Episodic memory retrieval is supported by rapid replay of episode content

Authors: G. Elliott Wimmer[1,2,*,†], Yunzhe Liu[1,2,*], Neža Vehar[1,2], Timothy E.J. Behrens[2,3], Raymond J. Dolan[1,2]

Affiliations:

[1]Max Planck University College London Centre for Computational Psychiatry and Ageing Research, University College London, London, UK.

[2]Wellcome Centre for Human Neuroimaging, University College London, London, UK.

[3]Wellcome Centre for Integrative Neuroimaging, Centre for Functional Magnetic Resonance Imaging of the Brain, University of Oxford, Oxford, UK.

* These authors contributed equally to this work.

† Corresponding author. Email: e.wimmer@ucl.ac.uk (G.E.W.)

Abstract

Our everyday experience shapes how we represent the structure of the world. Retrieval of these experiences from memory is fundamental for informing our future decisions. To uncover the fine-grained neurophysiological mechanisms that support such retrieval we studied participants who first experienced unique multi-component episodes, and subsequently completed cued memory retrieval of these whilst undergoing magnetoencephalography (MEG). Successful retrieval was supported by sequential replay of episodes, with a temporal compression factor greater than 60. This sequential replay was stronger in those participants with weaker overall memories. Replay direction, forward or backward, was dependent on whether a participant's goal was to retrieve elements of an episode that followed, or preceded, a retrieval cue. Our results demonstrate that memory-based decisions are supported by a rapid replay mechanism that flexibly shifts in direction.

Introduction

Although a subject of intense study, the fine-grained mechanisms underlying how we retrieve episodes of experience are unknown (*1*). Understanding the underlying neurophysiological processes can throw light on how episodes are represented in memory and subsequently retrieved to guide behavior (*2, 3*). Here we investigate whether episodes of experience are represented in a way that yields compressed sequential replay during retrieval, whether replay supports successful retrieval, and whether replay is flexibly influenced by internal goals.

Observations from animal studies have identified offline reactivation of sequences of hippocampal place cells that reflect past and future trajectories, thought to support memory consolidation, retrieval, and planning (*4-6*). Recently, animal studies have established a relationship between such replay strength and successful performance on spatial navigation tasks (*4, 5*). While untested, it is speculated that compressed replay might also support episodic memory retrieval in humans (*7*).

Human neuroimaging studies provide evidence for cue-elicited reactivation of stimulus associations at retrieval (*8-16*). A limitation of these studies is their inability to probe mechanisms supporting structured and temporally compressed reactivation, i.e. replay that proceeds at a rate faster than the original experience. An important advance in recent human neuroimaging research has been the identification of rapid sequential replay of internal state representations (*17, 18*). Here, we leverage these same methods to ask whether sequential replay supports memory decisions in humans.

We tested a hypothesis that episodic memory retrieval depends on rapid compressed replay of memory elements. Previous research demonstrating replay, which did not link replay to behavior, identified an approximately 40-50 ms lag between states (elements of a sequence) either during tasks involving lengthy planning periods or during rest periods (*17, 18*). Under similar conditions in rodents replay is known to occur preferentially during brief high-frequency sharp-wave ripple (SWR) events in the hippocampus (*19-21*). In contrast, slower theta-related sequence events are observed during active navigation and decision making in rodents (*19, 20, 22, 23*). Thus, during

memory retrieval, we expected that performance in the current experiment would be supported by replay events with a relatively longer lag between states.

Replay direction, forward or backward, is known to be influenced by conditions such as active movement and reward receipt (*18, 24, 25*), potentially serving different computational functions (*26*). Consequently, we expected replay direction may change flexibly based on internal states or task demands. In our study we predicted replay would switch direction depending on whether the current goal was to retrieve memory components that followed a cued episode versus having to retrieve memory components that preceded a cued episode. Inspired by evidence from rodents we predicted that replay onset, irrespective of directionality, would be coupled to increased low-to-mid-frequency power in the medial temporal lobe (MTL; *18*). Finally, we reasoned that the strength of encoding, as reflected in better memory performance, would relate to enhanced memory consolidation (*1, 7*). As greater experience is associated with less marked replay in rodents (*23, 27*), we expected a less dominant expression of replay in participants with near-ceiling memory performance. In these participants, theoretical considerations led us to predict performance could be supported by a form of clustered pattern completion for episode elements (*9, 28, 29*).

We designed a novel episodic memory task and combined this with our recently-developed MEG analytic methods (*17, 18*). In brief, on day 1 participants experienced temporally extended self-oriented episodes, where each single-exposure episode was composed of five discrete and unique picture stimuli assembled into a narrative story (**Fig. 1a** and **Fig S1**). Following overnight consolidation, we elicited cued retrieval of these episodes whilst at the same time obtaining MEG data to index fast neural dynamics supporting retrieval (**Fig. 1b**).

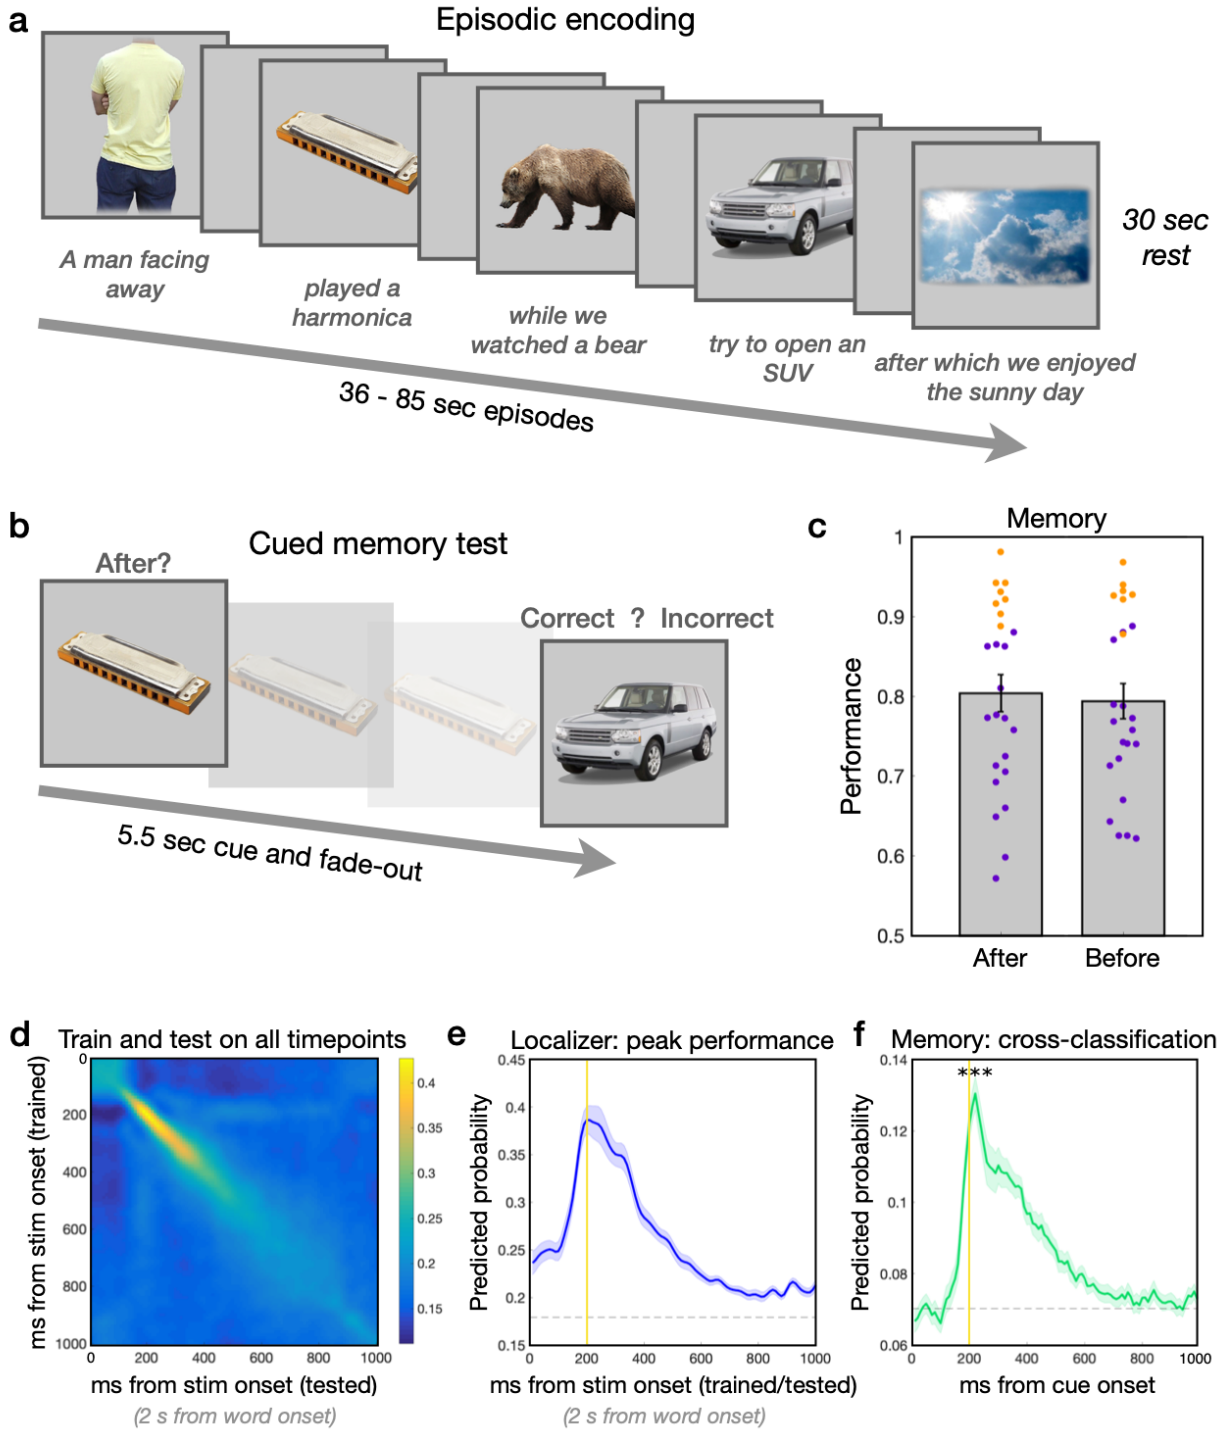

Fig. 1. Experimental design and decoding of the episode elements. (**a**) On day 1, in the episodic encoding phase we presented subjects with eight extended non-spatial episodes, with a single exposure per episode. Episodes contained five stimulus elements. The first four episode elements were selected from six distinct categories of

pictures. Participants were incentivised to encode the precise order of the episode elements. (**b**) On day 2, in the episodic memory test phase, participants were instructed to think about stimuli that followed a cued element from an individual episode, referred to as an 'after' condition trial. A test probe was then presented after 5.5 sec. The sequential order referred to any stimulus from the same episode that followed this cue; here, the depicted answer would be 'correct'. By contrast, in a 'before' condition trial, participants were instructed to think about what elements preceded a presented cue picture, followed by a test probe. (**c**) Mean memory performance in the after and before conditions respectively. Purple dots represent individual data points for regular participants with sufficient incorrect response (error) trials free from MEG artifacts for accuracy analyses (after, $n = 17$; before, $n = 18$); the remaining very high performance participants are shown in orange (see also **Fig. S1**). (**d**) Classifier performance for episode element categories presented during the localizer phase, training and testing at all time points, showing good discrimination of the 6 categories used to compose the first four episode elements. In localizer trials, note that a word naming the upcoming stimulus appeared 2 s before the stimulus, contributing to above-chance classification at 0 ms. (**e**) Peak classifier performance at 200 ms after stimulus onset in the localizer phase (see also **Fig. S2**). (**f**) Application of the trained classifier to cue onset in memory retrieval trials demonstrated above chance decoding of the current on-screen category during retrieval. (Error bars and shaded error margins represent standard error of the mean (SEM).)

Results

As a first step we confirmed that we could reliably identify neural patterns associated with individual episode elements, each drawn from one of six different stimulus categories. Note that the final element of each episode was not taken from a decoded category. A classifier trained on the localizer phase showed successful discrimination of the categories that made up the episodes with peak decoding at 200 ms after stimulus onset (**Fig. 1d-e**; **Fig. S3**), in line with previous reports (*17, 18*). This classifier generalized to the cued memory phase, showing significant across-phase classification

of cue category (peaking at 210 ms after the cue; compared versus chance at 200 ± 10 ms (the peak timepoint in localizer phase) $t_{(24)} = 9.80$, $p < 0.001$; **Fig. 1f**).

To test our specific predictions of a replay mechanism underlying episodic retrieval, we looked for compressed sequential reactivation of episode elements during the retrieval period. In this analysis, we first derived measures of category evidence – representing potential reactivation of memory elements – at each timepoint by applying the trained classifiers to the retrieval period MEG data. We then tested for lagged cross-correlations between episode element reactivations across the retrieval period, yielding a measure of 'sequenceness' in both forward and backward directions (*17, 18*) (**Fig. S2; Methods**). Following previous reports, to identify time lags showing potential sequenceness and examine the relationship to individual differences in memory performance, we tested for a difference between the forward and reverse direction components (*17, 18*). Our initial analyses focused on memory retrieval in the *after* condition, where participants are oriented to retrieve the episode in a forward sequential direction from the cue, as we expected that this condition would be easier and more naturalistic than the *before* condition. After identifying a time lag of interest, we examined both the relationship between sequenceness and individual differences in memory performance and finally the relationship between sequenceness and trial-by-trial memory retrieval success.

Our first sequenceness analysis was conducted to identify potential state-to-state time lags of interest and focused on correct trials, where we expected stronger sequenceness. In the *after* condition, we identified an overall dominance of reverse replay (backwards > forwards sequenceness) during correct trials, peaking between 100-120 ms (**Fig. 2a**). This difference did not survive correction for number of tests across lags, so it should not be interpreted on its own. However, it provides a time window for replay analyses, and all further analyses use this time-window unless otherwise noted. Notably, this time-window for rapid online retrieval is at a longer state-to-state time lag than the 40-50 ms lag found in other experiments reporting replay during extended planning or rest (*17, 18*). As in rodents, these resting replay events are associated with sharp-wave ripples in humans (*18*). However, rodents also show sequence events during ongoing behaviour that are slower and associated with ongoing

hippocampal theta rhythms (*22, 23*). Such online sequence events have not yet been identified in humans.

To provide an initial test of the relationship between replay and episodic retrieval, we examined the relationship between replay strength in correct trials, averaged across the 100-120 ms time lags, and overall memory performance. We found that differential sequenceness was correlated with mean memory performance (100-120 ms lag; r = 0.4254, p = 0.034; **Fig. 2b**). As sequenceness was on average negative – showing predominantly a reverse direction replay – this suggests that stronger reverse replay characterised individuals with weaker performance. This relationship between replay and memory strength is in line with the data from rodents showing stronger replay during initial versus late learning (*23, 27*).

