Closed-loop modulation of remote hippocampal representations with neurofeedback

Michael E. Coulter[1,2], Anna K. Gillespie[3], Joshua Chu[4], Eric L. Denovellis[1,2], Trevor T.K. Nguyen[1,5], Daniel F. Liu[1,2], Katherine Wadhwani[1,2], Baibhav Sharma[1,2], Kevin Wang[5], Xinyi Deng[6], Uri T. Eden[7], Caleb Kemere[4], Loren M. Frank[1,2]

[1]Kavli Institute and Department of Physiology UCSF, [2]Howard Hughes Medical Institute, [3]Departments of Biological Structure and Lab Medicine and Pathology, University of Washington, [4]Neuroengineering Initiative, Rice University, [5]SpikeGadgets Inc., [6]Dept. of Statistics, Beijing University of Technology, [7]Dept. of Mathematics and Statistics, Boston University

Correspondence: loren.frank@ucsf.edu

SUMMARY

Humans can remember specific events without acting on them and can influence which memories are retrieved based on internal goals. However, current animal models of memory typically present sensory cues to trigger retrieval and assess retrieval based on action[1–5]. As a result, it is difficult to determine whether measured patterns of neural activity relate to the cue(s), the retrieved memory, or the behavior. We therefore asked whether we could develop a paradigm to isolate retrieval-related neural activity in animals without retrieval cues or the requirement of a behavioral report. To do this, we focused on hippocampal "place cells." These cells primarily emit spiking patterns that represent the animal's current location (local representations), but they can also generate representations of previously visited locations distant from the animal's current location (remote representations)[6–13]. It is not known whether animals can deliberately engage specific remote representations, and if so, whether this engagement would occur during specific brain states. So, we used a closed-loop neurofeedback system to reward expression of remote representations that corresponded to uncued, experimenter-selected locations, and found that rats could increase the prevalence of these specific remote representations over time; thus, demonstrating memory retrieval modulated by internal goals in an animal model. These representations occurred predominately during periods of immobility but outside of hippocampal sharp-wave ripple (SWR)[13–15] events. This paradigm enables future direct studies of memory retrieval mechanisms in the healthy brain and in models of neurological disorders.

MAIN TEXT

Remembering is distinct from acting. Humans can remember places we have been and experiences we have had without overt behavioral signs that these memories have been retrieved. Remembering can also happen in the absence of specific external retrieval cues. Humans can retrieve specific memories based on internal goals[16], and in some cases memories can even seem to be retrieved spontaneously[17]. Thus, memory retrieval is a process that can be expressed in the brain separately from both external cues that trigger the memory and the decision to act based on the content of the memory.

Current animal models of memory are not designed to allow the direct study of the neural substrates of retrieval not confounded by sensory cues or the decision to act. First, current experimental approaches typically engage retrieval by presenting external cues such as those found in a particular spatial context[1,18–20]. These cues can strongly influence patterns of brain activity, including spiking in brain regions like the hippocampus[6] that contribute to memory formation and retrieval[19]. Thus, representations of the current sensory inputs can be difficult to disentangle from memory-related activity. Second, current approaches typically assess whether a memory has been retrieved based on a behavioral report. Widely used paradigms including contextual-fear conditioning[2–4], the Morris Water Maze[18], and many others place animals in a familiar context and quantify memory based on whether the animal's exhibit specific behaviors (freezing or directed navigation to a remembered goal location). Activity patterns measured in these paradigms are the result of a complex combination of sensory information processing, memory retrieval, a decision process, and action, and so it is difficult to isolate and study the retrieval process on its own.

We therefore asked whether we could develop a paradigm that would permit the study of deliberate memory retrieval in the absence of specific retrieval cues or behavioral outputs. We focused on hippocampal activity patterns as a substrate for retrieval, both because the hippocampus is critical for spatial and event memory retrieval for recent experiences[19] and because hippocampal activity patterns permit the experimental identification of activity consistent with retrieval. Specifically, as an animal explores an environment, each hippocampal place cell is predominately active when the animal occupies one or more locations in an environment (the cell's "place fields"). Sets of place cells can also "reactivate" when the animal is far outside the cells' place fields[6–13], reinstating a remote representation. These events can in turn be associated with activity across the brain consistent with the retrieval of a memory[12,15,21]. But whether animals can deliberately engage specific remote representations based on internal goals remains unknown. In addition, remote representations have been identified in association with two distinct local field potential (LFP) signatures: sharp-wave ripples (SWRs)[12,15] that occur during immobility and the late phases of the 8Hz theta rhythm[9–11] that is prominent during locomotion. Whether deliberately engaged representations are consistently associated with either of these signatures is also unknown.

Closed-loop hippocampal neurofeedback

To address these questions, we developed a behavioral task and a real-time neurofeedback system that rewarded animals for generating specific, experimenter-chosen remote representations. Rats were surgically implanted with tetrodes targeting hippocampus and following recovery, participated in a feedback-based task in a modified Y-maze equipped with lighted reward ports

(Fig. 1a, S1a, S2a). In each of three daily sessions, rats were first visually cued to explore the left and right arms of the environment (Fig. 1a, Task phase 1). Following this exploration period, the lighted ports in the outer arms were extinguished, and the feedback portion of the session began at a central reward port away from the arms (Fig. 1a, Task phase 2).

In the feedback portion of the task, for the first four days of training, rats received behavioral feedback based on head direction. If the rat turned its head in the specified direction (either left or right) and was near the central reward port, a reward-cue tone played. If the rat then quickly nose-poked in the port, reward was delivered (see Methods). Following the four days of head direction feedback, we switched to a neurofeedback task where neural activity triggered reward. In this version of the task, animals were rewarded only if they generated a hippocampal representation of a remote target region: the distal end of either the left or right arm of the track (see Methods and Fig. 1a). This remote representation was identified by a real-time system that continuously decoded (every 6 milliseconds) the hippocampal representation of space using an encoding model that related hippocampal spiking activity to the rat's position during the exploration phase of each session (Fig. 1b, S1a-e)[22–24]. If the decoder detected an extended representation of the target region and the animal was near the central reward port, the tone played; then, if the animal quickly nose-poked, reward was delivered (Fig. 1b, right). The target region was held fixed for three days (nine sessions) and then switched to the other arm. This switching continued for 6-18 days per animal depending on the quality of the recordings. In each session, the feedback period last for either 30 minutes or until animals received 75 rewards.

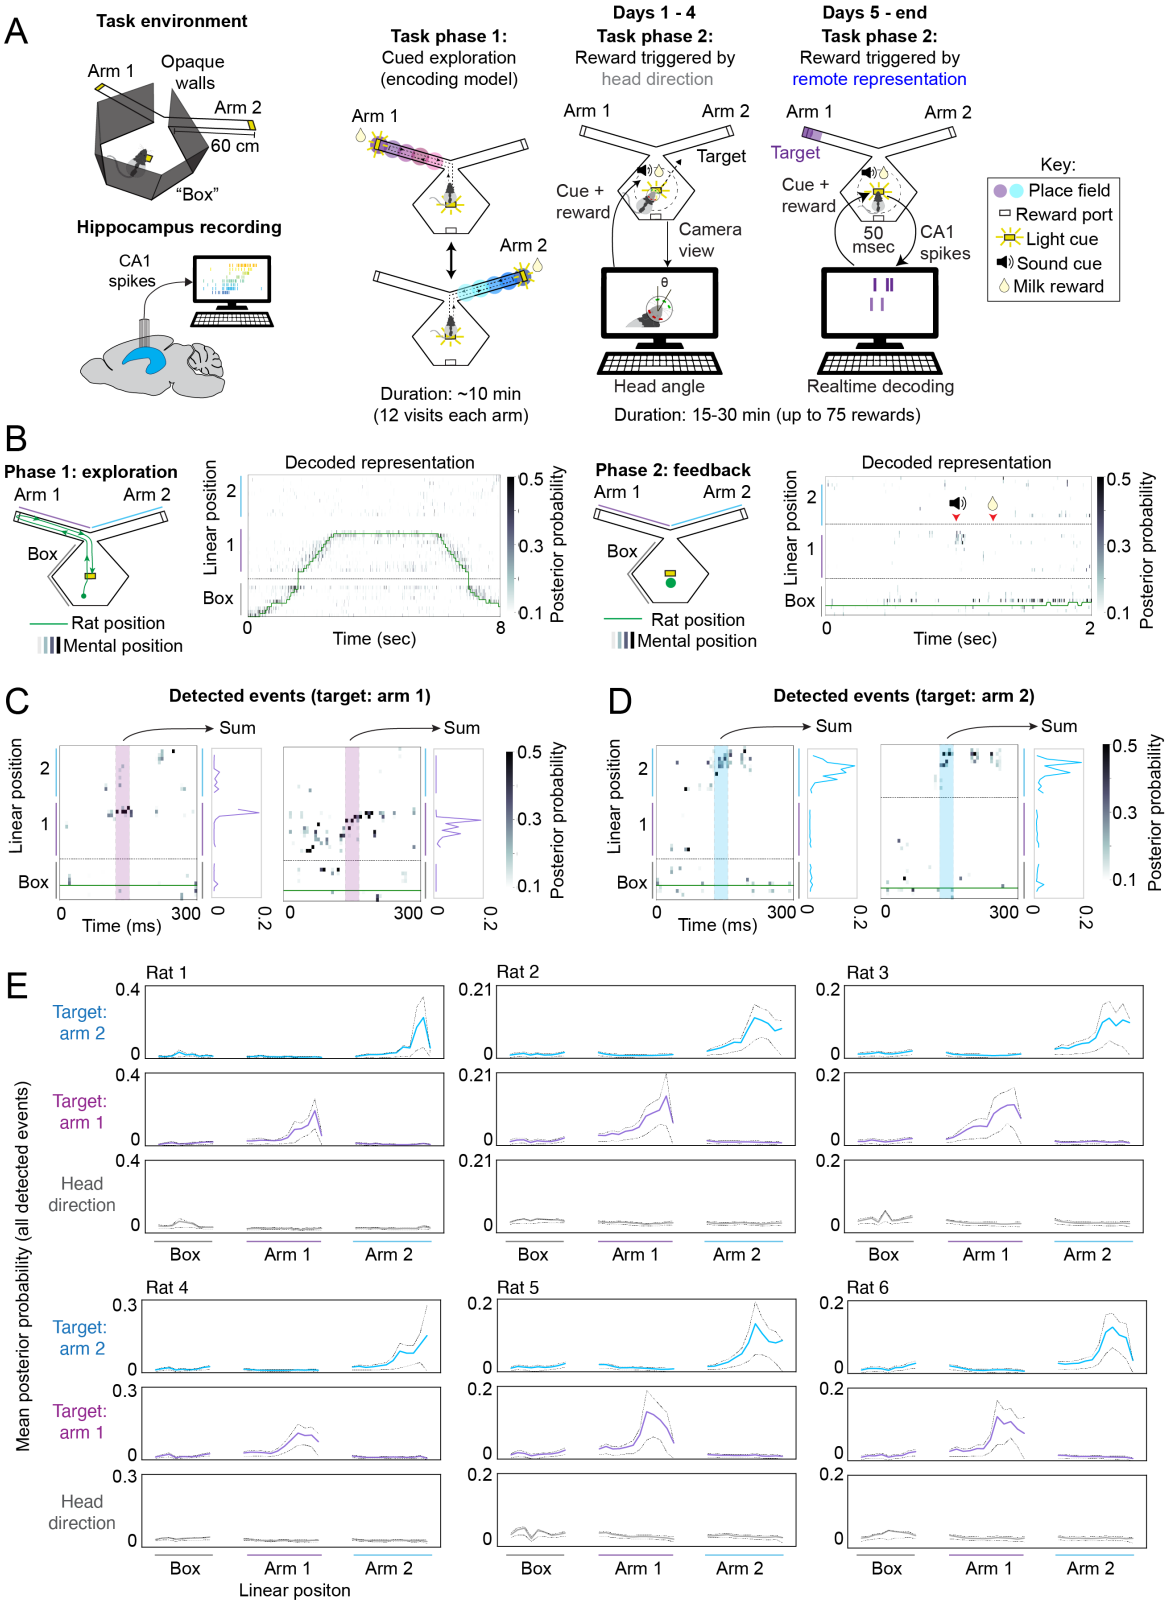
A
Task environment
Arm 1
Opaque walls
Arm 2
60 cm
"Box"
Hippocampus recording
CA1 spikes
Task phase 1:
Cued exploration
(encoding model)
Arm 1
Arm 2
Duration: ~10 min
(12 visits each arm)
Days 1 - 4
Task phase 2:
Reward triggered by
head direction
Arm 1
Arm 2
Target
Cue + reward
Camera view
Head angle
Days 5 - end
Task phase 2:
Reward triggered by
remote representation
Arm 1
Arm 2
Target
Cue + reward
50 msec
CA1 spikes
Realtime decoding
Duration: 15-30 min (up to 75 rewards)
Key:
Place field
Reward port
Light cue
Sound cue
Milk reward
B
Phase 1: exploration
Arm 1
Arm 2
Box
Rat position
Mental position
Decoded representation
Linear position
Box
Time (sec)
Posterior probability
Phase 2: feedback
Arm 1
Arm 2
Box
Rat position
Mental position
Decoded representation
Linear position
Box
Time (sec)
Posterior probability
C
Detected events (target: arm 1)
Sum
Linear position
Box
Time (ms)
Posterior probability
D
Detected events (target: arm 2)
Sum
Linear position
Box
Time (ms)
Posterior probability
E
Rat 1
Rat 2
Rat 3
Rat 4
Rat 5
Rat 6
Target: arm 2
Target: arm 1
Head direction
Mean posterior probability (all detected events)
Box
Arm 1
Arm 2
Linear positon

Figure 1. Closed-loop hippocampal feedback system.
(A) The task environment consisted of a central "box" area with a central reward port and two arms, each with a reward port at the end. The end of one of the two arms was used as the target location for neurofeedback in each session. Note that the walls surrounding the central box are opaque. Each task session contained two task phases, exploration and feedback. During the feedback phase, either specific head directions or remote target representation were detected and triggered a tone. A nosepoke at the center well within 3 seconds of the tone then triggered delivery of reward. **(B)** Clusterless decoding of hippocampal activity accurately tracked the rat's actual position during movement. During feedback, the decoder detected remote representations that triggered tone and reward. **(C,D)** The rats received reward when the decoded position matched the target location. Examples of rewarded representations for arm 1 (**C**) and arm 2 (**D**). Shaded areas contain time bins that contribute to detection (30 msec). The linearized sum of the representation is shown for each example (right plot). **(E)** Hippocampal representation averaged across all detected events. Organized by target region: arm 2 neurofeedback (blue), arm 1 neurofeedback (purple), or head direction feedback (grey). Posterior probability was averaged over 30 msec that triggered detection (as show in in (**C**)). Dashed lines represent 75% confidence interval.

