

# Hippocampal–prefrontal input supports spatial encoding in working memory

Timothy Spellman<sup>1</sup>, Mattia Rigotti<sup>2,3,4</sup>, Susanne E. Ahmari<sup>5,6</sup>, Stefano Fusi<sup>2,7</sup>, Joseph A. Gogos<sup>1,2</sup> & Joshua A. Gordon<sup>8,9</sup>

**Spatial working memory, the caching of behaviourally relevant spatial cues on a timescale of seconds, is a fundamental constituent of cognition. Although the prefrontal cortex and hippocampus are known to contribute jointly to successful spatial working memory, the anatomical pathway and temporal window for the interaction of these structures critical to spatial working memory has not yet been established. Here we find that direct hippocampal–prefrontal afferents are critical for encoding, but not for maintenance or retrieval, of spatial cues in mice. These cues are represented by the activity of individual prefrontal units in a manner that is dependent on hippocampal input only during the cue–encoding phase of a spatial working memory task. Successful encoding of these cues appears to be mediated by gamma–frequency synchrony between the two structures. These findings indicate a critical role for the direct hippocampal–prefrontal afferent pathway in the continuous updating of task-related spatial information during spatial working memory.**

Spatial working memory (SWM) is an essential feature of goal-directed action. Locating a resource, a threat, or even oneself within a dynamic or unfamiliar environment requires a cached representation of relevant spatial features that must be continuously updated, preserved and applied as needed to the execution of adaptive behaviours<sup>1</sup>. Despite long-standing interest in the neurobiological underpinnings of SWM, its multiple cognitive components, distributed anatomical constituents and distinct temporal phases have rendered its underlying circuit mechanisms elusive<sup>2–6</sup>. Nevertheless, an extensive body of work supports the idea that the prefrontal cortex (PFC) has a central role in the successful execution of tasks requiring SWM<sup>7</sup>. Moreover, the joint contribution of the medial prefrontal cortex (mPFC) and hippocampus (HPC) supports successful SWM in rodents<sup>8–12</sup>. It remains unclear, however, which phase(s) of SWM (encoding, maintenance, and/or retrieval) require the joint participation of the HPC and mPFC, what behaviourally relevant information is conveyed between the two structures, and through which anatomical pathway(s) they interact.

Direct HPC–mPFC connectivity is confined to a unidirectional projection from the CA1/subiculum of the ventral-most two-thirds of the hippocampus (vHPC)<sup>13–15</sup>. Cells in both the vHPC and mPFC exhibit location-specific firing<sup>16–19</sup>, and damage to the vHPC disrupts goal-related activity in the mPFC<sup>16</sup>, suggesting that the direct vHPC–mPFC projection may transmit critical location information during SWM tasks.

We used a rodent delayed non-match-to-place (DNMTP) task, known to require activity in the vHPC and mPFC, to test the role of vHPC–mPFC afferents in SWM. We applied a projection-specific, optogenetic silencing approach that afforded anatomical and temporal precision. Inhibiting vHPC–mPFC direct input disrupted encoding, but not maintenance or retrieval, of location cues necessary for task performance. Furthermore, we found that goal-selective firing in the mPFC was dependent on vHPC direct input exclusively during the encoding phase of each trial. Finally, we found evidence suggesting

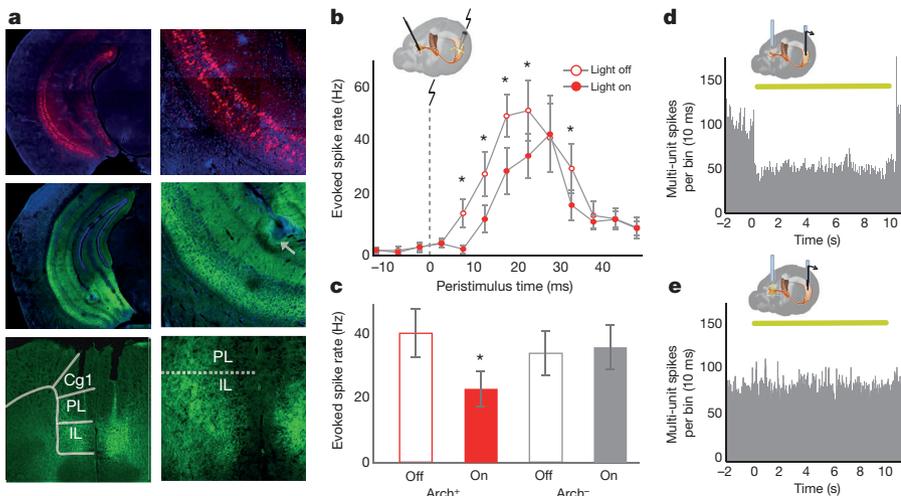
that the transmission of task-critical information through the vHPC–mPFC pathway is mediated by the synchronization of mPFC units to gamma oscillations in the vHPC. Together, these findings suggest that the direct vHPC–mPFC pathway enables the encoding of salient spatial cues during SWM performance.

## Optogenetic inhibition of vHPC terminals

To interfere specifically with vHPC–mPFC inputs, a projection-specific targeting approach was used. An adeno-associated virus vector (AAV2/5) engineered to express fluorescently labelled archaerhodopsin (eArch3.0–enhanced yellow fluorescent protein (eYFP))<sup>20,21</sup> was targeted to subfield CA1 of the vHPC in mice (Methods). Robust expression was seen in the dendrites and axons of the vHPC and in projection axons in the mPFC (Fig. 1a).

Arch-mediated hyperpolarization of distal axon terminals disrupted synaptic transmission *in vivo* without affecting spontaneous vHPC firing. The effect of terminal illumination on synaptic transmission was measured in acutely anaesthetized mice. Electrical stimuli were delivered via a bipolar stimulating electrode to ventral CA1 in both Arch-expressing (Arch<sup>+</sup>) mice and Arch-negative controls (Arch<sup>-</sup>). Postsynaptic multi-unit responses were observed in the mPFC (Fig. 1b). Light pulses delivered to the mPFC on interleaved trials reduced the evoked response by ~40% for Arch<sup>+</sup> animals but not Arch<sup>-</sup> animals (Fig. 1b, c). To measure the effects of terminal illumination on vHPC cell bodies, mice were implanted with optical-fibre-coupled stereotrodes in the vHPC and optical fibres in the mPFC. In the awake, resting state, light pulses delivered to the vHPC reduced local spontaneous multi-unit activity by ~50% (Fig. 1d), while illumination of terminal fields in the mPFC had no effect on multi-unit activity in the vHPC (Fig. 1e). These experiments demonstrated successful and specific inhibition of terminals *in vivo*, an approach that was then applied to a behavioural paradigm to examine what role the vHPC–mPFC projection has in spatial working memory.

<sup>1</sup>Department of Physiology and Cellular Biophysics, Columbia University, 630 West 168th Street, New York, New York 10032, USA. <sup>2</sup>Department of Neuroscience, Columbia University, 1051 Riverside Drive, New York, New York 10032, USA. <sup>3</sup>IBM T. J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, New York 10598, USA. <sup>4</sup>Italian Academy for Advanced Studies in America, Columbia University, 1161 Amsterdam Avenue, New York, New York 10032, USA. <sup>5</sup>Translational Neuroscience Program, Department of Psychiatry, University of Pittsburgh, 450 Techology Drive, Pittsburgh, Pennsylvania 15219, USA. <sup>6</sup>Center for Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, 200 Lothrop Drive, Pittsburgh, Pennsylvania 15261, USA. <sup>7</sup>Kavli Institute for Brain Sciences, Columbia University, 1051 Riverside Drive, New York, New York 10032, USA. <sup>8</sup>Department of Psychiatry, Columbia University, 1051 Riverside Drive, New York, New York 10032, USA. <sup>9</sup>Division of Integrative Neuroscience, New York State Psychiatric Institute, 1051 Riverside Drive, New York, New York 10032, USA.



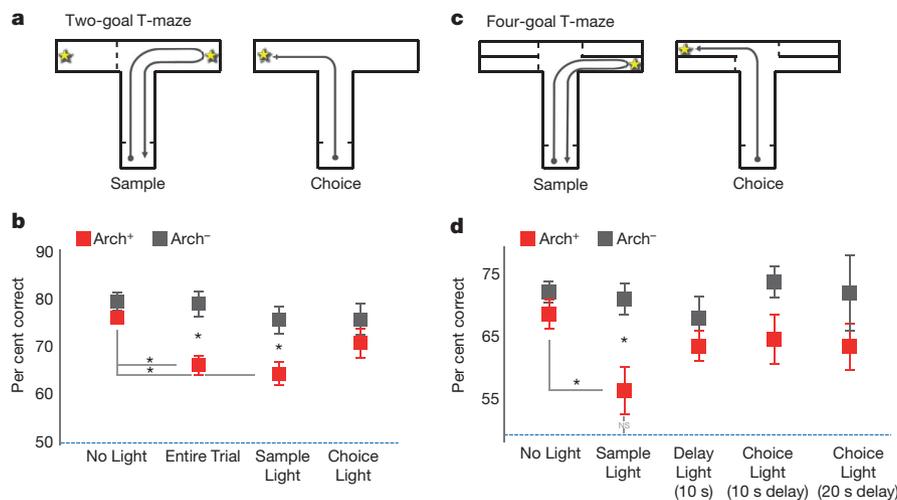
**Figure 1 | Optogenetic inhibition of vHPC–mPFC terminals *in vivo*.** **a**, Expression of Arch (middle row, green) and mCherry (top row, red) in ventral CA1. Arrow, lesion marking electrode location. Bottom row, Arch in terminals in the prelimbic (PL) and infralimbic (IL) mPFC. Cg1, anterior cingulate cortex. Original magnifications,  $\times 2$  (left),  $\times 20$  (right). **b**, mPFC multi-unit responses to vHPC stimulation in Arch<sup>+</sup> mice.  $n = 16$  sites from 3 animals, ANOVA  $F = 5.6$ ,  $P = 0.02$  for light effect,  $*P < 0.05$ , post-hoc  $t$ -test. Baseline rate =  $6.1 \pm 0.14$  Hz. **c**, Group mean evoked mPFC spike rate summed across 5–40 ms post-stimulus. ANOVA  $F = 31.4$ ,  $P = 4 \times 10^{-6}$  for virus-by-light interaction;  $n = 16$  sites from 3 mice,  $t = 6.68$ ,  $P = 0.0004$ ;  $n = 17$  sites from 3 mice,  $t = 1.57$ ,  $P = 0.3$ , for Arch<sup>+</sup> and Arch<sup>-</sup> mice, respectively. **d**, **e**, Multi-unit activity traces from vHPC of Arch<sup>+</sup> mice during somatic (**d**) and terminal field (**e**) illumination *in vivo*. Yellow bar, light on (throughout). Error bars show  $\pm$  standard error of the mean (s.e.m.).

### Terminal inhibition impairs encoding

To assay working memory performance, a T-maze DNMT task was employed. In this task, each trial is divided into three phases (Fig. 2a). In the sample phase, one of two goal locations is blocked by a wall, and the mouse is directed towards a food reward in the open location; during this phase the animal must encode the location of the sample goal. In the delay phase, the mouse returns to the start box and must maintain the sample goal in working memory during a variable delay (Methods). In the choice phase, the wall is removed, and the mouse must select the previously closed arm to receive a second reward. After successful task acquisition, light was delivered to vHPC–mPFC terminal fields in Arch<sup>+</sup> and Arch<sup>-</sup> mice during the entire trial (Entire Trial condition), the sample phase only (Sample Light), or the choice phase only (Choice Light). Trial types were randomly interleaved.

In Arch<sup>+</sup> but not Arch<sup>-</sup> mice, performance was impaired in Entire Trial and Sample Light conditions; Choice Light did not result in a statistically significant impairment (Fig. 2b). These data raised the possibility that vHPC–mPFC input is critical for the encoding of location cues associated with the sample goal but may not be required for the retrieval of such cues.