As an initial test of our prediction that internal goals – whether looking forward or backward in time through an experience – are important for retrieval and replay, we examined whether the relationship between replay and individual differences in performance changed from the *after* compared to the *before* condition. If a task goal affected replay, we would expect stronger forward sequenceness to be related to weaker performance. Indeed, in the before condition we found the degree of dominantly forward sequenceness negatively correlated with mean memory performance (100-120 ms lag; r = -0.4077; p = 0.0431; **Fig. 2c**). This relationship differed significantly from that in the after condition (z = 2.411; p = 0.0159; two-tailed, conservatively using the test for dependent correlations), providing initial support for the prediction that retrieval orientation influences the characteristics of replay that supports behaviour. Importantly, the direction of the results in both the after and before conditions indicated that replay was stronger in participants with lower overall performance, such that replay plays less of a role in retrieval with near-ceiling levels of performance.

We found no relationship between sequenceness and behavior in the shorter 40-50 ms state-to-state time lag as identified in previous studies (**Fig. S4**). In an exploratory analysis that examined evidence for sequences of episode elements present in any of the other 7 episodes (but not present in the current episode), we found a numerically negative sequenceness effect at 40 ms, but again found no relationship to memory performance (**Fig. S4**).

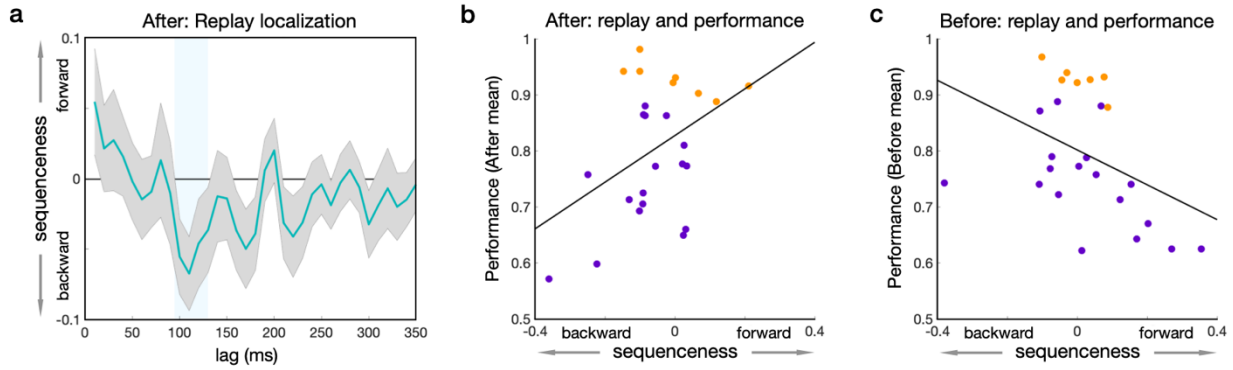

Fig. 2. Mean sequenceness (replay) in the *after* condition and the relationship between sequenceness and performance in the *after* and *before* memory retrieval conditions respectively. (**a**) In the *after* condition, mean forward minus backwards sequenceness for correct memory trials (when participants accurately answered the memory question). On correct trials, a peak of reverse sequenceness was observed in lags from 100-120 ms. This time window was used for subsequent analyses. (Shaded error margins represent SEM.) (**b**) In the *after* condition, stronger mean reverse sequenceness on correct trials negatively correlated with overall mean memory performance (percentage of correct trials). (As in Fig. 1c the data points for the regular performance participants are shown in purple; high performance participants are shown in orange.) (**c**) In the *before* condition, stronger forward sequenceness was related to lower performance. The results in the *after* and *before* conditions support a stronger role for replay in retrieving weaker memory traces. (*p < 0.01; **p < 0.01)

We next utilised analytic techniques that simultaneously examined the influence of forward and backward sequenceness on memory performance. First, we examined the relationship between sequenceness and individual differences in performance, which confirmed the above results: weaker memory performance related to stronger reverse replay in the after condition, while weaker memory related to stronger forward replay in the before condition (see **Supp. Results**).

To examine the relationship between trial-by-trial sequenceness and accuracy, we used multilevel regression analyses. In these analyses we necessarily focus on a subgroup of participants, excluding the very high performing participants who have too few incorrect trials to support reliable estimates. Critically, we found that reverse sequenceness from 100-120 ms was positively related to trial-by-trial accuracy (multilevel regression on accuracy in n = 17 participants with sufficient incorrect trials; forward β = -0.1503 [-0.310 -0.001]; z = -1.912; p = 0.0504; reverse β 0.180 [0.020 0.322]; z = 2.308; p = 0.0184; **Fig. 3a**). An example of a reverse sequence in the *after* condition for a single participant is shown in **Fig. 3c**. By contrast, in the before condition forward, but not reverse, sequenceness related positively to accuracy (regression in n = 18 participants with sufficient incorrect trials; forward β = 0.148 [0.007 0.293]; z = 2.051; p = 0.0400; reverse β = -0.051 [-0.189 0.097]; z = -0.690; p = 0.486; **Fig. 3A**). An example of a forward sequence in the before condition for a single participant is shown in **Fig. 3d**. As above, we found no relationship between sequenceness at a 40-50 ms lag and behavior (**Fig. S6**); we also did not find any relationship between the sequenceness measure derived from the alternative 7 episodes and behavior at either 40-50 or 100-120 ms lags (**Fig. S6**).

Importantly, the relationship between 100-120 ms replay and successful memory retrieval was significantly affected by the after versus before goal condition (condition by forward replay β = -0.162 [-0.275 -0.057]; z = -2.929; p = 0.0016; condition by reverse replay β = 0.12926 [0.023 0.236]; z = 2.331; p = 0.0160; **Fig. 3b**; n = 15 participants with sufficient incorrect trials in both the after and before conditions). The relationship between sequenceness and successful memory retrieval in both the after and before condition provides a clear link between sequenceness and behavior. Further, taking the results of the individual difference analyses together with these trial-by-trial analyses, we establish a double dissociation between replay direction and a participant's internal goal condition during retrieval. These findings demonstrate a flexibility in replay directionality that goes beyond previously reported effects of external events such as reward receipt (*18, 25*).

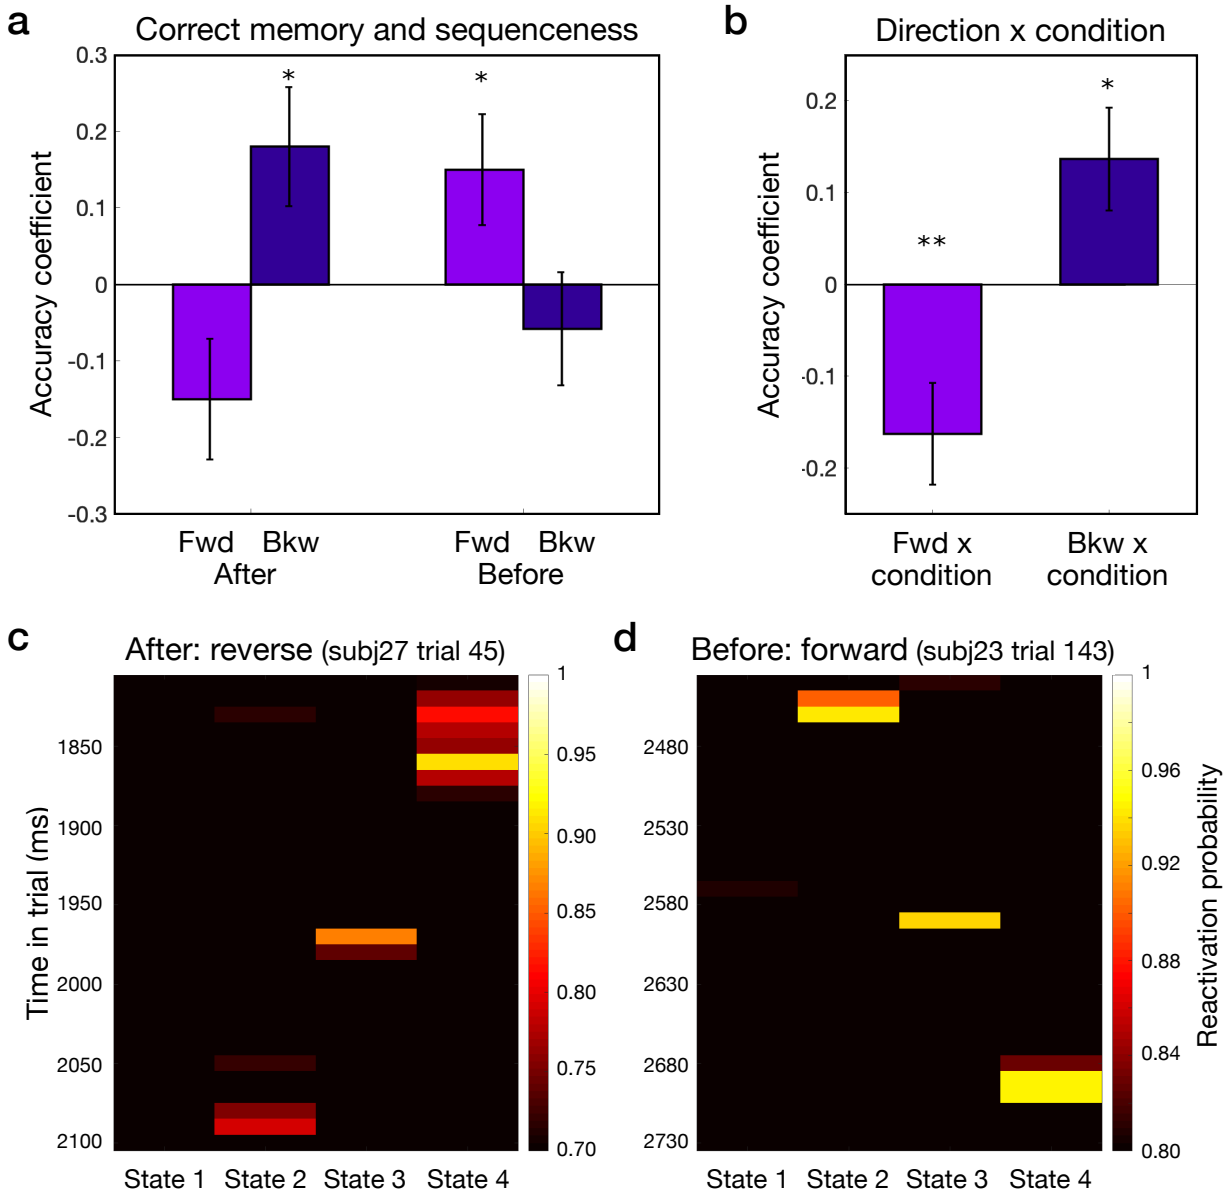

Fig. 3. Relationship between forward and backward sequenceness and memory retrieval, in the after and before conditions. (**a**) In the *after* condition (left), successful memory retrieval was supported by reverse sequenceness. In the *before* condition (right), retrieval was supported by forward sequenceness. (**b**) Interaction of replay direction (forward, backward) by condition (after, before) showing a stronger effect of forward replay on memory in the before condition and a stronger effect of backward replay on memory in the after condition. (The regular performance group in the combined sequenceness analysis included n = 15 participants common to the regular

performance group across the after and before conditions.) (**c**) Example of reverse sequenceness in the after condition. (**d**) Example of forward sequenceness in the before condition. (*p < 0.05; **p < 0.01; error bars represent standard error).

Inspired by neurophysiological studies showing that the hippocampus is a source for replay events, we next examined whether replay event onset related to power increases within the medial temporal lobe (*18*). Candidate replay onsets were identified by locating sequential reactivation events showing a 110 ms lag, applying a stringent threshold to these events, and using beamforming analysis to localize broadband power changes related to replay event onsets. For reverse replay events (in the *after* condition), and for forward replay events (in the *before* condition), this analysis localized activity at replay onset to the right anterior MTL, encompassing the hippocampus and entorhinal cortex (after: z = 3.72, p <0.001 whole-brain FWE; before: z = 3.73, p < 0.001 whole-brain FWE; **Fig. 4a, c**; **Table S2**), consistent with human fMRI results during rest in a cognitive paradigm (*30*). The increase in MTL power was selective to replay onset (**Fig. 4d**), in addition to a secondary peak in the after condition 1 lag later, at 110 ms. Replay onset also related to activity in two significant clusters in the right visual cortex in the after condition (**Table S2**; **Fig. S7**). Finally, we found evidence for increased power immediately preceding replay onset in the left anterior MTL in participants with lower performance (and stronger sequenceness; z = 3.82, p = 0.003 whole-brain FWE; **Fig 4b**; **Fig. S5**).

Replay onset was associated with broadband power increases of up to 50 Hz in both the after and before conditions (**Fig. S8**). However, we found no evidence for power increases in the high gamma frequency range that have been associated with replay events during rest (*18*) (**Fig. S8**). Interestingly, when looking at the relationship between time-frequency effects and accuracy, we found an increase in alpha frequency (8-12 Hz) during the retrieval period for correct versus incorrect trials (**Fig. S8**), a finding in the opposite direction of previous reports (*14, 31*). Notably, the alpha frequency range overlaps with the frequency expected from a 100-120 ms sequenceness lag.

