual-visual paired association learning but not configural learning in macaca monkeys. *Eur. J. Neurosci.* **9**, 1524–1529 (1997).
5. Farovik, A., Dupont, L.M., Arce, M. & Eichenbaum, H. Medial prefrontal cortex supports recollection, but not familiarity, in the rat. *J. Neurosci.* **28**, 13428–13434 (2008).
6. Sperling, R.A. *et al.* Encoding novel face-name associations: a functional MRI study. *Hum. Brain Mapp.* **14**, 129–139 (2001).
7. Kim, H. Neural activity that predicts subsequent memory and forgetting: a meta-analysis of 74 fMRI studies. *Neuroimage* **54**, 2446–2461 (2011).
8. Jones, M.W. & Wilson, M.A. Theta rhythms coordinate hippocampal–prefrontal interactions in a spatial memory task. *PLoS Biol.* **3**, e402 (2005).
9. Hyman, J.M., Zilli, E.A., Paley, A.M. & Hasselmo, M.E. Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. *Hippocampus* **15**, 739–749 (2005).
10. Siapas, A.G. & Wilson, M.A. Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron* **21**, 1123–1128 (1998).
11. Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M. & Tanila, H. The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron* **23**, 209–226 (1999).
12. Rainer, G., Rao, S.C. & Miller, E.K. Prospective coding for objects in primate prefrontal cortex. *J. Neurosci.* **19**, 5493–5505 (1999).
13. Colgin, L.L., Moser, E.I. & Moser, M.-B. Understanding memory through hippocampal remapping. *Trends Neurosci.* **31**, 469–477 (2008).
14. Schultz, W. Predictive reward signal of dopamine neurons. *J. Neurophysiol.* **80**, 1–27 (1998).
15. Matsumoto, M. & Hikosaka, O. Lateral habenula as a source of negative reward signals in dopamine neurons. *Nature* **447**, 1111–1115 (2007).

16. Abe, M. *et al.* Reward improves long-term retention of a motor memory through induction of offline memory gains. *Curr. Biol.* **21**, 557–562 (2011).
17. Bethus, I., Tse, D. & Morris, R.G.M. Dopamine and memory: modulation of the persistence of memory for novel hippocampal NMDA receptor–dependent paired associates. *J. Neurosci.* **30**, 1610–1618 (2010).
18. Wirth, S. *et al.* Single neurons in the monkey hippocampus and learning of new associations. *Science* **300**, 1578–1581 (2003).
19. Bunsey, M. & Eichenbaum, H. Conservation of memory function in rats and humans. *Nature* **379**, 255–257 (1996).
20. Murray, E.A., Gaffan, D. & Mishkin, M. Neural substrates of visual stimulus-stimulus association in rhesus monkeys. *J. Neurosci.* **13**, 4549–4561 (1993).
21. Sakai, K. & Miyashita, Y. Neural organization for the long-term memory of paired associates. *Nature* **354**, 152–155 (1991).
22. Erickson, C.A. & Desimone, R. Responses of macaque perirhinal neurons during and after visual stimulus association learning. *J. Neurosci.* **19**, 10404–10416 (1999).
23. Messinger, A., Squire, L.R., Zola, S.M. & Albright, T.D. Neuronal representations of stimulus associations develop in the temporal lobe during learning. *Proc. Natl. Acad. Sci. USA* **98**, 12239–12244 (2001).
24. Eichenbaum, H., Sauvage, M., Fortin, N., Komorowski, R. & Lipton, P. Towards a functional organization of episodic memory in the medial temporal lobe. *Neurosci. Biobehav. Rev.* **36**, 1597–1608 (2012).
25. Wirth, S. *et al.* Trial outcome and associative learning signals in the monkey hippocampus. *Neuron* **61**, 930–940 (2009).
26. Histed, M.H., Pasupathy, A. & Miller, E.K. Learning substrates in the primate prefrontal cortex and striatum: sustained activity related to successful actions. *Neuron* **63**, 244–253 (2009).
27. Jutras, M.J., Fries, P. & Buffalo, E.A. Gamma-band synchronization in the macaque hippocampus and memory formation. *J. Neurosci.* **29**, 12521–12531 (2009).
28. Düzel, E., Penny, W.D. & Burgess, N. Brain oscillations and memory. *Curr. Opin. Neurobiol.* **20**, 143–149 (2010).
29. Fell, J. & Axmacher, N. The role of phase synchronization in memory processes. *Nat. Rev. Neurosci.* **12**, 105–118 (2011).
30. Fries, P. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn. Sci.* **9**, 474–480 (2005).
31. Buschman, T.J. & Miller, E.K. Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* **315**, 1860–1862 (2007).
32. Siegel, M., Warden, M.R. & Miller, E.K. Phase-dependent neuronal coding of objects in short-term memory. *Proc. Natl. Acad. Sci. USA* **106**, 21341–21346 (2009).
33. Kopell, N., Ermentrout, G.B., Whittington, M.A. & Traub, R.D. Gamma rhythms and beta rhythms have different synchronization properties. *Proc. Natl. Acad. Sci. USA* **97**, 1867–1872 (2000).
34. Buzsáki, G. & Draguhn, A. Neuronal oscillations in cortical networks. *Science* **304**, 1926–1929 (2004).
35. Luu, P., Tucker, D.M. & Makeig, S. Frontal midline theta and the error-related negativity: neurophysiological mechanisms of action regulation. *Clin. Neurophysiol.* **115**, 1821–1835 (2004).
36. Buzsáki, G. & Moser, E.I. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat. Neurosci.* **16**, 130–138 (2013).
37. Engel, A.K. & Fries, P. Beta-band oscillations—signaling the status quo? *Curr. Opin. Neurobiol.* **20**, 156–165 (2010).
38. Dudek, S.M. & Bear, M.F. Homosynaptic long-term depression in area CA1 of hippocampus and effects of *N*-methyl-D-aspartate receptor blockade. *Proc. Natl. Acad. Sci. USA* **89**, 4363–4367 (1992).
39. Skaggs, W.E. *et al.* EEG Sharp waves and sparse ensemble unit activity in the macaque hippocampus. *J. Neurophysiol.* **98**, 898–910 (2007).