We first validated the accuracy of our real-time decoder in post-experimental analyses where we used the encoding model to decode the actual location of the animal during the cued exploration period (Task phase 1; Fig. 1b, left; Fig. S2b,c; Methods); a small number of lower quality decoding sessions were excluded from further analyses. We then examined the representations that triggered reward in the neurofeedback sessions to validate the real-time system. Inspection of individual decoded probabilities preceding detected events revealed a variety of patterns, including representations largely confined to the target region as well as representations that appeared to progress sequentially from the animal's location in the box to the target region (Fig. 1c,d). As expected, these representations were specific to the chosen remote target region (averaged across rewarded representations, Fig. 1e). By contrast, during head direction feedback sessions, the representation before reward was consistent with the rat's current position (Fig. 1e). Thus, the real-time system correctly identified and rewarded periods of remote spatial representations consistent with memory retrieval.

Behavioral results showed that all six rats were able to rapidly and consistently maximize the number of rewards per session with behavioral feedback based on head direction. Further, all six rats also reached the maximal number of rewards per session during remote representation neurofeedback, although there was substantial variability across sessions and rats; some animals achieved high performance in many sessions while others achieved high performance in a smaller fraction of sessions (Fig. S3).

Increased representation of the target location

Our system established relatively strict criteria for a remote spatial representation (see Methods), and a detailed inspection of the decoded representations present during the neurofeedback periods revealed numerous instances of apparent target representations that did not meet these criteria. Therefore, to determine whether target representation was consistently enriched through neurofeedback (which would provide evidence that animals can control representation content), we calculated the fraction of time that at least 40% of the representation (probability mass) was within the target region while the rat was at the center port (see Methods; Fig. 2a,b). During these periods, the majority of the representation was almost always within the target arm (Fig. S4).

We first compared remote representation during head direction feedback sessions to remote representation during neurofeedback sessions. Because reward itself can drive patterns of hippocampal activity that express remote representations[25,26], we only compared remote representation from sessions with similar reward amounts (>90% of possible rewards). We found that remote representation neurofeedback drove a substantial increase in the generation of

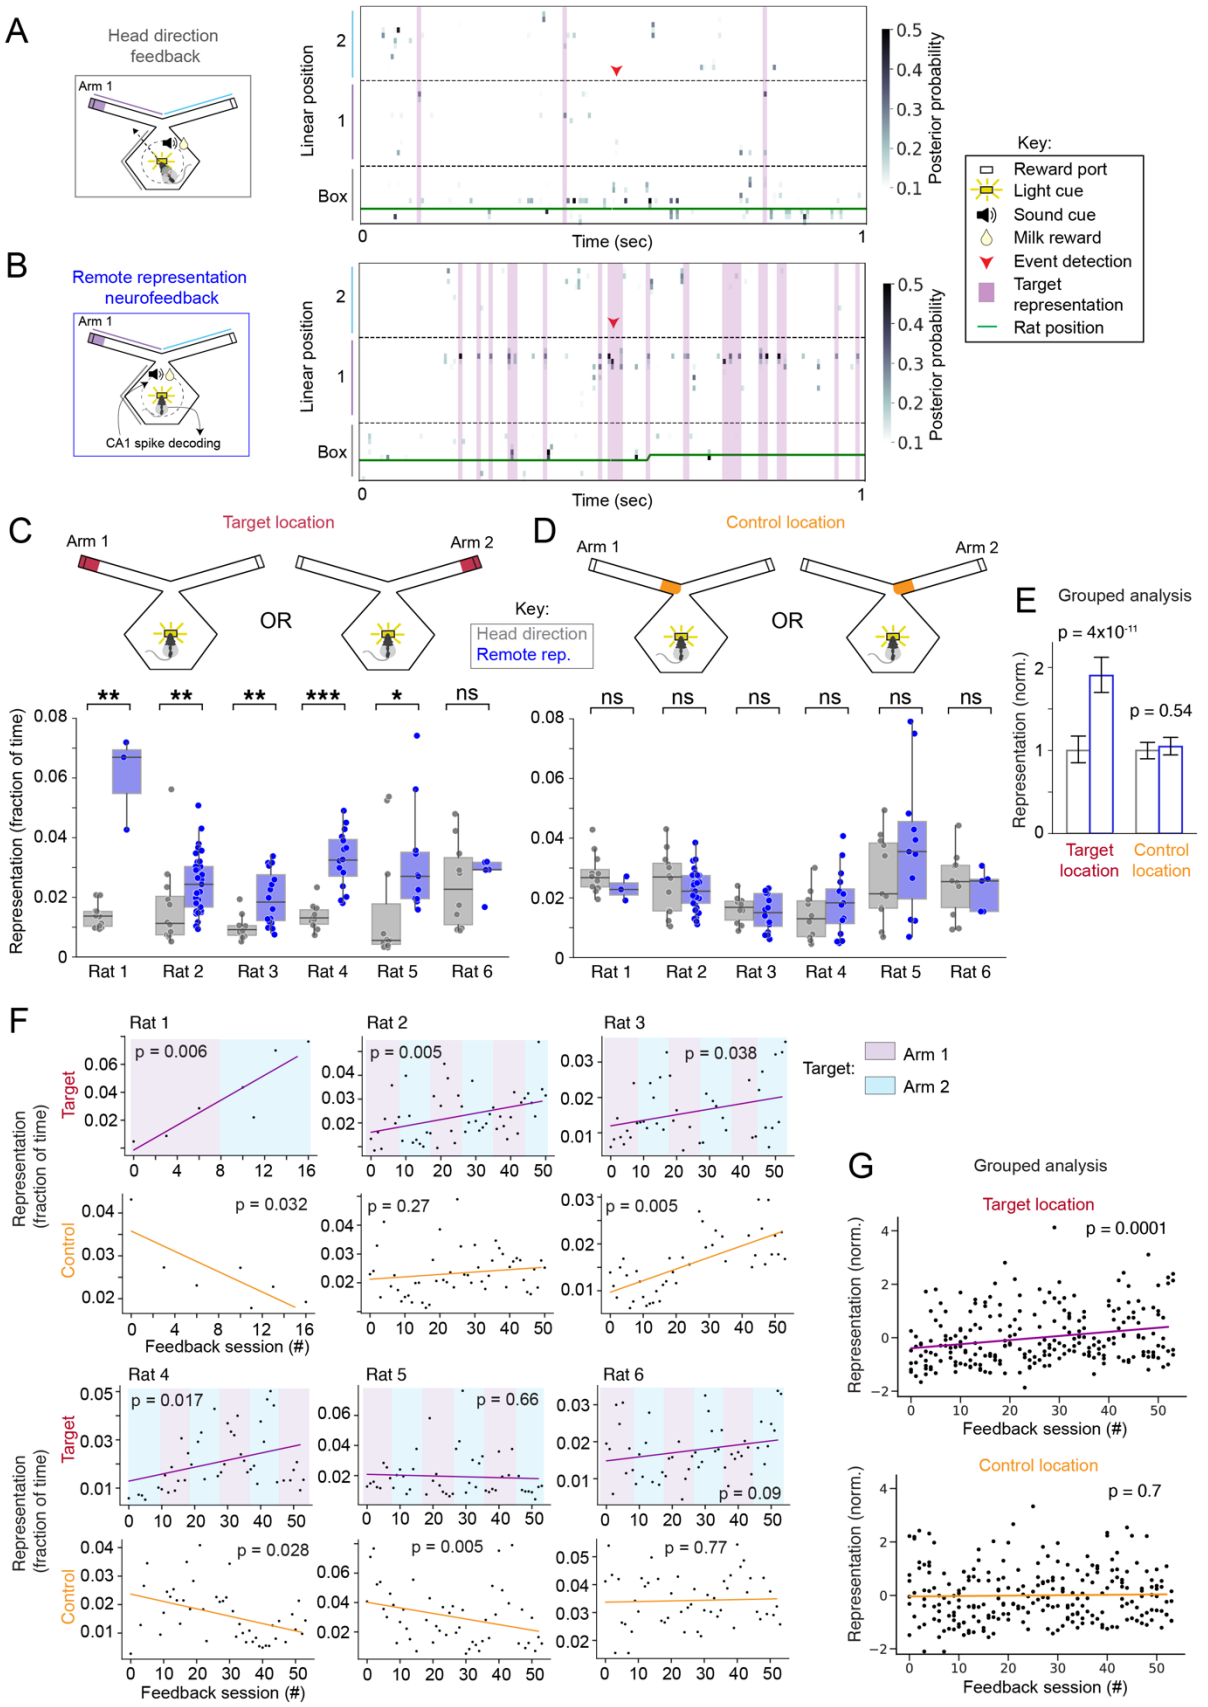
A
Head direction feedback
Arm 1
Linear position
Box
Time (sec)
Posterior probability
B
Remote representation neurofeedback
CA1 spike decoding
Key:
Reward port
Light cue
Sound cue
Milk reward
Event detection
Target representation
Rat position
C
Target location
Arm 1
Arm 2
OR
Key:
Head direction
Remote rep.
D
Control location
Representation (fraction of time)
Rat 1
Rat 2
Rat 3
Rat 4
Rat 5
Rat 6
E
Grouped analysis
p = 4x10-11
p = 0.54
Representation (norm.)
Target location
Control location
F
Target:
Arm 1
Arm 2
Target
Control
Feedback session (#)
G
Grouped analysis
Target location
p = 0.0001
Control location
p = 0.7

Figure 2. Modulation of hippocampal representations through neurofeedback.
(A) Head direction feedback (schematic at left) and example of decoded mental position around the time of a detected correct head direction event (right). **(B)** Remote representation feedback session (schematic at left) and example of decoded mental position around the time of a remote representation detection event (right). **(C)** Target location representation prevalence during high-reward sessions for head direction feedback vs remote representation feedback sessions for each rat. **(D)** Control location (non-rewarded) representation prevalence during high-reward sessions for head direction feedback vs. remote representation feedback sessions for each rat. **(E)** Grouped (LME) analysis across all 6 rats. **(F)** Target location representation prevalence across all remote representation feedback sessions (red) and for control location representations (orange). Colors in top plots represent designated target arm. Line shows linear fit; p-value corresponds to the slope of the linear fit. **(G)** Grouped analysis (linear regression) for normalized data from all 6 rats for target (top) and control locations (bottom). *: $p<0.05$, Mann-Whitney test, **: $p<0.01$, ***: $p<0.001$.

representations of the target region as compared to the head direction feedback. Representations of the target region were approximately two-fold more prevalent during neurofeedback sessions compared to head direction feedback sessions (Fig. 2c,e, individually significant in 5 of 6 rats ($p < 0.05$, Mann-Whitney); grouped analysis (linear mixed effects model, LME): all data: $p = 3.9e-11$, tone-triggering representation removed: $p = 3.6e-11$; Fig. S5a). Importantly, this did not reflect a non-specific increase in remote representations: representations of a physically closer location (the base of the target arm) were not different in neurofeedback sessions compared to head direction feedback sessions (Fig. 2d,e, non-significant in each rat, grouped analysis: $p = 0.54$, LME). Likewise, representation of the end of the opposite (non-target) arm was not enriched (Fig. S6a, grouped analysis: $p = 0.067$, LME).