This native version of the T-maze task, however, is not optimized to discriminate between encoding and retrieval of the sample goal location, as the animal could begin forming a motor action plan (that is, ‘go into the opposite arm’) any time after it encounters the sample goal. To segregate better the encoding, maintenance, and retrieval phases within each trial, a modified, four-goal T-maze was constructed (Fig. 2c). Here, as in the two-goal task, a single goal was made available for retrieval of the sample reward. During the choice run, the sample goal and one of the other three arms were open. This design



**Figure 2 | Inhibition of vHPC–mPFC terminals impairs encoding.** **a**, Two-goal DNMT task. **b**, Effect of mPFC illumination on performance in the two-goal task differed by trial phase and virus type (ANOVA, virus-by-light interaction,  $F = 5.92$ ,  $P = 0.02$ ). Post-hoc tests revealed effects of Entire Trial ( $t = 3.96$ ,  $P = 0.002$ ) and Sample Light ( $t = 2.98$ ,  $P = 0.011$ ) but not Choice Light ( $t = 1.1$ ,  $P = 0.29$ ) conditions in Arch<sup>+</sup> ( $n = 8$ ) versus Arch<sup>-</sup> ( $n = 6$ ) animals. Entire Trial and Sample Light but not Choice Light performance differed from the No Light condition in Arch<sup>+</sup> ( $t = 4.9$ ,  $P = 0.002$ , Entire Trial;  $t = 4.5$ ,  $P = 0.003$ , Sample Light;  $t = 1.7$ ,  $P = 0.125$ , Choice Light) but not Arch<sup>-</sup> animals ( $t = 0.18$ ,  $P = 0.87$ ;  $t = 1.1$ ,  $P = 0.3$ ;  $t = 1.3$ ,  $P = 0.25$ ,

respectively). **c**, Four-goal DNMT task. **d**, Effect of illumination of mPFC terminal fields on performance in the four-goal task (Arch<sup>+</sup>,  $n = 7$  Arch<sup>-</sup>,  $n = 6$ , ANOVA  $F = 3.1$ ,  $P = 0.03$  for virus-by-light interaction). Impairment was restricted to Sample Light trials ( $t = 3.1$ ,  $P = 0.0093$ ; and  $t = 1.1$ ,  $P = 0.29$ ;  $t = 1.0$ ,  $P = 0.34$ ;  $t = 1.91$ ,  $P = 0.08$ ;  $t = 1.2$ ,  $P = 0.24$ , for No Light, Delay Light, Choice Light 10 s and Choice Light 20 s, respectively). No Light and Sample Light performance were significantly reduced in Arch<sup>+</sup> ( $t = 2.5$ ,  $P = 0.04$ ) but not Arch<sup>-</sup> mice ( $t = 0.35$ ,  $P = 0.73$ ). Performance of Arch<sup>+</sup> mice during Sample Light runs was not significantly above chance ( $t = 1.9$ ,  $P = 0.11$ ). NS, not significant. Error bars show s.e.m.

prevents the mouse from formulating a spatially directed action plan until the choice phase, when, at the end of its centre arm run, it is presented with two of the four goals as options (Supplementary Video). Thus, selection of the choice goal was temporally restricted to the choice phase, separating out encoding from retrieval.

Again, performance was impaired in the Sample Light condition, confirming the requirement for vHPC–mPFC input during encoding (Fig. 2d). In the Choice Light condition (repeated here with either 10 or 20 s delays), a slight trend towards impairment did not reach significance, and post-hoc power analysis revealed that an  $n$  of 37 animals per group would have been required to detect a paired difference given the observed effect size. There was also no significant impairment with terminal illumination during the delay period (Delay Light; Fig. 2d). Although we cannot conclusively rule out an effect of terminal inhibition during choice, these experiments indicate that input from the vHPC to the mPFC is critical for encoding of spatial cues.

### Effect on spatial representation in mPFC

The behavioural findings suggested that task-related spatial locations may be represented by firing rates in the mPFC in a vHPC-input-dependent manner. Therefore, recordings of multiple single units were obtained from the prelimbic region of the mPFC in mice performing the four-goal task (Extended Data Fig. 1). A total of 792 well-isolated single units were obtained from 9 mice. Forty-four per cent of mPFC single units displayed selectivity for one or both spatial dimensions that distinguished goal arms (left/right, back/front), and/or their interaction, as assayed by two-way analysis of variance (ANOVA) (Extended Data Fig. 2a–d). Given the inhomogeneous and distributed nature of the representation of spatial information among recorded mPFC units, a maximum margin linear classifier<sup>22</sup> was used to decode the sample goal location from binned population firing rate vectors and to quantify the strength and reliability of the neural representation (Fig. 3). The classifier was cross-validated by training on data from half the trials and testing the model's performance on data from the remainder (see Methods).

The sample goal was decoded from spike histograms aligned to multiple trial events (Fig. 3a). Sample goal identity was decoded from the mPFC population at accuracies well above chance from the time

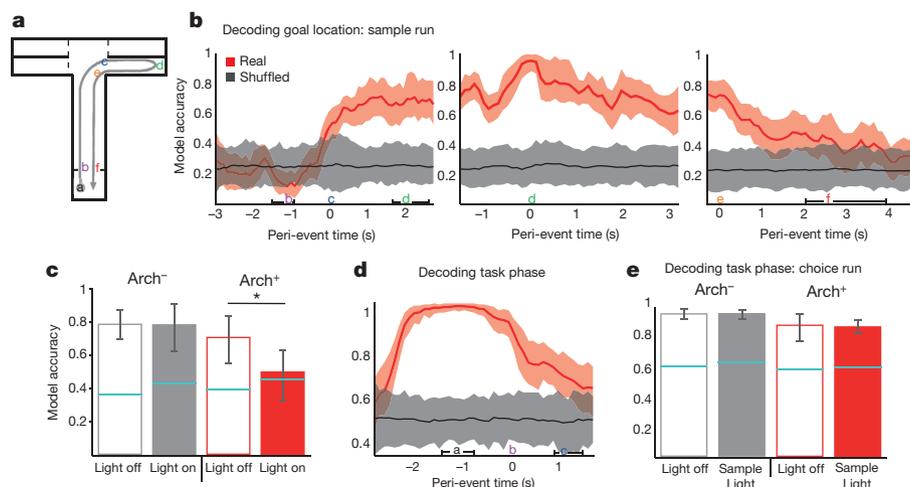
the animals entered the T-intersection, and peaked at 96% upon arrival at the reward port in the sample goal (Fig. 3b).

Sample goal representation at goal arrival time was then assessed during inhibition of vHPC input by training the model on firing rates from non-Sample Light trials (No Light, Delay Light and Choice Light trials combined) and testing on rates from Sample Light trials. For Arch<sup>-</sup> animals, the sample goal was decoded equally well with or without mPFC illumination, while in Arch<sup>+</sup> animals accuracy was reduced to chance by mPFC illumination (Fig. 3c). This result demonstrates that vHPC input is critical for the representation of the sample goal among mPFC units during encoding.

To assess the impact of terminal inhibition on encoding of non-spatial, task-relevant cues, the same classifier was trained to decode the task phase (sample versus choice) at the time the start-box doors opened immediately before running down the centre arm. Because this epoch was behaviourally equivalent in sample and choice runs, an accurate representation of the task phase at this time point must rely upon a memory of the preceding task phase. Task phase was decoded with near perfect accuracy (0.98) at the time bin corresponding with the opening of the doors, revealing a memory trace for the preceding task phase (Fig. 3d). When the decoder was trained on firing rates from this epoch in No Light trials and tested on choice runs from Sample Light trials, model accuracy was not affected (Fig. 3e), suggesting that vHPC–mPFC terminal inhibition does not generally interfere with the encoding of task-relevant information.

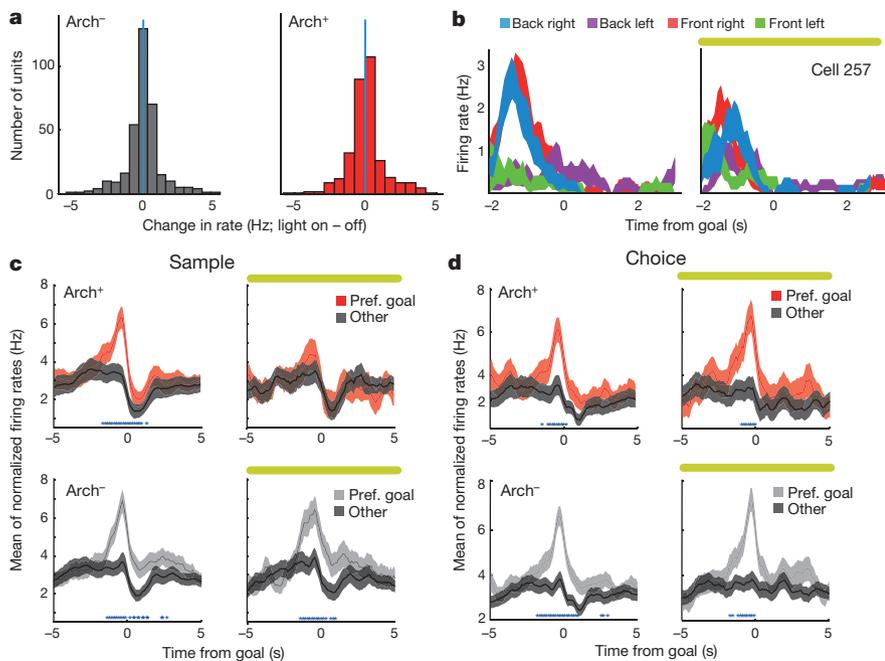
The finding of vHPC-input-dependent location coding in the mPFC leaves open the question of how vHPC input influences mPFC neurons. Overall firing rates of mPFC units were unaffected by terminal illumination in both Arch<sup>-</sup> and Arch<sup>+</sup> animals, whether all were considered together (Fig. 4a) or putative pyramidal cells and interneurons were separately classified on the basis of waveform features (Extended Data Fig. 4). Nonetheless, consistent effects of vHPC inputs might be revealed by a finer-grained analysis.

Each unit's preferred goal (that is, the goal in which firing rate was most different from the mean rate across the other three goals) was identified using weights generated by the classifier (Methods). Although this separation criterion allowed for the possibility that any given unit might represent location by an increase or decrease in rate, the difference was observed as an elevation in mean firing rate



**Figure 3** | mPFC units require vHPC input to encode location but not task phase. **a**, Schematic of sample run. **b**, Accuracy of goal decoding during sample run with light off. Solid lines, mean decoder accuracy; shaded areas, 95% confidence intervals.  $n = 727$  units from 9 mice. **c**, Decoding accuracy for sample goal upon arrival at the reward port (location 'd' in **a**) in the presence (filled bars) and absence (open bars) of illumination of vHPC–mPFC terminals (269 units from 4 Arch<sup>+</sup> mice, 285 units from 5 Arch<sup>-</sup> mice;  $n = 100$  permutations; ANOVA  $F = 1,978$ ,  $P = 5 \times 10^{-5}$  for virus-by-light interaction;

$t = 0.48$ ,  $P = 0.64$  for Arch<sup>-</sup>;  $*t = 161.2$ ,  $P = 1.2 \times 10^{-121}$  for Arch<sup>+</sup>). Error bars represent 95% confidence intervals; blue lines represent upper bounds of 95% confidence intervals for shuffled data. **d**, Decoding of task phase (sample versus choice) as a function of time relative to departure from the start box. Conventions as in **b** (792 units from 9 mice). **e**, Decoding of task phase at door opening (location 'a' in **a**) as a function of trial type (ANOVA,  $F = 1.94$ ,  $P = 0.17$  for virus-by-light interaction). Error bars show s.e.m.



**Figure 4 | Location selectivity requires vHPC input during encoding but not retrieval.**

**a**, Differences in firing rate with terminal illumination ( $n = 433$  units, sign rank  $z = -0.23$ ,  $P = 0.82$ ;  $n = 359$ ,  $z = -0.97$ ,  $P = 0.33$ , for Arch<sup>+</sup> and Arch<sup>-</sup> animals, respectively). **b**, Mean binned spike rates for an example mPFC unit from an Arch<sup>+</sup> animal, aligned to arrival at each of the four sample goals (colour coded) without (left) or with (right) terminal illumination. **c**, Peri-event firing rates during the sample phase for all location-selective units as they approach the preferred (Pref.; red/grey) and non-preferred (Other; black) goals (blue asterisks, time points with Bonferroni-corrected significance) during light on (right) and light off (left) trials.  $n = 67$  units from 7 Arch<sup>+</sup> mice, and 78 units from 6 Arch<sup>-</sup> mice. **d**, Same as c, but during the choice phase.  $n = 145$  units from 7 Arch<sup>-</sup> mice, and 77 units from 6 Arch<sup>+</sup> mice.

relative to other goals, when averaged across all units (Extended Data Fig. 5), and especially those units with high goal selectivity (Fig. 4b–d). The effect in Sample Light conditions was striking. Inhibition of terminals during the sample phase prevented the increase in firing rate that occurred in units' preferred goals, without affecting baseline rates (Fig. 4c, top). The same results were obtained when including all units in the analysis (Extended Data Fig. 5). These results reveal an excitatory role for vHPC input in the location-selective firing rate enhancement seen in mPFC units.