ONLINE METHODS

Subjects. All experiments were performed in two adult (~8–10 years old) rhesus macaques (*Macaca mulatta*), one male and one female, weighing 9 and 7.5 kg, respectively. The monkeys were pair-housed under a 12-h light/dark cycle (7 a.m. to 7 p.m. light), with experiments performed around the middle of their light cycle. Both monkeys were experimentally naive at the start of this study. Each monkey was implanted under general anesthesia with a titanium post for head restraint and two cylindrical 20-mm diameter titanium recording chambers. Chambers were stereotaxically placed over PFC and HPC in the left hemisphere based on coordinates from structural MRI scans in each monkey. All procedures followed the guidelines of the MIT Animal Care and Use Committee and the US National Institutes of Health.

Neurophysiological methods. Up to 16 microelectrodes in PFC, and up to 4 in HPC, were acutely inserted into the brain and removed at the end of each daily experiment. All recordings from PFC, and most from HPC, were performed with epoxy-coated tungsten microelectrodes (FHC). Some HPC recordings used 24-channel linear probes with 300-µm spacing between adjacent platinum-iridium recording contacts (U-Probes, Plexon). For PFC, electrodes were lowered through the dura using a custom-built screw microdrive assembly. For HPC, electrodes were inserted through a 25 gauge transdural cannula using a motorized microdrive system (NAN-S4, NAN Instruments).

Neural activity was amplified, filtered, digitized and stored using an integrated multichannel recording system (Multichannel Acquisition Processor, Plexon). The signal from each electrode was amplified by a high input–impedance headstage (HST/8o50-G1, Plexon), then separately filtered to extract spiking activity (250–8,000 Hz) and local field potentials (LFPs, 0.7–300 Hz). Both signals were referenced to ground, rather than to one of the electrodes, eliminating the possibility of artifactual synchrony due to neural signals measured by the reference itself. The spiking signal was threshold-triggered to separate neuronal spikes from background noise, and individual spike waveforms were digitized at 40 kHz and sorted offline into isolated neurons and multi-units using a combination of waveform shape and amplitude features (Offline Sorter, Plexon). To minimize any sampling bias of neural activity, we did not prescreen activity for responsiveness or task selectivity. Instead, electrodes were advanced until one or more neurons could be isolated, and then data collection began. Neurons were included in analyses only for the extent of time during which they were well isolated from background noise and other neurons. LFPs were recorded continuously at 1 kHz, and corrected offline for filtering-induced phase shifts (FPAlign Utility, Plexon)[40]. Only LFPs from electrodes recording at least one neuron (isolated or multiunit) were used for all analyses, to ensure they were in the appropriate cell layer.