Next, we examined target representation across all included remote representation neurofeedback sessions (6 or 18 days) and found that the prevalence of these target representations increased with time, consistent with a learning process. Even though the target representation switched every three days, the overall prevalence of the target representation increased over time (Fig. 2f,g, individually significant in 4 of 6 animals; grouped analysis: linear regression of z-scored values,

all data: p = 0.0001, tone-triggering representation removed: p = 0.0002; Fig. S5b). By contrast, there was no consistent change in the representation of the base of the target arm (Fig. 2f,g, significant increase in only one animal; grouped analysis: linear regression of z-scored values, p =0.7). Interestingly, representation of the opposite arm end also increased over time (Fig S6b, linear regression of z-scored values, p = 0.004), perhaps because the target switched back and forth between the two arms every three days. Overall, the longitudinal increases in representation were specific to rewarded locations.

Remote representations engage hippocampal cell assemblies

The real-time system and the associated analyses presented above used all CA1 spikes above an amplitude threshold to assess the structure of hippocampal spatial representations. This approach provides more accurate decoding than using only spikes that can be confidently associated with single neurons[27,28,29], but limits the conclusions that can be reached regarding single neuron activity. As memory retrieval is thought to engage the coordinated activity of ensembles of neurons[30], we also performed spike sorting and analyzed the resulting putative single neuron data to determine whether and how ensembles were engaged during periods of remote representation.

Our results demonstrate engagement of coordinated ensembles of single neurons during remote spatial representations. We often found that multiple single neurons had place fields (increased firing rate) in the target region and that these individual neurons were active around the time of target representation detection (Fig. 3a). These included neurons that showed a pronounced increase in activity at detection of the reward-triggering representation (cells 1 and 2) and others

that were active before, during, and after the detection of the reward-triggering representation (cells 3 and 4). We quantified this engagement by identifying coordinated activity of groups of neurons (cell assemblies[31]; see Methods) and segregating them into two groups based on whether or not their activity represented the target location during the exploration period of each session (target and non-target assemblies; Fig. 3b,c). We then computed the activity of these assemblies at the times when remote representations were expressed.

We first focused on times when the target representation was detected (and the reward cue was triggered) during the feedback period of the neurofeedback sessions. During these periods, when the animal was in the box, far from the actual target location, we nevertheless found strong and specific activation of target-representing assemblies: these assemblies were about 90 times more active at detection times compared to random times (Fig. 3d,e, each rat $p<0.001$, Mann-Whitney test) and were about 25 times more active than non-target assemblies (Fig. 3d,e, each rat $p<0.001$, Mann-Whitney test).

We also observed enhanced activation of target-representing assemblies during periods of target representation that did not overlap with target detection events (Fig. 3f). During these periods, target assemblies were about 4 times more active than at random times (Fig. 3g, 5 of 6 rats $p<0.001$, 1 rat, $p=0.16$, Mann-Whitney test; grouped analysis of all rats, $p = 3.2e-27$, LME) and were about 1.6 times more active than non-target assemblies at these same times (Fig. 3g, 4 of 6 rats, $p<0.001$, 2 rats: $p=0.4$, 0.44; grouped analysis, $p = 9.5e-18$, LME). We note that target assembly activity was not detected during all times of remote representation; this is likely because our system used

clusterless decoding (decoding without cell clustering using all spikes above a voltage threshold), which incorporates many spikes that are not clustered into single units during spike sorting[32].

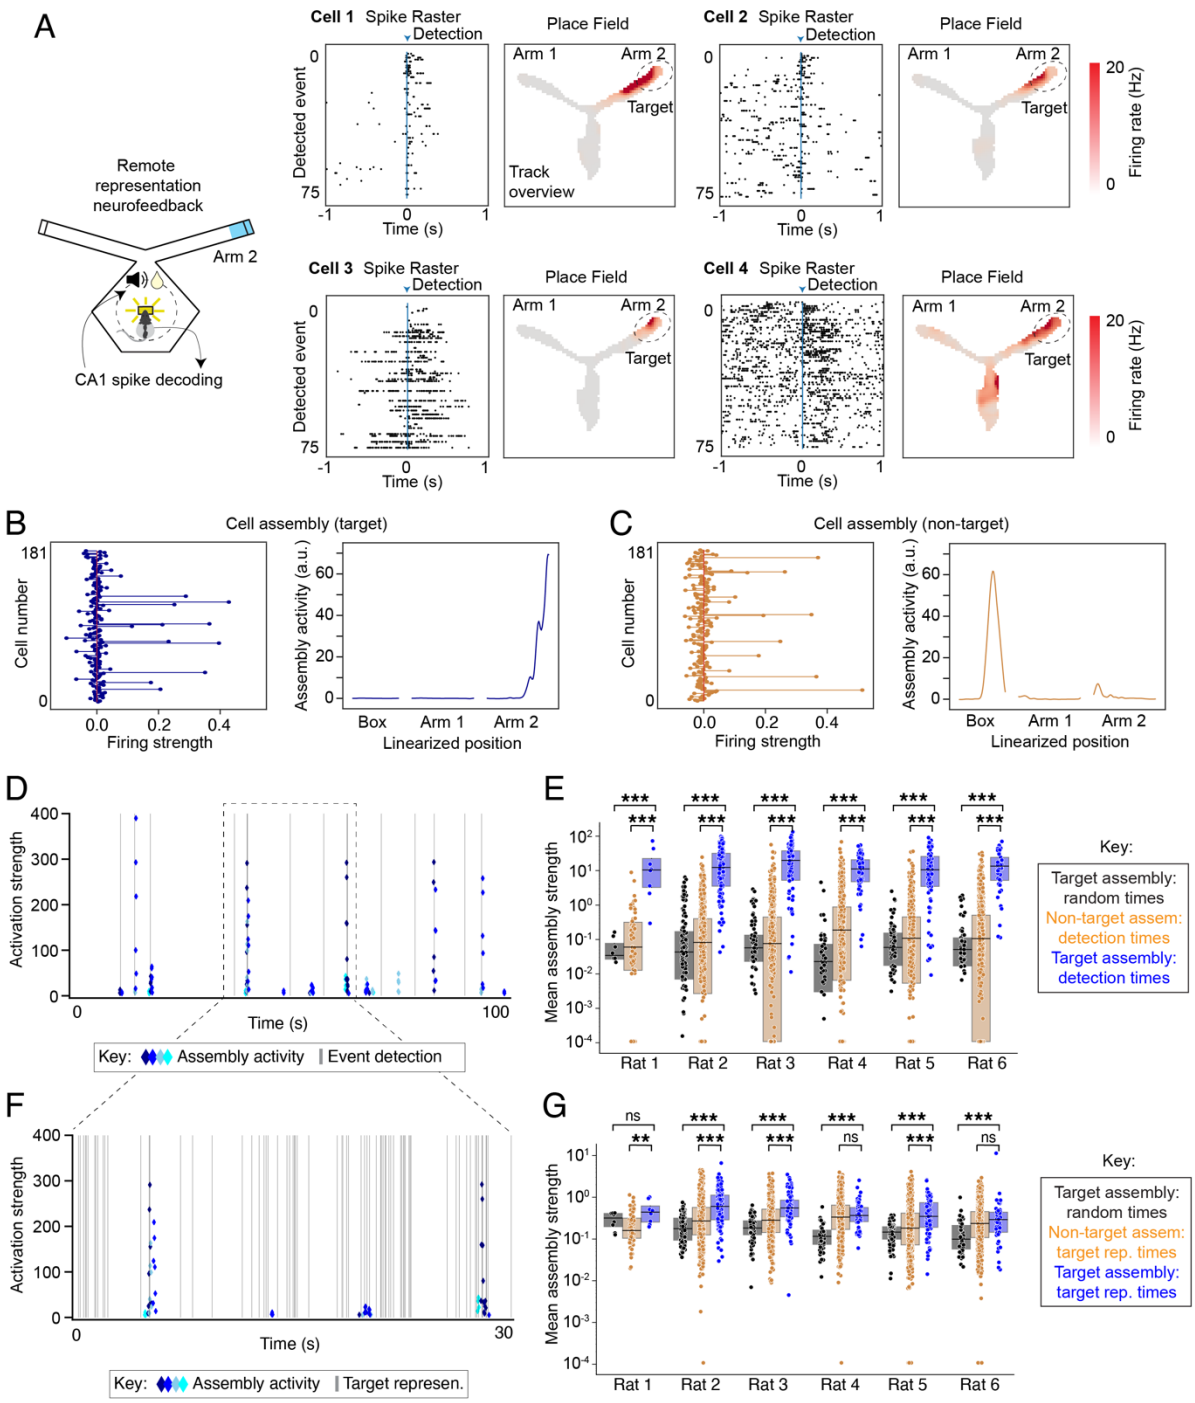

Figure 3. Cell assemblies are activated at the time of remote representations.
(A) Individual neurons active at remote representation detection times. For each cell, left: raster plot showing spikes surrounding each detected remote representation (blue arrow); right: place field computed during exploration phase. **(B)** Target cell assembly. Left: individual cell weights; right: combined location-specific assembly activity. **(C)** Same plot as **(B)** for a non-target cell assembly. **(D)** Example of activation strength for four target assemblies at the time of remote representation detection. Blue diamonds show activity of each target assembly and grey lines indicate detection times. **(E)** Target assembly activity at time of remote representation detection (blue) compared to random times (black) and to non-target assemblies at detection times (orange) across all sessions. **(F)** Zoom-in on plot from **(D)** showing target representation times (grey lines) and assembly activity (blue diamonds). **(G)** Target assembly activity at time of remote representation outside of detection events (blue) compared to random times (black) and to non-target assemblies at remote representation times (orange) across all sessions. ***: $p < 0.001$, Mann-Whitney test. In **(E)** and **(G)**, for plotting only, any values less than 1e-4 were set to 1e-4.

Finally, to verify that each detection event engaged multiple neurons (illustrated in an example, Fig. S7a), we counted the number of high-strength cells in target assemblies that were active before each target detection event (Fig. S7b, Methods). Across all sessions, rats had an average of 1.8 – 8 cells active immediately preceding each detection event, which corresponded to 10 – 15% of high assembly strength cells (Fig. S7c). This is comparable to the number of neurons expressing remote representations during movement periods as reported in previous papers[10,11], and given that we are recording a tiny fraction of the total number of neurons in the hippocampus, suggests that many neurons are engaged during each of these remote representation events.

Brain state during remote representation events

Previous work has identified remote spatial representations in the hippocampus in the context of two distinct physiological states. First, remote representations can be expressed during sharp-wave ripple (SWR) events. During waking, SWRs occur primarily during immobility, and spiking during SWRs is most strongly associated with memory consolidation processes[13–15]. Second, remote representations can be expressed in moving animals in the context of the ~8 Hz theta

rhythm, where late phases of theta are most strongly associated with remote activity[11,33]. Remote activity during theta is hypothesized to contribute to memory retrieval in the context of deliberative decision making[11,34]. We therefore asked whether the remote spatial representations occurred during SWRs, theta, or in a different brain state[35] (Fig. 4a).

Strikingly, remote representations were primarily expressed during stillness (head speed <4 cm/sec) but outside of clearly identifiable SWR events: only 10-20% of remote representation were associated with SWRs (Fig. 4b). Furthermore, target representations outside SWRs consistently increased in prevalence with neurofeedback (Fig. 4c, $p < 0.05$ in 5 of 6 individual rats, Mann-Whitney; grouped analysis: $p = 1.3e-11$, LME), while representations associated with SWRs did not consistently increase across animals (Fig. 4c, $p < 0.05$ in 1 of 6 rats, Mann-Whitney; grouped: $p = 0.056$, LME). There was no obvious LFP signature of the remote representations that occurred during stillness and outside of SWRs (although some rats had increased power >100 Hz, potentially reflecting high gamma activity or spiking, Fig. S8b). Indeed, and as expected, multiunit spiking activity peaked at the time of remote representations (Fig. S8c). Those remote representations that were expressed during movement showed characteristics generally similar to those seen in previous reports[3]. These representations were associated with specific theta phases in all rats ($p < 0.0001$ for each rat, Fig. S8a), and in 4 of 6 rats, remote representations were more prevalent in late theta phases.