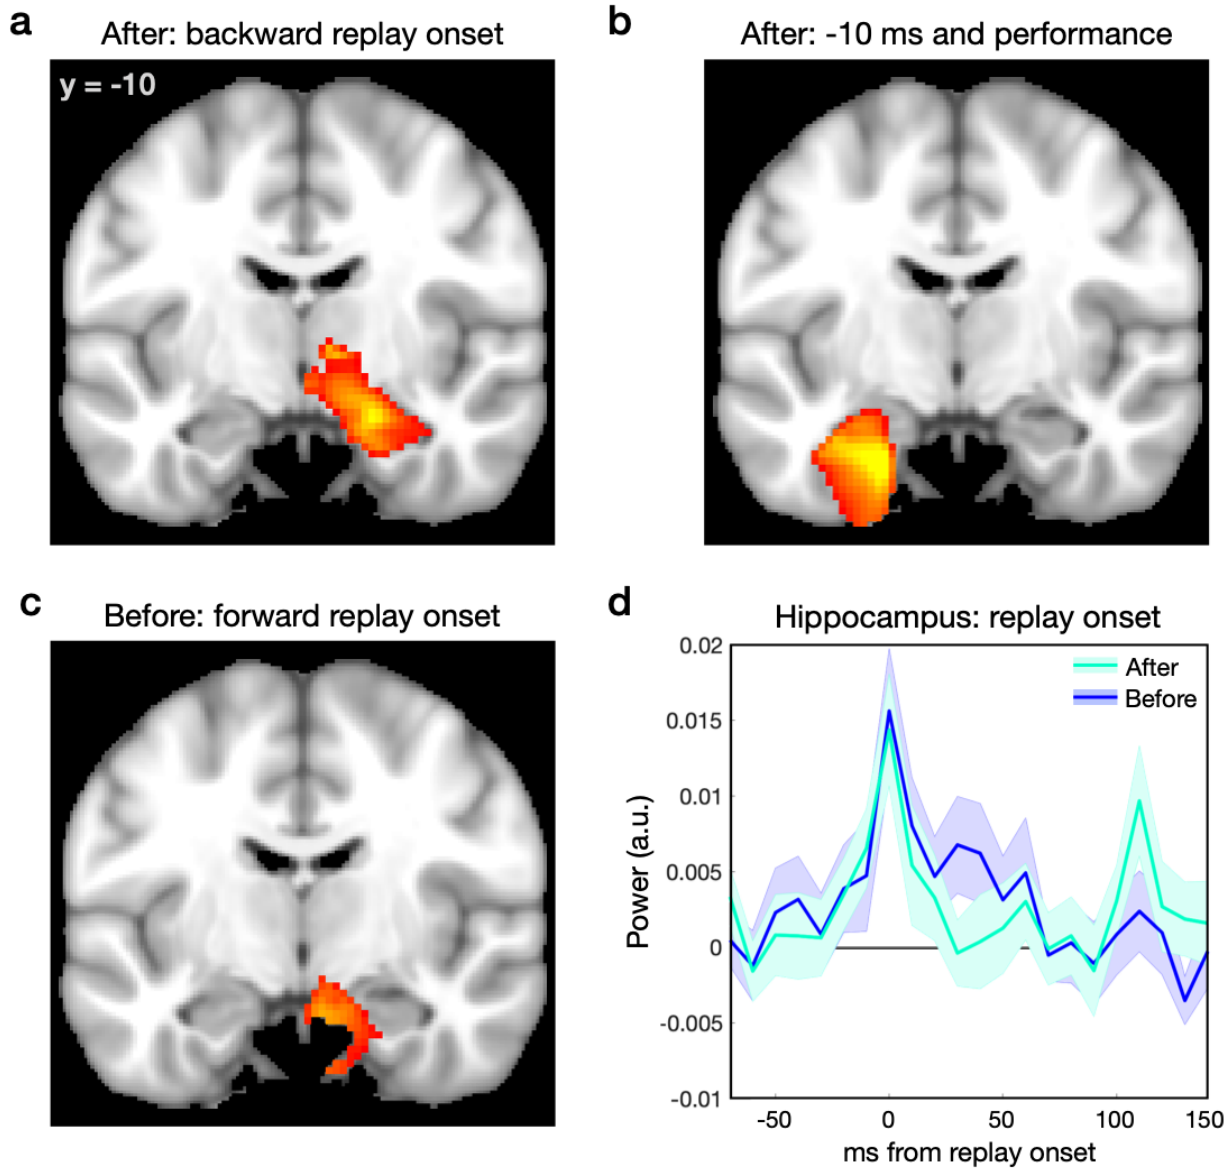

Fig. 4. Beamforming analysis of power increases at the onset of sequenceness events. (**a**) In the after condition, power in the right anterior MTL increased at onset of reverse sequenceness events (n = 25). (**b**) Power in the left MTL 10 ms before the onset of reverse sequenceness events correlated with performance, such that lower performing participants showed the strongest increase in power (**Fig. S5**). (**c**) In the before condition, power in the right anterior MTL increased at the onset of forward sequenceness events (n = 25). (**d**) Timecourse of power changes relative to replay onset in the anterior hippocampus in the after (cyan) and before (blue) conditions.

(Statistical maps thresholded at $p < 0.001$ uncorrected, for display; shaded error margins represent SEM.)

Finally, as very high performing participants did not show a relationship between replay and performance, we examined whether retrieval for strongly encoded memories was based on clustered pattern completion. Across all participants, rapidly after cue onset we found evidence for significant reactivation of within-episode elements compared to other-episode elements, none of which were displayed on the screen (average across timepoints showing best classification of on-screen element (**Fig. 1f**): 210 ±10 ms post-cue $t_{(24)} = 3.978$, $p < 0.001$; **Fig. 5a**). A reactivation event from a single participant is shown in **Fig. 5b**.

To examine the relationship between the cue-evoked reactivation effect and memory in very high performance participants, instead of a contrast of correct versus incorrect trials, we used a measure of mean performance for the episode cued on the current trial (a graded measure from 0 to 1). Cue-evoked reactivation from 200-250 ms positively related to performance on a given episode in very high performing participants ($n = 10$; $\beta = 0.0798$ [0.0368 0.1250]; $t = 3.452$; $p < 0.001$; **Fig. 5c**), an effect stronger in high compared to regular performance participants (regular $\beta = -0.0442$ [-0.1354 0.0050]; $t = -0.920$; $p = 0.3864$; difference $\beta = 0.1245$ [0.003 0.243]; $t = 2.029$; $p = 0.0432$; **Fig. 5c**). Additionally, although based on a very low number of trials, in the very high performing participants we found that incorrect trials were related to lower cue-evoked reactivation as compared to correct trials ($\beta = 2.470$ [0.401 4.683]; $t = 2.562$; $p = 0.020$). We found no significant relationship between cue-evoked responses and accuracy in regular performance participants (regular $\beta = 0.638$ [-0.322 1.596]; $t = 1.30$; $p = 0.1712$; difference $\beta = 1.78198$ [-0.390 3.972]; $t = 1.572$; $p = 0.1264$). Importantly, in regular performers, replay but not cue-evoked responses were related to performance: the trial-by-trial relationship between replay and accuracy in both the after and before conditions remained significant when including cue-evoked reactivation in the same model while effects for cue-evoked reactivation were not significant (**Fig. S8**). Finally, while we could also identify putative simultaneous reactivation events during the

subsequent retrieval period, we found no relationship between these events and performance in regular performers (**Fig. S8**), supporting the importance of sequential reactivation for successful episodic memory retrieval.

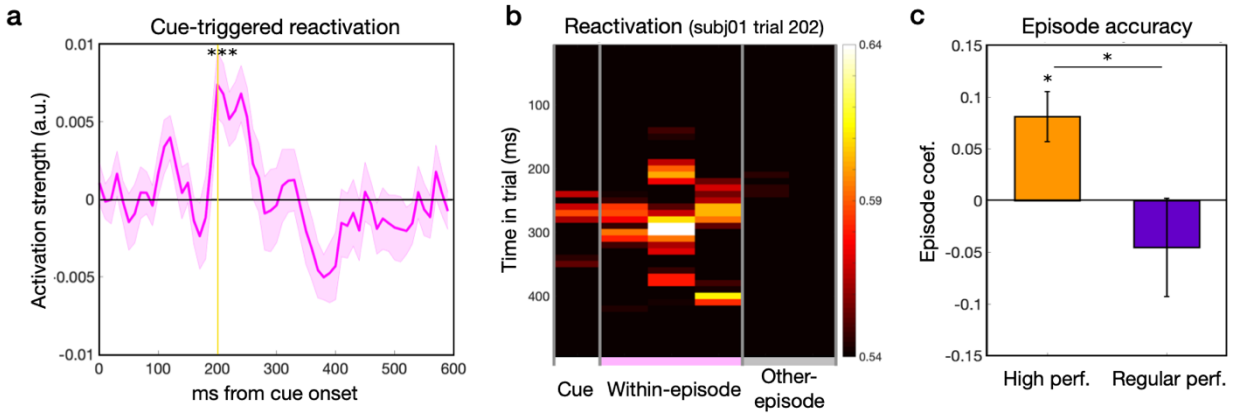

Fig. 5. The relationship between cue-evoked reactivation and performance. (**a**) Across the *after* and *before* conditions, we found evidence for cue-evoked reactivation of the elements present in the episode, peaking 200-250 ms after cue onset. (Shaded error margins represent SEM.) (**b**) Example of cue-evoked reactivation of within-episode elements in a single trial in a single participant. (**c**) Cue-evoked reactivation related to mean performance in a given episode for the high performance participants, but not regular performance participants (group breakdown based on number of incorrect trials across both the after and before conditions; high $n = 10$; regular $n = 15$; error bars represent standard error.)

Discussion

During episodic memory retrieval in humans, we show that a rapid sequential replay of episode elements relates to individual differences in memory performance, trial-by-trial retrieval success, and increased power in the anterior MTL. Importantly, as is the case with real-life episodes outside of the lab, in our experiment memory episodes were only

experienced a single time. Our results combined indicate an important role for replay with an element-to-element lag of 100-120 ms in online memory retrieval, establishing a novel connection between replay and ongoing behaviour in humans that has only recently been demonstrated in animal research (*4, 5, 27*).

Replay events spanned a temporal horizon of multiple seconds during retrieval, in contrast to a single instance of clustered pattern completion (*9, 28*). The latter was seen in very high performing participants alone, where there was a cue-evoked reactivation that closely resembled pattern completion. The absence of sequential replay in very high performing participants could reflect a difficulty in detection due to a sparse distribution or rapid decay of replay event frequency. Alternatively, when episodes are strongly encoded during the experience itself, different representations could begin to form, where order information is no longer represented by sequential replay but instead by clustered reactivation as we observed. A potentially related finding of a decreasing expression of replay with increasing experience has been reported in rodents (*23, 27*). Here we speculate that more strongly encoded and consolidated representations are enhanced by spontaneous reactivation and replay during post-learning rest and sleep (*6, 32-34*) and that these representations may be differentially supported by cortical systems (*28, 29, 35, 36*).

Replay in the current experiment showed an element-to-element lag of approximately 110 ms, representing a temporal compression factor of 60 to 150. This compression is in line with, or exceeds, the degree reported in offline place cell sequences in rodents (*37, 38*). Previous MEG research on replay in humans has reported a shorter 40-50 ms lag between replayed elements for very well-learned sequences (*17, 18*). These studies allowed for tens of seconds of planning or involved acquisition over minutes of rest; moreover, replay during rest was related to putative SWR events (*18*). This contrasts with our current experiment where there was a requirement for relatively rapid ‘online’ decisions. These different effects, influenced by task demands, parallel well-established results in animals. Here, slower theta-related sequence events are predominantly found during active navigation, while more compressed replay events are found during rest and sleep (including SWRs; *19, 20-23*).

Individual episodes of experience are important building blocks for creating a representation of the structure of the world (*2*). Episodic representations that support replay are likely to be important for how we successfully navigate spatial, social, and abstract environments (*3, 6, 39-43*). In turn, memory closely interacts with decision making (e.g. *10, 42*). The ability to reactivate episodes in a highly compressed manner provides a novel mechanism for very rapid retrieval and replay of previous experiences during choice (*44-46*), and our results reported here can motivate new directions of research into memory encoding, consolidation and decision making. Further, the flexible direction of episodic retrieval replay events that we identify may affect choice dynamics. We speculate that sequential replay flexibility and strength might serve as markers for impaired associative binding between memory elements caused by negative emotional events. Impaired, or pathologically disturbed, memory organization has a strong negative impact on well-being and behaviour, and future human research into memory replay may provide novel insights into memory disturbances seen in psychiatric disorders such as post-traumatic stress disorder and schizophrenia (*47, 48*).

References

1. H. Eichenbaum, N. J. Cohen, *From Conditioning to Conscious Recollection: Memory Systems of the Brain*. (Oxford University Press, New York, 2001).

2. G. Buzsaki, E. I. Moser, Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat Neurosci* **16**, 130 (Feb, 2013).

3. H. Eichenbaum, The role of the hippocampus in navigation is memory. *J Neurophysiol* **117**, 1785 (Apr 1, 2017).

4. B. E. Pfeiffer, D. J. Foster, Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74 (May 2, 2013).

5. S. P. Jadhav, C. Kemere, P. W. German, L. M. Frank, Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**, 1454 (Jun 15, 2012).

6. H. F. Olafsdottir, D. Bush, C. Barry, The Role of Hippocampal Replay in Memory and Planning. *Current biology : CB* **28**, R37 (Jan 8, 2018).

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

8. S. M. Polyn, V. S. Natu, J. D. Cohen, K. A. Norman, Category-specific cortical activity precedes retrieval during memory search. *Science* **310**, 1963 (Dec 23, 2005).

9. A. J. Horner, J. A. Bisby, D. Bush, W. J. Lin, N. Burgess, Evidence for holistic episodic recollection via hippocampal pattern completion. *Nat Commun* **6**, 7462 (Jul 2, 2015).

10. G. E. Wimmer, D. Shohamy, Preference by association: how memory mechanisms in the hippocampus bias decisions. *Science* **338**, 270 (Oct 12, 2012).

11. M. E. Wheeler, S. E. Petersen, R. L. Buckner, Memory's echo: vivid remembering reactivates sensory-specific cortex. *Proc Natl Acad Sci U S A* **97**, 11125 (Sep 26, 2000).

12. Z. Kurth-Nelson, G. Barnes, D. Sejdinovic, R. Dolan, P. Dayan, Temporal structure in associative retrieval. *Elife* **4**, (Jan 23, 2015).

13. S. Michelmann, B. P. Staresina, H. Bowman, S. Hanslmayr, Speed of time-compressed forward replay flexibly changes in human episodic memory. *Nature Human Behaviour* **3**, 143 (Feb, 2019).

14. S. Michelmann, H. Bowman, S. Hanslmayr, The Temporal Signature of Memories: Identification of a General Mechanism for Dynamic Memory Replay in Humans. *PLoS Biol* **14**, e1002528 (Aug, 2016).

15. T. I. Brown *et al.*, Prospective representation of navigational goals in the human hippocampus. *Science* **352**, 1323 (Jun 10, 2016).

16. G. E. Wimmer, C. Buchel, Learning of distant state predictions by the orbitofrontal cortex in humans. *Nature Communications* **10**, (Jun 11, 2019).

17. Z. Kurth-Nelson, M. Economides, R. J. Dolan, P. Dayan, Fast Sequences of Non-spatial State Representations in Humans. *Neuron* **91**, 194 (Jul 06, 2016).

18. Y. Liu, R. J. Dolan, Z. Kurth-Nelson, T. Behrens, Human replay spontaneously reorganises experience. *Cell*, (2019).

19. A. E. Papale, M. C. Zielinski, L. M. Frank, S. P. Jadhav, A. D. Redish, Interplay between Hippocampal Sharp-Wave-Ripple Events and Vicarious Trial and Error Behaviors in Decision Making. *Neuron* **92**, 975 (Dec 7, 2016).

20. G. Pezzulo, F. Donnarumma, D. Maisto, I. Stoianov, Planning at decision time and in the background during spatial navigation. *Curr Opin Beh Sci* **29**, 69 (2019).

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

22. A. M. Wikenheiser, A. D. Redish, Hippocampal theta sequences reflect current goals. *Nat Neurosci* **18**, 289 (Feb, 2015).

23. A. Johnson, A. D. Redish, Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci* **27**, 12176 (Nov 7, 2007).

24. K. Diba, G. Buzsaki, Forward and reverse hippocampal place-cell sequences during ripples. *Nat Neurosci* **10**, 1241 (Oct, 2007).

25. R. E. Ambrose, B. E. Pfeiffer, D. J. Foster, Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. *Neuron* **91**, 1124 (Sep 7, 2016).

26. D. J. Foster, M. A. Wilson, Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680 (Mar 30, 2006).

27. A. C. Singer, M. F. Carr, M. P. Karlsson, L. M. Frank, Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron* **77**, 1163 (Mar 20, 2013).

28. J. L. McClelland, B. L. McNaughton, R. C. O'Reilly, Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* **102**, 419 (Jul, 1995).

29. D. Marr, Simple memory: a theory for archicortex. *Philos Trans R Soc Lond B Biol Sci* **262**, 23 (Jul 1, 1971).

30. N. W. Schuck, Y. Niv, Sequential replay of nonspatial task states in the human hippocampus. *Science* **364**, (2019).

31. B. P. Staresina *et al.*, Hippocampal pattern completion is linked to gamma power increases and alpha power decreases during recollection. *Elife* **5**, (Aug 10, 2016).

32. A. Tambini, L. Davachi, Persistence of hippocampal multivoxel patterns into postencoding rest is related to memory. *Proc Natl Acad Sci U S A* **110**, 19591 (Nov 26, 2013).

33. A. Ben-Yakov, Y. Dudai, Constructing realistic engrams: poststimulus activity of hippocampus and dorsal striatum predicts subsequent episodic memory. *J Neurosci* **31**, 9032 (Jun 15, 2011).