To address whether direct vHPC input is necessary for all mPFC spatial representations, firing rate in preferred and non-preferred goals was examined during choice runs. As in sample runs, firing rate for highly goal-selective units during choice runs was also higher in the preferred goal than in non-preferred goals. However, unlike in sample runs, this elevation was unaffected by vHPC–mPFC terminal inhibition (Fig. 4d). In this finding, physiology mirrored the behavioural result, indicating that goal selectivity after encoding is no longer dependent on vHPC input.

### vHPC gamma organizes mPFC spike timing

Activity in distant brain regions can be coordinated by gamma oscillations<sup>23</sup>, and long-range gamma synchrony in both the cortex and hippocampus has been linked to spatial learning and memory<sup>24,25</sup>. Therefore, long-range synchrony was quantified using the magnitude (pairwise phase consistency<sup>26</sup> (PPC)) and significance (Rayleigh's test<sup>9</sup>) of phase non-uniformity of spike times in one brain region relative to local field potential (LFP) oscillations in another. To determine the temporal directionality of synchronous activity, lag analysis was performed, in which phase-locking was calculated at various temporal shifts; preferential phase-locking at a non-zero lag indicates a predictive relationship between oscillatory phase and spike timing.

A subset of mPFC units were significantly phase-locked to vHPC gamma (Fig. 5b, c), and the percentage of significantly phase-locked units was greatest at lags in which vHPC gamma preceded mPFC spiking (Fig. 5c). Moreover, the mean strength of mPFC unit phase-locking to vHPC gamma was maximal at lags in which vHPC led (Fig. 5d–f). These findings suggest the possibility that gamma-frequency inputs from the vHPC influence mPFC spike timing. Consistent with this suggestion, inhibition of vHPC–mPFC terminals reduced the overall strength of gamma phase-locking (Fig. 5g, h), indicating that the observed synchrony is mediated by direct

vHPC–mPFC input. Importantly, this directionality was specific to gamma oscillations, as lag analysis of phase-locking in the theta range revealed an opposite directionality; mPFC led vHPC (but not dorsal hippocampus (dHPC)) activity in the theta range, and theta synchrony was unaffected by terminal inhibition (Extended Data Fig. 7).

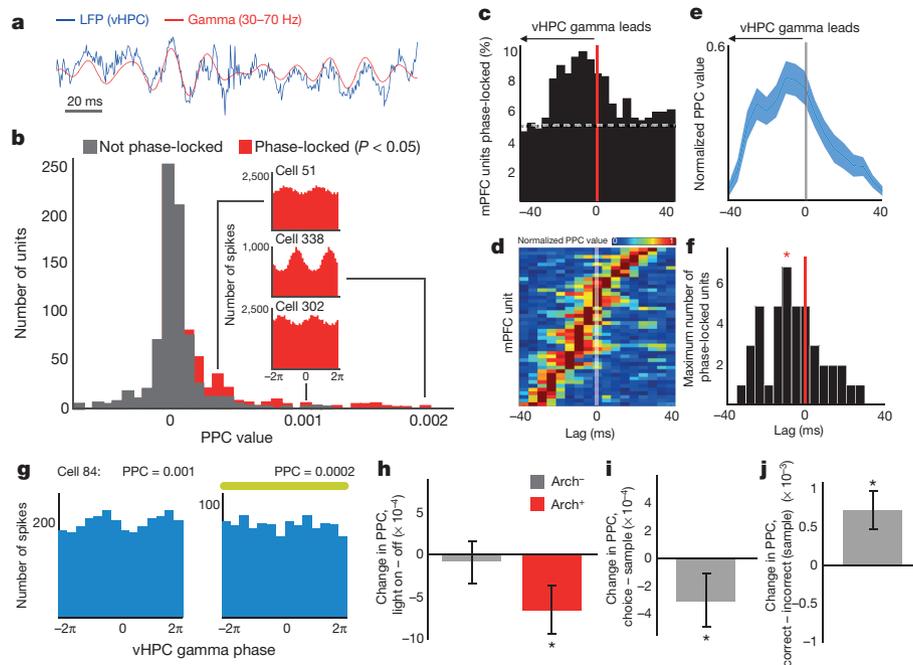
vHPC–mPFC gamma synchrony correlated with behaviour in two key ways. PPC values were higher during sample than choice runs, demonstrating that stronger gamma phase-locking is associated with the encoding phase of the trial (Fig. 5i). Additionally, phase-locking was stronger during sample runs of correct trials than of incorrect trials (Fig. 5j), suggesting that it may support effective encoding of location cues.

### Discussion

We leveraged a temporally precise, projection-specific manipulation to test the role of the vHPC–mPFC afferent pathway in spatial working memory. We found that direct vHPC–mPFC input is essential for successful encoding of task-related cues, both behaviourally and at the level of neural representation within the mPFC. vHPC–mPFC gamma synchrony correlated with successful cue encoding and was also disrupted by vHPC terminal inhibition. These findings point to a role for the vHPC–mPFC afferent pathway as a conduit for the updating of task-critical location cues.

The finding of a dependence of task performance on vHPC input only during sample runs, replicated across both behavioural paradigms used in this study, provides a strong argument for the importance of vHPC–mPFC afferent input during the trial phase in which relevant spatial cues are encoded. The role of vHPC–mPFC input in encoding goal location is further supported by the effect of terminal inhibition on the neural representation of goal location in the mPFC. Moreover, while contemporaneous location was robustly represented as a goal-selective enhancement of firing in preferred goals during both sample and choice runs, this representation depended on vHPC input only during the sample run. Here the physiology agrees with the behavioural observation that vHPC–mPFC input is critical for encoding, not retrieval, of task-relevant location cues.

Interestingly, we find no evidence of retrospective location coding in mPFC activity (Fig. 3 and Extended Data Fig. 3). Within the parameters of this study, this argues against persistent firing within mPFC neurons as a means of maintaining the stimulus representation between encoding and retrieval. Previously, retrospective and pro-



**Figure 5 | Task-dependent modulation of mPFC spiking by vHPC gamma.**

**a**, Example raw and gamma-filtered vHPC LFP. **b**, Distribution of phase-locking values for all units from spikes recorded at all times, coloured by significance (Rayleigh's test,  $P < 0.05$ ). Insets, vHPC gamma phase histograms from example mPFC units (Cell 51:  $z = -3.24$ ,  $P < 0.001$ ,  $\text{PPC} = 0.0003$ ; Cell 338:  $z < -6$ ,  $P < 0.0001$ ,  $\text{PPC} = 0.002$ ; Cell 324:  $z = -2.8$ ,  $P = 0.002$ ,  $\text{PPC} = 0.001$ ). **c**, Percentage of mPFC units significantly phase-locked to vHPC gamma across a range of lags. Dashed line, chance. **d**, Pseudocolour plot of normalized PPC values, sorted by lag of maximal phase-locking, for mPFC units with Bonferroni-corrected significance ( $P < 0.0029$ ). **e**, Mean normalized PPC value

by lag. **f**, Distribution of lags at peak phase-locking strength; shifted towards a vHPC lead ( $n = 43$  units, sign rank,  $z = -2.2$ ,  $P = 0.014$ ). Asterisk indicates mean lag. **g**, Distribution of gamma phases for spikes from an example mPFC unit from an Arch<sup>+</sup> animal during all light off runs (Rayleigh's  $P = 0.03$ ) and light on runs (Rayleigh's  $P = 0.3$ ). **h–j**, Change in phase-locking comparing light on versus off ( $n = 140$  units from 7 Arch<sup>+</sup> mice,  $z = -3.9$ ,  $P = 8.7 \times 10^{-5}$ ; and  $n = 222$  units from 6 Arch<sup>-</sup> mice,  $z = -1.83$ ,  $P = 0.07$ ) (**h**); choice versus sample phases ( $n = 458$  units,  $z = -3.2$ ,  $P = 0.0016$ ) (**i**); and correct versus incorrect trials ( $n = 270$  units,  $z = -4.2$ ,  $P = 3.5 \times 10^{-5}$ ) (**j**). Significance by sign rank. Error bars shows s.e.m.

spective location representations have been seen in mPFC cells of rodents during delays in spatial alternation tasks in which past and future location are intertwined<sup>27,28</sup>, while in primates there is ample evidence for stimulus-specific delay-period firing in the dorsolateral PFC<sup>29,30</sup>. The primate dorsolateral PFC is functionally analogous to, but anatomically distinct from, the rodent prelimbic region of the mPFC<sup>30,31</sup>. But retrospective place firing has been shown to be absent in the rodent mPFC when past and future location are independent<sup>17</sup>. While the delays used here were relatively brief (10–20 s), multiple previous studies have demonstrated a requirement for the mPFC at similar delays<sup>11,32–35</sup>.

The combined results that vHPC–mPFC input is necessary only during cue encoding, that the mPFC lacks retrospective location representation during maintenance and retrieval, and that the mPFC representation of contemporaneous location during retrieval is not dependent on direct vHPC input, suggest that the mPFC processes sample goal location transiently and that some downstream structure(s) may maintain the information thereafter. Likely candidate structures would include the dHPC and the thalamic nucleus reuniens. We cannot, however, rule out the possibility that retrospective location is persistently represented in the mPFC in a form not detectable using the analytic techniques presented here, such as in temporary changes in synaptic weights, or transient reactivation that is too brief to be reliably decoded from binned spike histograms.

We find that vHPC–mPFC gamma synchrony is correlated with successful location encoding, and that inhibition of this input disrupts long-range gamma-frequency but not theta-frequency synchrony. vHPC gamma oscillations entrain local output in a phase-coherent manner (Extended Data Fig. 6), which influences mPFC spiking at putatively monosynaptic delays. These findings demonstrate that entrainment of mPFC spikes to vHPC gamma oscillations is a physio-

logical signature of task-critical long-range signal propagation. This phenomenon is subtle—at the lag with the greatest phase-locking, fewer than 10% of mPFC units were significantly phase-locked to vHPC gamma. Nevertheless, the observed correlation between vHPC–mPFC gamma synchrony and successful encoding suggests that gamma synchrony could be a behaviourally relevant marker of effective long-range functional connectivity<sup>25</sup>.

The finding that mPFC theta activity leads vHPC theta, and that theta synchrony between the two structures did not appear to depend upon vHPC–mPFC afferents, ran counter to our initial hypotheses. Previous work had shown that dHPC theta leads mPFC theta<sup>10,36</sup>, theta waves travel from the dHPC to vHPC<sup>37,38</sup>, and silencing of vHPC activity affects dHPC–mPFC theta synchrony<sup>39</sup>, supporting the idea that the vHPC theta-patterned activity might directly entrain mPFC theta. The current findings regarding theta synchrony suggest that there may be an alternative explanation for the previous finding of reduced dHPC–mPFC synchrony with vHPC silencing. It is possible that pharmacological inactivation of vHPC affected downstream targets in the dHPC, disrupting dHPC theta activity and thus dHPC–mPFC theta synchrony. This possibility is supported by the finding that vHPC silencing with muscimol reduces dHPC theta power<sup>39</sup>.

Our findings point to a role for the vHPC–mPFC afferent pathway as a conduit for the updating of task-critical location cues, extending previous work that implicates the vHPC and mPFC in contextual learning to short timescales<sup>40–42</sup>. Future work in this area should seek to implicate intermediary and/or upstream structures in the transmission of theta-patterned activity between the HPC and mPFC, to determine what (if any) role the mPFC has in maintenance and/or retrieval of task-related cues, and to identify which of the other major inputs to the mPFC serve to mediate the performance of spatial working memory.