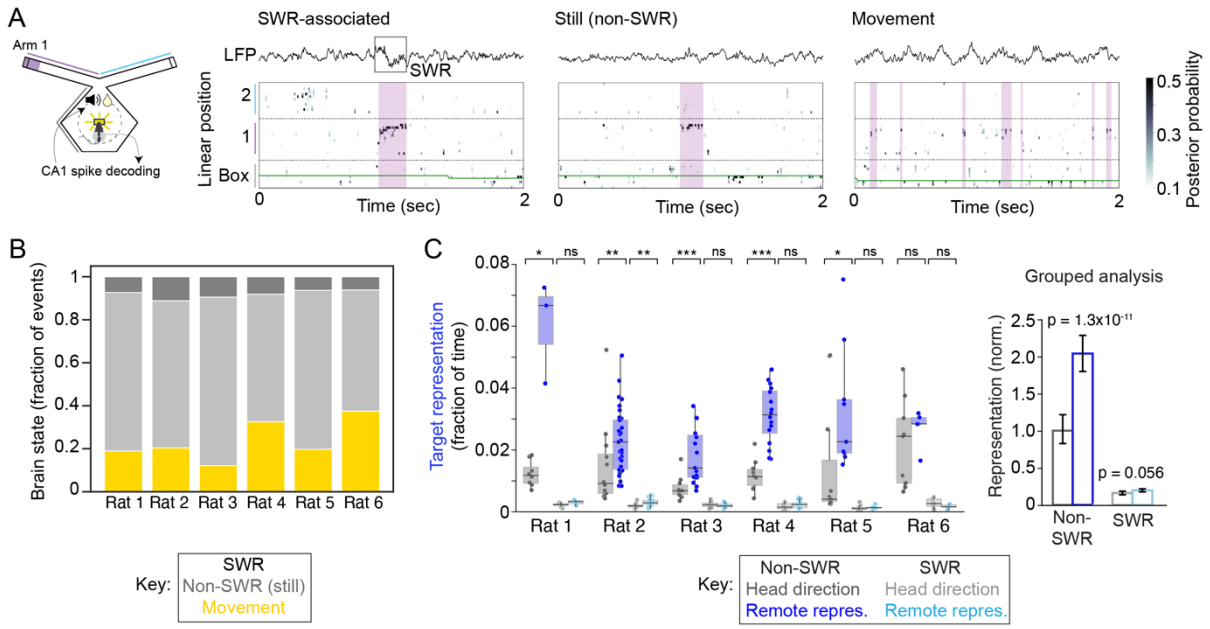

Figure 4. Brain state during remote representation.
(A) Example remote representations during SWR, stillness outside of SWR, and movement. Left: schematic. Right: example plots; top: LFP trace, bottom: decoded linear position. **(B)** Summary of brain state during remote representations. **(C)** Change in target representation during SWR-associated and non-SWR times for neurofeedback sessions compared to head direction feedback sessions. Grouped analysis (LME) for 6 rats. *: $p<0.05$, Mann-Whitney test, **: $p<0.01$, ***: $p<0.001$.

DISCUSSION

We developed a closed-loop neurofeedback system for hippocampal representations and used this system to reward rats for generating specific remote spatial representations. Although no retrieval cues were presented, rats learned to activate specific representations corresponding to experimenter-chosen spatial locations. The rats were also able to generate different representations at different times in response to changes in the target location, demonstrating flexibility. These remote representations engaged one or more cell assemblies and multiple neurons within each assembly, consistent with the reactivation of a population of neurons representing the remote

location. Our work establishes a model for studying how the brain can deliberately reinstate representations related to previous experience in the absence of cues and specific behavioral outputs.

Enrichment of specific neural activity patterns through operant conditioning has been repeatedly demonstrated in motor and pre-motor cortex[36–38], but very rarely in hippocampus. One study reported conditioning of the activity of a single hippocampal cell[39], but this cell was not associated with any specific past experience. Here we demonstrated conditioning of population-level activity (cell assembly analysis, Fig. 3) associated with coherent spatial representations (clusterless decoding, Fig. 1). And recently, Lai et al. reported a hippocampal brain machine interface (BMI) that demonstrated volitional control of hippocampal population activity in rats for mental navigation in a virtual reality environment[40]. Our work complements and extends those findings in two ways. First, in Lai et al., hippocampal conditioning occurred in a continuously updating virtual reality environment that provided visual cues to guide neural activity. By contrast, our system occurs in a physical environment where the rat receives no cues or feedback until it generates the target representation, and so more closely models spontaneous memory retrieval. Second, Lai et al. used recurrent neural networks (RNNs) to decode spatial representations from hippocampal activity, whereas our system uses an explicit model that provides a direct understanding of how the detected CA1 spiking activity generates a remote spatial representation.

Our work also provides insights into the brain states and activity patterns that can support the generation of remote representations consistent with memory retrieval. Most remote representations in this task occurred during stillness but did not overlap with SWRs, and remote

representations that did occur during SWRs were not enriched despite many days of training. Thus, while it is possible to train animals to increase the rate of SWR generation[41], these findings suggest that representations within SWRs are less amenable to control. This possibility is consistent with the idea that SWRs are critical for "offline" memory consolidation and updating functions rather than "online" retrieval to inform immediate decisions[13,12,15].

We also found that while some of the events occurred during movement, most occurred during periods of stillness and were not associated with an obvious local field potential signature. This state has also been associated with immobility-related place cell activity[42–44], and our results indicate that in rats, the state is also consistent with a more "generative" function associated with remote representations[35]. Specifically, the lack of movement and associated sensory input may be conducive to the generation of internally driven patterns of activity associated with specific past experiences.

The specific mechanisms that support the generation of these remote representations are unknown, but previous studies of remote hippocampal representations in the context of SWRs[12,15] and movement[9,21] highlight the engagement of multiple cortical regions including the prefrontal cortex before, during, and after remote hippocampal events. It seems plausible that related networks are engaged when animals deliberately activate specific remote representations, but the mechanisms that might differentiate more automatic from more deliberate activation are not understood. More broadly, the generation of patterns of activity associated with past experience is central to the retrieval of memories[45], and also to related cognitive functions like planning and imagination[34]. The ability to train animals to generate such patterns in the absence of specific cues and specific

behavioral outputs thereby provides a potentially powerful tool for understanding how the brain can support these remarkable functions.

SUPPLEMENTARY INFORMATION

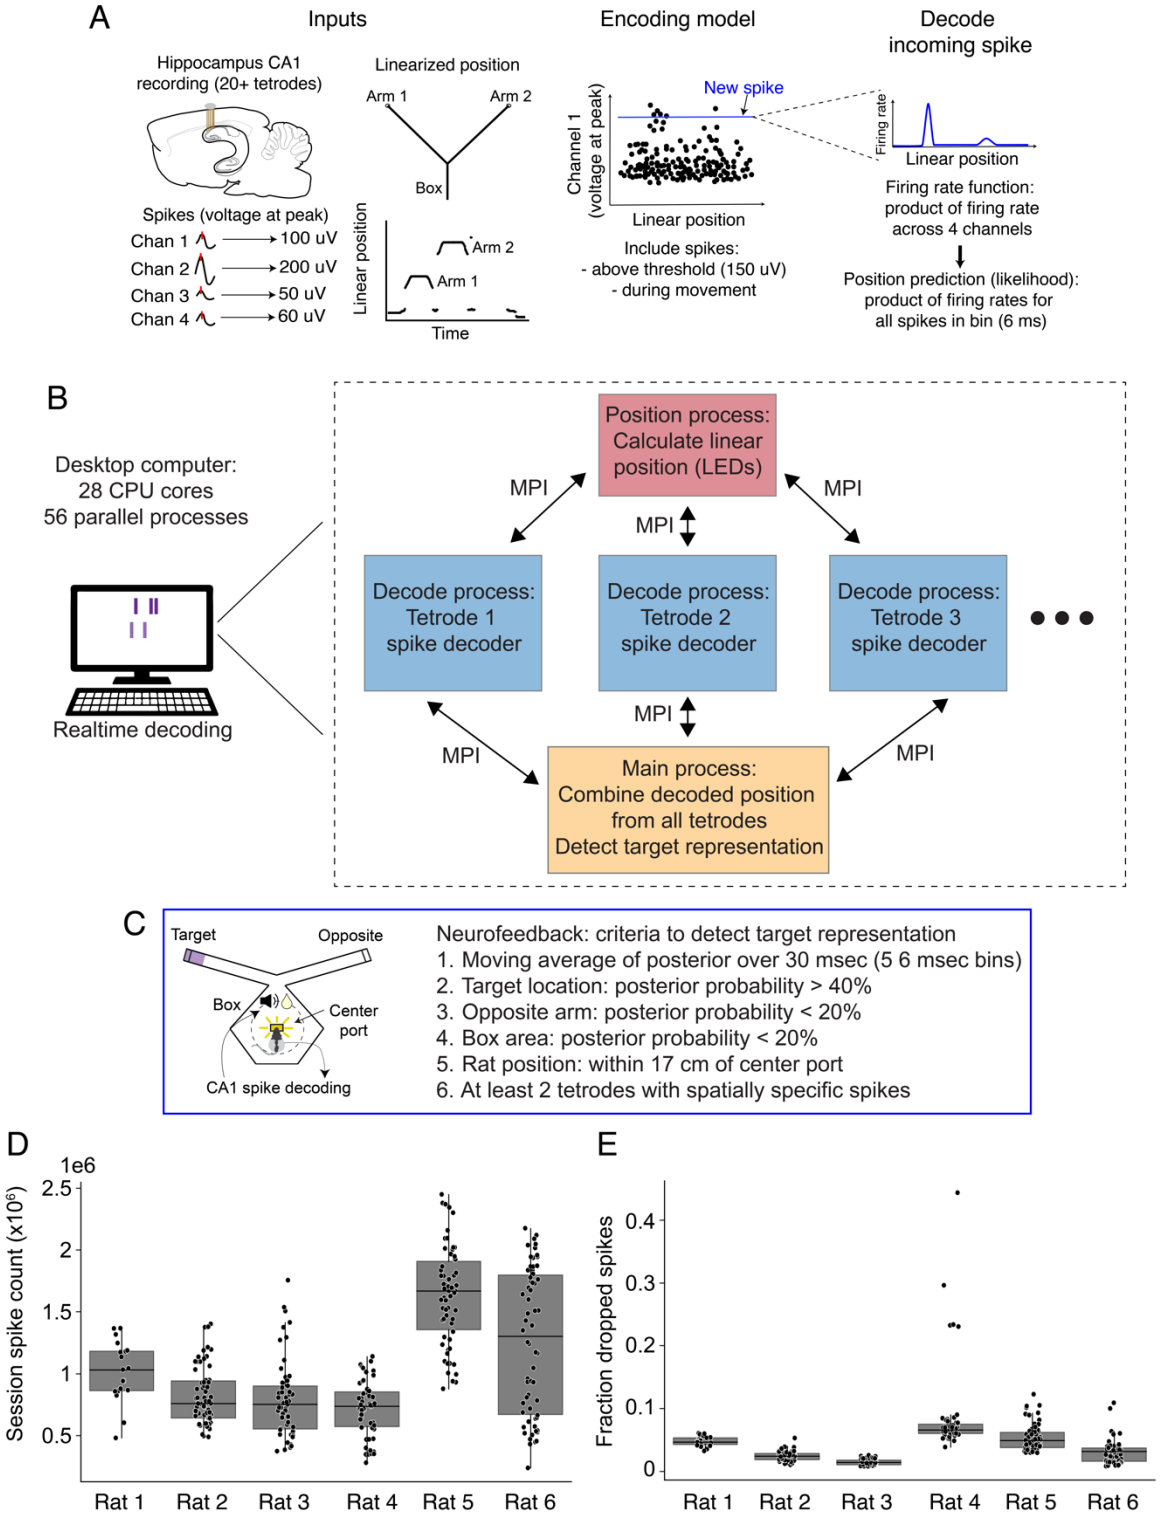

Figure S1. (A) Schematic of clusterless decoding algorithm. (B) Schematic of software architecture for real-time decoder implementation. MPI: message passing interface. (C) Criteria for detection of remote representation to trigger reward cue. (D) Total recording session spike counts for each rat. (E) Fraction dropped spikes for each rat (not decoded because of decoder latency).

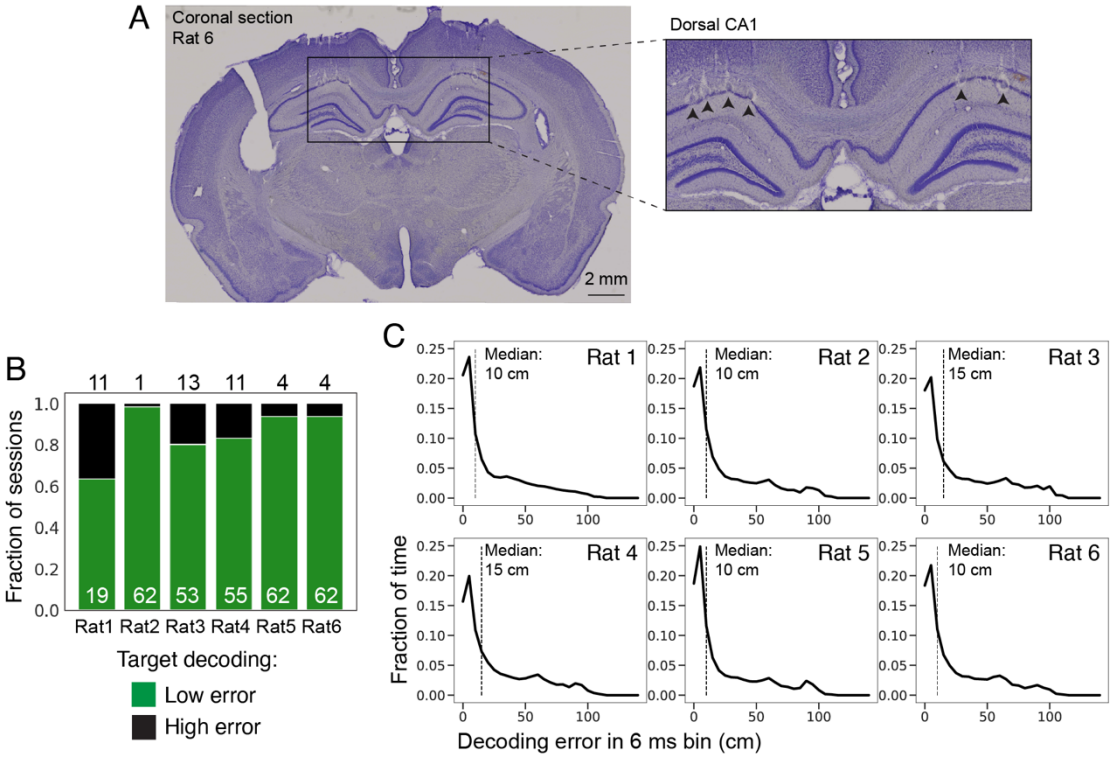

Figure S2. (A) Histology (Nissl stain) showing recording tetrodes located in dorsal CA1 of hippocampus. (B) Fraction of recording sessions for each rat with low decoding error vs. high decoding error (low: <35% decoding error in target region). (C) Decoding error during movement. Histogram of error (distance between rat's real position and decoded mental position) for each 6 msec time bin while the rat was moving (>4cm/sec). Median is marked with dashed vertical line.