34. L. Deuker *et al.*, Memory consolidation by replay of stimulus-specific neural activity. *J Neurosci* **33**, 19373 (Dec 4, 2013).

35. D. Kumaran, D. Hassabis, J. L. McClelland, What Learning Systems do Intelligent Agents Need? Complementary Learning Systems Theory Updated. *Trends Cogn Sci* **20**, 512 (Jul, 2016).

36. D. Tse *et al.*, Schemas and memory consolidation. *Science* **316**, 76 (Apr 6, 2007).

37. A. K. Lee, M. A. Wilson, Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* **36**, 1183 (Dec 19, 2002).

38. Z. Nadasdy, H. Hirase, A. Czurko, J. Csicsvari, G. Buzsaki, Replay and time compression of recurring spike sequences in the hippocampus. *J Neurosci* **19**, 9497 (Nov 1, 1999).

39. E. C. Tolman, Cognitive maps in rats and men. *Psychol Rev* **55**, 189 (Jul, 1948).

40. T. E. J. Behrens *et al.*, What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron* **100**, 490 (Oct 24, 2018).

41. R. Kaplan, N. W. Schuck, C. F. Doeller, The Role of Mental Maps in Decision-Making. *Trends Neurosci* **40**, 256 (May, 2017).

42. D. Shohamy, N. D. Daw, Integrating memories to guide decisions. *Curr Opin Behav Sci* **5**, 85 (Oct, 2015).

43. R. J. Dolan, P. Dayan, Goals and habits in the brain. *Neuron* **80**, 312 (Oct 16, 2013).

44. M. N. Shadlen, D. Shohamy, Decision Making and Sequential Sampling from Memory. *Neuron* **90**, 927 (Jun 1, 2016).

45. A. Bakkour, A. Zylberberg, M. N. Shadlen, D. Shohamy, Value-based decisions involve sequential sampling from memory. *bioRxiv*, (2018).

46. G. E. Wimmer, C. Buechel, Reactivation of reward-related patterns from single past episodes supports memory-based decision making. *J Neurosci* **36**, 2868 (2016).

47. C. R. Brewin, J. D. Gregory, M. Lipton, N. Burgess, Intrusive images in psychological disorders: characteristics, neural mechanisms, and treatment implications. *Psychol Rev* **117**, 210 (Jan, 2010).

48. J. Suh, D. J. Foster, H. Davoudi, M. A. Wilson, S. Tonegawa, Impaired hippocampal ripple-associated replay in a mouse model of schizophrenia. *Neuron* **80**, 484 (Oct 16, 2013).

49. E. K. Braun, G. E. Wimmer, D. Shohamy, Retroactive and graded prioritization of memory by reward. *Nat Commun* **9**, 4886 (Nov 20, 2018).

50. B. D. Van Veen, W. Van Drongelen, M. Yuchtman, A. Suzuki, Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Transactions on biomedical engineering* **44**, 867 (1997).

Acknowledgments: The authors thank Zeb Kurth-Nelson for helpful discussions. This work was supported by a Wellcome Trust Investigator Award (098362/Z/12/Z) to R.J.D. Y.L. is supported by a UCL Graduate Research Scholarship and an Overseas Research Scholarship. The Max Planck University College London Centre is a joint initiative supported by University College London and the Max Planck Society. The Wellcome Centre for Human Neuroimaging is supported by core funding from the Wellcome Trust (203147/Z/16/Z).

Author contributions: G.E.W., Y.L., and N.V. designed the experiment. G.E.W. and N.V. collected the data. G.E.W and Y.L. wrote the analysis code, analyzed, and interpreted the data. T.E.J.B. and R.D. contributed to data interpretation. G.E.W. wrote the paper with input from Y.L., N.V., T.E.J.B., and R.J.D.

Competing interests: Authors declare no competing interests.

Methods

Twenty-eight healthy volunteers participated and completed both sessions of the experiment. Participants were recruited from the UCL Institute of Cognitive Neuroscience Subject Database. Data from three participants were excluded due to poor memory performance (described below) leaving data from 25 participants for analyses (14 female; mean age 24 (range 18-32). Participants were required to meet the following criteria: age between 18-35, fluent English speaker, normal or corrected-to-normal vision, without current neurological or psychiatric disorders, no non-removable metal, and no participation in an MRI scan in the two days preceding the MEG session. The study was approved by the University College London Research Ethics Committee (Approval ID Number: 9929/002). All participants provided written informed consent before the experiment. Participants were paid for their time, for their memory performance (up to £10 based on percent correct performance above chance), and a bonus for localizer phase target detection performance (up to £2).

Participants were excluded from analysis if two of the following three criteria were met: (1) accuracy below 50 % on the cued retrieval task on the second day, (2) accuracy below 50 % in the episode component reordering task on the second day, and (3) indication on the post-experiment questionnaire that the participant had mentally reordered the episodes from their original day 1 order. In the current sample, no participants were excluded based on MEG decoding performance, specifically, the classification of the 6 categories in the MEG localizer phase data.

Experimental Task

We designed our memory experiment to investigate the neural processes supporting retrieval of episodic experiences where the original episodes were only experienced once, similar to many experiences outside the lab; this is in explicit contrast to paradigms with many repetitions of the same (sequence of) stimuli. Retrieval was also separated from encoding by approximately 24 hours, again to increase ecological validity. To allow for many unique episodes with multiple while at the same time

maintaining reasonable decoding performance, episodes were designed such that they were made up of elements from 6 different categories. On the first day, participants experienced 8 different temporally extended episodes with one exposure per episode in a testing room (**Fig. 1a**). Episodes were composed of 5 discrete picture elements and an accompanying story written in first-person perspective. On the following day, participants returned for the MEG scanning session where they completed a cued retrieval phase and a category localizer phase during the acquisition of MEG data (**Fig. 1b**). Behavioral piloting in a separate sample of participants was used to optimize the design and ensure that memory retrieval performance on day 2 was both reliably above chance but below ceiling in the majority of participants.

Episodic encoding session procedure

On the first day, participants completed the episodic encoding phase. This phase presented eight episodes each composed of five unique sequential picture components. Episode components were accompanied with a text segment of a story to encourage the maintenance of the true episode order in memory. The story was written in first-person perspective to better align with the perspective of veridical personal episodic memories. The first four elements of each episode were taken from 6 potential categories of stimuli: faces, buildings, body parts, objects, animals, and cars. The final element in each episode was not taken from these categories; instead, it represented a unique ending element. Participants were instructed to try to remember the order of the episodes and further instructed that a performance bonus would be tied to their performance on questions testing memory for the sequential order of the episode elements. A practice episode was presented first after which the participants were asked to type in the name of the 1st stimulus element presented in the episode, then the 2nd, 3rd, 4th, and 5th elements.

In each episode, participants were presented with the initial picture element along with the piece of story text shown below (**Fig. 1a**; **Table S1**). A grey screen background was used for all experimental phases. The stimulus faded in for 0.5 sec and was then presented with the story text for 2 sec. The text then disappeared and for the

remaining 2.5 sec, participants performed a target detection task, pressing the '1' key whenever they saw a small grey square appear at any location over the stimulus (mean of 1 target per stimulus). The stimulus then faded out for 0.5 sec. Total stimulus duration including fade-in and fade-out was ~ 5.5 sec. A grey 'bokeh' image faded in as the stimulus faded out. After the stimulus disappeared, participants responded with the 'up arrow' key to a series of 1-3 arrow indicators ('^ ^ ^') in order to progress to the next element of the episode. If participants did not respond to an arrow within 6 sec, a warning appeared instructing the participant to respond faster. The mean inter-stimulus interval was 6.5 sec (1 sec for short duration episodes; 12 sec for long duration episodes). For the final component of the episode, a white square initially occluded the stimulus. Participants pressed the 'space' key to reveal the stimulus and story text. After the final component of the episode, a delay of 2 sec was followed by the text "Positive ending: you won +£1.00!" or "Negative ending: you lost - £0.50!" depending on whether the story ended in a positive or negative way. Participants were then presented with a probe requiring them to type in the name of a particular episode element (selected pseudo-randomly from elements 1-4). A 30 sec rest period followed each episode. After the completion of the 8 episodes, participants were instructed not to explicitly rehearse the episodes or to record the episodes in any way.

Episodes were constructed from a pseudo-random combination of category elements in addition to a final component that was not taken from these categories. The stimuli consisted of 40 photographs taken from the internet and previous studies from our group in the following categories: human faces (6), buildings (6), body parts (5), objects (5), animals (5), automobiles (5), and eight final component pictures (4 negative and 4 positive). Brief story text connected the sequence of stimuli into a short story (**Table S1**). As noted above, half of the episodes were of a longer duration, achieved via manipulating the inter-stimulus-interval (1 sec or 12 sec). The story in half of the episodes ended in a positive element and half ended in a negative element (**Table S1**). The ordering of long versus short and positive versus non-positive episodes was pseudo-randomized in two counterbalance orders.

After a 5 min break to decrease the potential influence of temporal proximity on performance for the last episodes, participants completed a short cued retrieval phase

that tested recall of the order of the elements presented in each episode. The memory test was brief to minimize additional exposure to the episode stimuli. Following a practice trial (using stimuli from the practice episode), participants completed 8 trials in the "after" condition and then 8 trials in the "before" condition. Each mini-block of 8 trials was preceded by text indicating the current condition. Participants were shown a picture cue and instructed to retrieve the associated episode in order to make a response about the sequential order of the subsequent answer stimulus. In the after condition, participants attempted to remember what came after (at any point) the cue in the same episode (**Fig. 1b**). For the example in **Fig. 1**, if the participant was cued with the bear and shown an answer stimulus of the SUV or sunny blue sky, the answer would be correct. If the answer was the man's back, the harmonica, or a stimulus from any other episode, the answer would be incorrect. Answers were 'correct' for any position after the cue, not just immediately after. In the before condition, participants attempted to remember what came before (at any point) the cue in the same episode. In both conditions, when the answer picture was presented, participants were shown the response options "Correct" and "Incorrect" in text below the picture. Cues in this memory test were only taken from the second state 2 (of 5 total episode states) in the after condition or the fourth state in the before condition. The answer on half of the trials was correct.

On each cued retrieval trial, the cue picture was presented in full opacity for 0.5 sec and then faded to 0 % opacity across the remaining 5 sec of the retrieval period (**Fig. 1**). Then the answer picture was presented. The answer text indicated the mapping between key responses and answers, e.g. "Correct (1)" and "Incorrect (2)"; the left and right text locations were randomly selected on each trial. There was no time constraint on the answer period. After the answer was recorded, following a brief 0.1 sec pause, a 2-level confidence scale ("High" and "Low") was presented, with the left and right location of options randomized. After a 0.1 sec pause, a fixation period of mean 1.5 sec followed (randomly sampled from the values [1.0, 1.5, 2.0]).

MEG session procedure

Participants returned for the MEG scan on the following day. After initial setup in the MEG room, participants were reminded of the instructions for the cued memory phase and completed 4 practice questions (based on the practice episode from the previous day). During scanning, the memory response period was time-constrained. This limit was added to encourage participants to retrieve as much information from memory as possible during the cue period, to facilitate later MEG analysis of neural processes underlying successful retrieval. Participants were instructed to try to retrieve the episodes as well as possible during the presentation of the cue picture and that in this way, they could respond faster (and avoid missed responses) when the answer appeared. Participants were also reminded of the performance bonus based on memory accuracy.

As described above for the memory test on the first day, on each cued retrieval trial, the cue picture was presented in full opacity for 0.5 sec and then faded to 0 % opacity across the remaining 5 sec of the retrieval period (**Fig. 1**). The gradual fade of the cue across the retrieval period was designed to avoid any sharp stimulus offset effects which could negatively affect MEG decoding. Then the answer stimulus was displayed. The text indicating the key response, e.g. “Correct (1)” and “Incorrect (2)”, was randomly presented on the left and right of the screen. If a response was not made in this time period, the warning “Please try to respond more quickly!” was presented for 2 sec. The answer picture was presented for 1-3 sec with the duration based on the recent rate of missed trials in the past 20 trials. If participants made no response on more than 14 % of recent trials, the answer period was increased in duration by 0.25 sec (with a ceiling of 3 sec). If participants made no response on less than 5 % of recent trials, the answer period was decreased in duration by 0.25 sec (with a floor of 1 sec). After the answer period, following a brief 0.1 sec pause, a 2-level confidence scale (“High” and “Low”) was presented, with the left and right location of options randomized. If a response was not made in time, the warning “Please try to respond more quickly!” was presented for 2 sec. After a 0.1 sec pause, a fixation period of mean 1.5 sec followed (randomly sampled from the values [1.0, 1.5, 2.0]).

In each of 5 blocks in the cued retrieval phase, the trials in the after and before conditions were separated into mini-blocks of 10-12 trials. Each mini-block was

preceded by an instruction screen: “Next: What picture came after (before)?” along with the instruction to press the ‘1’ key to continue. At the mid-point of each block, participants were given a 30 sec pause, followed by a reminder of the current condition and an instruction to press the ‘1’ key to continue. Each of the five blocks of cued retrieval included 43 trials and lasted for approximately 8 minutes. Brief rest breaks were inserted between blocks. In the cued retrieval phase, we collected ~ 27 trials per episode and ~ 43 trials per state (episode positions 1 to 5) for a total of 215 trials. All trials with a cue from state 1 were after condition trials. All trials with a cue from state 5 were before condition trials. Trials with a cue from state 3 were composed of equal numbers of after and before condition trials, while trials with a cue from state 2 and state 4 were a weighted mixture of after and before condition trials. The presented answer was correct on ~ 39 % of trials. On ~ 9 % of trials, a ‘lure’ answer was presented that was from the same episode but in the incorrect direction as the current condition. For example, in an after condition trial where the cue was from state 3, a picture from that episode in state 1 was presented as the answer. Trials were presented in a pseudo-random order with the constraint that no episode was queried on sequential trials.

The cued retrieval phase was followed by a functional localizer to derive participant-specific sensor patterns that discriminated each of the 6 categories that made up the episodes by repeatedly presenting each of the 32 unique stimuli. The localizer design followed a design used previously (Kurth-Nelson et al. 2016; Liu et al. 2019). Participants were instructed to read a word shown on the screen, pay attention to the picture that followed, and respond if any grey square targets appeared superimposed over the picture. The instructions were followed by 4 practice trials.

In a localizer trial, participants were presented with a brief name corresponding to one of the pictures, presented in text on the center of the screen for 2 sec. Participants were instructed to imagine the corresponding picture. The text then disappeared and the named picture appeared on the screen for 0.75 sec. During picture presentation, participants performed a target detection task, responding with a ‘1’ button press if the picture contained a small grey square. Targets were rare events, appearing on 15.4 % of trials. A mean 0.75 sec fixation ITI followed (range 0.25 – 1.25) during which responses were still recorded. If performance on the target detection task fell below

70 % correct (across missed responses and false alarms), a warning was presented: “Please improve your detection of the grey squares!” Finally, as in the cued retrieval phase, a mid-block rest of 30 sec was inserted during each block. After each localizer block, participants were shown yellow ‘stars’ on the screen, ranging from 0-4, depending on their target detection accuracy in the preceding block.

The stimulus pictures were presented in a pseudo-random order, with the constraint that no category repeat in subsequent trials. Each picture from a given category was presented an equivalent number of times, with 78 repetitions per picture category. The localizer was presented in 5 blocks, with 94 trials in the first four blocks and 92 trials in the last block for a total of 468 trials.

