Title: Hippocampal ripples initiate cortical dimensionality expansion for memory retrieval

Casper Kerrén[1*], Sebastian Michelmann[2] & Christian F Doeller[1,3]

Affiliations:

[1] *Max Planck Institute for Human Cognitive and Brain Sciences; Leipzig, Germany*

[2] *Department of Psychology, New York University; New York City, New York*

[3] *Kavli Institute for Systems Neuroscience, Centre for Neural Computation, Egil and Pauline Braathen and Fred Kavli Centre for Cortical Microcircuits, Jebsen Centre for Alzheimer's Disease, NTNU Norwegian University of Science and Technology; Trondheim, Norway*

*Corresponding author. Email: kerren@cbs.mpg.de

Abstract

How are past experiences reconstructed from memory? Learning is thought to compress external inputs into low-dimensional hippocampal representations, later expanded into high-dimensional cortical activity during recall. Hippocampal ripples—brief high-frequency bursts linked to retrieval—may initiate this expansion. Analysing intracranial EEG data from patients with pharmaco-resistant epilepsy during an episodic memory task, we found that cortical dimensionality increased following ripple events during correct, but not incorrect, retrieval. This expansion correlated with faster reaction times and reinstatement of the target association. Crucially, hippocampal theta and cortical gamma phase–amplitude coupling emerged after ripples but before cortical expansion, suggesting a mechanism for ripple-driven communication. Ripple events also marked the separation of task-relevant variables in cortical state space, revealing how hippocampal output reshapes the geometry of memory representations to support successful recall.

Main Text

Episodic memory allows us to store detailed records of past experiences and consciously reconstruct those experiences at later points in time (*1*). Like any computational system though, the human brain operates with finite resources (*2*). To cope with these constraints, efficient encoding and retrieval of memories is thought to rely on compression and expansion of neural representations (*3*–*6*). During encoding, environmental information flows through cortical and subcortical regions to the hippocampus, where memories are initially stored (*7*, *8*). During retrieval, internal or external cues trigger the hippocampus to detect matches with stored traces (*9*). A partial match initiates pattern completion, leading to memory reactivation and reconstruction in cortical networks (*10*, *11*). Yet, how hippocampal pattern completion gives rise to cortical reinstatement remains poorly understood.

Recent studies have begun to conceptualise this process as shifts in the geometric relationships of points in neural state space, a framework that has found broader applications in studies of decision-making and working memory (*12*–*19*). We hypothesise that memory retrieval might likewise involve a transformation in dimensionality, such that low-dimensional hippocampal representations are expanded into a higher-dimensional cortical state, allowing mnemonic information to be decoded for successful recall (*5*, *10*, *20*, *21*). Preliminary evidence supports this hypothesis, showing a shift from semantic to perceptual representations along the ventral visual stream during retrieval (*22*–*26*). However, whether and how the hippocampus might initiate this dimensionality expansion is unclear.

One potential means to drive cortical representational expansion is through hippocampal ripples. Ripples are known to coordinate the transfer of compressed representations and changes of brain-wide functional connectivity during offline periods in rodents (*27*–*31*). In humans, an increase of hippocampal ripple rates precedes episodic memory recall (*32*–*37*) and neocortical reinstatement of previously encoded memories consistently follows hippocampal ripple events (*32*–*35*). In parallel, theoretical accounts and empirical evidence have highlighted cross-frequency interactions—particularly theta-gamma-phase–amplitude coupling (TG-PAC) —as a mechanism for coordinating long-range communication (*38*–*44*). While TG-PAC has been widely studied during mnemonic processing particularly within hippocampus (*45*–*53*), its role during retrieval, and specifically in mediating ripple-initiated cortical transformations, remains largely unexplored.

We hypothesise that the reconstruction of information in cortical regions is supported by a ripple-based mechanism, where compressed representations are expanded via PAC-based

connectivity between hippocampus and cortex. To test this hypothesis, we analysed intracranial data from 12 patients with pharmaco-resistant epilepsy as they performed an associative memory task (Fig. 1, A and B, fig. S1). We asked whether ripple events could trigger both an increase in cortical representational dimensionality and TG-PAC between hippocampus and cortex—linking local hippocampal dynamics to global cortical transformations during successful episodic retrieval (Fig. 1C).