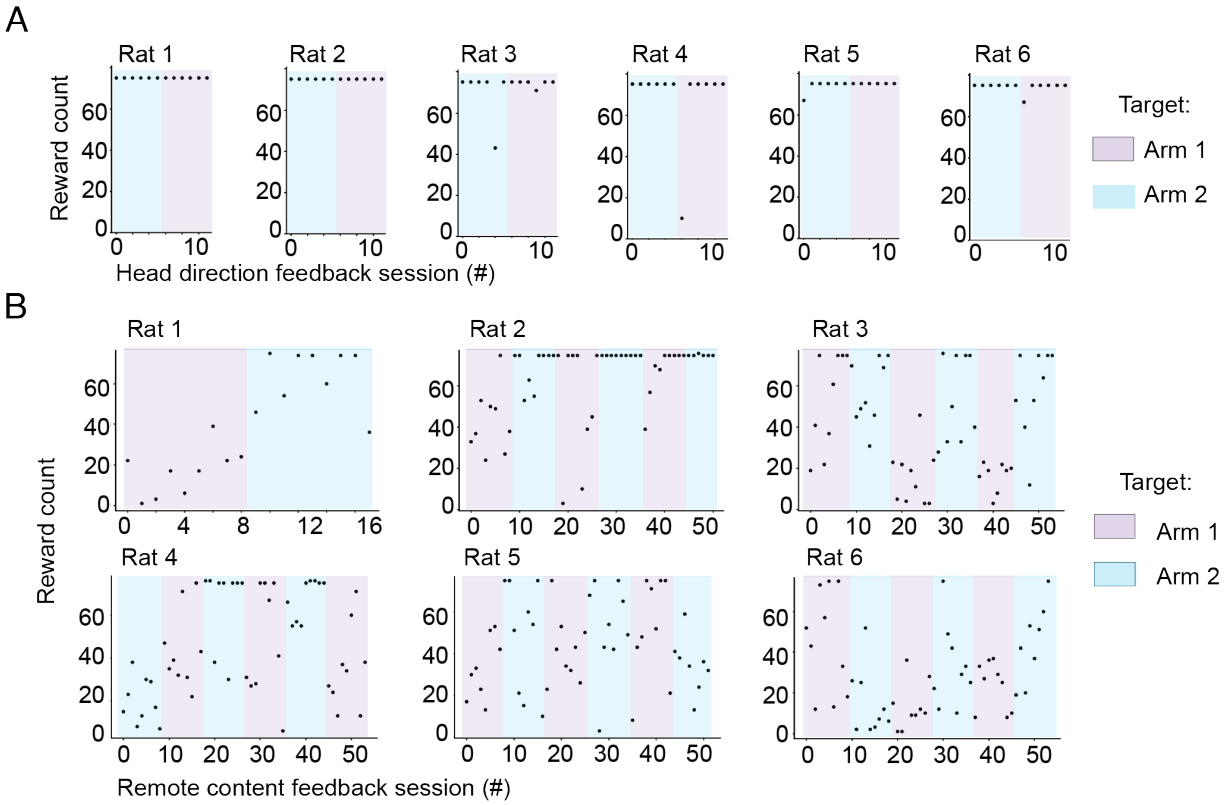

Figure S3. (A) Reward counts for all head direction feedback sessions. (B) Reward counts for all remote representation neurofeedback sessions. Target arm is indicated by color of plot background. Maximum rewards per session was 75.

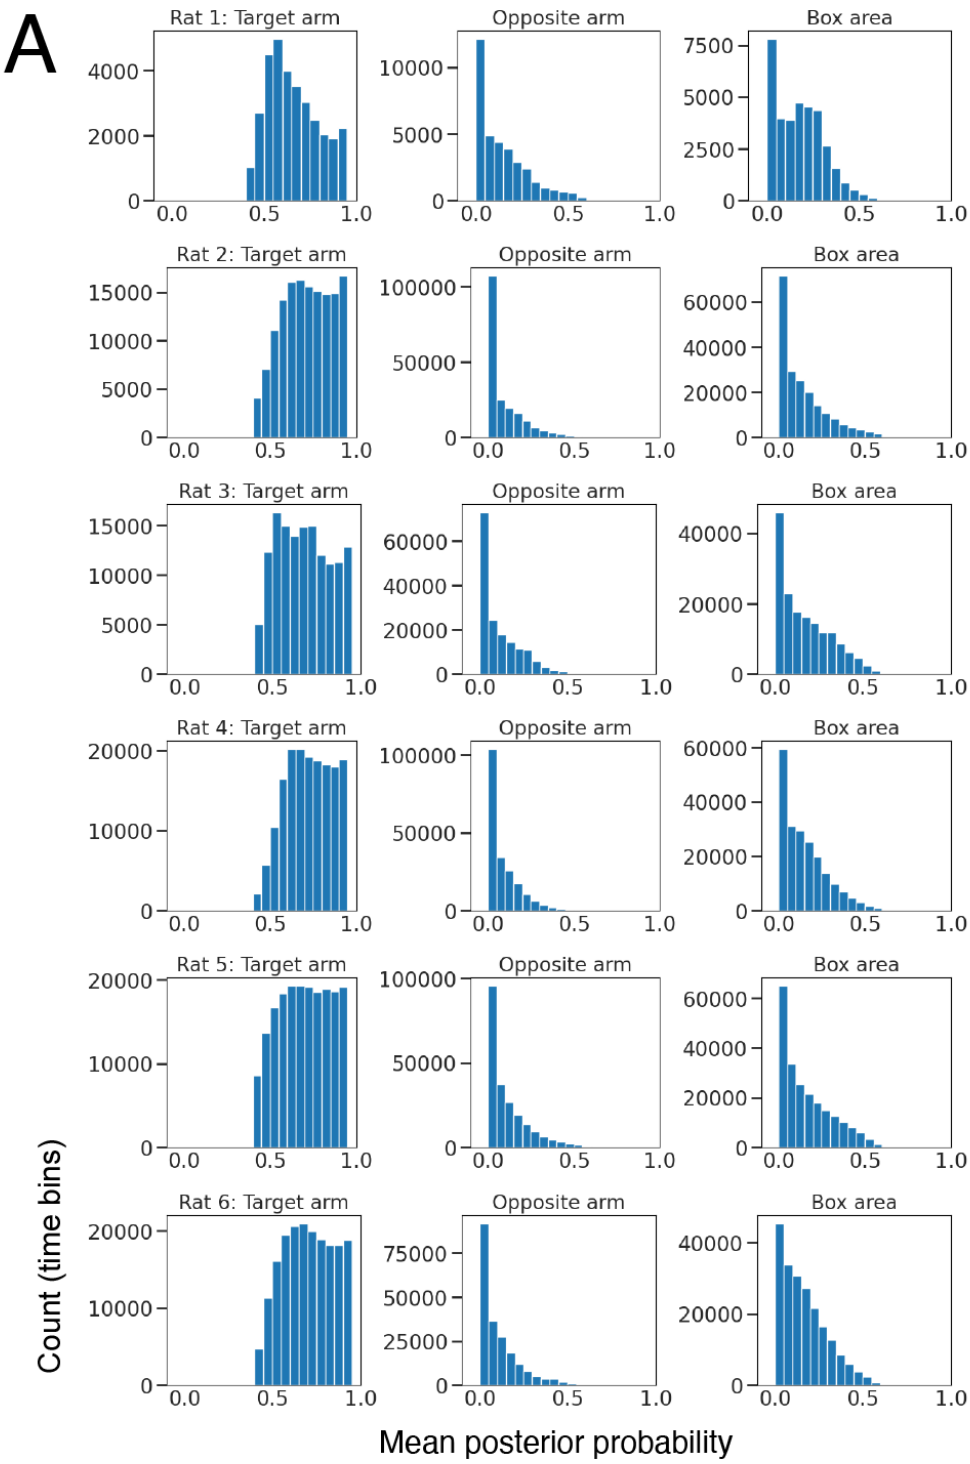

Figure S4. (A) Decoded posterior probability at target representation times. This shows that for each rat, nearly all times of target representation had at least 50% representation in the target arm and much less representation of either the opposite arm or the local area (box).

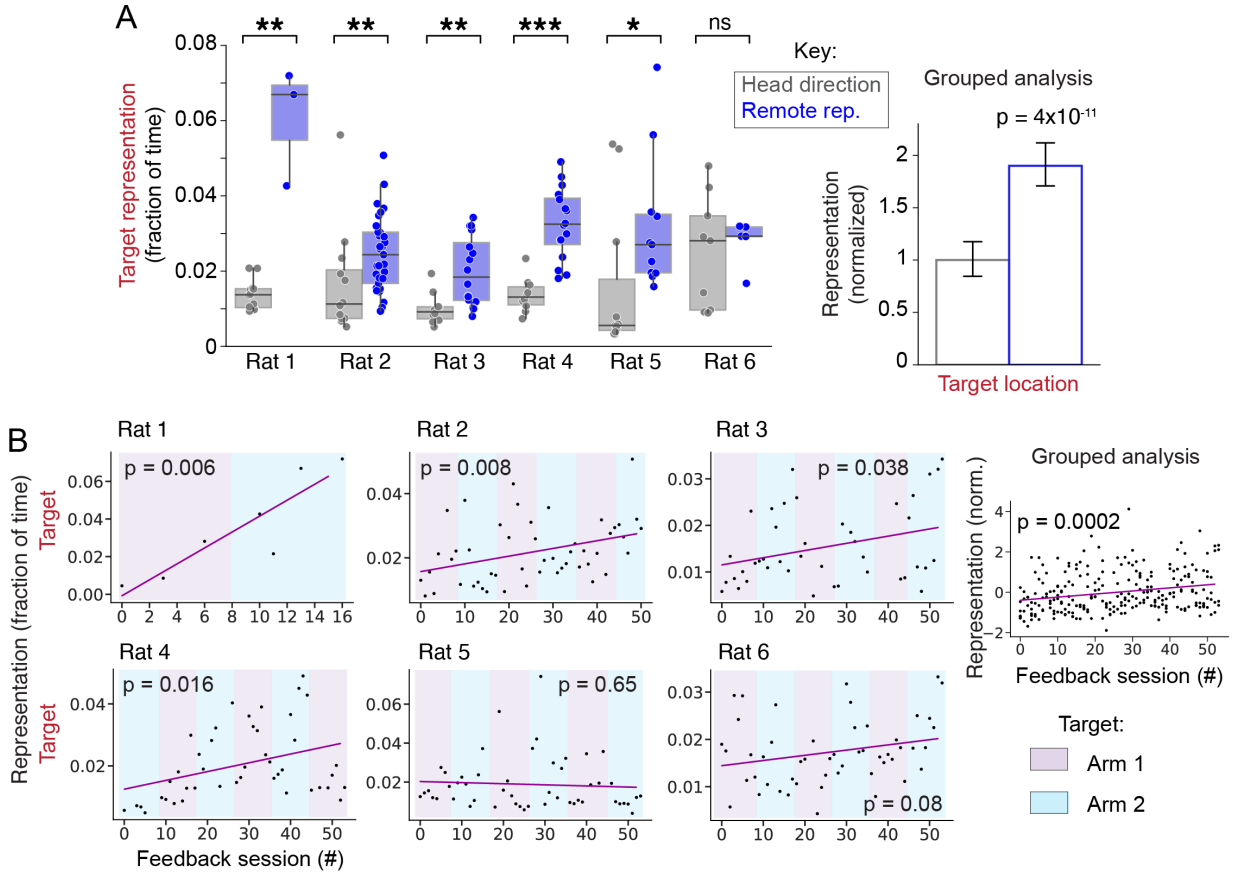

Figure S5. (A) Same plots as in Figure 2C with triggering representation removed. (B) Same plots as in Figure 2D with triggering representation removed. *: p<0.05, Mann-Whitney test, **: p<0.01, ***: p<0.001.