**Online Content** Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 29 November 2014; accepted 25 March 2015.

Published online 8 June 2015.

- Baddeley, A. & Hitch, G. in *Recent Advances in Learning and Motivation* Vol. 8 (ed. Bower, G. A.) 47–90 (Academic, 1974).
- Klauer, K. C. & Zhao, Z. Double dissociations in visual and spatial short-term memory. *J. Exp. Psychol. Gen.* **133**, 355–381 (2004).
- Andrade, J. *Working Memory in Perspective* (Psychology Press, 2001).
- Miyake, A. & Shah, P. *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (Cambridge Univ. Press, 1999).
- Baddeley, A. Working memory: looking back and looking forward. *Nature Rev. Neurosci.* **4**, 829–839 (2003).
- de Zubicaray, G. I., McMahon, K., Wilson, S. J. & Muthiah, S. Brain activity during the encoding, retention, and retrieval of stimulus representations. *Learn. Mem.* **8**, 243–251 (2001).
- Curtis, C. E. & D'Esposito, M. The effects of prefrontal lesions on working memory performance and theory. *Cogn. Affect. Behav. Neurosci.* **4**, 528–539 (2004).
- Hyman, J. M., Zilli, E. A., Paley, A. M. & Hasselmo, M. E. Working memory performance correlates with prefrontal-hippocampal theta interactions but not with prefrontal neuron firing rates. *Front. Integr. Neurosci.* **4**, 2 (2010).
- Jones, M. W. & Wilson, M. A. Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial working memory task. *PLoS Biol.* **3**, e402 (2005).
- Sigurdsson, T., Stark, K. L., Karayiorgou, M., Gogos, J. A. & Gordon, J. A. Impaired hippocampal-prefrontal synchrony in a genetic mouse model of schizophrenia. *Nature* **464**, 763–767 (2010).
- Lee, I. & Kesner, R. P. Time-dependent relationship between the dorsal hippocampus and the prefrontal cortex in spatial memory. *J. Neurosci.* **23**, 1517–1523 (2003).
- Wang, G. W. & Cai, J. X. Disconnection of the hippocampal-prefrontal cortical circuits impairs spatial working memory performance in rats. *Behav. Brain Res.* **175**, 329–336 (2006).
- Hoover, W. B. & Vertes, R. P. Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Struct. Funct.* **212**, 149–179 (2007).
- Jay, T. M. & Witter, M. P. Distribution of hippocampal CA1 and subicular afferents in the prefrontal cortex of the rat studied by means of anterograde transport of Phaseolus vulgaris-leucoagglutinin. *J. Comp. Neurol.* **313**, 574–586 (1991).
- Oh, S. W. A mesoscale connectome of the mouse brain. *Nature* **508**, 207–214 (2014).
- Burton, B. G., Hok, V., Save, E. & Poucet, B. Lesion of the ventral and intermediate hippocampus abolishes anticipatory activity in the medial prefrontal cortex of the rat. *Behav. Brain Res.* **199**, 222–234 (2009).
- Jung, M. W., Qin, Y., McNaughton, B. A. & Barnes, C. L. Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cereb. Cortex* **8**, 437–450 (1998).
- Kjelstrup, K. B. *et al.* Finite scale of spatial representation in the hippocampus. *Science* **321**, 140–143 (2008).
- Royer, S., Sirota, A., Patel, J. & Buzsáki, G. Distinct representations and theta dynamics in dorsal and ventral hippocampus. *J. Neurosci.* **30**, 1777–1787 (2010).
- Chow, B. Y. *et al.* High-performance genetically targetable optical neural silencing by light-driven proton pumps. *Nature* **463**, 98–102 (2010).
- Gradinaru, V. *et al.* Molecular and cellular approaches for diversifying and extending optogenetics. *Cell* **141**, 154–165 (2010).
- Rigotti, M. *et al.* The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).
- Engel, A. K., Konig, P., Kreiter, A. K. & Singer, W. Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science* **252**, 1177–1179 (1991).
- Carr, M. F., Karlsson, M. P. & Frank, L. M. Transient slow gamma synchrony underlies hippocampal memory replay. *Neuron* **75**, 700–713 (2012).
- Yamamoto, J., Suh, J., Takeuchi, D. & Tonegawa, S. Successful execution of working memory linked to synchronized high-frequency gamma oscillations. *Cell* **157**, 845–857 (2014).
- Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P. & Pennartz, C. The pairwise phase consistency: a bias-free measure of rhythmic neuronal synchronization. *Neuroimage* **51**, 112–122 (2010).
- Baeg, E. H. *et al.* Dynamics of population code for working memory in the prefrontal cortex. *Neuron* **40**, 177–188 (2003).
- Horst, N. K. & Laubach, M. Working with memory: evidence for a role for the medial prefrontal cortex in performance monitoring during spatial delayed alternation. *J. Neurophysiol.* **108**, 3276–3288 (2012).
- Funahashi, S., Bruce, C. & Goldman-Rakic, P. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* **61**, 331–349 (1989).
- Goldman-Rakic, P. S. Cellular basis of working memory. *Neuron* **14**, 477–485 (1995).
- Seamans, J. K., Laphs, C. C. & Durstewitz, D. Comparing the prefrontal cortex of rats and primates: insights from electrophysiology. *Neurotox. Res.* **14**, 249–262 (2008).
- Rogers, D. C. *et al.* Photothrombic lesions of the frontal cortex impair the performance of the delayed non-matching to position task by rats. *Behav. Brain Res.* **49**, 231–235 (1992).
- Shaw, C. & Aggleton, J. The effects of fornix and medial prefrontal lesions on delayed non-matching-to-sample by rats. *Behav. Brain Res.* **54**, 91–102 (1993).
- Sloan, H. L., Good, M. & Dunnett, S. B. Double dissociation between hippocampal and prefrontal lesions on an operant delayed matching task and a water maze reference memory task. *Behav. Brain Res.* **171**, 116–126 (2006).
- Izaki, Y., Takita, M. & Akema, T. Specific role of the posterior dorsal hippocampus-prefrontal cortex in short-term working memory. *Eur. J. Neurosci.* **27**, 3029–3034 (2008).
- Siapas, A. G., Lubenov, E. V. & Wilson, M. A. Prefrontal phase locking to hippocampal theta oscillations. *Neuron* **46**, 141–151 (2005).
- Lubenov, E. V. & Siapas, A. G. Hippocampal theta oscillations are travelling waves. *Nature* **459**, 534–539 (2009).
- Patel, J., Fujisawa, S., Berényi, A., Royer, S. & Buzsáki, G. Traveling theta waves along the entire septotemporal axis of the hippocampus. *Neuron* **75**, 410–417 (2012).
- O'Neill, P. K., Gordon, J. A. & Sigurdsson, T. Theta oscillations in the medial prefrontal cortex are modulated by spatial working memory and synchronize with the hippocampus through its ventral subregion. *J. Neurosci.* **33**, 14211–14224 (2013).
- Ruediger, S., Spirig, D., Donato, F. & Caroni, P. Goal-oriented searching mediated by ventral hippocampus early in trial-and-error learning. *Nature Neurosci.* **15**, 1563–1571 (2012).
- Fanselow, M. A. & Dong, H. W. Are the dorsal and ventral hippocampus functionally distinct structures? *Neuron* **65**, 7–19 (2010).
- Komorowski, R. W. *et al.* Ventral hippocampal neurons are shaped by experience to represent behaviorally relevant contexts. *J. Neurosci.* **33**, 8079–8087 (2013).

**Supplementary Information** is available in the online version of the paper.

**Acknowledgements** The authors would like to thank M. Topiwala for technical assistance, M. Kheirbek for advice and assistance with the design of fibre optics, and M. Shapiro for advice with regards to designing the four-arm T-maze. This work was supported by grants from the National Institutes of Health (MH096274 and MH081968), the Hope for Depression Research Foundation, the International Mental Health Research Organization, the Gatsby Charitable Foundation and the Swartz Foundation.

**Author Contributions** T.S., J. A. Gogos and J. A. Gordon designed the experiments. T.S. performed the experiments and analysed the data. M.R. and S.F. developed the linear classifier, adapted it for use with the T-maze data set, and provided guidance on its implementation. S.E.A. participated in the design of optogenetic experiments. T.S., S.F., J. A. Gogos and J. A. Gordon interpreted the results. T.S. and J. A. Gordon wrote the paper.

**Author Information** Reprints and permissions information is available at [www.nature.com/reprints](http://www.nature.com/reprints). The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to J. A. Gordon ([jg343@columbia.edu](mailto:jg343@columbia.edu)).

## METHODS

**Subjects.** Male C57BL/6 mice (Jackson Labs) were used for all experiments, aged 8–12 weeks at first use. Mice were housed in a New York State Psychiatric Institute satellite facility and were maintained on a 12-h light–dark cycle. Except when food-restricted for the purpose of behavioural training and testing, all mice were given *ad libitum* access to food and water. Pre-surgical mice were group-housed with littermates, while mice with chronic recording implants were singly housed in divided cages with visual, auditory and olfactory contact with another implanted mouse. Six mice were used in the acute anaesthetized experiment (3 Arch<sup>+</sup>, 3 Arch<sup>-</sup>), 14 in the two-goal T-maze experiment (8 Arch<sup>+</sup>, 6 Arch<sup>-</sup>), and 13 in the four-goal T-maze experiment (7 Arch<sup>+</sup>, 6 Arch<sup>-</sup>). The effect size for behavioural impairment in the Arch<sup>+</sup> group in the four-goal experiment was 1.48, and the probability of observing a significant effect (statistical power) was 90% for an *N* of 7, 83% for an *N* of 6 (Arch<sup>+</sup> and Arch<sup>-</sup> group sizes, respectively). Mice were randomized to a given viral type. All procedures were approved by Columbia University and the New York State Psychiatric Institute Institutional Animal Care and Use Committees.

**Surgical preparation.** Animals were placed inside a flow box and anaesthetized with isoflurane gas (2%) until sedated, at which point they were placed in a stereotax and maintained on 0.5% isoflurane for the duration of the surgery. Craniotomies were made bilaterally above the mPFC, dHPC and vHPC (coordinates below), and skull screws placed over the cerebellum and olfactory bulb served as ground and reference, respectively. In the acute anaesthetized experiment, surgical depth was maintained with isoflurane for the duration of the experiment.

**Viral transduction.** AAV2/5 of titre exceeding 10<sup>12</sup> vg ml (K. Deisseroth via UNC Vector Core and UPenn Vector Core) was used to package the virus. In the acute stimulation/silencing experiment and the two-goal T-maze experiment, a CamKII $\alpha$ -eArch3.0-eYFP sequence was used to express the opsin and CamKII $\alpha$ -mCherry was used as an opsin-negative control. For the four-goal experiment, hSyn-eArch-eYFP and hSyn-eYFP were used for opsin and control, respectively. The hSynapsin promoter was chosen for the four-goal experiment to account for possible long-range GABAergic vHPC–mPFC projections and to avoid potential toxicity effects resulting from opsin expression under the stronger CamKII $\alpha$  promoter. Virus was targeted to multiple targets within the stratum pyramidale of ventral CA1 (two mediolateral rows at anteroposterior (AP) 2.95 and 3.25, with sites at mediolateral (ML)/dorsoventral (DV): 2.65/4.5, 3.0/4.3, 3.35/3.9, 3.7/3.3–2.9. An additional row was made at AP 3.1, with ML/DV sites at 2.8/1.55 and 3.15/1.7. All coordinates are reported in mm, all AP and ML coordinates are with respect to bregma, DV coordinates with respect to brain surface. A 200 nl bolus was delivered to each site via glass micropipette (20–40  $\mu$ M diameter) at a rate of 100 nl min<sup>-1</sup> continuous infusion, with a wait time of 5 min between infusion and retraction.