Following scanning, the participants completed a post-experiment questionnaire that assessed memory strategy and potential mental reordering of the episodes, and also asked participants to try to write down a brief version of each story. The re-ordering question asked “Did you change the order of the stories to make your own story order? 1= never, 5=always”. Participants who responded with a 4 or 5 were considered for exclusion, in conjunction with performance on the memory and sequence memory test. We observed a negative correlation in the full group (prior to exclusions) between response to this question and memory performance in the MEG session.

Finally, participants completed a computerized sequence memory test where they attempted to place the stimuli from a given episode in the correct order. In this phase, the stimuli from an episode were presented in a random order on the left side of the computer screen. Participants then moved each stimulus from the left side (starting from the top) into one of 5 empty boxes spread from the left to the right across the screen. Stimuli were moved using the left and right arrow keys; the space bar confirmed placement. Accuracy was measured as the mean rate of correct replacement across each location across all episodes.

MEG acquisition

The participants were scanned while sitting upright inside an MEG scanner located at the Wellcome Centre for Human Neuroimaging at UCL. A whole-head axial gradiometer

MEG system (CTF Omega, VSM MedTech) recorded data continuously at 600 samples per second, utilizing 273 channels (2 original channels of the 275 channels are not operational). Three head position indicator coils were used to locate the position of participant's head in the three-dimensional space with respect to the MEG sensor array. They were placed on the three fiducial points: the nasion and left and right pre-auricular areas. The coils generate a small magnetic field which is used to localize the head and enable continuous movement tracking. We also used an Eyelink eye-tracking system to monitor participant's eye movements and blinks. The task was projected onto a screen suspended in front of the participants. The participants responded during the task using a 4-button response pad to provide their answers (Current Designs), responding with self-selected digits to the first and second buttons.

MEG Pre-processing

MEG data were processed using MATLAB packages SPM12 (Wellcome Trust Centre for Neuroimaging) and FieldTrip. The CTF data were imported using OSL (the OHBA Software Library, from OHBA Analysis Group, OHBA, Oxford, UK) and down-sampled from 600 Hz to 100 Hz (yielding 10 ms per sample) for improved signal to noise ratio and to conserve processing time. Slow drift was removed by applying a first order IIR high-pass filter at 0.5 Hz.

Preprocessing was conducted separately for each block. An initial preprocessing step in OSL identified potential bad channels whose characteristics fell outside the normal distribution of values for all sensors. Then independent component analysis (FastICA, http://research.ics.aalto.fi/ica/fastica) was used to decompose the sensor data for each session into 150 temporally independent components and associated sensor topographies. Artifact components were classified by automated inspection of the combined spatial topography, time course, kurtosis of the time course, and frequency spectrum for all components. For example, eye-blink artifacts exhibited high kurtosis (>20), a repeated pattern in the time course and consistent spatial topographies. Mains interference had extremely low kurtosis and a frequency spectrum dominated by 50 Hz line noise. Artifacts were then rejected by subtracting them out of the data. All

subsequent analyses were performed directly on the filtered, cleaned MEG signal, in units of femtotesla.

In the cued retrieval blocks, an 8.5 second epoch was extracted for potential analysis for each trial, encompassing 500 ms preceding cue onset and continuing past the answer response. In the analyses below, we analyzed the first two-thirds of the cued retrieval period. Given the speeded response demands to the answer stimulus, the end of the period would involve increasing response preparation that may decrease the ability to detect sequenceness events. We also excluded the initial 160 ms following cue presentation to allow time for early stimulus processing. Thus, our retrieval period analysis window focused on 160 - 3667 ms of the full 5500 ms period. In the localizer blocks, a 4.5 second epoch was extracted for potential analysis for each trial, encompassing 500 ms preceding text onset through the end of the picture presentation period. In both the retrieval and localizer blocks, preceding the analysis steps below, we further excluded time periods within individual channels that exhibited extreme outlier events (determined by values > 7x the mean absolute deviation).

MEG data decoding and cross-validation

Lasso-regularized logistic regression models were trained for each category. Methods followed previous studies (Kurth-Nelson et al. 2016 ; Liu et al. 2019). Only the sensors that were not rejected across all scanning sessions in the preprocessing step were used to train the decoding models. A trained model k consisted of a single vector with length of good sensors n consisting of 1 slope coefficient for each of the sensors together with an intercept coefficient. Decoding models were trained on MEG data elicited by direct presentations of the visual stimuli.

For each category we trained one binomial classifier. Positive examples for the classifier were trials on which that category was presented. Negative examples consisted of two kinds of data: trials when another category was presented, and data from the fixation period before the text pre-cue appeared. The null data were included to reduce the correlation between different classifiers – enabling all classifiers to report low probabilities simultaneously. Prediction accuracy was estimated by treating the highest

probability output among all classifiers as the predicted category. Sensor distributions of beta estimates are shown in **Fig. S2** and prediction performance of classifiers trained on 200 ms on left-out trials in functional localizer task are shown in **Fig. S3**.

Sequenceness measure

The decoding models described above allowed us to measure spontaneous reactivation of task-related representations during memory retrieval. We next defined a 'sequenceness' measure, which describes the degree to which these representations were reactivated in a well-defined sequential order (Kurth-Nelson et al. 2016; Liu et al. 2019). Here we utilized an updated general linear model approach (Liu et al. 2019). This analysis approach is illustrated in **Fig. S2**.

First, we applied each of the six category decoding models to the cued retrieval period MEG data. This yielded six timeseries of reactivation probabilities for each trial, each with length N, where N is the number of time samples included in the retrieval period analysis window. Below, we use the term "stimulus" for simplicity to refer to the category-level information.

We then used a linear model to ask whether particular sequences of stimulus activations appeared above chance in these timeseries. For each stimulus *i*, at each possible time lag Δt, we estimated a separate linear model:

$$\mathrm{Y}_i = \mathrm{X}(\Delta t) * \beta_i(\Delta t)$$

The predictors $\mathrm{X}(\Delta t)$ were time-lagged copies of the six reactivation timeseries. The model predicted Y_i, the reactivation of stimulus *i*. The linear model had N rows, with each row a time sample. We estimated $\beta_i(\Delta t)$, a vector of coefficients that described the degree to which stimulus *i*'s reactivation was predicted by activation of each other stimulus at time lag Δt. By repeating this procedure for each stimulus *i*, we obtained $\beta_i(\Delta t)$, a 6x6 matrix that can be viewed as an empirical transition matrix between the six stimuli (categories) at lag Δt.

Specifically:

$$\mathrm{Y}_i = \sum_{j=1}^{s} \mathrm{X}_j(\Delta t)\beta_{ij}(\Delta t)$$

Where $\mathrm{X}_j(\Delta t)$ are time-lagged copies of Y_j, *s* is the number of states, and therefore:

$$\mathrm{Y}_i(t) = \sum_{j=1}^{s} \mathrm{Y}_j(t-\Delta t)\beta_{ij}(\Delta t)$$

The matrix $\beta_i(\Delta t)$ is obtained by solving the following set of equations for each stimulus *i*, up to state *s*.

$$\mathrm{Y}_{i=1}(t) = \sum_{j=1}^{s} \mathrm{Y}_j(t-\Delta t)\beta_{ij}(\Delta t)$$

$$\mathrm{Y}_{i=2}(t) = \sum_{j=1}^{s} \mathrm{Y}_j(t-\Delta t)\beta_{ij}(\Delta t)$$

$$\mathrm{Y}_{i=s}(t) = \sum_{j=1}^{s} \mathrm{Y}_j(t-\Delta t)\beta_{ij}(\Delta t)$$

We next asked whether the $\beta_i(\Delta t)$ was consistent with a specified 6x6 transition matrix by taking the Frobenius inner product between these two matrices (the sum of element-wise products of the two matrices). This resulted in a single number $Z_{\Delta t}$, which pertained to lag Δt. Finally, differential forward – backward sequenceness was defined as $Z_{f\Delta t} - Z_{b\Delta t}$. In our initial analyses and individual differences analyses, we used the difference between correlations in the forward ($Z_{f\Delta t}$) and backward ($Z_{b\Delta t}$) direction in order to remove common autocorrelation which would otherwise add significant variance. In the analyses testing for a relationship between sequenceness and trial-by-trial accuracy, we entered the separate forward ($Z_{f\Delta t}$) and backward ($Z_{b\Delta t}$) sequenceness measures into the regression analyses. As our analysis was on trial-based data and not rest, we did not need to control for alpha rhythm (Liu et al. 2019).

The transition matrix was defined as the stimulus (category) order in each episode. Our primary results focus on comparisons of sequenceness on correct versus incorrect retrieval trials; as such, we do not conduct comparisons to a null value. Here, as category orders were pseudo-randomly shuffled across episodes, we did not conduct permutation tests. To ensure that the results were not overfit to the regularization parameter of the logistic regression, all results were obtained with the lasso regularization parameter that yielded the strongest mean decoding in the localizer (l1 = 0.002). The decoding models used to evaluate sequenceness were trained on functional localizer data taken from 200 ms following stimulus onset. The 200 ms time point exhibited the strongest decoding accuracy during the localizer; notably, this time point of category decoding was also consistent with the individual stimulus decoding findings of Kurth-Nelson et al. (2016) and Liu et al. (2019).

Identifying Replay Onsets

Replay onsets were defined as moments when a strong reactivation of a stimulus was followed by a strong reactivation of the next (or preceding) stimulus in the sequence from an episode (Liu et al. 2019). Specifically, we first found the stimulus-to-stimulus time lag Δt at which there was maximum evidence for sequenceness (as described above), time shifted the reactivation matrix *X* up to this time lag Δt, obtaining $\mathrm{X}(\Delta t)$. We then multiplied *X* by the transition matrix *P*, corresponding to the unscrambled sequences: $X \times P$. Next, we elementwise multiplied $\mathrm{X}(\Delta t)$ by $X \times P$. The resulting matrix had a column for each stimulus, and a row for each time point in the cue period for each trial. We then summed over columns to obtain a long vector *R*, with each element indicating the strength of replay at a given moment in time. Finally, we thresholded *R* at its 95th percentile to only include high-magnitude putative replay onset events. We also imposed a constraint that a replay onset event had 100 ms of preceding replay-free time.

Specifically:

$$Proj = X(\Delta t)$$

Matrix *Proj* is obtained by time shifting the reactivation matrix *X* to time lag Δt.

$$Orig = X \times P$$

Matrix *Orig* is obtained by matrix multiplication between reactivation matrix *X* and transition matrix *P*.

$$R_t = \sum_i^s Orig_{ti} * Proj_{ti}$$

Vector *R* is obtained by elementwise multiplication between matrix *Orig* and *Proj*, and then summing over columns.

Cue-triggered reactivation analyses

In the cued retrieval period, we tested for cue-triggered reactivation of episode elements. This analysis compared evidence for categories present in a cued episode versus categories not present in a cued episode. The analysis utilized the raw classifier evidence vectors (*n* categories by *t* trial timepoints) to investigate differential activity near the peak stimulus response at ~ 200 ms. For each episode, the within-episode categories that were not presented as a cue were averaged to derive a measure of reactivation of within-episode elements. In the after condition, there were 3 within-episode categories; in before condition, trials where the cue came from state 5 had 4 categories entered into the within-episode analysis. The 2 categories that were not members of the cued episode were averaged to derive a measure of other-episode reactivation. The timepoints showing the strongest difference between these two measures were averaged for each trial to derive trial-by-trial reactivation measures representing relative within- versus other-element activity. These values were subsequently entered into multilevel regression analyses. We examined a relationship between the trial-by-trial reactivation measure and mean episode accuracy: the average performance across trials for the episode cued on a given trial. We also examined the relationship to trial-by-trial accuracy, but this analysis was under-powered in the very

high performing participants. The reactivation analyses collapsed across the after and before conditions.

Time-frequency analyses

A frequency decomposition (wavelet transformation) was computed for the memory retrieval period in every trial. We examined whether the log power for correct versus incorrect trials was related to accuracy. The analysis focused on alpha (8-12 Hz), as previous research has found alpha decreases predictive of successful memory retrieval (*14, 31*). Statistical comparison focused on the first two-thirds of the retrieval period, excluding the first 600 ms after the cue to allow for an initial cue-evoked alpha peak to return approximately to baseline. Results are displayed in **Fig. S8**.

Zero-lag correlation analysis

In a supplemental analysis, we examined the relationship between reactivation of within-episode elements compared to other-episode elements with a zero time lag. This measure was a basic correlation between the time series of category evidence: the average of 3 correlations for the within-episode elements and 2 correlations for the other-episode elements. We did not find a greater correlation between within-episode elements than between other-episode elements. Through thresholding of the category evidence time series, we found that correlations were driven by increases in evidence and that these increases were brief (**Fig. S9**). However, we found no relationship between the correlation of within-episode elements across the retrieval period and behavior (**Fig. S9**).

Multilevel modelling

We conducted all pre-processing of behavioral and MEG data for multilevel modelling in Matlab. Multilevel models were implemented in R, following previous procedures (*49*). We used a multi-level logistic regression model (glmer, in the lmer4 package) to predict

correct memory responses. A correct response in the cued retrieval phase was an answer stimulus correctly identified as coming after the cue in a given episode, an answer stimulus correctly rejected as coming after the cue in a given episode, etc. All missed response trials (where no response was recorded within the response time window) were excluded from analysis.

In the main sequenceness analyses, we fit separate intercept, forward sequenceness, and backward sequenceness effects for each participant. In the model, we also included control variables representing performance in neighboring trials. These variables were included because we found that performance 1 and 2 trials in the past and performance 1 and 2 trials in the future was positively related to current trial performance, an effect similar to what we have observed in previous memory studies. In analyses of continuous variables such as mean correct performance for the episode cued on the current trial, we used multi-level regression (lmer).

For all models, to ensure convergence, models were run using the bobyqa optimizer set to 10^6 iterations. We estimated confidence intervals using the confint.merMod function and p-values using the bootMer function (both from the lmer4 package) using 2500 iterations. All reported p-values are two-tailed.

MEG Source Reconstruction

All source reconstruction was performed in SPM12 and FieldTrip utilizing OAT. Forward models were generated on the basis of a single shell using superposition of basis functions that approximately corresponded to the plane tangential to the MEG sensor array.

Linearly constrained minimum variance beamforming (*50*) was used to reconstruct the epoched MEG data to a grid in MNI space, sampled with a grid step of 5 mm. The sensor covariance matrix for beamforming was estimated using data in broadband power across all frequencies. The baseline activity was the mean power activity averaged over -100 ms to -50 ms relative to replay onset. All non-artifactual trials were baseline corrected at source level. We looked at the main effect of the

initialization of replay. This analysis was conducted separately to investigate backward replay events in the after condition and forward replay events in the before condition.

The statistical significance of clusters identified in the beamforming analysis was calculated using SPM12. An initial cluster-forming threshold of $p < 0.001$ was applied and regions exceeding $p < 0.05$ whole-brain family-wise-error corrected (FWE) at the cluster level are reported. The timepoint preceding replay onset (- 10 ms) was additionally investigated to explore whether individual differences in memory performance related to differential MTL power preceding replay onset.