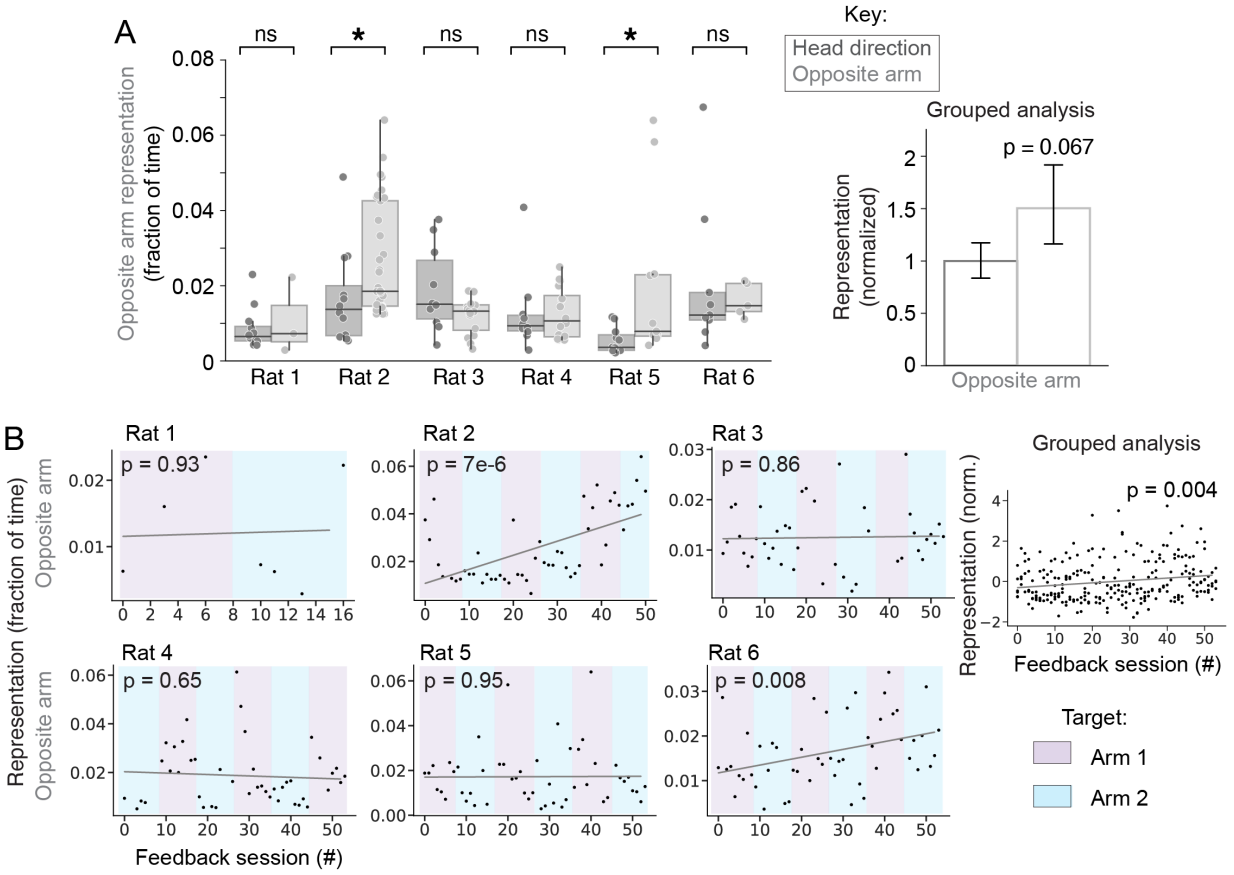

Figure S6. (A) Same plots as in Figure 2C showing representation of opposite arm end (non-target). (B) Same plots as in Figure 2D showing representation of opposite arm end. *: p<0.05, Mann-Whitney test.

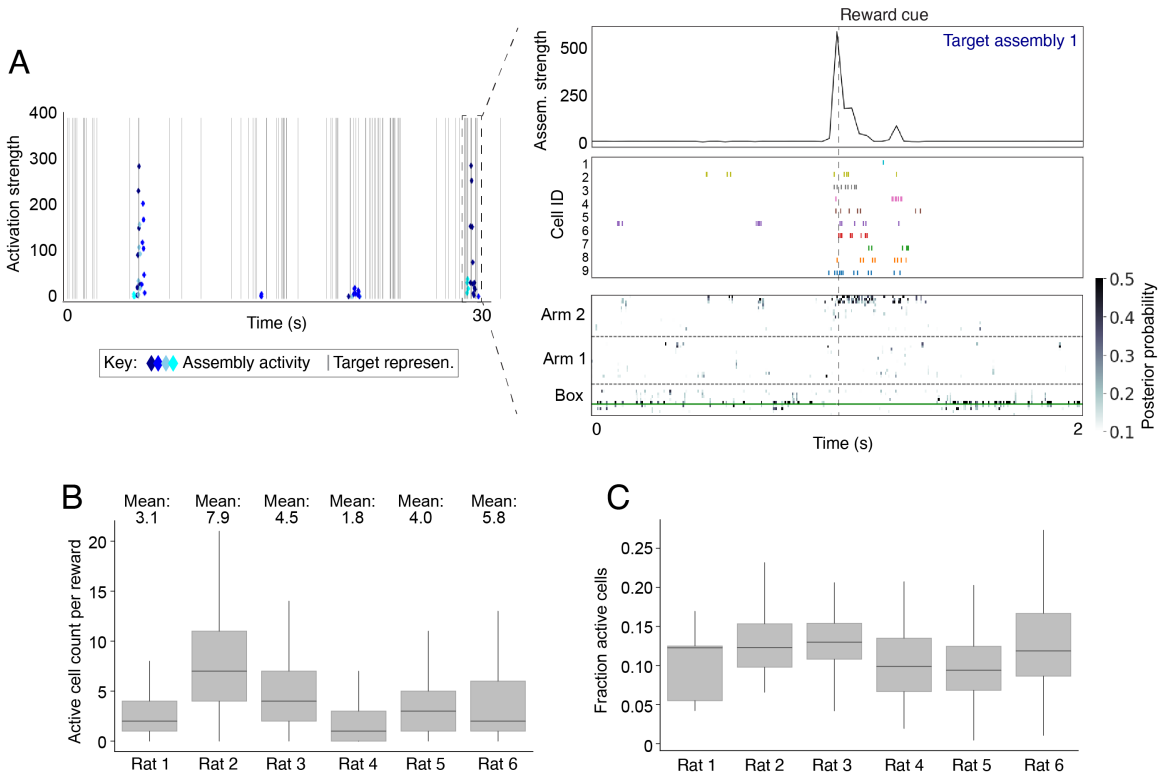

Figure S7. (A) Left: same panel from Figure 3F, showing target assembly activity (blue diamonds) and remote representation (grey lines). Right: zoom in on a detected target representation. During this representation, target assembly activity is high (top) and there are spikes from several high-strength cells within this assembly during the representation (middle). Clusterless decoding showing the target representation at the end of arm 2 (bottom). (B) We counted the number of high-strength cells with spiking activity immediately before target representations and found that, on average, all rats have more than one spiking cell (within 100 ms). (C) This plot shows the fraction of high-strength cells active before each target representation. The fraction is similar across rats showing that much of the variability in cell count in (B) is the result of variations in the number of high-strength cells across rats, an indicator of recording quality.

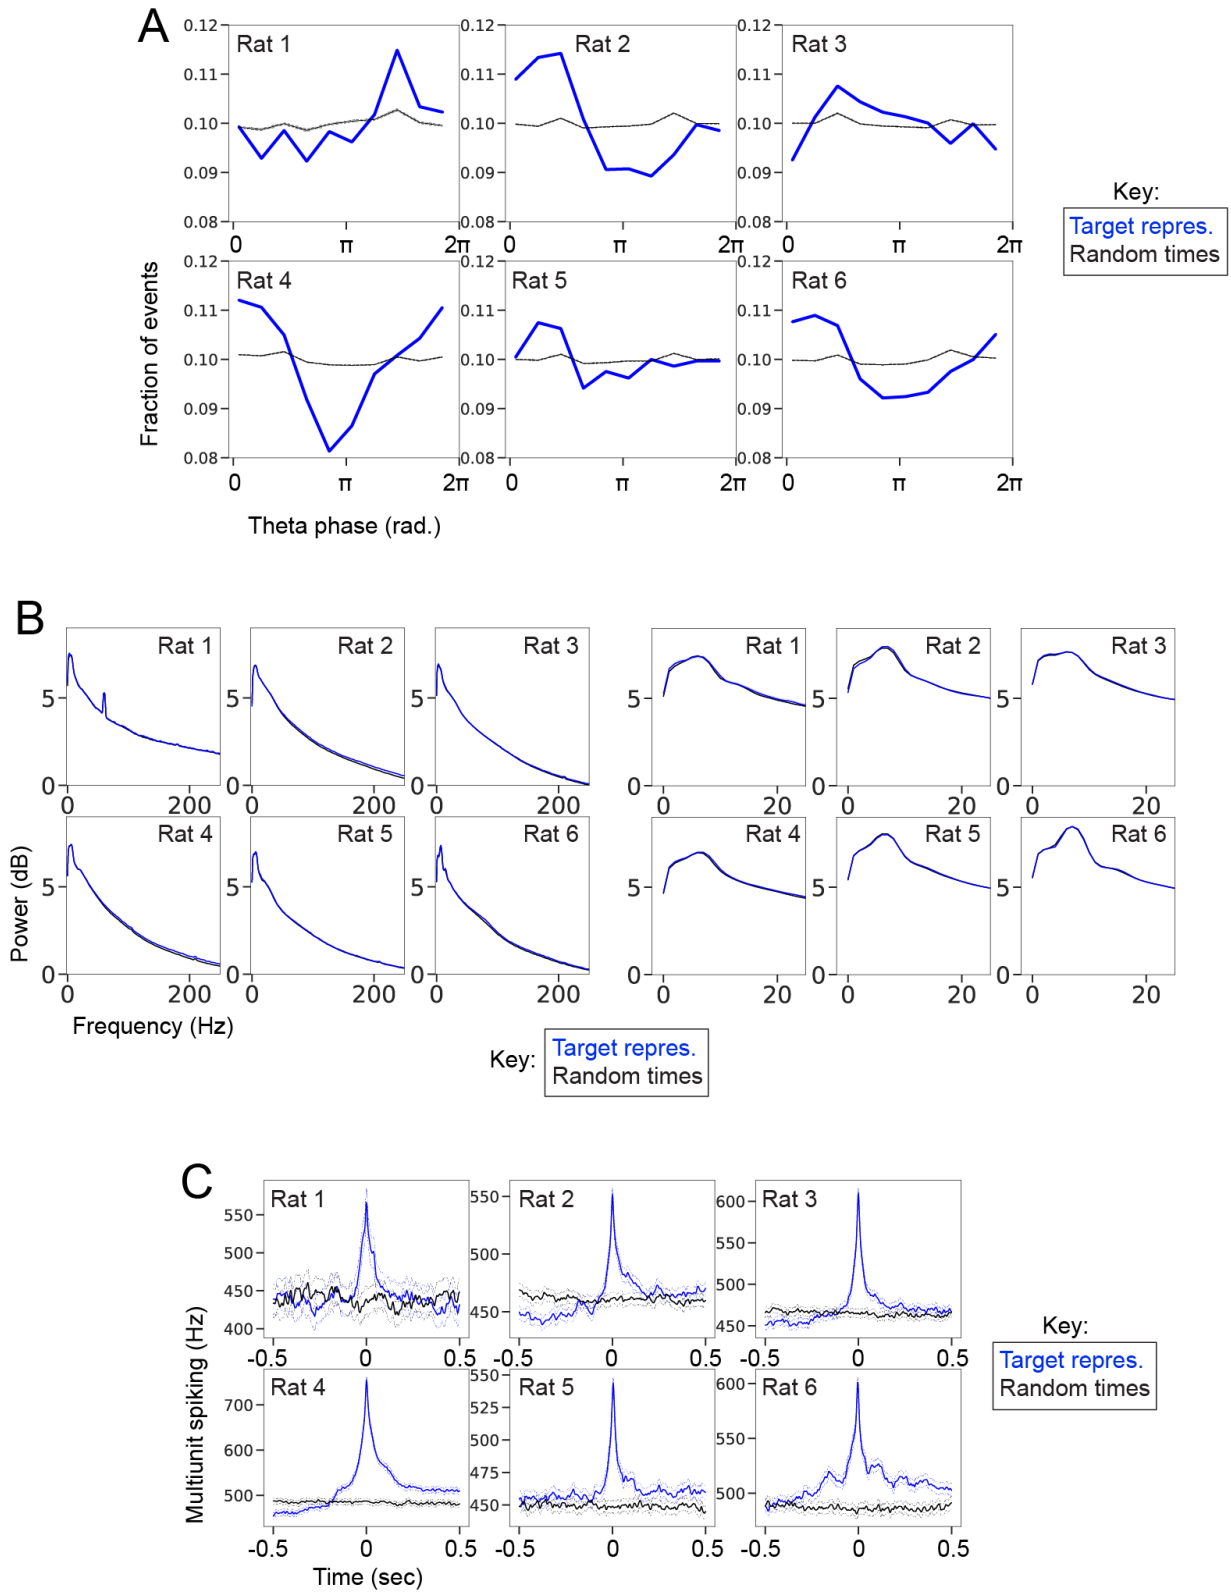

Figure S8. (A) Remote representation times (blue) are modulated by theta phase. Black line is 95% confidence interval for theta phase of random times. All possible theta phase values were

grouped into 10 equal bins and the fraction of remote representation events in each bin is plotted. (B) LFP spectrogram of remote representation times. Left: tetrode in CA1 cell layer (referenced), right: tetrode above cell layer (unreferenced). Theta power is most accurately measured above the cell layer with an unreferenced tetrode, and so right plot is zoomed in to low frequency range. (C) Multiunit spiking at times of remote representation. Blue: remote representation, black: random times. Dashed lines: 95% confidence intervals of the mean.