**Electrode and fibre implantation.** LFPs were recorded using 50- $\mu$ M-diameter tungsten wire, while spikes and LFPs were recorded using stereotrodes (mPFC) and tetrodes (vHPC) made from 13- $\mu$ M-diameter tungsten fine wire. For the two-goal T-maze experiment, stereotrodes were coupled to ferrule-bound optical fibres (Thorlabs, 200- $\mu$ M-diameter core, 0.39 NA) positioned 300–500  $\mu$ M dorsal to the stereotrode tips, which were arrayed semi-circularly around the lateral edge of the fibre. Fibre-coupled stereotrode bundles were then implanted bilaterally in the mPFC (1.8 mm anterior, 0.4 mm lateral, 1.4 mm ventral), while LFP wires were implanted bilaterally in the dHPC (1.85 mm posterior, 1.25 mm lateral, 1.45 mm ventral) and the vHPC (3.1 mm posterior, 3.0 mm lateral, 3.9 mm ventral). For the four-goal experiment, 9 of the 13 animals were implanted as described earlier, while 4 animals (2 Arch<sup>+</sup> and 2 Arch<sup>-</sup>) were instead implanted with tetrodes over the vHPC (3.7 mm ventral), which were advanced until spikes with putative pyramidal-cell waveforms were detected (Extended Data Fig. 4). mPFC optical fibres in these 4 animals were coupled to stationary bilateral LFP wires. All movable microdrives were advanced at a rate of 40  $\mu$ M per day across all recording days. Recording sites were histologically confirmed by visual examination of electrothermolytic lesions made before killing and perfusing implanted animals. Lesions were induced by passing current through an electrode at each implanted site (50  $\mu$ A, 20 s). Perfused and fixed tissue was then sectioned, and DNA was stained using 4',6-diamidino-2-phenylindole (DAPI) Fluoromount-G mounting medium (Southern Biotech). vHPC LFP wires used for phase-locking analysis were located in the stratum pyramidale, stratum radiatum or stratum lacunosum-moleculare of ventral CA1; no differences in gamma power, peak frequency or synchrony with mPFC units were seen for the different laminar locations.

**Recording and spike sorting.** Recordings were amplified, band-pass filtered (1–1,000 Hz LFPs, 600–6,000 Hz spikes), and digitized using the Neuralynx Digital Lynx system. LFPs were collected at a rate of 2 kHz, while spikes were detected by online thresholding and collected at 32 kHz. Units were initially clustered using Klustakwik (Ken Harris), sorted according to the first two principal components,

voltage peak and energy from each channel. Clusters were then accepted, merged or eliminated based on visual inspection of feature segregation, waveform distinctiveness and uniformity, stability across recording session, and inter-spike interval distribution.

**Behavioural training.** Mice undergoing behavioural testing were given 5 weeks to recover from surgery and allow for expression and transport of the opsins, at which time they were placed on a food-restricted diet consisting of 1.5–2.5 g of food per day as needed to maintain 85% of their post-recovery bodyweight. The researcher responsible for training and running the mice was blinded as to viral group by another laboratory member, who randomly assigned opsin and control viruses. Mice were then given 2 days of habituation to the maze, which consisted of 10 min free exploration and foraging (plugged into optical fibres and recording tether) with all doors open, followed by 5 min of interleaved laser light pulses in the start box with doors closed (5–10 s on, 30 s off, 5 min). On the subsequent 2 days mice underwent behavioural shaping consisting of 10 min of running to baited goal arms in alternating directions. Mice then underwent training on the T-maze task until criterion performance (consisting of 70% correct trials on 2 out of 3 consecutive days) was achieved, in the absence of illumination. All animals that underwent successful surgeries met criterion. Inter-trial delay was 20 s. On the two-goal experiment, reward consisted of dustless pellets (Bio-Serv). Also in the two-goal experiment only, to minimize across-session performance drift and ceiling effects, intra-trial delay was adjusted according to the preceding day's performance, beginning at 10 s and progressively increasing delay by an additional 5 s following days on which performance was above 80% (as in ref. 12). For the four-goal experiment, reward consisted of sweetened condensed milk (~50  $\mu$ l, 3:1 dilution). To allow for direct comparison of delay within animals in this task, intra-trial delay was fixed at 10 s, with the addition of a single session of 20 s delays as the final session for each animal. Light stimulation was 532 nm, 10 mW in all cases. Task performance was lower when sample and choice goals matched in the left/right dimension than when they differed (*n* = 13 mice; mean = 0.65  $\pm$  0.01 and 0.71  $\pm$  0.02, respectively; *t* = -3.91, *P* = 0.001), but performance on these trials remained above chance (*t* = 4.0  $\times$  10<sup>3</sup>, *P* = 3.1  $\times$  10<sup>-38</sup>).

**Statistics.** All effects presented as statistically significant exceeded an  $\alpha$ -threshold of 0.05. All independence tests were two-tailed. All independence testing of paired values (that is, changes across conditions) used paired *t*-tests or (where stated) signed rank tests. All ANOVA tests involving multiple observations per subject/unit (every figure except Extended Data Fig. 2d) were done as repeated measures group-by-condition tests. All *t*-tests and rank tests performed with more than two groups were done post-hoc to ANOVA tests except where Bonferroni correction for multiple comparisons is specifically cited.

**Spike analysis using a linear classifier.** A maximum margin linear classifier was used to decode sample goal location from binned population firing rate vectors and to quantify the strength and reliability of the neural representation. This type of classifier had the advantage of integrating comparisons between each pair of goal locations into a single inference, thus capturing all possible schemes of location discrimination in a way that can be applied across a diverse neural population. Matlab scripts are available for sharing upon request.

Analyses were performed on all units for which there was at least one training and one testing trial per condition.

Analysis using the population decoder was performed on binned spike vectors (500 ms bins, 100 ms increments) of all recorded units from sessions with at least two trials for each feature class under consideration. As in previous work with spike data from the PFC, we therefore took advantage of the accumulated data to decode task-relevant features by treating all units as a single, pseudosimultaneously recorded population<sup>43</sup>. This approach assumes equivalent spike statistics across animals and recording sessions, and ignores potential contributions to population coding from correlated spike variability that would be observed if all units were indeed simultaneously recorded<sup>44</sup>.

Model training was performed using constrained quadratic programming<sup>45,46</sup> that employed a maximal margin perceptron<sup>46–48</sup>. The training samples to this algorithm are generated by averaging the recorded spike counts within the relevant conditions that need to be decoded, and across a specified training set of trials. This procedure gives a mean activity vector per condition, in which each component of the vector represents the trial-averaged activity of a given neuron at a given condition. The quadratic programming procedure then aims at finding a set of readout weights that maximally separate the mean activity vectors corresponding to the conditions that have to be discriminated. Multi-class discrimination problems are reduced to a set of binary discrimination problems involving all pairwise combinations of conditions<sup>49</sup>.

The model was then tested by cross-validating its performance on a test set of recorded trials. Given a task condition, a test vector is generated by sampling its components from the distribution of spike counts recorded from the corresponding neuron during the test trials. At each test phase 100 test vectors per condition

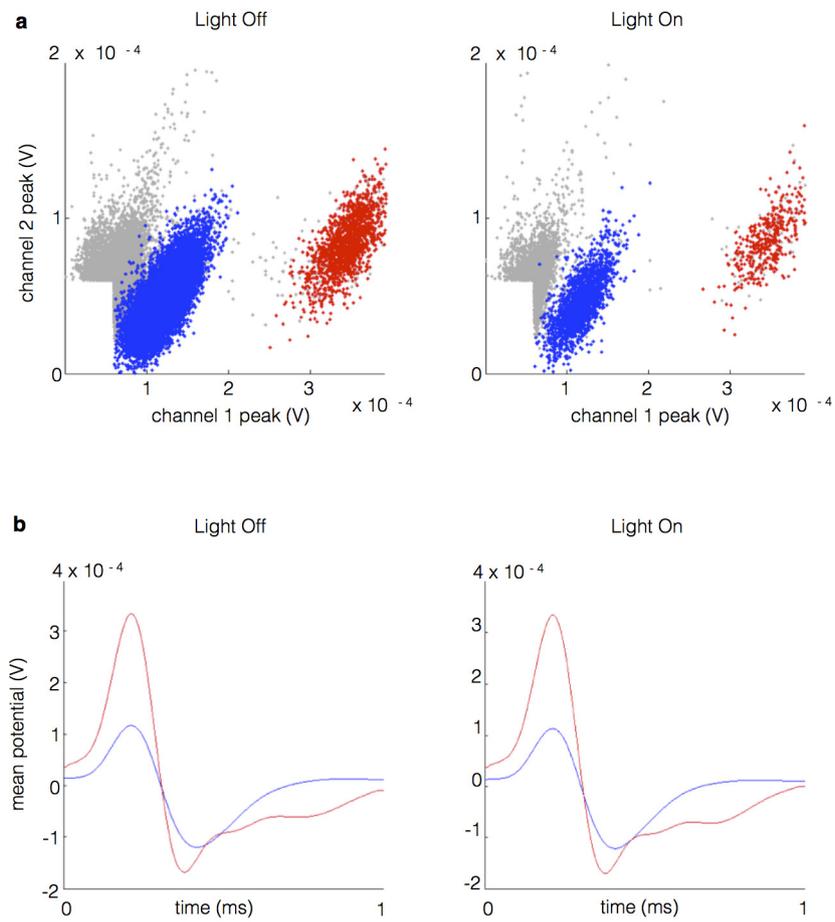
are randomly resampled with replacement, and the performance of the model is quantified as the average accuracy in classifying them. For each time bin, model training and testing was performed 100 times (at which point estimates of model accuracy approached asymptote) on non-overlapping subsets of trials (half of trials to train, half to test, random subsampling without replacement), with subsets constrained to include at least one trial corresponding with each feature class under consideration. For training and testing across separate trial conditions (Fig. 4), the same trial number requirement was applied across training and testing sets to ensure equal population sizes and equal representation of feature classes in the two trial sets.

For identification of each unit's preferred goal, absolute values of model weights for each goal's three binary classifications (comparison with each other goal) were summed, and the goal with the highest summed value was judged to be the preferred goal. To segregate units with high and low goal selectivity, the absolute model weights for all six binary classifications (each goal-goal comparison) were summed; units in the upper and lower quintiles were judged to have high and low goal selectivity, respectively.

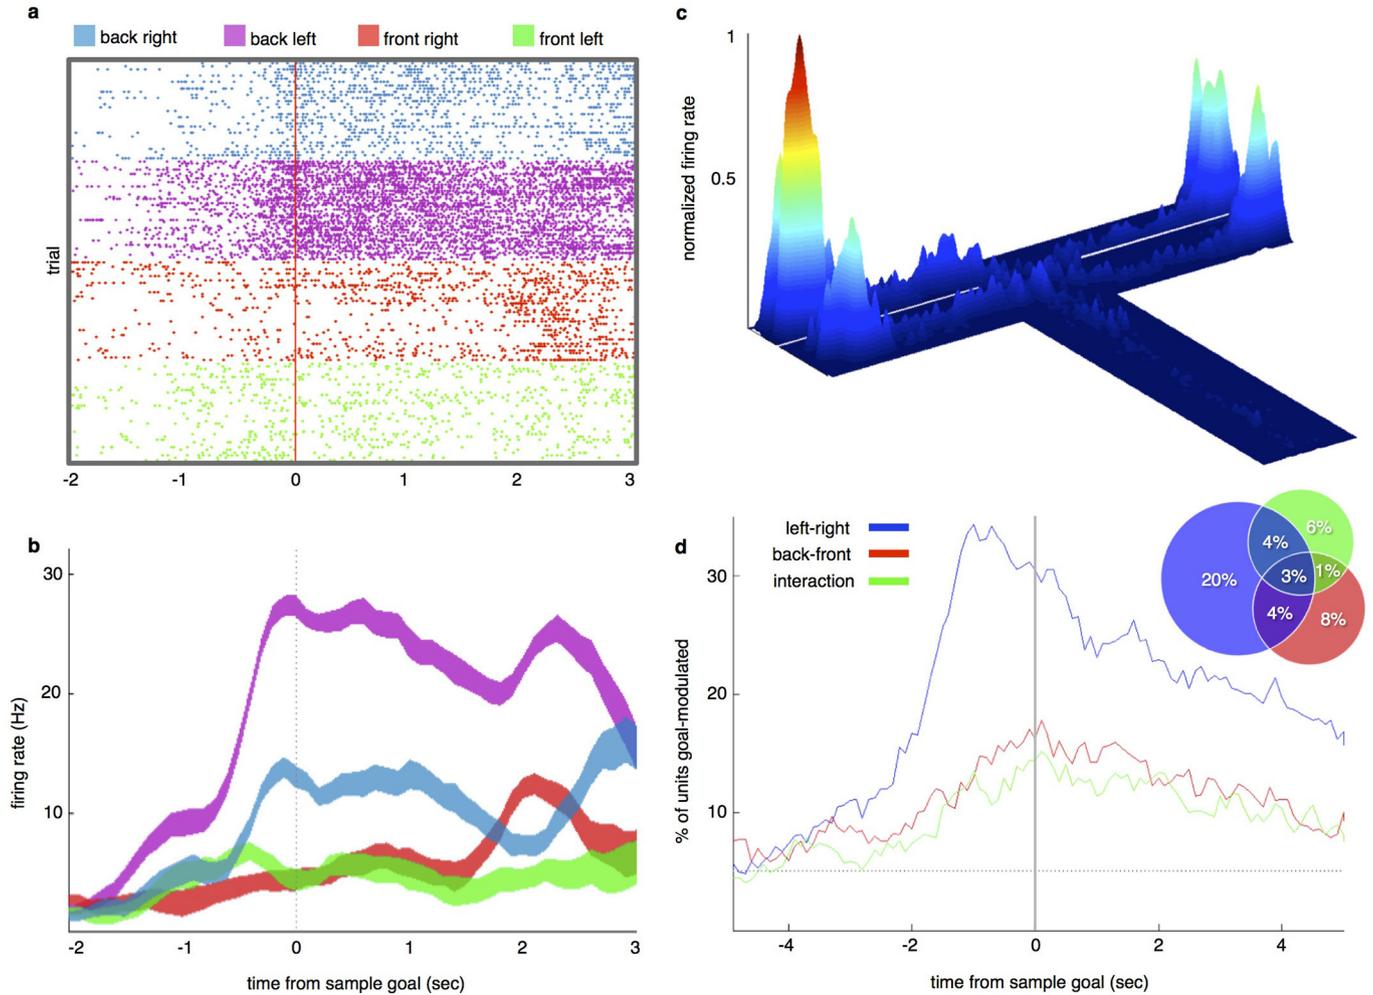
**Phase-locking analysis.** Phase-locking of spikes to the oscillatory phase of LFPs was performed using pairwise phase consistency, which, unlike other commonly used measures of phase-locking, is unbiased by spike number<sup>26</sup>. Nevertheless, to ensure a representative estimate of spike phase, we set a threshold of 100 spikes for