Individual differences

We tested for a relationship between MEG measures of sequenceness and mean memory performance in the after and before conditions. For sequenceness, we used differential (forward-backward) sequenceness given the strong decaying autocorrelation evident in the raw forward and backward sequenceness estimates (Kurth-Nelson et al. 2016; Liu et al. 2019). In a supplemental analysis, we estimated the relationship between replay and memory performance using a regression, separately entering forward and backward sequenceness as predictor variables. These analyses used Pearson correlations, reporting two-tailed p-values. A statistical comparison of the correlations between of sequenceness and behavior in the after condition and the before condition was conducted using a test for the difference between two dependent correlations. This test is conservative, as the performance measures in the after condition and the before condition were not identical, while the test assumes full dependence.

Data and materials availability

Complete behavioral data will be publicly available on the Open Science Framework. The full MEG dataset will be publicly available on openneuro.org.

Supplementary Materials for

Episodic memory retrieval is supported by rapid replay of episode content

G. Elliott Wimmer*, Yunzhe Liu*, Neža Vehar, Timothy E.J. Behrens, Raymond J. Dolan

This file includes:

Supplementary Results

Figs. S1 to S9

Tables S1-S3

Supplemental Results

Replay and individual differences in memory performance

The primary analysis of the relationship between sequenceness and individual differences in memory performance utilized the differential sequenceness measure (fwd – bkw sequenceness; **Fig. 2c-d**). This measure provides a summary of the overall evidence for sequenceness and finds that the same sequenceness direction is important as the trial-by-trial analysis of accuracy. However, given the specific relationship between backwards versus forwards sequenceness in the trial-by-trial analysis of accuracy, we verified that the individual difference relationship was also selective. In the after condition, we found that reverse sequenceness was negatively related to average performance (fwd $t_{(23)}$ = 2.265, p = 0.0337; bkw $t_{(23)}$ = -2.9111, p = 0.0081). In the before condition, we found that forward sequenceness was related to average performance (fwd $t_{(23)}$ = -2.2419, p = 0.0354; bkw $t_{(23)}$ = 1.0456, p = 0.3071). Results from both the after and before conditions show stronger sequenceness in lower-performing participants (reverse sequenceness in the after condition; forward sequenceness in the before condition). These analyses give qualitatively similar results as those reported in the main analysis which used differential forward-backward sequenceness.

Supplementary Figures and Tables

Figure S1. Memory performance as a function of episode length and whether the episode ended in a positive or negative element and performance on final episode re-ordering test. (**a** and **b**) Memory did not differ in the after condition by length ($t_{(24)}$ = -1.389; p = 0.178) or the before condition by length ($t_{(24)}$ = 0.661; p = 0.515). (**c** and **d**) Memory did not differ in the after condition by end valence ($t_{(24)}$ = -0.068; p = 0.946) or the before condition by end valence ($t_{(24)}$ = 0.1478; p=0.88). Given the null behavioral differences, primary MEG analysis collapsed across these variables. (**e**) Performance on the post-scan episode sequence memory re-ordering test (n=24 participants with sequence test data). Individual scores were the average of accurate placements of each element within each episode. (Error bars represent SEM.)

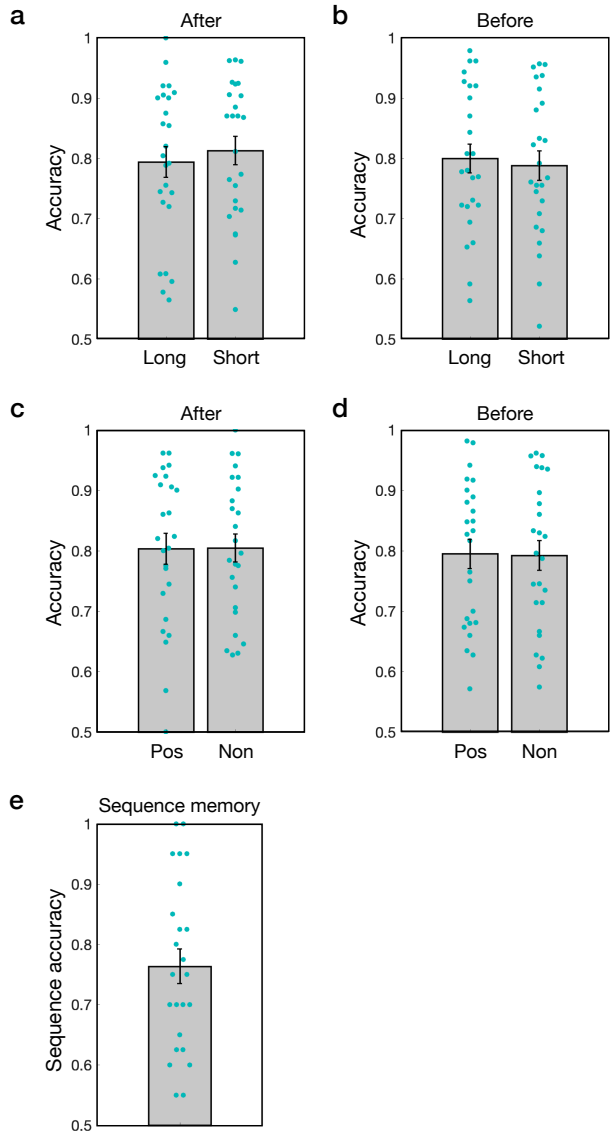

Figure S2. Sequenceness analysis schematic and classifier sensor weighting. (**a**) Classifiers were trained on the 6 categories that made up the episodes. The mean weighting (approximate importance) of each sensor for a given category, minus the mean across all other categories, for illustration only. (Anterior = top; posterior = bottom.) (**b**). Mean sensor weighting across all categories. (**c**) Illustration of how the trained classifiers are applied to the MEG data timeseries for each cued retrieval period, where state 1 - 4 represents episode components 1-4 from **Fig. 1a**. (**d**) The sequenceness analysis detects systematic time shifts (T) in category evidence. A forward sequence illustration is shown on the left; a backward sequence illustration is shown on the right.

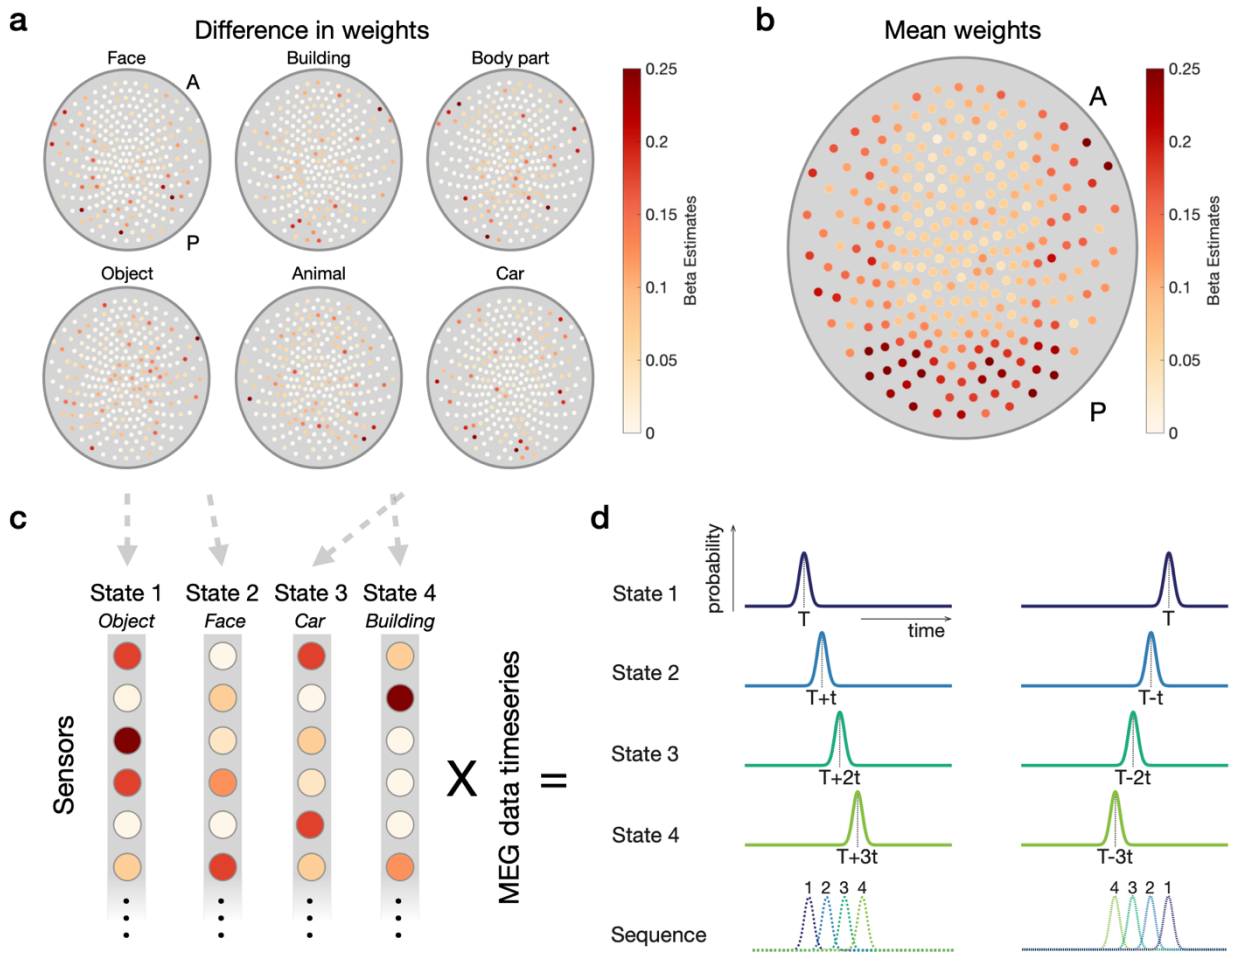

Figure S3. Illustration of localizer performance for the six stimulus categories that made up the first 4 components of episodes (face, building, body part, object, animal, and car). The peak response at approximately 200 ms in each plot represents the results for training on a given category of stimuli at the 200 ms time point and testing on the same category. Other lines indicate performance for training on the correct category and testing on alternate categories. Dashed line indicates classifier baseline performance estimated by shuffling labels. (Shaded error margins represent SEM.)

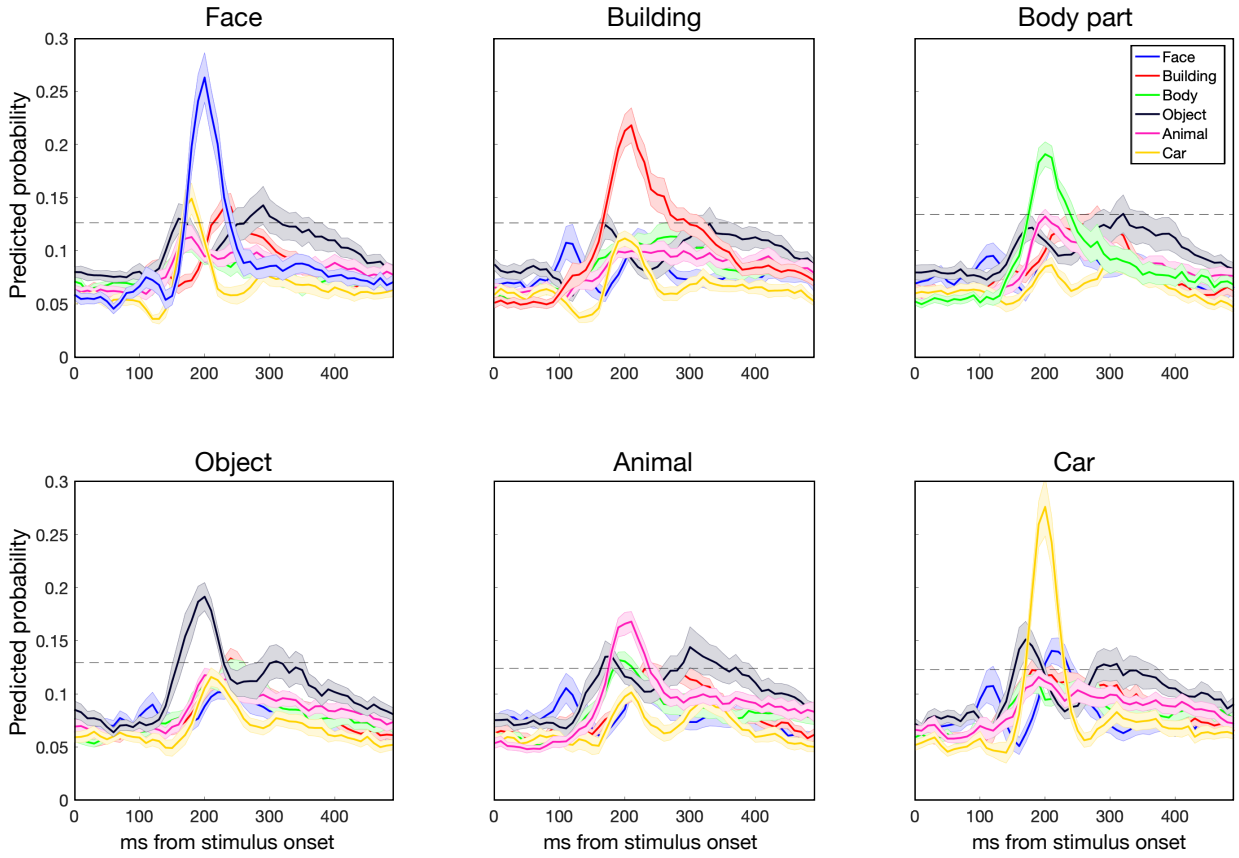

Figure S4. Differential sequenceness for the current episode and across all other episodes in the after and before conditions. (**a**) Differential sequenceness (forward - backward) in the after condition for regular performance participants as in **Fig. 2a**, here showing the subset of n = 17 participants with sufficient incorrect trials. (**b**) Differential sequenceness (forward - backward) in the before condition for regular performance participants (subset of n = 18 participants with sufficient incorrect trials in the before condition). (**c**) Differential sequenceness for all other episodes (excluding the current episode) in the after condition for all participants (n=25). (**d**) Differential sequenceness for all other episodes in the before condition for all participants (n=25). (Shaded error margins represent SEM.)