ACKNOWLEDGMENTS

We thank all members of Loren Frank's lab for thoughtful feedback throughout the project. We thank Vanessa Bender, Massimo Scanziani, and Andrew M Lee for feedback on the manuscript.

MEC was supported by F32MH123003 from the NIMH, LMF is supported by Howard Hughes Medical Institute (investigator).

AUTHOR CONTRIBUTIONS

MEC designed and performed the experiments, analyzed the data, and wrote the manuscript. AKG helped design the experiments. JC, ELD, TTKN, DLF, KW, and XD contributed to the real-time decoding software. KW and BS helped perform the experiments. UTE and CK helped supervise the project. LMF supervised the project, helped design the experiments, and wrote the manuscript.

COMPETING INTERESTS

The authors have no competing interests to disclose.

METHODS

RESOURCE AVAILABILITY

Materials Availability

Not applicable.

Data and Code Availability

Raw data will be available on DANDI.

Analysis code is available on a GitHub repository called:
https://github.com/orgs/LorenFrankLab/hippocampus_content_neurofeedback

Realtime spatial decoder code is available on GitHub repositories called:
Current version of code:
https://github.com/orgs/LorenFrankLab/realtime_decoder
Version of code used in the manuscript:
https://github.com/MichaelCoulter/realtime_decoder_v0

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Adult wild type Long Evans rats were obtained from Charles River Labs. Housed at UCSF following all IACUC guidelines. All surgical procedures done following IACUC guidelines.

METHOD DETAILS

Behavioral training

Adult wild type Long Evans rats (6-9 months of age) were food restricted to 85% of their free body weight. Before neural implant surgery, rats were trained to nosepoke in reward ports to receive a milk reward on a linear track. Rats were then exposed to the task environment and trained in the task structure, in which they had to visit lit reward ports to receive a milk reward. The task environment consisted of a central area with a milk reward port and two attached arms with a reward port at each arm end. The task structure consisted of two phases during 40-45 minute recording sessions. In the first 10-15 minutes (exploration phase), the rat explored the Y-shaped environment (center box area, arm 1, and arm 2) and received milk reward by nosepoking in each lit reward port (300 uL of sweetened condensed milk was delivered automatically via the Trodes recording software). The light cues directed the rat to visit each arm 12 times and return to the center after each arm visit (arm 1 and arm 2 visits were randomly ordered). In the second phase (operant conditioning), lasting 20-30 minutes, the rat could only receive reward at the center port. During the conditioning phase, a sound cue was played and then a milk reward was available at the center port. In pretraining, the rat learned the association between the sound cue and reward: when the sound cue played, the rat had to nosepoke within 3 seconds to receive reward. During initial training the sound cue occurred randomly (centered on intervals of 5-30 seconds). For head direction feedback, the sound cue was triggered when the rat turned its head in a specific direction For remote representation neurofeedback, the sound cue was triggered by online detection of a specific remote hippocampal representation (as in pretraining, the rat had to nosepoke within 3 seconds of the sound cue to receive the milk reward).

Neural implant

Rats that completed the pretraining were surgically implanted with a Microdrive containing 64 individually moveable tetrodes targeting dorsal hippocampus[46]. This surgery was performed as described previously[13] and in compliance with the UCSF IACUC animal safely protocols. Following surgery, tetrodes were lowered to the dorsal CA1 pyramidal cell layer over 2-3 weeks.

Clusterless decoding of hippocampal representation

Once tetrodes were lowered to the pyramidal cell layer, hippocampal spatial representations were decoded from recordings using a decoding algorithm described previously[27]. To decode position from action potential firing (spiking) of CA1 neurons, this algorithm builds a marked point-process model to relate all CA1 spikes (above a set threshold) to the rat's position and then applies this model to estimate the position of each new spike. Position estimates for all spikes in a small time window (6 msec) are combined to generate a prediction of the spatial position represented in

hippocampus. A key feature of this decoding algorithm is that it does not require spike sorting ("clusterless").

Real-time spatial decoding

The clusterless decoding algorithm was implemented for fast operation using a desktop computer with 28 cores and parallelization of the decoding algorithm using message passing interface (MPI). During each recording session, the system decoded ~1 million spikes. With a decoding latency of 30 msec, 90-99% of incoming spikes were decoded. The details of the real-time decoder are described in an accompanying manuscript[24].

Head direction feedback

After surgery recovery and tetrode adjusting, each rat started with 4 days of head direction conditioning. Each training day included 3 recording sessions. The recording sessions followed the task structure described above. During the feedback part of the task, the sound cue and reward was triggered by the rat's head direction. Reward cue was triggered by a head turn to either 30 degrees to left (arm 1) or to the right (arm 2) when facing forward towards the base of the outer arms. The angular accuracy required to trigger the reward cue increased over the first 25 rewards of each feedback session (from +/- 20 degrees to +/- 3 degrees). The feedback phase lasted for 30 minutes or until the rat received 75 rewards, whichever occurred first. For the first 2 days, the target direction was pointing towards one arm and then the target switched to the opposite arm for the next 2 days.

Remote representation neurofeedback

Following head direction feedback, the rat switched to remote representation neurofeedback. Each day had 3 recording sessions and the target location (either the distal 25 cm of arm 1 or arm 2) switched every 3 days. Based on recording quality rats were trained for 6 days (rat 1) or 18 days (rats 2-6). During the exploration phase at the beginning of each session, the recorded CA1 spikes were used to build an encoding model that associated spiking activity with spatial locations in the environment. Then, during the feedback phase, when the decoder detected a coherent representation of the target location, the sound cue was triggered and reward was available at the center port. The following requirements had to be met to trigger the reward cue: 30 msec running average of decoded position, >40% in target arm end, <20% in opposite arm, <20% in box area, at least 2 tetrodes with spatially specific spikes during 30 msec window, and rat physical distance to center port < 17cm. The feedback phase lasted for 30 minutes or until the rat received 75 rewards, whichever occurred first. In rat 1, only 1 tetrode with spatially specific spikes was required.

Data collection and processing

Hippocampal tetrode recordings were collected using SpikeGadgets hardware and Trodes recording software (sampling rate, 30 kHz). The rat's physical position was recorded using a Manta camera (frame rate 30 Hz) and red and green LEDs attached to the headstage were tracked using Trodes. These LEDs allowed calculation of head direction. Raw electrical signal, position tracking, and reward port nosepokes were extracted from the Trodes recording file and converted to an NWB file (https://github.com/orgs/LorenFrankLab/rec_to_nwb). The NWB file was inserted into a database organized by the Spyglass package for reproducible analysis of neural data

(https://github.com/orgs/LorenFrankLab/spyglass)[47]. All subsequent analysis was done on data within the database.

Real-time decoder output
The decoded mental position as estimated by the real-time decoder was saved each session in a separate recording file. This file included the rat's actual position and decoded position for every 6 msec time bin and each detected reward event (head direction or remote representation).

Data inclusion
Recording sessions were included with high quality spatial decoding. High quality sessions were defined as: during the exploration period, when the rat was running outward from the box and in the end of the target arm, the posterior maximum of the decoded position was in the target arm in at least 65% of time bins.

Spike Sorting
The raw tetrode recording was band-pass filtered from 600-6000 Hz and spikes were detected and sorted using MountainSort4[32]. During subsequent curation, noise clusters and multiunit clusters were removed automatically. Manual inspection was used to merge clusters that were inappropriately split. The remaining clusters are considered "isolated single units."

LFP analysis
For LFP analysis, the raw tetrode recordings were band-pass filtered from 0 – 600 Hz. Previously described methods were used to detect sharp wave ripples[13]. Theta phase was calculated from one tetrode in each rat located above the CA1 pyramidal cell layer (stratum oriens).

Cell assembly analysis
Cell assemblies were identified using ICA with methods described previously[31]. The spike times of all isolated single units for a single recording session were binned into 30 msec bins and then co-firing cell assemblies were detected.

Histology
Rat brains were fixed with 4% PFA, serially sectioned, and then treated with Nissl staining to visualize hippocampal cellular structure and electrolytic lesions (made after the experiment) at the tetrode recording sites.

QUANTIFICATION AND STATISTICAL INFORMATION

A few analyses could not be performed on all recording sessions because of errors with data saving and preprocessing. The missing sessions are listed here.

Figure 1, 2: no missing data

Figure 3: rat 1 (2 missing sessions), rat 2 (2 missing), rat 3 (1 missing), rat 4 (1 missing), rat 5 (4 missing), rat 6 (3 missing).

Figure 4: rat 1 (4 missing head direction, 2 missing neurofeedback), rat 2 (0 missing), rat 3 (1 missing neurofeedback), rat 4 (0 missing), rat 5 (4 missing neurofeedback), rat 6 (2 missing neurofeedback).

Figure 1d
The decoded mental position that triggered reward (30 msec) was averaged across all sessions of each type for each rat (head direction, arm 1 representation, or arm 2 representation). These 3 average representations are plotted for each rat. Dashed lines are 75% confidence intervals of the data.

Figure 2c
Spatial representation was calculated by counting the number of 6 msec time bins while the rat was near the center reward port (within 17cm) and > 40% of the decoded mental position (posterior mass) matched the specified location. For target representation (left), the location was the farthest 25 cm of the target arm from the box area ("arm end"), for non-rewarded representation (right), the location was the closest 25 cm of target arms to the box area ("arm base"). To calculate the fraction of time, the number of time bins meeting these criteria was divided by the total number of time bins at the reward port during the feedback part of the session. To match the amount of reward for head direction and remote representation conditioning, only sessions when the rats received >90% of the maximum rewards were included. Each dot represented one session, and significance was calculated with the Mann-Whitney test. Grouped analysis: a linear mixed model was created using the values just described grouped by rat. The tested variable was session type: either head direction or remote representation.
n: Rat 1: Head direction: n=12 sessions, neurofeedback: n=3. Rat 2: Head direction: n=12, neurofeedback: n=33. Rat 3: Head direction: n=10, neurofeedback: n=16. Rat 4: Head direction: n=10, neurofeedback: n=16. Rat 5: Head direction: n=11, neurofeedback: n=11. Rat 6: Head direction: n=9, neurofeedback: n=5.

Figure 2d
In each rat, linear regression was performed on the fraction of target representation for each neurofeedback session. Top panels show change in target representation over time and bottom panels show change in non-target representation. Grouped analysis: fraction of target representation was z-scored within each rat and then data from all rats was combined and linear regression was performed on the combined data.
n: Rat 1: n=7 sessions, Rat 2: n=50. Rat 3: n=42. Rat 4: n=44. Rat 5: n=50. Rat 6: n=51.

Figure 3a
Place field for each cell was calculated as the occupancy normalized firing rate in the 2D task environment. Spatial bins are 3cm x 3cm.

Figure 3b and 3c
Left plot shows the weights of each cell for an example cell assembly. Right plot shows occupancy-normalized assembly activity at each location in the task environment.

Figure 3e

For each session, the firing of assemblies that represented the target location (highest assembly strength for either box, arm 1 end, or arm 2 end, with strength > 100) was compared at times of representation detection (reward cue) to randomly selected times while the rat was at the center reward port. Target assembly at reward cue times was also compared to non-target assembly (all other assemblies) at reward cue times. Significance calculated with Mann-Whitney test.
n: Rat 1: target: n=8 cell assemblies, random: n=8, non-target: n=72. Rat 2: target: n=126, random: n=126, non-target: n=694. Rat 3: target: n=81, random: n=81, non-target: n=527. Rat 4: target: n=57, random: n=57, non-target: n=568. Rat 5: target: n=85, random: n=85, non-target: n=578. Rat 6: target: n=53, random: n=53, non-target: n=769.

Figure 3g
For each session, the firing of assemblies that represented the target location was compared at times of target representation (>40% posterior mass in target location) to randomly selected times while the rat was at the center reward port. Target assembly at target representation times was also compared to non-target assembly (all other assemblies) at target representation times. Significance calculated with Mann-Whitney test.
n: Rat 1: target: n=8 cell assemblies, random: n=8, non-target: n=72. Rat 2: target: n=126, random: n=126, non-target: n=694. Rat 3: target: n=81, random: n=81, non-target: n=527. Rat 4: target: n=57, random: n=57, non-target: n=568. Rat 5: target: n=85, random: n=85, non-target: n=578. Rat 6: target: n=53, random: n=53, non-target: n=769.

Figure 4b
Each recording session was divided into three states: movement (rat speed > 4cm/sec), stillness (rat speed < 4cm/sec without SWR) and SWR (still with detected SWR). Decoding time bins (6 msec) with target representation (>40% posterior mass in target location) were assigned to each state across all recording sessions for each rat. The bar plot shows the fraction of these assignments for each rat.