all analysis; in comparisons of phase-locking across conditions (Fig. 5h-j and Extended Data Fig. 7h), only units that fired 100 spikes in each condition were included. LFP signal was digitally band-pass filtered (4–12 Hz for theta, 30–70 Hz for gamma) using a zero-phase-delay filter (filter provided by K. Harris and G. Buzsaki, order = sample frequency). The phase component was calculated by a Hilbert transform, and a corresponding phase was assigned to each spike.

43. Meyers, E. M., Freedman, D. J., Kreiman, G., Miller, E. K. & Poggio, T. Dynamic population coding of category information in inferior temporal and prefrontal cortex. *J. Neurophysiol.* **100**, 1407–1419 (2008).
44. Abbott, L. F. & Dayan, P. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* **11**, 91–101 (1999).
45. Barak, O. & Rigotti, M. A simple derivation of a bound on the perceptron margin using singular value decomposition. *Neural Comput.* **23**, 1935–1943 (2011).
46. Anlauf, J. K. & Biehl, M. The AdaTron: an adaptive perceptron algorithm. *Europhys. Lett.* **10**, 687–692 (1989).
47. Rosenblatt, F. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms* (Spartan, 1962).
48. Krauth, W. & Mezard, M. Learning algorithms with optimal stability in neural networks. *J. Phys. A Math. Gen.* **20**, L745–L752 (1987).
49. Dietterich, T. G. & Ghulum, B. Error-correcting output codes: a general method for improving multiclass inductive learning programs. *Proc. AAAI* **91**, 572–577 (1991).

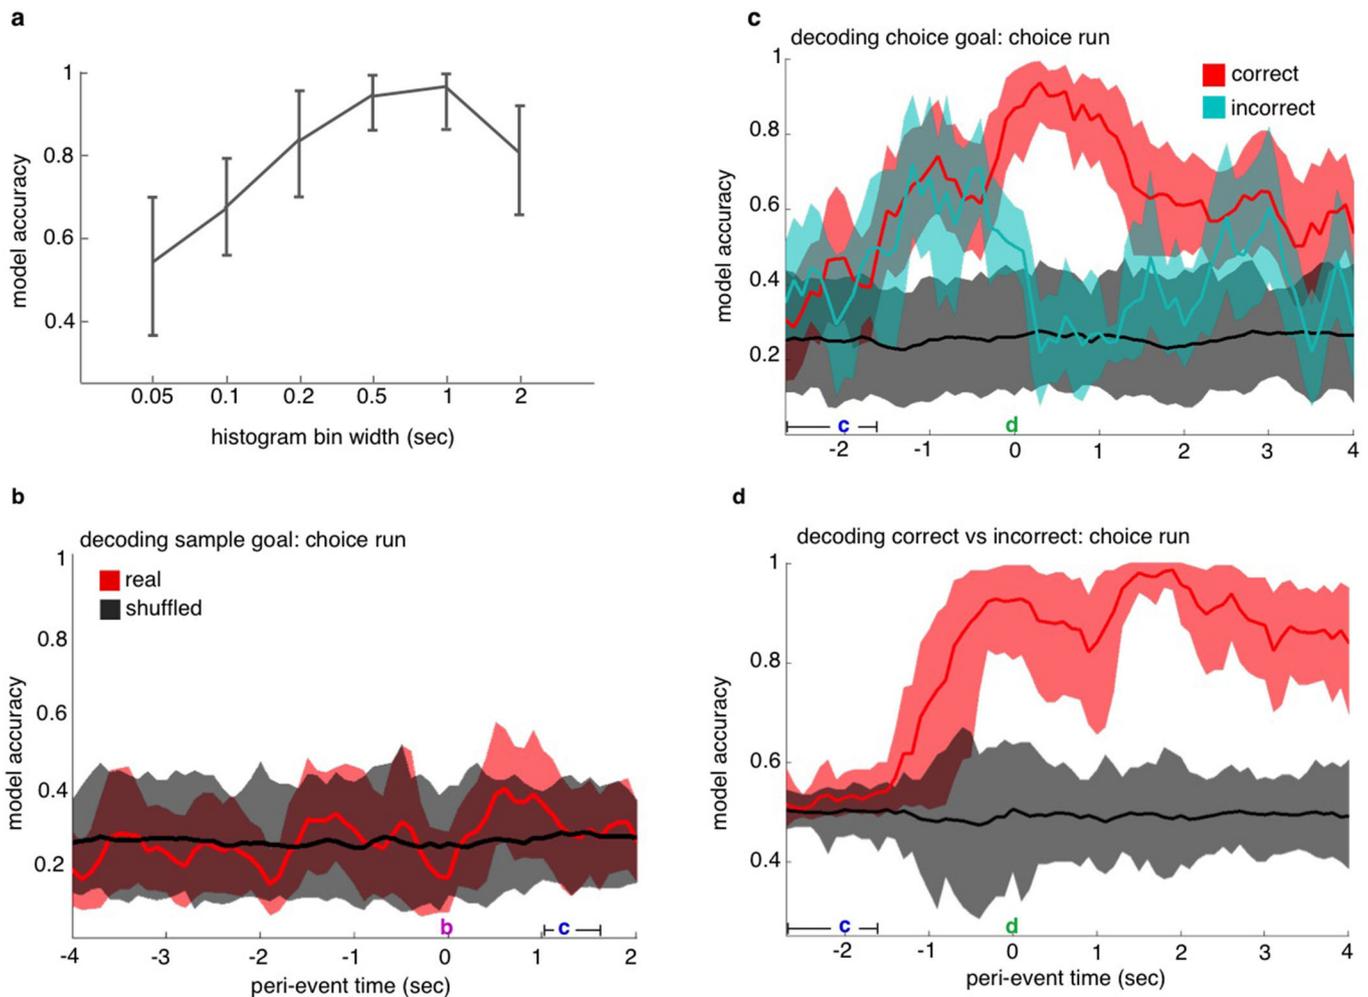


**Extended Data Figure 1 | Individual mPFC units clustered from fibre-coupled stereotrodes. a**, Multiple individual units clustered from stereotrode recordings in the mPFC in the absence and presence of illumination. **b**, Mean waveforms of extracellular potentials from example units in **a**.



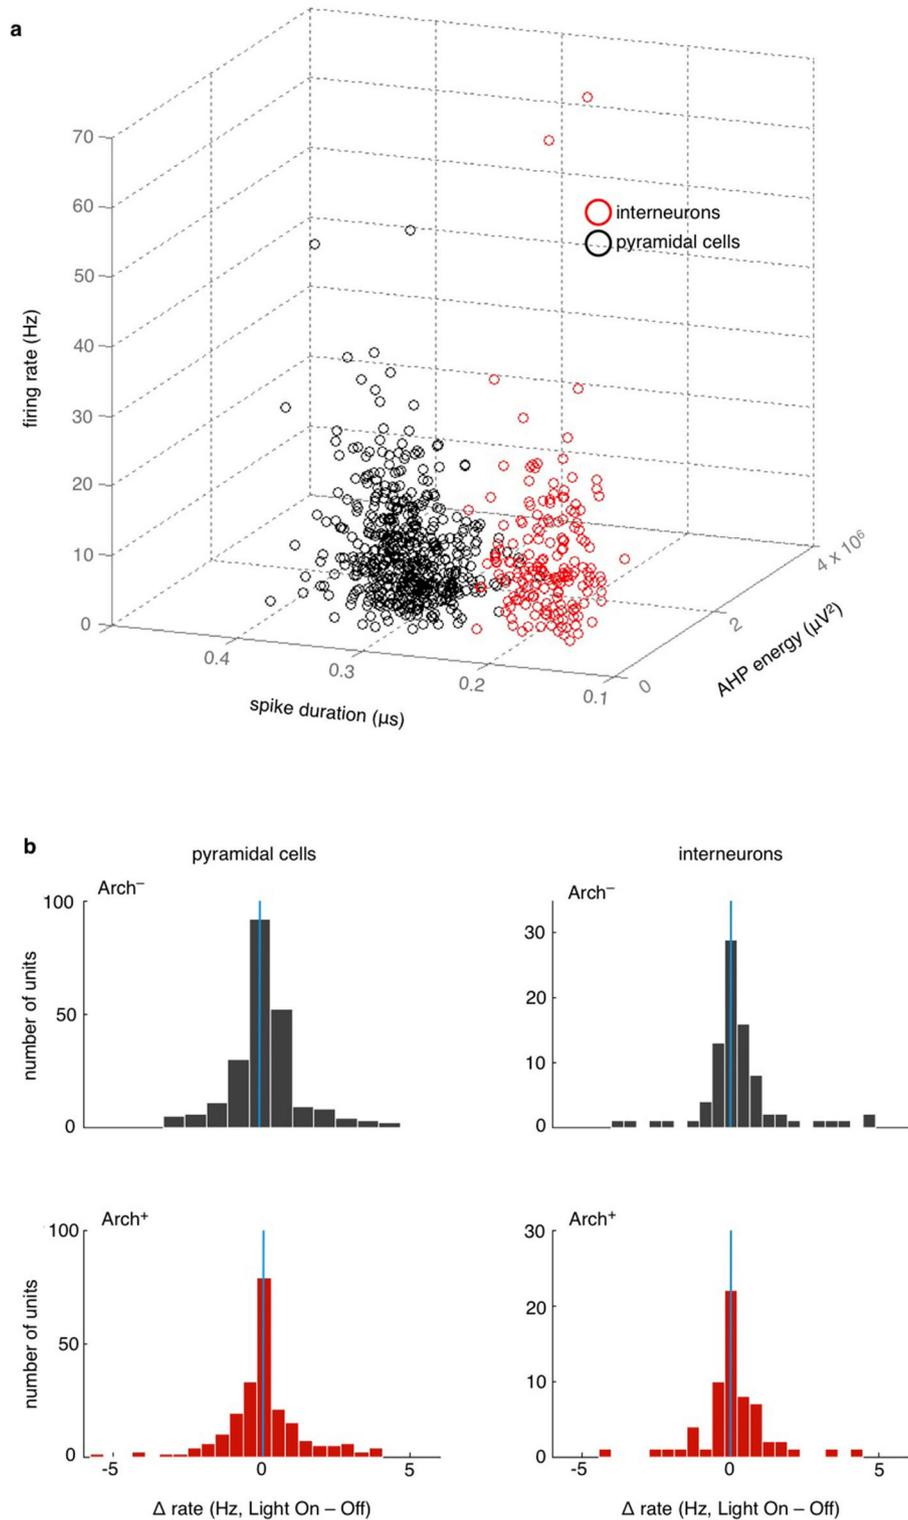
**Extended Data Figure 2 | mPFC cells encode goal location both categorically and globally.** **a**, A raster plot of spikes fired by an example single unit across trials, sorted by sample goal, temporally aligned to arrival at sample goal. **b**, Traces of firing rates averaged across trials by sample goal location, for the unit from **a**. This unit shows location selectivity, firing preferentially in the back left goal. Traces are mean  $\pm$  s.e.m. **c**, Spatial map of firing rates for the same unit for the full recording session. Goal-selective units tended to fire more at the preferred goal than at the other goals, and more

at all goals than in the rest of the environment. **d**, Percentage of units that were goal-selective as a function of time from sample goal, according to two-way repeated measures ANOVAs performed on binned spike rates. Units were identified as having selectivity for left/right (blue), back/front (red), and/or combined spatial dimensions (green). Dashed line represents chance ( $P = 0.05$ ). Inset, percentage of units having each type and/or combination of selectivity at time zero (arrival at sample goal). Percentages are out of 792 recorded units.