Electrodes were targeted using custom MATLAB software that co-registered each monkey's implanted recording chambers and structural MRIs in stereotaxic coordinates, and resliced the MRIs along the electrode's path. The sequence of distinct neurophysiological compartments (gray matter, white matter and sulcus/ventricle) along the electrode's path to HPC was compared online to these images. HPC was also identified by its characteristic high-amplitude LFPs, low spike rates and bursty spiking patterns[39] (**Fig. 2c**). Recordings targeted dorsolateral and ventrolateral PFC (parts of areas 46, 45 and 8; **Fig. 2a**), and all subregions (dentate gyrus/CA4, CA3, CA2, CA1, and subiculum) of the anterior ~3/4 of the hippocampal formation (**Fig. 2b**). For hippocampal subregion analyses, recordings were pooled across the dentate gyrus, CA3, and CA2 (HPC input/local-processing), and across CA1 and the subiculum (HPC output). A total of 496 PFC and 270 HPC neurons (156 from locally projecting and 111 from output subregions) were recorded across all sessions.

To restrict analysis to only those sessions with successful learning, we used a learning criterion of 32 correct responses over the final 50 trials of each association ($p \approx 0.01$, binomial test). Only sessions where all four associations were learned to criterion were included in the reported analyses (61 of 87 sessions, including a total of 319 PFC and 199 HPC neurons (104 from locally projecting and 93 from output subregions)). No statistical methods were used to pre-determine sample sizes, but our sample sizes are similar to those generally employed in the field.

Behavioral task. Monkeys performed an object paired associate learning task that required them to rapidly learn arbitrary associations between pairs of objects. For each daily recording session, six objects never before seen by the monkey were chosen from an image database (Hemera Photo-Objects). Four were randomly designated as cue objects and the remaining two as associate objects, and each cue was randomly paired with an associate. The resulting many-to-one (4-to-2) mapping from cues to associates distinguished neural activity related to the cue from retrieval of its associate (see below), and encouraged prospective recall of the associate before its appearance[12]. The monkeys' task was to learn, through trial-and-error guessing, which associate was paired with each cue object.

Each trial started when the monkey acquired fixation of a white dot at the center of the stimulus screen. After a blank fixation baseline (500 ms), a cue object (foveal, 3° wide) was presented (500 ms), followed by a blank delay interval (750 ms) in which the monkey was expected to recall the paired associate object from memory. The two associate objects were then presented in a randomly-ordered series, with each presentation (500 ms) followed by a brief delay (100 ms) and a response target (250 ms, 7.5° to the left or right of fixation). The monkey was required to make a saccadic response to the target immediately following the correct paired associate for the given cue. Response to the correct associate was rewarded with juice and a short wait (3 s) until the subsequent trial. Response to the incorrect associate was punished by withholding reward, displaying a red "error screen" as secondary reinforcement (1.5 s), and a longer wait (6 s) for the subsequent trial. The location (left versus right) of the response target after each associate was randomized and unrelated to task performance, so a specific motor plan could not be formed until the target was shown, and striatal-mediated procedural (stimulus-response) learning[2] could not be used to correctly perform the task. Saccadic responses were used because their stereotypy obviates the possibility that changes in motor performance might be confounded with associative learning (**Supplementary Fig. 1c**).

Each session began with a short block of 36 trials where the cue and associate objects for that day were passively presented to the monkey under fixation control (3 objects per trial at 500 ms each, 750-ms blank inter-stimulus interval), and a block of 96 identity match-to-sample trials in which each object was matched to itself, rather than to an arbitrary associate. These trials familiarized the monkeys with the stimuli, and eliminated any contribution of novelty-based or familiarity-based memory processes[41–43] to our results.

Eye movements and pupil size were monitored and recorded at 1 kHz using an infrared eye tracking system (EyeLink II, SR Research). Fixation was required to be maintained within a 1.5° window around the fixation dot through the entire trial, until the response period; fixation breaks terminated the trial without reward. Behavioral monitoring and visual stimulus presentation were handled by the NIH CORTEX real-time control system, and displayed on a CRT monitor with a 100-Hz refresh rate.

Data analysis. For all behavioral and neural analyses, only trials where the monkey made a valid response to the correct or incorrect associate were analyzed, thus excluding trials where the monkey broke fixation or failed to respond.