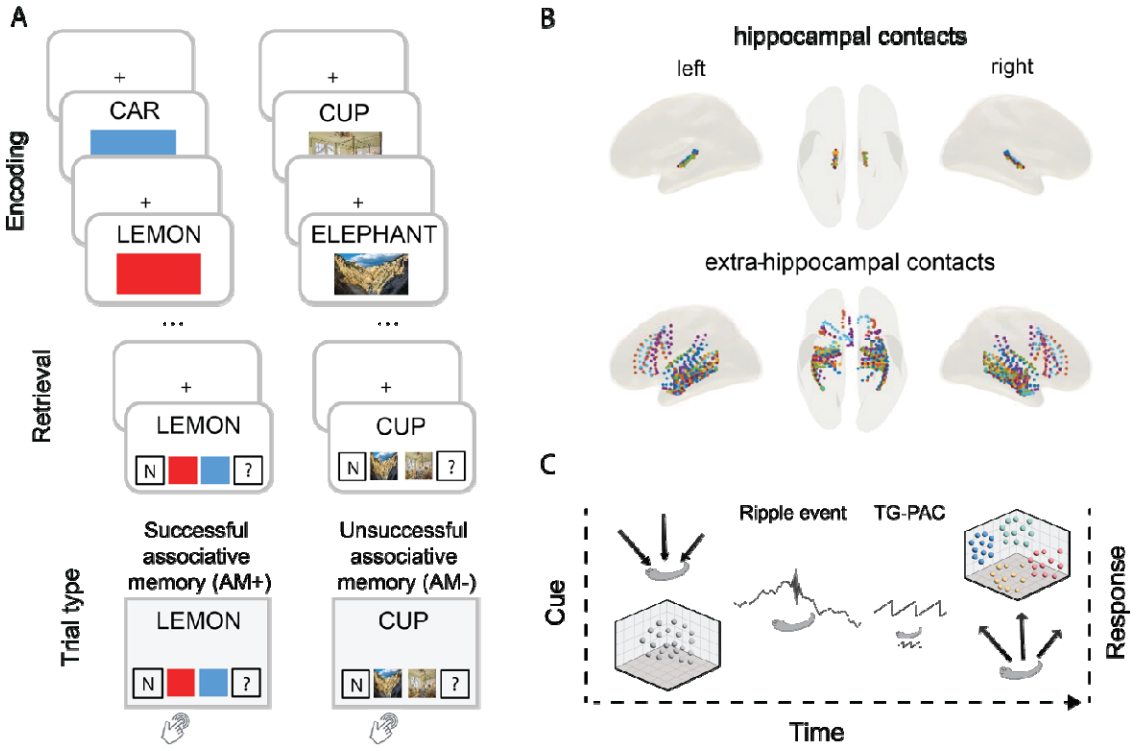

Fig. 1. Paradigm, implantation scheme, hypothesis

(A) During encoding, participants associated a noun either with one of two colours (*left*) or one of two scenes (*right*; alternating across runs). During retrieval, participants were prompted with the noun alongside response options, yielding either successful associative memory (AM+) or unsuccessful associative memory (AM-; combining incorrect and 'don't know' responses). **(B)** Hippocampal (*top*) and extra-hippocampal (*bottom*) iEEG contacts included in the analyses, highlighted in different colours for each participant. **(C)** Schematic depiction of the hypothesis: During encoding, the representation of an event is stored in hippocampus as a lower-dimensional pointer to the higher-dimensional representation in cortical areas. During retrieval, hippocampal ripples initiate a process that initiates a hippocampal-cortical connectivity through theta-gamma phase-amplitude coupling (TG-PAC), leading to the low-dimensional representation being expanded in cortical areas, which is linked to behavioural memory performance.

Behavioural results

The experiment yielded a balanced amount of successful and unsuccessful associative memory trials (AM+: M = 49.95%, SE = 5.12%; AM-: M = 50.05%, SE = 5.12%; t(1,11) = -.01,

p = .992). Response latencies for AM+ trials were faster (M = 1.90 sec, SE = .12 sec) than for AM- trials (M = 2.06 sec, SE = .12 sec; t(1,11) = -2.68, p = .021).

Greater ripple-density for successful retrieval

On average, 787.33 hippocampal ripples (SE = 145.97) per participant were detected during all retrieval trials, after excluding false positives, with a spectral mean peak at 89.17 Hz (SE = .57 Hz) (Fig. 2A; fig. S1 for participant-specific ripple plots). Previous research has shown that ripple density increases prior to successful free recall performance (e.g., (*32*)). To assess whether this finding extends to associative recognition as employed here, we extracted the ripple density for each trial, i.e., the number of ripples normalised by the trial's reaction time, and averaged across trials in each condition. Ripple density was higher during AM+ (M = .24 Hz, SE = .03 Hz) compared to AM- (M = .19 Hz, SE = .02 Hz), t(11) = 3.58, p < .01 (Fig. 2B).