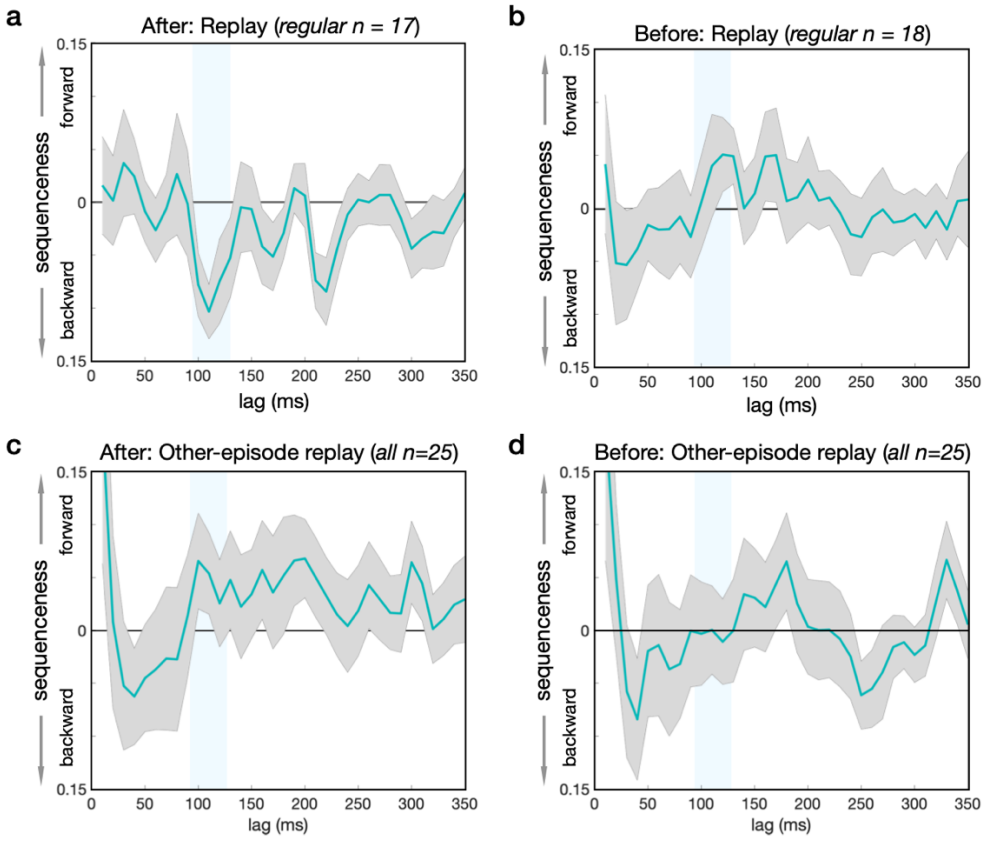

Figure S5. No significant relationship between sequenceness and trial-by-trial behavior at other time lags and in the other episode sequenceness analysis. (**a**) In the after condition, mean replay strength (forward-backward) with a 40-50 ms lag did not relate to overall mean memory performance (percentage of correct trials). As in **Fig. 1c** the data points for the regular performance participants are shown in purple; high performance participants are shown in orange. (**c**) As in panel a, here for the before condition. (**c**) In the after condition, 40-50 ms sequenceness for other episode transitions (excluding the current episode) did not relate to mean memory performance. (**d**) As in panel c, here for the before condition. (**e**) As in panel c, here for a 100-120 ms lag. (**f**) As in panel e, here for the before condition.

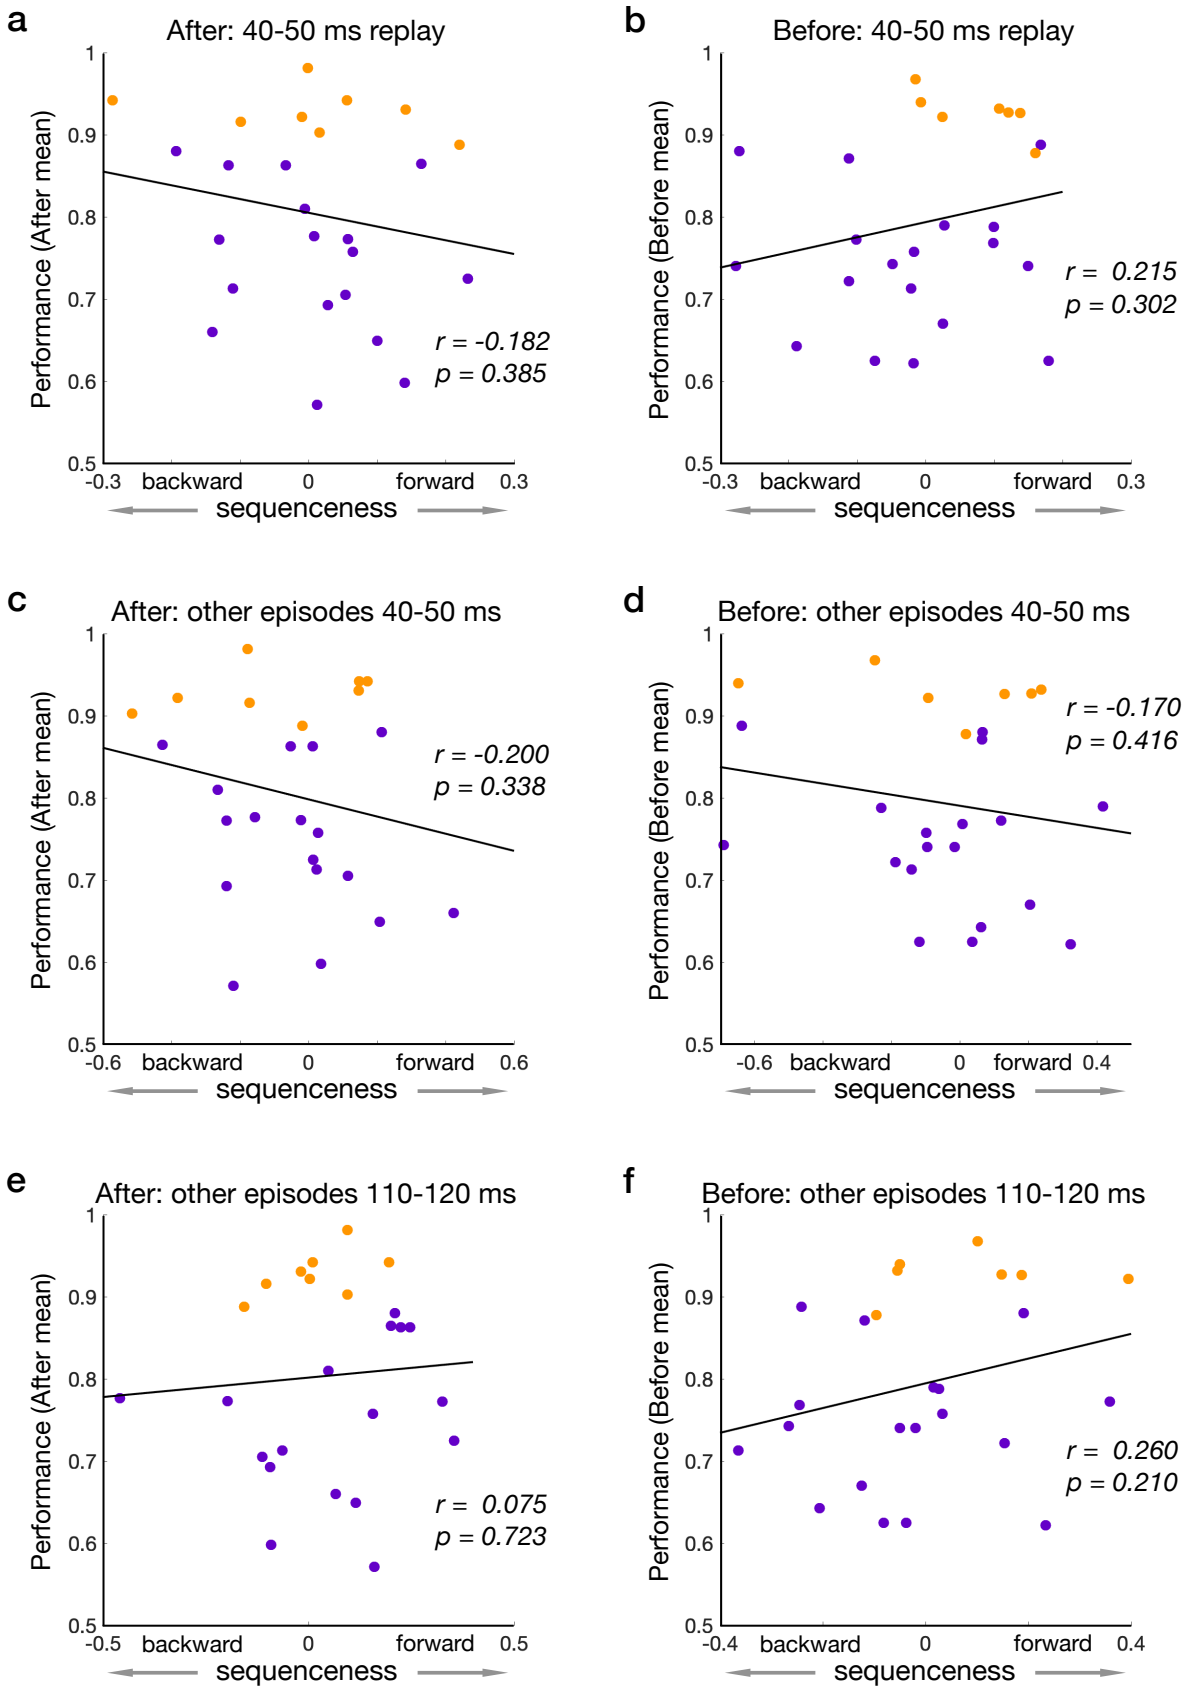

Figure S6. No significant relationship between sequenceness and behavior at other time lags and in the other episode sequenceness analysis. (**a**) In the after condition, trial-by-trial memory retrieval success did not relate to forward or backward replay. (**b**) In the after condition, 40-50 ms sequenceness for other episode transitions (excluding the current episode) did not relate to mean memory performance. (**c**) As in panel c, here showing the results for the other episode sequenceness measure for a 100-120 ms lag. (Error bars represent standard error.)

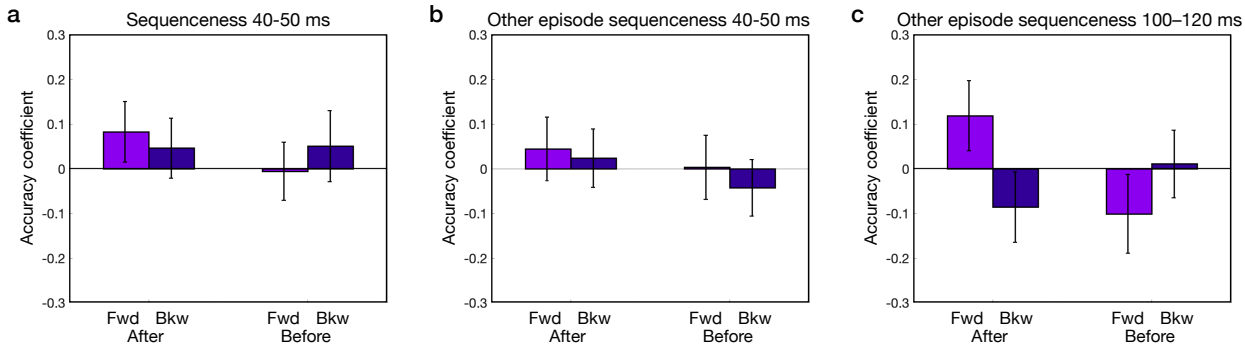

Figure S7. Additional beamforming results in the after condition. (**a**) Power in the right anterior hippocampus increased 10 ms before replay onset in participants with lower memory performance. Data are for visualization purposes only and represent the peak coordinate as in **Fig. 4b**. High performance participants in orange; regular performance participants in purple. (**b**) Power in the right visual cortex at replay onset in the after condition, displaying a different view of the whole-brain results shown in a coronal section in **Fig. 4a**. (Statistical maps thresholded at $p < 0.001$ uncorrected, for display.)

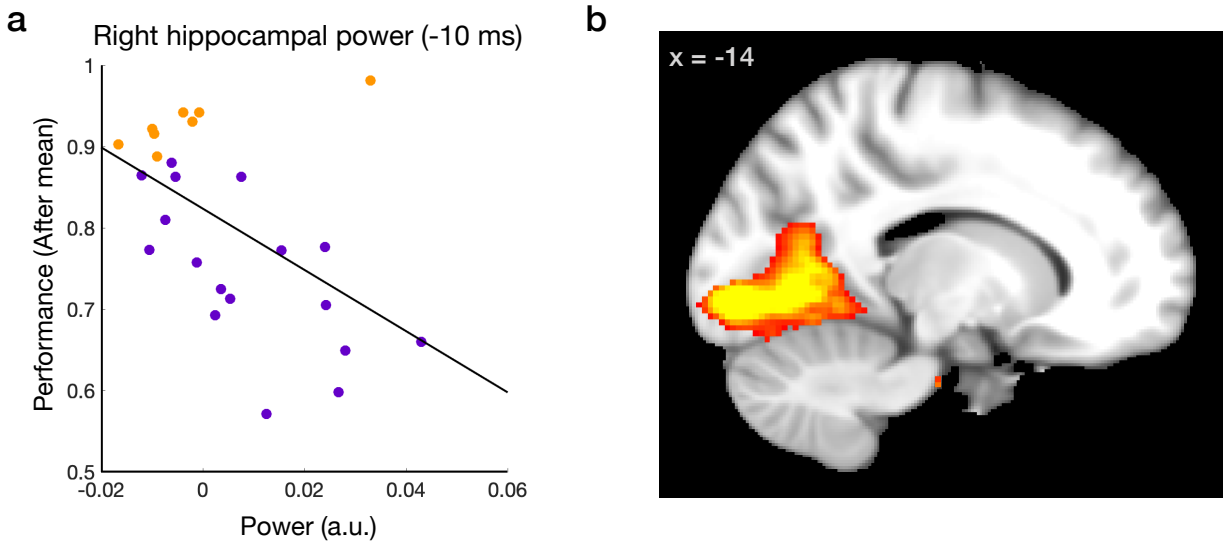

Figure S8. Time-frequency analysis of replay onsets and the relationship between memory accuracy and alpha power (8-12 Hz). (**a**) Time-frequency analysis showing power increases at replay onset in the after condition using the primary downsampled (100 Hz) data which yields frequencies up to ~ 50 Hz. 0 ms represents the onset of putative replay events. (Average across all n=25 participants in correct trials.) (**b**) Time-frequency analysis as in panel A, here in the before condition. (**c**) Time-frequency analysis of high frequencies in the after condition, showing power increases at replay onset in the after condition using data that was not downsampled (600 Hz). (**d**) Time-frequency analysis of high frequencies as in panel C, here in the before condition. (**e**) Time-frequency difference between correct and incorrect trials across frequencies up to ~50 Hz in the after and before conditions (n = 15 participants with enough incorrect trials in both the after and before conditions). (**f**) Extracted timecourse of power in alpha (8 – 12 Hz) for correct (green) and incorrect (red) trials. Bracketed line indicates the time period of interest for statistical comparison. (**g**) Difference between correct and incorrect trials from the end of the cue-induced alpha peak at 600 ms to 3670 ms ($t_{(14)}$ = 2.95, p = 0.015; error bars and error margin represent SEM). These results are in the opposite direction of previous findings which instead report alpha decreases associated with successful memory retrieval (*14, 31*). (Error bars and shaded error margins represent SEM.)