Changes in behavior and neural activity across learning were measured by performing analyses independently in sliding trial windows, each defined by a percentile of the total number of session trials to normalize for differences in session length. Analyses focused on contrasts between (a) the recalled associate objects (signals reflecting associative learning), and (b) correct and incorrect trial outcomes (signals likely used to guide learning). For (a), neural signals were temporally aligned to retrieval cue onset, and session trials were finely sampled via sliding trial windows (**Fig. 3a**). For (b), neural signals were aligned to the outcome feedback (reward versus no-reward) on each trial. Because the proportion of correct and incorrect trials changes with learning, we balanced their trial numbers for all analyses of trial outcome, which necessitated coarser sampling of session trials (**Figs. 3b**,**c** and **5**). First, for a given session, we found the smallest number of correct or incorrect trials in any trial window (usually incorrect trials late in learning). We then randomly sampled only this number of trials from each outcome (correct or incorrect) within each trial window, calculated the statistic of interest for this subsample of trials, and averaged the resulting statistic across several random samplings to obtain a robust estimate. Note that correct and incorrect trial outcomes entail a number of perceptual differences that might affect neural activity (for example, auditory click of reward solenoid, tactile/gustatory responses to juice reward). These distinctions remain constant across learning trials, however, so we focus our analysis on changes in activity with learning.

Statistics. Significance testing was conducted using random resampling methods that make no assumptions about the data distribution[44]. To test hypotheses that a statistic is distinct from a specified value (as in a one-sample *t* test), we used a bootstrap where the distribution of the statistic is estimated empirically by recalculating it repeatedly from random resamples-with-replacement of the observed data. The *P* value is the proportion of resampled values less than the specified comparison value (for example, zero). To test hypotheses that a statistic takes different values for different groups (as in a one-way ANOVA), we used permutation tests in which the null distribution reflecting no actual difference is estimated by repeatedly shuffling data values between groups. The *p* value is the proportion of shuffled values larger than the actual observed value. For comparisons involving two factors and their interaction (as in a two-way ANOVA), data values were shuffled across observations while maintaining the set of factor labels for each observation, and therefore any correlations between factors. The *p* value for each factor and interaction was calculated as above. At least 10,000 iterations were performed for each test.

Code availability. All analyses were performed using custom code written in MATLAB (MathWorks), available from the authors upon request. In addition, as mentioned below, some analyses used functions available in open-source MATLAB code: the Torrence & Compo wavelet[45] (http://paos.colorado.edu/research/wavelets/), BSMART[46] (http://www.brain-smart.org/), FieldTrip[47] (http://fieldtrip.fcdonders.nl/), and Chronux[48] (http://www.chronux.org/) toolboxes.

Behavioral analyses. Behavioral learning curves were estimated in two ways. A sliding window analysis (**Fig. 1c**) measured the percent of correct responses within a window of width equal to 10% of all trials in a given session, stepped in 2.5% increments from the start to the end of the session. This permits identical trial sampling for behavioral and neural data analysis, but underestimates learning rate due to the smoothing inherent in averaging across several abrupt, laterally-shifted learning curves[49]. To estimate the true learning rate, we also fit binary (correct/incorrect) outcomes across trials with a bounded logistic curve (**Fig. 1c**)

$$p\left(\text{correct response on trial } x\right) = a + \frac{b-a}{1+\exp\left(-\frac{x-\mu}{\sigma}\right)}$$

where the probability *p* of a correct response on each trial is estimated as a sigmoidal learning curve with center μ, width σ (inversely related to learning rate), initial guess rate *a* (~0.5 for our two-choice task), and post-learning asymptote *b*. These four parameters were fitted for each learned association using nonlinear least-squares estimation (MATLAB lsqnonlin function) with reasonable parameter bounds based on the data.

Spike preprocessing. For analyses of spiking activity, spike times were converted into smoothed rate (spike density) functions via convolution with a Hann window:

$$y(t) = \begin{cases} 0.5\left(1+\cos\left(\pi t/a\right)\right) \text{ if } -a \le t \le a \\ 0 \qquad\qquad\qquad \text{otherwise} \end{cases}$$

with width parameter *a* = 175 ms (nearly identical to a 70 ms s.d. Gaussian, but with finite spread of $\pm a$). For summary analyses, spike rates were instead calculated within time epochs designed to capture either primarily transient (100–500 ms after retrieval cue or outcome feedback onset) or sustained (600–1,350 ms) neural responses during the delay or inter-trial interval periods.