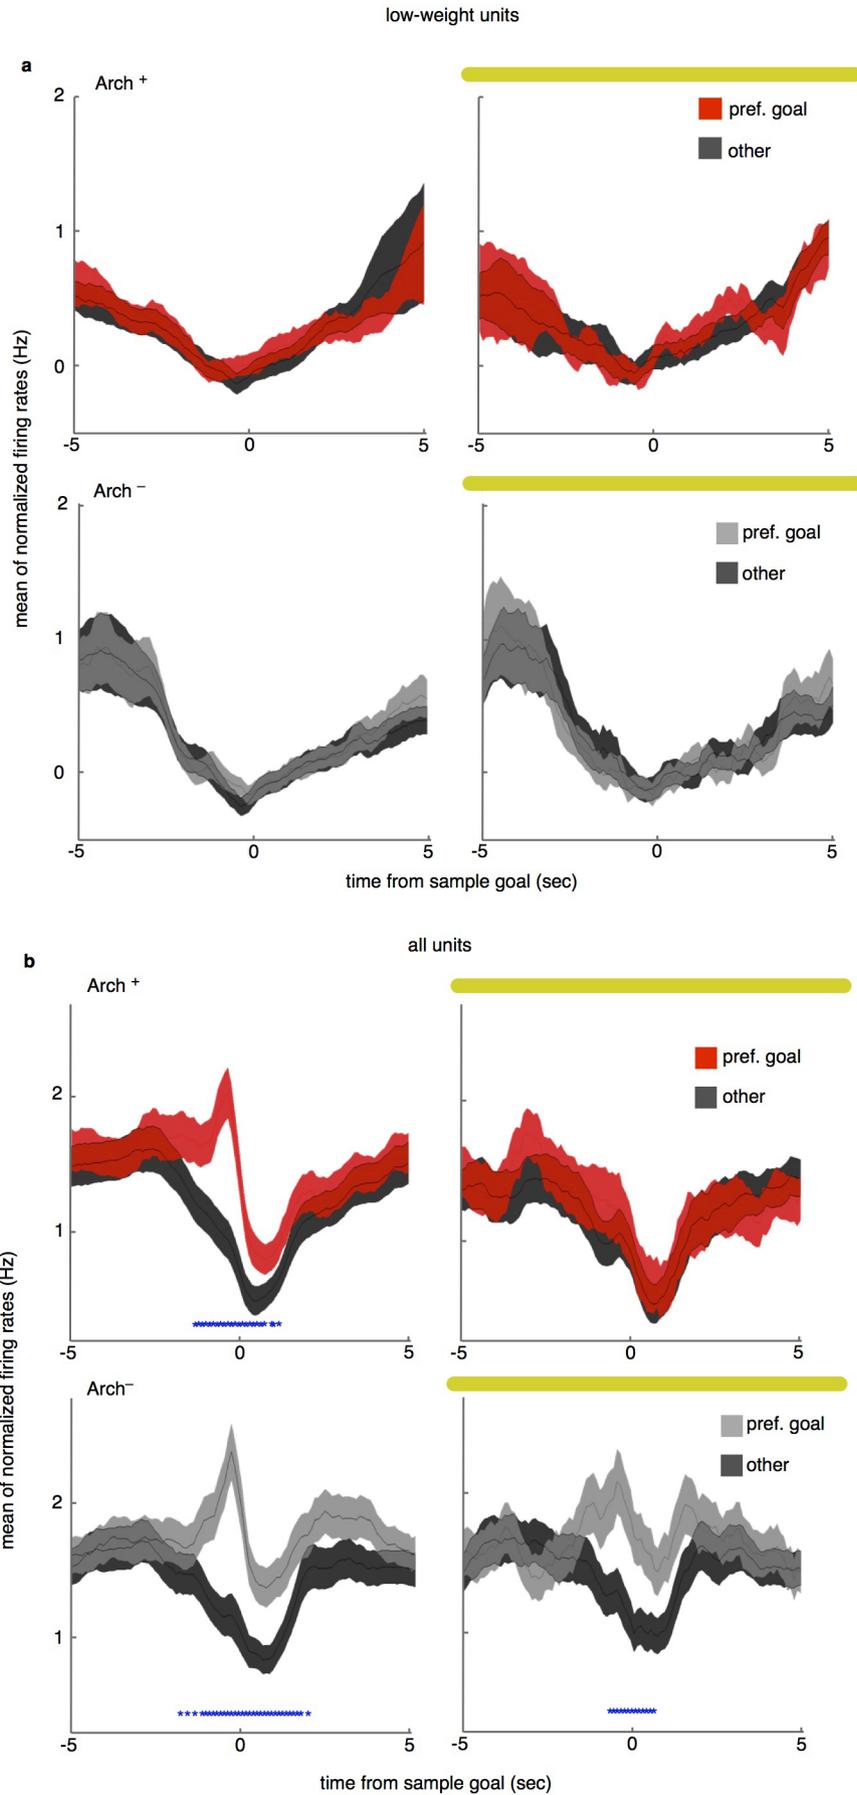
**Extended Data Figure 3 | mPFC units represent choice goal location, not sample goal location, during choice runs.** **a**, Model accuracy at the time bin corresponding with arrival at the sample goal port during the four-goal task was highest for spike histograms with time bins of 500 ms and 1,000 ms. Five-hundred-millisecond time bins were used for spike analyses. **b**, Decoding sample goal location during subsequent choice run during the four-goal task. Using the linear decoder, previously visited location was not detectable above chance accuracy. Ten- and twenty-second delay trials were combined. **c**, Decoding choice goal during choice run, correct versus incorrect trials during the four-goal task. Location decoded for this analysis was chosen goal (that is,

the mouse's current location) rather than correct goal. Model accuracy reached 0.93 upon arrival at the goal on correct trials. On incorrect trials, model accuracy exceeded chance during goal approach but dropped to chance levels upon reaching the goal. Ten- and twenty-second delay trials were combined. **d**, Decoding choice accuracy (correct versus incorrect) during choice trials. Model accuracy peaked at 0.99 at 1.9 s after arrival at the goal. **b–d**, Histograms were aligned to departure from start box. Ten- and twenty-second delay trials were combined. Data show mean  $\pm$  95% confidence intervals for **b** and **d**; s.e.m. for **c** and **e**.



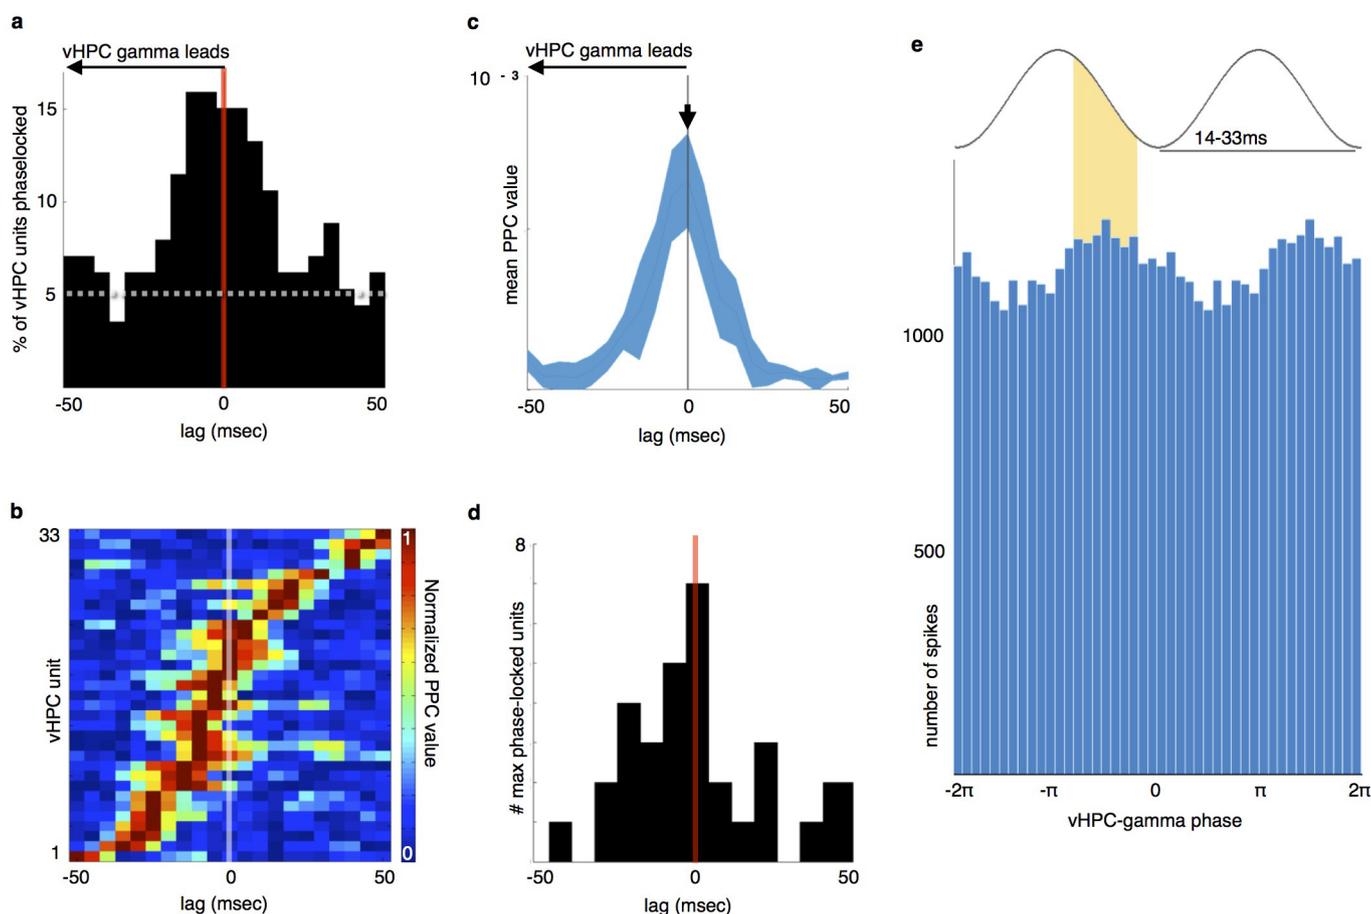
**Extended Data Figure 4 | vHPC-mPFC terminal inhibition does not alter mPFC spike rate.** **a**, Waveform features used to separate putative cell types. Spike duration was defined as the peak-to-trough time, while afterhyperpolarization (AHP) energy was taken as the area over the curve after the second zero-crossing. Spike duration yielded the clearest separation.