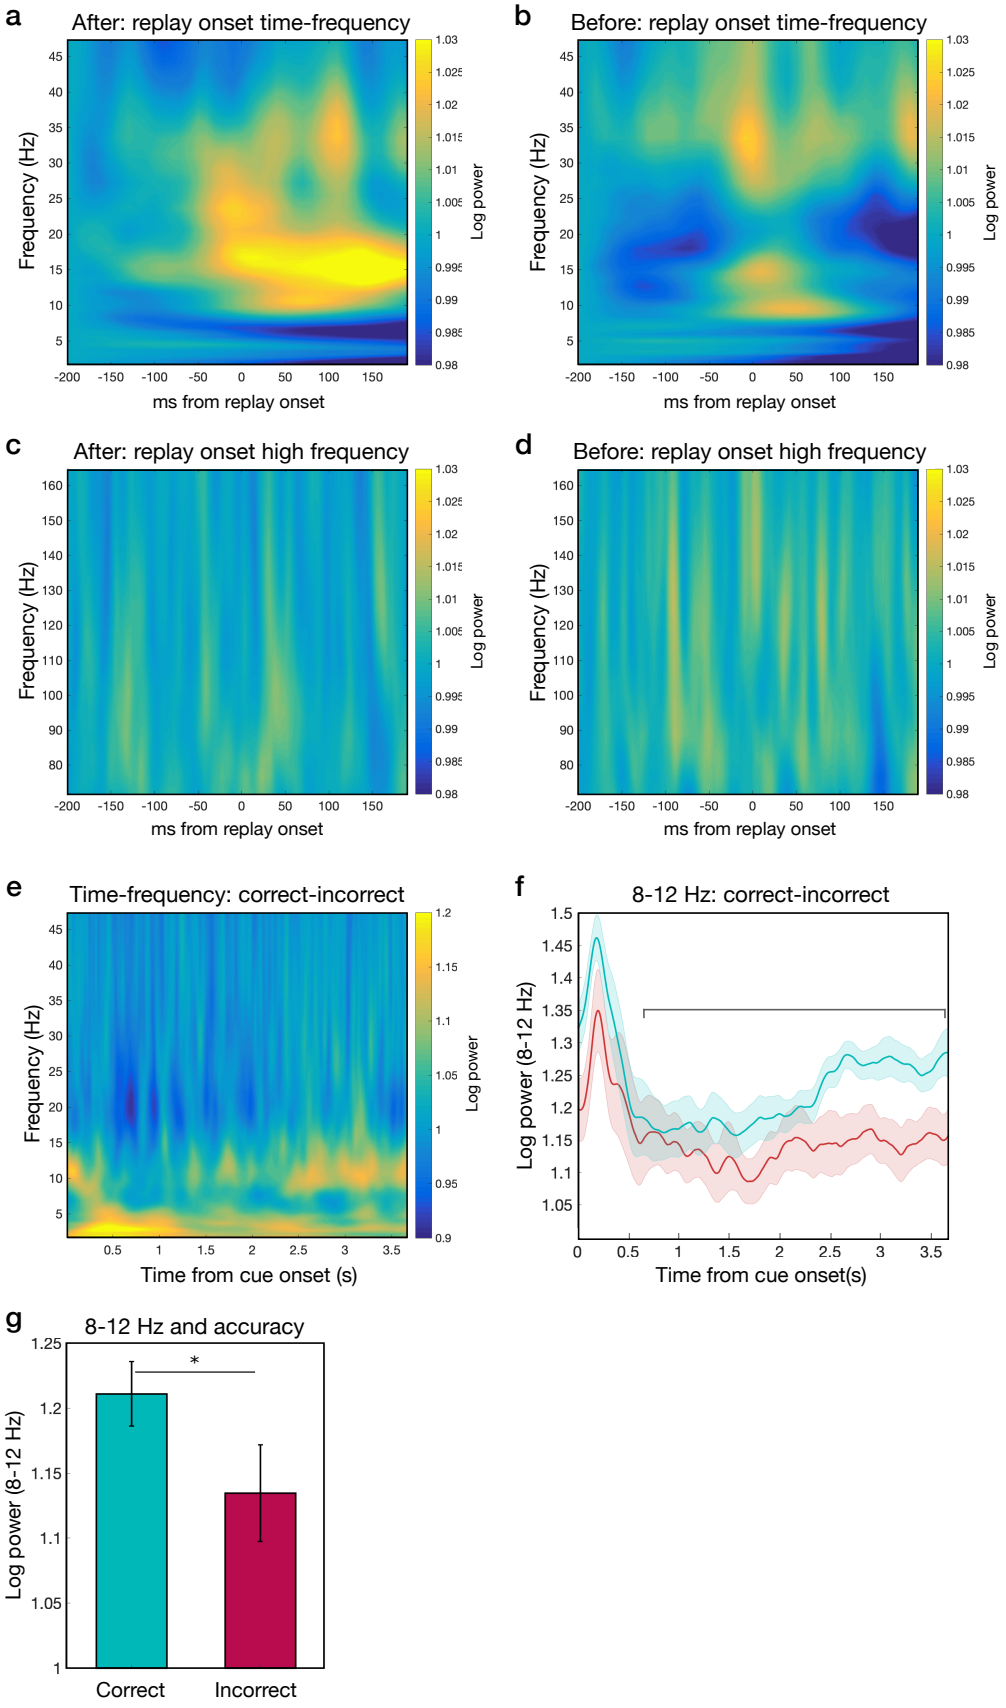

Figure S9. Relationship between accuracy, cue-evoked reactivation, and zero-lag correlation between within-episode category evidence during the retrieval period. (**a-b**) Cue-evoked reactivation of within-episode elements minus other-episode elements from 200-250 ms in the after condition (**a**) and before condition (**b**), included in the regression model with forward and backward sequenceness. The effect of cue-evoked reactivation is non-significant (after: 0.878 ± 0.682; z = 1.287, p = 0.198; before: 0.578 ± 0.714; z = 0.809, p = 0.418); the effects for sequenceness are unaffected. (**c-d**) The correlation between evidence for within-episode categories minus the correlation between all other pairings (zero-lag correlation) across the 160 ms – 3667 ms cue period of analysis (**c**) is not related to trial-to-trial accuracy in very high or regular performance participants: High performance (-0.534 ± 0.644; z = -0.829, p = 0.407); regular performance (-0.093 ± 0.354; z = -0.263, p = 0.792). (**d**) The correlation between within-episode category evidence is driven by high-magnitude events (>= 95 % of mean), and activity for these events peaks and falls rapidly. The purple line represents the mean across participants in the after condition. (Error bars represent SEM.) (**e-f**) The zero-lag correlation between evidence for within-episode categories minus the correlation between all other pairings included in the regression model with forward and backward sequenceness in the after condition (**e**) and in the before condition (**f**). The effect of clustered reactivation is non-significant (after: 0.137 ± 0.463; z = 0.296, p = 0.767; before: -0.721 ± 0.494; z = -1.460, p = 0.144); the effects for sequenceness are unaffected. (Error bars represent standard error.)

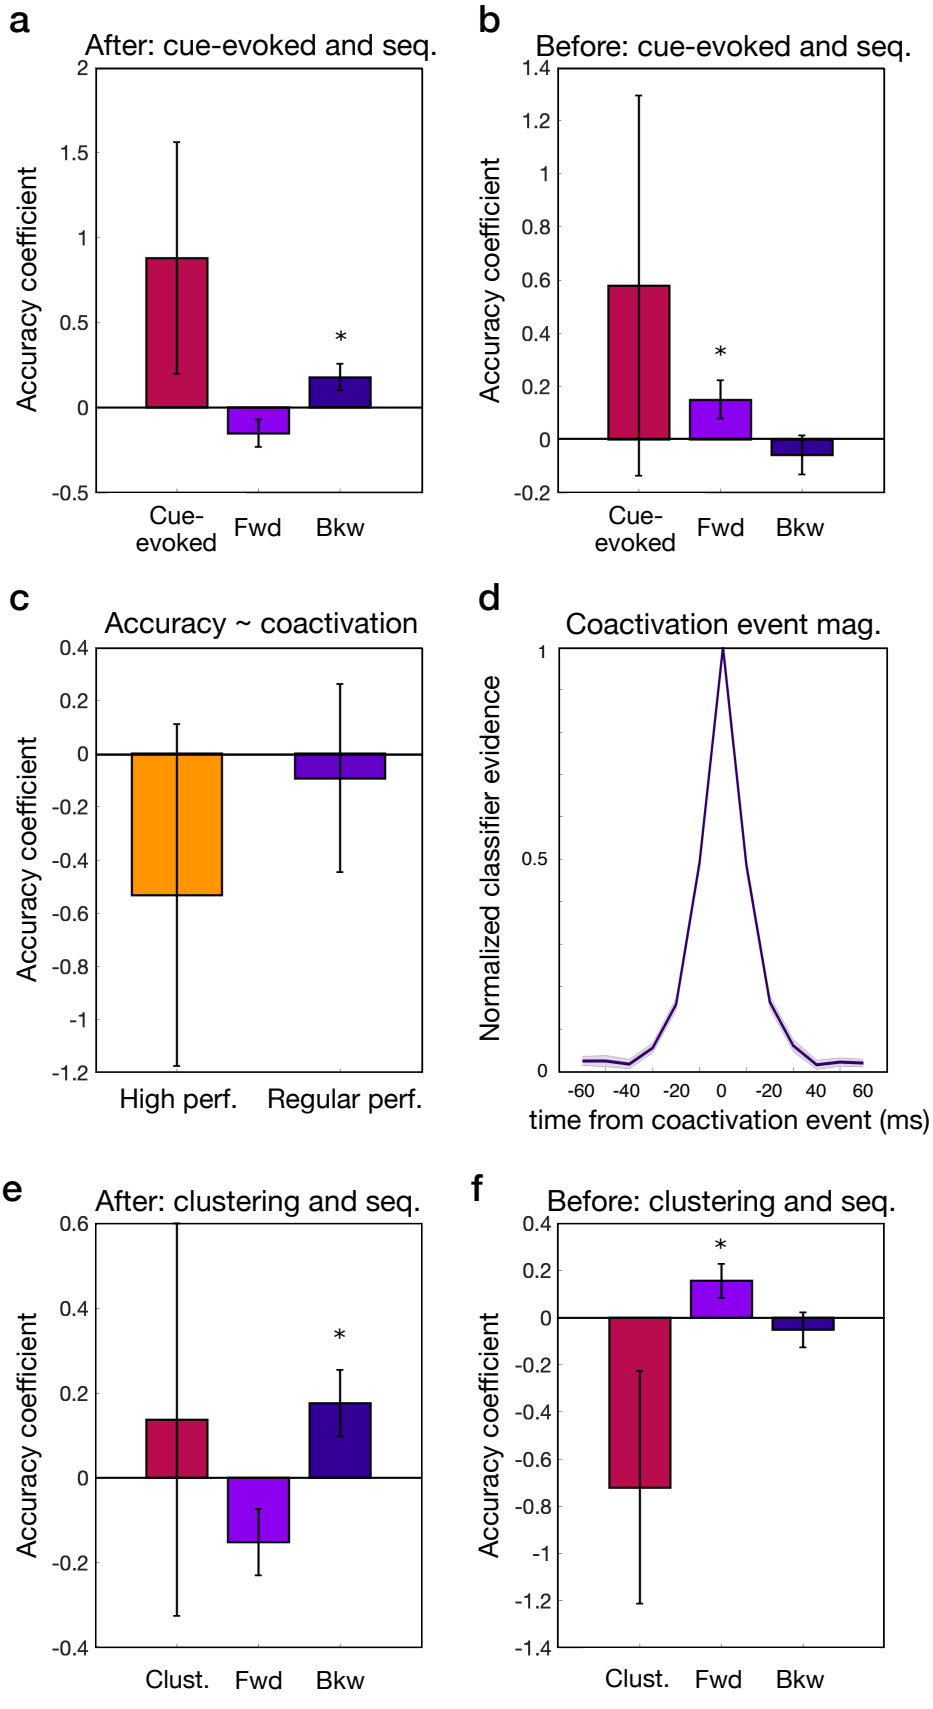

Table S1. Story text example used in the episodic memory encoding phase on the first day. The stimuli for the first 4 components were taken from the categories: face, building, body party, object, animal, and car. The alternative counterbalance order changed component 5 across episodes from positive to negative.

Episode	Component 1	Component 2	Component 3	Component 4	Component 5
1	I had a big elephant	and guided it to the barn	a freckled woman was waiting there	and cleaned it with a toothbrush	then we all had birthday cake.
2	A man facing away	played a harmonica	while we watched a bear	try to open an SUV	after which we enjoyed the sunny day.
3	I was sitting outside the stone house	trying to fix a computer mouse	when a sports car pulled up	and a young Asian man	gave me a pile of gold coins for my work.
4	I found the key I needed	to get into the warm cabin	an Asian woman	massaged my sore shoulder	and we celebrated her graduation with balloons.
5	I was playing with a girl	outside her big white mansion	when a turtle appeared	and walked over her feet	but we had to hide from the thunderstorm.
6	I called for a taxi	to give my tired knees a rest	and rode with a guy in my class	to go look at a deer	but then I had to go study for an exam.
7	I was using scissors	to trim the man's beard	then we took a mini car	to the pastel hotel	but I slipped and fell on some marbles.
8	At the greenhouse	my friend pressed her hand to the glass	then I saw a pickup truck drive by	and noticed a horse getting groomed outside	but then we had to drag out the trash.

Table 2. Multilevel modeling results for the interaction between sequenceness and episode length (long, short) or episode end valence (positive, negative).

After condition: Length				
Variable	**coef.**	**ste**	**z-stat**	**p-value**
Intercept	0.8262	0.1172	7.053	<0.0001
SeqFwd	-0.1498	0.0792	-1.892	0.052
SeqBkw	0.1877	0.0782	2.403	0.012*
Length	-0.106	0.0712	-1.278	0.1528
SeqFwd X Length	0.0065	0.0781	0.083	0.9096
SeqBkw X Length	0.0125	0.0778	0.161	0.876
Before condition: Length				
Intercept	0.9295	0.1172	7.934	<0.0001
SeqFwd	0.1484	0.0728	2.038	0.0392*
SeqBkw	-0.0628	0.0739	-0.850	0.3872
Length	0.0811	0.0711	1.142	0.2600
SeqFwd x Length	-0.0811	0.0722	-1.124	0.2744
SeqBkw X Length	0.0429	0.074	0.580	0.5536
After condition: End valence				
Intercept	0.8477	0.1202	7.054	<0.0001
SeqFwd	-0.1416	0.0785	-1.804	0.0728
SeqBkw	0.1741	0.0786	2.214	0.0304*
Reward	0.0081	0.0761	0.107	0.8968
SeqFwd X Valence	0.1131	0.0785	1.440	0.1352
SeqBkw X Valence	-0.1966	0.0792	-2.483	0.0096*
Before condition: End valence				
Intercept	0.9235	0.1167	7.912	<0.0001
SeqFwd	0.1451	0.0730	1.987	0.0520
SeqBkw	-0.0557	0.0742	-0.751	0.4352
Reward	0.0340	0.0827	0.411	0.6809
SeqFwd X Valence	-0.0006	0.0726	-0.008	0.9968
SeqBkw X Valence	0.0014	0.0740	0.018	0.9952

Table S3. Whole-brain beamforming MEG results for replay onset in the after and before conditions. Clusters significant whole-brain FWE-corrected after an initial threshold of $p < 0.001$ to provide interpretable clusters.

Contrast	Regions	Cluster size	x	y	z	Peak z stat
After backward replay onset	L Lingual Gyrus	5495	-14	-86	-4	4.69
	L Lingual Gyrus		-22	-70	2	
	L Middle Occipital Gyrus		-32	-84	6	
	R Calcarine Sulcus	1198	26	-58	20	3.9
	R Parietal Lobe		36	-60	32	
	R Calcarine Sulcus		28	-46	4	
	R Anterior Hippocampus	2151	20	-10	-18	3.72
	R Ventral Thalamus		4	-20	-6	
	R Anterior Hippocampus		20	-2	-22	
Before forward replay onset	R Midbrain	1707	2	-32	-18	3.73
	R Parahippocampal Gyrus		14	0	-34	
	R Entorhinal Cortex		14	-2	-24	
After backward -10 ms & performance	L Entorhinal Cortex	1046	-22	-8	-32	3.82
	L Entorhinal Cortex		-18	-16	-26	
	L Anterior Hippocampus		-34	-12	-24	