To ensure our results were not affected by any slow fluctuations in spike rate unrelated to task factors, we removed them before further analysis. Slow trends were estimated at each time point/epoch by convolving the spike rate across trials with a Gaussian broad enough to blur out rate differences related to individual conditions (s.d. = 32 trials = 8 repetitions of all associations). This estimate was subtracted from the individual-trial spike density functions before further analysis. Results were similar, but with weaker signal-to-noise, without this detrending.

Spike rate selectivity. To measure the strength of spike rate signals reflecting each task factor of interest (i.e., cue and associate object identities, trial outcome), we calculated the percent of explained variance (PEV) in the smoothed rate functions by the task factor at each time point (**Fig. 3a**,**b**). For analysis of trial outcome, this was calculated via a one-way ANOVA (two levels: correct, incorrect). For cue and associate identity, we used a nested two-way ANOVA[50], which measures the effect of a nesting factor (associate: A_1, A_2) while partialing out the effect of a nested factor (cue: $C_{1.1/1.2}$ nested within associate A_1, $C_{2.1/2.2}$ within associate A_2; **Fig. 1a**). To examine PEV across learning, these analyses were performed independently within sliding trial windows (for cue/associate identity, window width = 10% and step = 2.5% of total number of session trials; for trial outcome, both = 33%, due to necessity of trial-balancing). Because the traditional formulation of the PEV statistic

$$\text{PEV}_{\eta^2} = \text{SS}_{\text{Groups}} / \text{SS}_{\text{Total}}$$

is biased toward positive values, we instead used the bias-corrected formulation[51]

$$\text{PEV}_{\omega^2} = \frac{\text{SS}_{\text{Groups}} - \text{df}_{\text{Groups}}\text{MS}_{\text{Error}}}{\text{SS}_{\text{Total}} + \text{MS}_{\text{Error}}}$$

where $\text{SS}_{\text{Groups}}$ and SS_{Total} are the between-groups (task conditions) and total sums of squares, $\text{df}_{\text{Groups}}$ is the groups degrees of freedom, and MS_{Error} is the mean squared error. This resulted in an unbiased metric, with an expected value of zero when there is no difference between conditions. Similar results were obtained using shuffle-corrected Shannon mutual information or area under ROC curve analyses (**Supplementary Fig. 11**).

As formulated, another unwanted signal might still contribute to associate-object PEV. Given that it is computed as a contrast between pairs of cue-object conditions (that is, $A_1 = [C_{1.1}$ or $C_{1.2}]$ versus $A_2 = [C_{2.1}$ or $C_{2.2}]$), some portion of associate PEV might be due to selective activation of a neuron by random combinations of cue objects. We reasoned that if this were the case, activation to combinations of cues not paired with same associate (for example, $C_{1.1}$ and $C_{2.2}$) would be just as likely as to cues paired with same associate (for example, $C_{1.1}$ and $C_{1.2}$). To control for this possibility, we recalculated associate-object PEV using all mispaired cue-associate nesting relationships, and subtracted the average of these from the PEV calculated from the actual, proper condition grouping; the resulting corrected statistic is plotted in **Figure 3a**. Simulations showed this statistic to have an expected value >0 for neurons active for a specific associate object, but ≤0 for neurons activated by one or more cues not paired with the same associate.

To measure population neural bias toward representing one condition over another (that is, correct more than incorrect outcomes or vice versa), we calculated a signed version of the PEV (**Fig. 3c**), where the PEV at each time point is multiplied by the sign of the difference in rate between the two conditions of interest

$$\text{PEV}_{\text{signed}}(t) = \text{PEV}(t) * \text{sign}\left(\text{rate}_{\text{cond1}}(t) - \text{rate}_{\text{cond2}}(t)\right)$$

This metric has an expected value of zero if the two conditions are represented equally across the neural population.

Finally, for display purposes only, spike selectivity plots were smoothed with a two-dimensional Gaussian with s.d. (5% of session length, 50 ms) and interpolated to a finer sampling grid.