Figure 4c
This is the same analysis as Figure 2c, except that times during SWRs have been separated from the remainder of the recording session. Fraction of target representation was calculated as described during SWRs times and non-SWRs times. Grouped analysis same as Figure 2c.
n: Rat 1: Head direction: n=8 sessions, neurofeedback: n=3. Rat 2: Head direction: n=12, neurofeedback: n=33. Rat 3: Head direction: n=10, neurofeedback: n=15. Rat 4: Head direction: n=10, neurofeedback: n=16. Rat 5: Head direction: n=11, neurofeedback: n=9. Rat 6: Head direction: n=9, neurofeedback: n=5.

Figure S1c
Total number of spikes detected by the real-time decoder was counted for each included recording session.
n: Rat 1: n=19 sessions, Rat 2: n=62. Rat 3: n=53. Rat 4: n=55. Rat 5: n=62. Rat 6: n=63.

Figure S1d
Dropped spikes were defined as spikes that were decoded with a latency >30 msec. Fraction dropped spikes: dropped spike count / total spike count for each session.
n: Rat 1: n=19 sessions, Rat 2: n=62. Rat 3: n=53. Rat 4: n=55. Rat 5: n=62. Rat 6: n=63.

Figure S2b
High quality sessions ("good") were defined as: during the exploration period, when the rat was running outward from the box and in the end of the target arm, the posterior maximum of the decoded position was in the target arm in at least 65% of time bins.

Figure S2c
Decoding error was calculated as the difference between the rat's actual linear position and the posterior maximum for each 6 msec time bin during movement. The decoding error for all included sessions was combined for each rat and the plot is a histogram of error distance. Dashed line is median for all movement bins.
n: Rat 1: n=19 sessions, Rat 2: n=62. Rat 3: n=53. Rat 4: n=55. Rat 5: n=62. Rat 6: n=63.

Figure S4a
Posterior for all position bins was averaged across all time bins with >40% representation of the target location. The full posterior was split into 3 non-overlapping categories: target arm, opposite arm, and box area. Plots show histogram of the posterior probability for each category across all included time bins.

Figure S5
This is the same plot as Figure 2c and 2d except that the remote representations that triggered reward are excluded.
n (A): Rat 1: Head direction: n=12 sessions, neurofeedback: n=3. Rat 2: Head direction: n=12, neurofeedback: n=33. Rat 3: Head direction: n=10, neurofeedback: n=16. Rat 4: Head direction: n=10, neurofeedback: n=16. Rat 5: Head direction: n=11, neurofeedback: n=11. Rat 6: Head direction: n=9, neurofeedback: n=5.
n (B): Rat 1: n=7 sessions, Rat 2: n=50. Rat 3: n=42. Rat 4: n=44. Rat 5: n=50. Rat 6: n=51.

Figure S6
Same plot as Figures 2c and 2d except for opposite (non-target) arm end.
n (A): Rat 1: Head direction: n=12 sessions, neurofeedback: n=3. Rat 2: Head direction: n=12, neurofeedback: n=33. Rat 3: Head direction: n=10, neurofeedback: n=16. Rat 4: Head direction: n=10, neurofeedback: n=16. Rat 5: Head direction: n=11, neurofeedback: n=11. Rat 6: Head direction: n=9, neurofeedback: n=5.
n (B): Rat 1: n=7 sessions, Rat 2: n=50. Rat 3: n=42. Rat 4: n=44. Rat 5: n=50. Rat 6: n=51.

Figure S7b
Cell number active at reward times was calculated by taking the strongest cells from each target assembly (strength > 5 s.d. above mean) and counting how many of these cells had spikes in the 100 msec preceding each reward.
n: Rat 1: n=322 reward times, Rat 2: n=2999. Rat 3: n=1970. Rat 4: n=1737. Rat 5: n=2037. Rat 6: n=1176.

Figure S7c

For this plot, the number of active cells at reward times was divided by the total number of high strength cells from reward assemblies, thus giving a fraction of measured cells active before each reward.
n: Rat 1: n=322 reward times, Rat 2: n=2999. Rat 3: n=1970. Rat 4: n=1737. Rat 5: n=2037. Rat 6: n=1176.

Figure S8a
Theta phase during time bins of remote representation (>40% posterior mass in target location). Only movement times were considered (rat speed > 4cm/sec). To create null distribution (grey line with 95% confidence interval dashed lines) theta phase was calculated at random time bins matching the number of remote representation time bins. The random sampling was done 1000 times and the distribution was used to calculate the 95% confidence interval.
n: Rat 1: n=13822 representation times. Rat 2: n=42027. Rat 3: n=21747. Rat 4: n=76860. Rat 5: n=37556. Rat 6: n=71226.

Figure S8b
Power spectrum of time (1 sec) surrounding remote representation. Spectra were calculated for each remote representation time bin outside of SWRs when the rat was still (speed < 4cm/sec) at the center well and then averaged. As a control, power spectra were calculated for random times when the rat was at the center reward port. Left plots: average power spectrum calculated from tetrode within pyramidal cell layer. Right plots: average power spectrum from tetrode above cell layer (used for theta band analysis).
n: Rat 1: n=1071 representation times (1 sec non-overlapping bins, rat still, non-SWR). Rat 2: n=16032. Rat 3: n=13605. Rat 4: n=17164. Rat 5: n=15025. Rat 6: n=14241.

Figure S8c
Combined population spiking at time of remote representations (surrounding 1 sec). Combined spiking was calculated by summing spiking of all isolated single units in 10 msec time bins. As a control, random times were selected when the rat was at the central reward port. Dashed lines are 95% confidence intervals of the mean.
n: Rat 1: n=1840 representation times (1 sec non-overlapping bins). Rat 2: n=18211. Rat 3: n=22150. Rat 4: n=21021. Rat 5: n=23581. Rat 6: n=23092.

REFERENCES

1. Kim, J. J. ;Fanselow, M. S. Modality-specific retrograde amnesia of fear. *Science* **256**, 675–677 (1992).

2. Maren, S., Phan, K. L. & Liberzon, I. The contextual brain: implications for fear conditioning, extinction and psychopathology. *Nat Rev Neurosci* **14**, 417–28 (2013).

3. Anagnostaras, S. G., Gale, G. D. & Fanselow, M. S. Hippocampus and contextual fear conditioning: recent controversies and advances. *Hippocampus* **11**, 8–17 (2001).

4. Wang, M. E. *et al.* Differential roles of the dorsal and ventral hippocampus in predator odor contextual fear conditioning. *Hippocampus* **23**, 451–66 (2013).
5. Fanselow, M. S. Conditioned and unconditional components of post-shock freezing. *Pavlov. J. Biol. Sci.* **15**, 177–182 (1980).
6. Muller, R. A quarter of a century of place cells. *Neuron* **17**, 813–822 (1996).
7. Karlsson, M. P. & Frank, L. M. Awake replay of remote experiences in the hippocampus. *Nat Neurosci* **12**, 913–918 (2009).
8. Johnson, A. & Redish, A. D. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci* **27**, 12176–89 (2007).
9. Yu, J. Y. & Frank, L. M. Prefrontal cortical activity predicts the occurrence of nonlocal hippocampal representations during spatial navigation. *PLOS Biol.* **19**, e3001393 (2021).
10. Wikenheiser, A. M. & Redish, A. D. Hippocampal theta sequences reflect current goals. *Nat Neurosci* **18**, 289–94 (2015).
11. Kay, K. *et al.* Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. *Cell* **180**, 552-567.e25 (2020).
12. Buzsaki, G. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus* **25**, 1073–188 (2015).
13. Gillespie, A. K. *et al.* Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice. *Neuron* **109**, 3149-3163 e6 (2021).
14. Diba, K. & Buzsaki, G. Forward and reverse hippocampal place-cell sequences during ripples. *Nat Neurosci* **10**, 1241–2 (2007).
15. Joo, H. R. & Frank, L. M. The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation. *Nat Rev Neurosci* **19**, 744–757 (2018).

16. Tarder-Stoll, H., Jayakumar, M., Dimsdale-Zucker, H. R., Günseli, E. & Aly, M. Dynamic internal states shape memory retrieval. *Neuropsychologia* **138**, 107328 (2020).
17. Mildner, J. N. & Tamir, D. I. Spontaneous Thought as an Unconstrained Memory Process. *Trends Neurosci.* **42**, 763–777 (2019).
18. Morris, R. Developments of a water-maze procedure for studying spatial learning in the rat. *J Neurosci Methods* **11**, 47--60 (1984).
19. Eichenbaum, H. & Cohen, N. J. *From Conditioning to Conscious Recollection*. (Oxford University Press, New York, 2001).
20. Toader, A. C. *et al.* Anteromedial thalamus gates the selection and stabilization of long-term memories. *Cell* **186**, 1369-1381.e17 (2023).
21. Tang, W., Shin, J. D. & Jadhav, S. P. Multiple time-scales of decision-making in the hippocampus and prefrontal cortex. *eLife* **10**, e66227 (2021).
22. Deng, X., Liu, D. F., Kay, K., Frank, L. M. & Eden, U. T. Clusterless Decoding of Position from Multiunit Activity Using a Marked Point Process Filter. *Neural Comput* **27**, 1438–60 (2015).
23. Denovellis, E. L. *et al.* Hippocampal replay of experience at real-world speeds. *eLife* **10**, (2021).
24. Joshua P Chu *et al.* RealtimeDecoder: A fast software module for online clusterless decoding. *bioRxiv* 2024.05.03.592417 (2024) doi:10.1101/2024.05.03.592417.
25. Singer, A. C. & Frank, L. M. Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron* **64**, 910–921 (2009).
26. Ambrose, R. E., Pfeiffer, B. E. & Foster, D. J. Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. *Neuron* **91**, 1124–36 (2016).

27. Deng, X., Liu, D. F., Kay, K., Frank, L. M. & Eden, U. T. Decoding position from multiunit activity using a marked point process filter. *BMC Neurosci.* **16**, 1–1 (2015).

28. Kloosterman, F., Layton, S. P., Chen, Z. & Wilson, M. A. Bayesian Decoding using Unsorted Spikes in the Rat Hippocampus. *J Neurophysiol* (2013) doi:10.1152/jn.01046.2012.

29. Eden, U. T. & Brown, E. N. Continuous-Time Filters for State Estimation from Point Process Models of Neural Data. *Stat Sin* **18**, 1293–1310 (2008).

30. Buzsaki, G. Neural syntax: cell assemblies, synapsembles, and readers. *Neuron* **68**, 362–85 (2010).

31. Lopes-dos-Santos, V., Ribeiro, S. & Tort, A. B. L. Detecting cell assemblies in large neuronal populations. *J Neurosci Methods* **220**, 149–166 (2013).

32. Chung, J. E. *et al.* A Fully Automated Approach to Spike Sorting. *Neuron* **95**, 1381-1394.e6 (2017).

33. Skaggs, W. E., McNaughton, B. L., Wilson, M. A. & Barnes, C. A. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149–72 (1996).

34. Comrie, A. E., Frank, L. M. & Kay, K. Imagination as a fundamental function of the hippocampus. *Philos. Trans. R. Soc. B Biol. Sci.* **377**, 20210336 (2022).

35. Kay, K. & Frank, L. M. Three brain states in the hippocampus and cortex. *Hippocampus* (2018) doi:10.1002/hipo.22956.

36. Wessberg, J. *et al.* Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* **408**, 361–365 (2000).

37. Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M. & Shenoy, K. V. High-performance brain-to-text communication via handwriting. *Nature* **593**, 249–254 (2021).

38. Carmena, J. M. *et al.* Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.* **1**, E42 (2003).

39. Ishikawa, D., Matsumoto, N., Sakaguchi, T., Matsuki, N. & Ikegaya, Y. Operant conditioning of synaptic and spiking activity patterns in single hippocampal neurons. *J Neurosci* **34**, 5044–53 (2014).

40. Lai, C., Tanaka, S., Harris, T. D. & Lee, A. K. Volitional activation of remote place representations with a hippocampal brain-machine interface. *Science* **382**, 566–573 (2023).

41. Gillespie, A. K., Astudillo Maya, D. A., Denovellis, E. L., Desse, S. & Frank, L. M. Neurofeedback training can modulate task-relevant memory replay rate in rats. Preprint at https://doi.org/10.7554/eLife.90944.2 (2024).

42. Kay, K. *et al.* A hippocampal network for spatial coding during immobility and sleep. *Nature* **531**, 185–90 (2016).

43. Yu, J. Y. *et al.* Distinct hippocampal-cortical memory representations for experiences associated with movement versus immobility. *Elife* **6**, (2017).

44. Tabuchi, E., Mulder, A. B. & Wiener, S. I. Reward value invariant place responses and reward site associated activity in hippocampal neurons of behaving rats. *Hippocampus* **13**, 117–132 (2003).

45. Goode, T. D., Tanaka, K. Z., Sahay, A. & McHugh, T. J. An Integrated Index: Engrams, Place Cells, and Hippocampal Memory. *Neuron* **107**, 805–820 (2020).

46. Voigts, J., Newman, J. P., Wilson, M. A. & Harnett, M. T. An easy-to-assemble, robust, and lightweight drive implant for chronic tetrode recordings in freely moving animals. (2019) doi:10.1101/746651.

47. Lee, K. H. *et al.* Spyglass: a framework for reproducible and shareable neuroscience research. *BioRxiv Prepr. Serv. Biol.* 2024.01.25.577295 (2024) doi:10.1101/2024.01.25.577295.