**b**, Putative fast-spiking (FS) and non-FS cells, sorted by spike width, showed no effect of terminal illumination on spike rate (Arch<sup>-</sup> non-FS: sign rank  $z = -1.7$ ,  $P = 0.095$ ; Arch<sup>-</sup> FS:  $z = -1.6$ ,  $P = 0.11$ ; Arch<sup>+</sup> non-FS:  $z = -2.7$ ,  $P = 0.79$ ; Arch<sup>+</sup> FS:  $z = -0.49$ ,  $P = 0.62$ ).



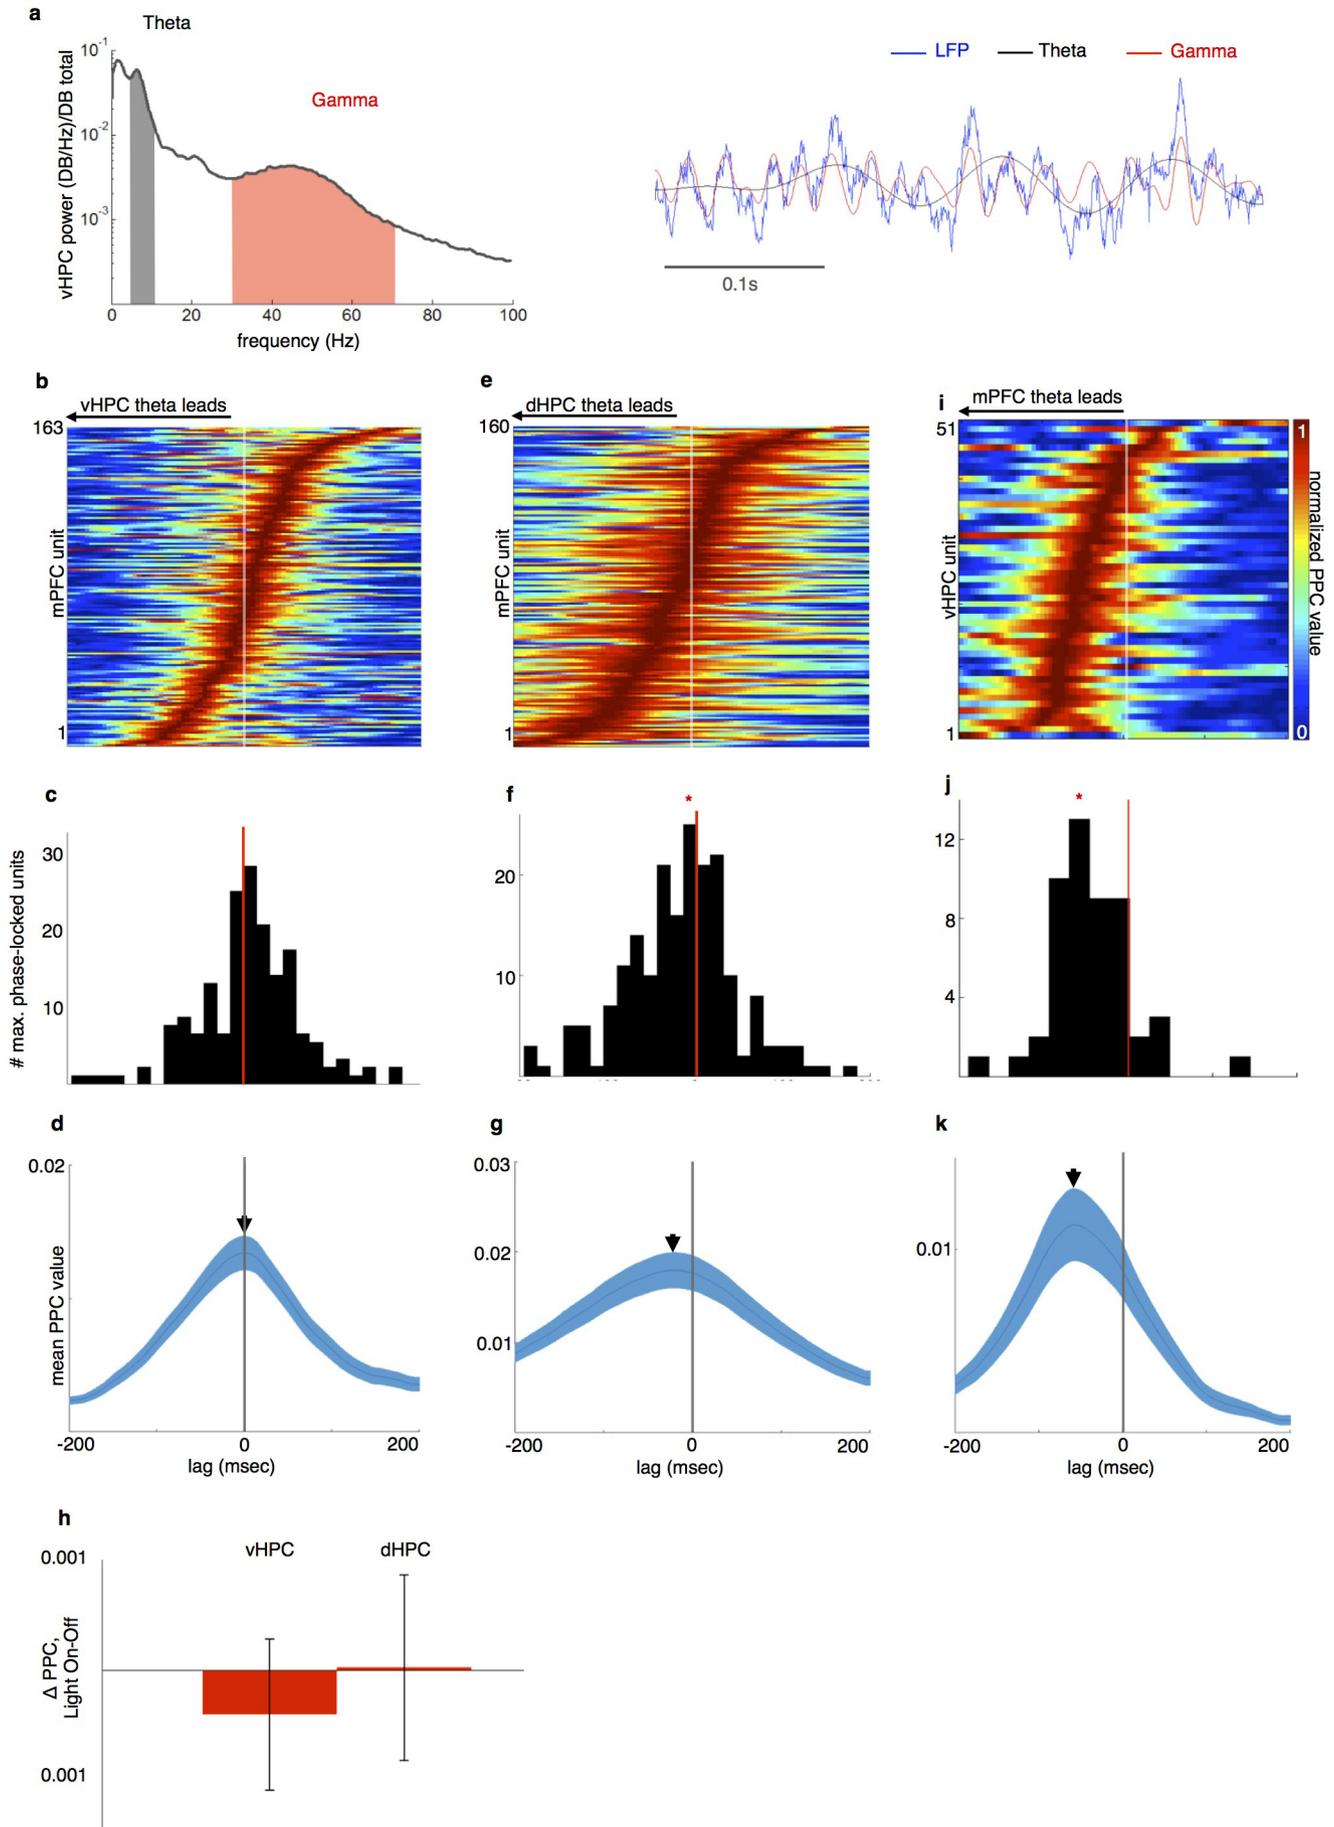
**Extended Data Figure 5 | Effect of mPFC illumination on goal-selective firing in the mPFC.** **a**, Low-weighted units, as identified using the classifier, show no difference in firing between the goal with the highest weight relative to the other goals. In the sample goal these units fire at rates not different than their session mean rates. Traces indicate mean  $\pm$  s.e.m. of normalized firing rate (bin FR – session FR). **b**, Terminal inhibition eliminates firing rate differences in preferred (Pref.) versus non-preferred (Other) goal during

encoding across all units. On sample runs with no light, units from both Arch<sup>-</sup> (bottom left) and Arch<sup>+</sup> animals (top left) had elevated firing rates in preferred goal relative to non-preferred goal (red asterisks mark time points with Bonferroni-corrected significance). In Sample Light runs, units from Arch<sup>-</sup> animals maintain elevated firing in the preferred goal (bottom right), while units from Arch<sup>+</sup> animals show no significant firing rate difference (top right;  $N = 358$  Arch<sup>-</sup> units, 325 Arch<sup>+</sup> units, sign rank  $P < 0.0005$ ).



**Extended Data Figure 6 | vHPC gamma modulates vHPC output.** **a**, vHPC units phase-lock maximally to the vHPC gamma rhythm at a lag of zero ( $P$  value from Rayleigh's test  $< 0.05$ , dashed line indicates chance rate). **b**, Normalized PPC values, sorted by lag of maximal phase-locking, for significantly phase-locked vHPC units. Units with Bonferroni-corrected significance within the  $-40$  to  $40$  ms lag window (Rayleigh test,  $P < 0.0029$ ) were included. **c**, Mean normalized PPC value for the population shown in **b**. Shading is s.e.m. **d**, Histogram of units with maximum PPC value at each lag.

Units maximally phase-locked at a lag of zero, with no net difference from zero across the population. **e**, vHPC units share a common preferred gamma phase. Pooled spikes from significantly phase-locked vHPC units were modulated by vHPC gamma phase at zero-lag ( $N = 26,303$  spikes, Rayleigh's  $z = 17.6$ ,  $P = 2.2 \times 10^{-8}$ , PPC value = 0.002), with peak spiking in the descending phase of the gamma cycle. (Note that spikes and LFPs were both recorded from stereotrodes in the stratum pyramidale and that this gamma phase would probably differ from that recorded in SLM, as in Fig. 5).



**Extended Data Figure 7 | mPFC theta activity follows dHPC and leads vHPC during the task.** **a**, Example vHPC LFP (blue, right) and spectrogram (left) demonstrating robust theta (grey, 4–12 Hz) and gamma (red, 30–70 Hz) components during all runs towards goals. **b**, Pseudocolour plot of relative strength of mPFC unit phase-locking to vHPC theta at lags from –200 ms to 200 ms, for units with Bonferroni-corrected significance in at least one lag. Warmer colours indicate stronger phase-locking. **c**, Distribution of lags at peak phase-locking strength for significantly phase-locked mPFC units. Distribution centred at 0 ( $N = 189$  units,  $z = 2.05$ ,  $P = 0.98$ ). **d**, Mean  $\pm$  s.e.m. PPC value of mPFC units and vHPC theta, as a function of lag. **e–g**, Phase-locking of

mPFC units to dHPC theta as a function of lag, as in **b–d**. Distribution of lags at peak phase-locking is significantly shifted towards a dHPC lead ( $N = 160$  units, sign rank  $z = -4.4$ ,  $P = 6 \times 10^{-6}$ ). **h**, No difference in strength of phase-locking of mPFC units to vHPC (left) and dHPC (right) theta in light on versus light off trials. Mean and s.e.m. shown for each ( $N = 140$  units, sign rank  $z = -1.3$ ,  $P = 0.2$ ;  $z = -1.4$ ,  $P = 0.12$ ). **i–k**, Phase-locking of vHPC units to mPFC theta as a function of lag, as in **b–d**. Distribution of lags at peak phase-locking is significantly shifted towards an mPFC lead ( $N = 51$  units,  $z = -5.03$ ,  $P = 2.4 \times 10^{-7}$ ).