Spike rate classification. To assess whether the learning-related change in HPC trial outcome bias reflected a switch in the preference of individual neurons, or simply a relative modulation of neurons with consistent preference through learning, we performed a population classification analysis. A logistic regression (linear) classifier was trained to discriminate correct versus incorrect trials based on the ITI-epoch spike rates of all 199 HPC neurons (with neurons recorded in different sessions concatenated into a 'pseudo-population'[52]). From each neuron, we randomly selected 20 trials (10 correct, 10 error) to train the classifier and 10 distinct trials (5 correct, 5 error) to test its predictive accuracy; results were averaged across 100 random trial selections. We asked whether classifier weights trained on trials selected from the early-learning stage could transfer to accurately predicting outcome in late-learning trials—the preference-flip model predicts they would not generalize, while the consistent-preference model suggests they would. As a baseline for comparison, we examined the accuracy of classifiers

doi:10.1038/nn.3954

both trained and tested on independent sets of late-learning trials. A very similar pattern of results was obtained when classifying PFC instead of HPC activity (but with lower overall accuracy), and when using Fisher linear discriminant, linear-kernel support vector machine (SVM), or Poisson naive Bayes classifiers.

LFP preprocessing. Any power line noise, and its harmonics, was estimated by fitting sinusoids in 5-s sliding windows, and subtracted them from the raw data (rmlinesmovingwinc function in Chronux toolbox[48]). To remove the contribution of signal components phase-locked to trial events (event-related evoked potentials), these were measured as the across-trial mean of the raw LFP and subtracted off each individual trial's LFP before further analysis[53]. This correction was performed separately for each trial window and condition, to account for any possible changes in evoked activity across these factors. For the multivariate autoregression-based causality analysis (see below), we additionally normalized each trial by the across-trial s.d., and downsampled the LFP signals to 200 Hz[54].

For summary analyses, LFP metrics (power, synchrony, etc.) were pooled within spectrotemporal regions with time ranges based on epochs used for spike analyses, but extended by 50% to account for the longer duration LFP responses (100–700 and 600–1,725 ms), and with frequency ranges based on our empirical results (2–6 and 9–16 Hz, respectively).

LFP-LFP synchrony analysis. For measures of LFP-LFP synchrony, LFPs were transformed into the time-frequency domain using complex Morlet wavelets[45] (wavenumber = 6, evaluated at 0.25 octave intervals from 1–256 Hz), from which we extracted their phase. The strength of neural synchrony was quantified by the phase-locking value (PLV)[55], the length of the across-trial vector average of cross-electrode differences in phase φ, for a given time point t and frequency f

$$\mathrm{PLV}(f,t)=\left|\frac{1}{nTrials}\sum_{trl=1}^{nTrials}\exp\left(i\left[\varphi_{trl,elec1}(f,t)-\varphi_{trl,elec2}(f,t)\right]\right)\right|$$

PLV measures the degree to which LFP pairs maintain the same phase relationship—independent of their absolute phases and amplitudes—across repeated trials. This analysis was performed separately for each electrode pair, trial window, and condition. To quantify the difference in synchrony between task conditions, we subtracted the PLV(f,t) spectrogram for one condition from another (for example, correct – incorrect outcome). This was normalized to a z-score-like statistic by subtracting the mean, and dividing by the s.d., of the between-condition PLV difference calculated across 50 random permutations of the condition labels across trials. Similar results were obtained using classical coherence or pairwise phase consistency[56] instead of the PLV.

For display purposes only, LFP synchrony plots were smoothed with a two-dimensional Gaussian with s.d. (0.15 octaves, 100 ms) and interpolated to a finer sampling grid.

LFP-LFP causality analysis. We measured directional influences between the PFC and HPC in two ways. First, we calculated the mean phase lag between LFPs on each pair of electrodes

$$\varphi_{lag}(f,t)=\mathrm{angle}\left(\frac{1}{nTrials}\sum_{trl=1}^{nTrials}\exp\left(i\left[\varphi_{trl,elec1}(f,t)-\varphi_{trl,elec2}(f,t)\right]\right)\right)$$

To test whether the distribution of phase lags across all electrode pairs was significantly different from zero—suggestive of a directionality between the two signals—we used a circular bootstrap test[57].

Though consistent phase lags much smaller than a full oscillatory cycle (**Fig. 5a**) are suggestive of directional influences, they are in principle ambiguous because of the cyclic nature of the signals. We therefore also computed the GPDC[58] between pairs of LFPs in each area. GPDC is a frequency-domain analog of Granger causality[59], which measures putative causality in terms of the degree to which one signal (LFP) can be predicted by past values of another signal (LFP from a distinct electrode), with its own past history factored out. This approach is based on a multivariate autoregressive (MVAR) model fit to pairs of LFP time series

$$\mathbf{X}(t)=\sum_{k=1}^{p}\mathbf{A}_k\mathbf{X}(t-k)+\varepsilon(t)$$

where $\mathbf{X}(t)$ is the data vector of the pair of LFP signals at time t, $\mathbf{A}_k$ the 2×2 matrix of autoregressive coefficients at the k^{th} time lag, p is the maximum number of lags (model order), and $\varepsilon(t)$ the residual prediction error. We evaluated the Bayesian Information Criterion on a few representative LFP pairs to select a fixed model order of 17 (max lag 85 ms). MVAR models were fit separately on LFP data within each sliding time window (500-ms width, 250-ms step between successive windows), trial window, and task condition. The parameters of the MVAR model in each window were estimated using Morf's modification of the Levinson-Wiggins-Robinson algorithm[46,60]. The fitted MVAR parameters were then transformed from the time domain into the frequency domain

$$\bar{\mathbf{A}}(f)=\mathbf{I}-\sum_{k=1}^{p}\mathbf{A}_k\exp\left(-2\pi ikf/f_{samp}\right)$$

where $\mathbf{I}$ is the $p\times p$ identity matrix, f_{samp} is the LFP sampling rate, and the spectral coefficients $\bar{\mathbf{A}}(f)$ were evaluated at 1-Hz steps from 1–64 Hz. The GPDC reflecting the directional influence of LFP_1 on LFP_2 is then calculated as

$$GPDC_{2\leftarrow 1}=\frac{\frac{1}{\sigma_2}\left|\bar{A}_{2\leftarrow 1}(f)\right|}{\sqrt{\frac{1}{\sigma_2^2}\left|\bar{A}_{2\leftarrow 1}(f)\right|^2+\frac{1}{\sigma_1^2}\left|\bar{A}_{1\leftarrow 1}(f)\right|^2}}$$

where σ_k^2 is the variance of the prediction error for channel k. Similar results were obtained using classical spectral Granger causality[61].

For display purposes only, LFP causality plots were smoothed with a two-dimensional Gaussian with s.d. (1 Hz, 100 ms) and interpolated to a finer sampling grid.

Controls for other behavioral changes. To examine whether observed changes in neural activity with learning might be related to changes in arousal across the learning session (**Supplementary Fig. 1e**), we measured pupil size, which has been strongly linked to global arousal and associated noradrenergic modulation[62–64]. Pupil diameter was calculated during the delay period, when it is least influenced by external factors[63].

To examine whether neural changes might be attributed to ingestion-related orofacial movements (**Supplementary Fig. 1f**), we made electromyographic (EMG) recordings from the dorsal lip muscles[65] (orbicularis oris) in separate *post hoc* experiments. Since the original animals used for all other reported results were no longer available, EMG signals were obtained from two animals performing a working memory–guided saccade task[66] (4 sessions) or a visuomotor associative learning task[67] (6 sessions). Standard surface EMG methods were used[68] (monopolar recording from 6 mm Ag/AgCl disc electrodes; filtered 10–250 Hz; full-wave rectified).

Power-stratification controls. To ensure that the observed inter-area synchrony and causality effects were not due to differences in within-area LFP power, we repeated these analyses with power balanced across the relevant conditions (**Supplementary Fig. 9**). Power balancing was performed using a stratification method[47,69] that trims trials with extreme power values from each condition until the histogram of trial power values is closely matched between compared conditions (that is, correct and incorrect outcomes).

A **Supplementary Methods Checklist** is available.

40. Nelson, M.J., Pouget, P., Nilsen, E.A., Patten, C.D. & Schall, J.D. Review of signal distortion through metal microelectrode recording circuits and filters. *J. Neurosci. Methods* **169**, 141–157 (2008).
41. Miller, E.K., Gochin, P.M. & Gross, C.G. Habituation-like decrease in the responses of neurons in inferior temporal cortex of the macaque. *Vis. Neurosci.* **7**, 357–362 (1991).

42. Xiang, J.-Z. & Brown, M.W. Neuronal responses related to long-term recognition memory processes in prefrontal cortex. *Neuron* **42**, 817–829 (2004).
43. Yanike, M., Wirth, S., Smith, A.C., Brown, E.N. & Suzuki, W.A. Comparison of associative learning-related signals in the macaque perirhinal cortex and hippocampus. *Cereb. Cortex* **19**, 1064–1078 (2009).
44. Manly, B.F.J. *Randomization, Bootstrap, and Monte Carlo Methods in Biology* (Chapman & Hall/CRC, 2007).
45. Torrence, C. & Compo, G. A practical guide to wavelet analysis. *Bull. Am. Meteorol. Soc.* **79**, 61–78 (1998).
46. Cui, J., Xu, L., Bressler, S.L., Ding, M. & Liang, H. BSMART: a Matlab/C toolbox for analysis of multichannel neural time series. *Neural Netw.* **21**, 1094–1104 (2008).
47. Oostenveld, R., Fries, P., Maris, E. & Schoffelen, J.-M. FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput. Intell. Neurosci.* **2011**, 1–9 (2011).
48. Bokil, H., Andrews, P., Kulkarni, J.E., Mehta, S. & Mitra, P.P. Chronux: a platform for analyzing neural signals. *J. Neurosci. Methods* **192**, 146–151 (2010).
49. Gallistel, C.R., Fairhurst, S. & Balsam, P. The learning curve: implications of a quantitative analysis. *Proc. Natl. Acad. Sci. USA* **101**, 13124–13131 (2004).
50. Zar, J.H. *Biostatistical Analysis* (Prentice-Hall/Pearson, 2010).
51. Olejnik, S. & Algina, J. Generalized eta and omega squared statistics: measures of effect size for some common research designs. *Psychol. Methods* **8**, 434–447 (2003).
52. Meyers, E.M. & Kreiman, G. Tutorial on pattern classification in cell recording. in *Visual Population Codes* (eds. Kriegskorte, N. & Kreiman, G.) 517–538 (MIT Press, 2012).
53. Kalcher, J. & Pfurtscheller, G. Discrimination between phase-locked and non-phase-locked event-related EEG activity. *Electroencephalogr. Clin. Neurophysiol.* **94**, 381–384 (1995).
54. Ding, M., Bressler, S.L., Yang, W. & Liang, H. Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment. *Biol. Cybern.* **83**, 35–45 (2000).
55. Lachaux, J.-P., Rodriguez, E., Martinerie, J. & Varela, F.J. Measuring phase synchrony in brain signals. *Hum. Brain Mapp.* **8**, 194–208 (1999).
56. Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P. & Pennartz, C.M.A. The pairwise phase consistency: A bias-free measure of rhythmic neuronal synchronization. *Neuroimage* **51**, 112–122 (2010).
57. Fisher, N.I. *Statistical Analysis of Circular Data* (Univ. Press, 1995).
58. Baccalá, L.A., Sameshima, K. & Takahashi, D.Y. Generalized partial directed coherence. *15th Int. Conf. Digit. Signal Process* 163–166 (2007).
59. Granger, C.W. Investigating causal relations by econometric models and cross-spectral methods. *Econ. J. Econ. Soc.* **37**, 424–438 (1969).
60. Morf, M., Vieira, A., Lee, D.T. & Kailath, T. Recursive multichannel maximum entropy spectral estimation. *Geosci. Electron. IEEE Trans.* **16**, 85–94 (1978).
61. Geweke, J. Measurement of linear dependence and feedback between multiple time series. *J. Am. Stat. Assoc.* **77**, 304–313 (1982).
62. Hess, E.H. & Polt, J.M. Pupil size as related to interest value of visual stimuli. *Science* **132**, 349–350 (1960).
63. Kennerley, S.W. & Wallis, J.D. Reward-dependent modulation of working memory in lateral prefrontal cortex. *J. Neurosci.* **29**, 3259–3270 (2009).
64. Nassar, M.R. *et al.* Rational regulation of learning dynamics by pupil-linked arousal systems. *Nat. Neurosci.* **15**, 1040–1046 (2012).
65. Shepherd, S.V., Lanzilotto, M. & Ghazanfar, A.A. Facial muscle coordination in monkeys during rhythmic facial expressions and ingestive movements. *J. Neurosci.* **32**, 6105–6116 (2012).
66. Funahashi, S., Bruce, C.J. & Goldman-Rakic, P.S. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* **61**, 331–349 (1989).
67. Cromer, J.A., Machon, M. & Miller, E.K. Rapid association learning in the primate prefrontal cortex in the absence of behavioral reversals. *J. Cogn. Neurosci.* **23**, 1823–1828 (2011).
68. Merletti, R. & Di Torino, P. Standards for reporting EMG data. *J. Electromyogr. Kinesiol.* **9**, 3–4 (1999).
69. Schoffelen, J.-M. Neuronal coherence as a mechanism of effective corticospinal interaction. *Science* **308**, 111–113 (2005).