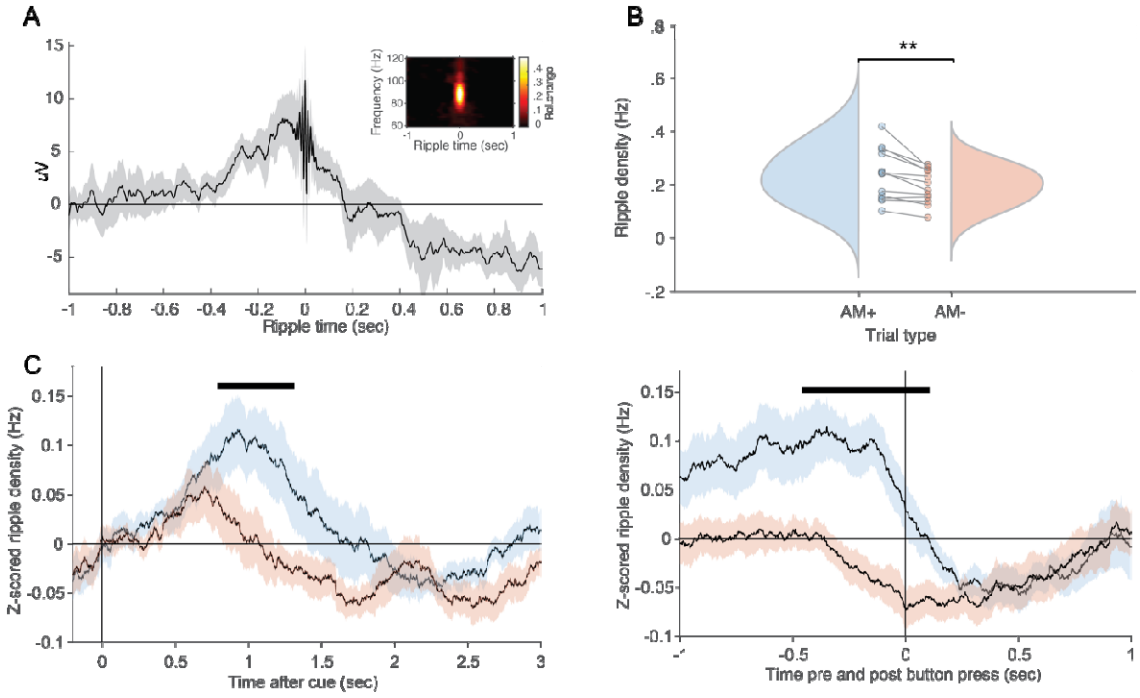

Fig. 2. Hippocampal ripple density increases during successful memory retrieval

(**A**) Grand average hippocampal ripple for all participants, aligned to ripple peak at time 0, with spectral grand average as inset, showing a peak at 89 Hz. (**B**) Ripple density significantly differed between AM+ and AM-. Circles represent participants. (**C**) Ripple density dynamics across time. *Left*: aligned to retrieval cue onset. *Right*: aligned to response time. Horizontal black lines denote significant differences between AM+ and AM- (p < .05, corrected for multiple comparisons across time).

Additionally, we examined whether the increased ripple rate occurred at specific time points during the successful retrieval and observed that ripple density was increased 800-1400ms post cue (for cue-aligned data; Fig. 2C, left) and leading up to the retrieval response (for response aligned data; Fig. 2C, right). These findings are consistent with previously reported latencies of episodic memory processes (*54*).

Neocortical reinstatement is linked to hippocampal ripple events

Building on these findings, we next examined the relationship between hippocampal ripple events and the reinstatement of information from encoding. Specifically, we assessed whether cortical patterns present during encoding re-emerged around ripple events during retrieval.

Our results confirm that ripple-locked memory reinstatement underlies successful associative recognition memory (AM+). The successful decoding of memory content was closely yoked to ripple peaks in our data and continued until the end of the analysis epoch (see figure S2 for various control analysis confirming that reinstatement was specific to the ripple latency of the current trial). This extends previous findings of the role of hippocampal ripples in the replay of encoding activity during free recall of verbal material (*34*) and suggests a generalised role of ripples in the reinstatement of information-rich cortical representations.

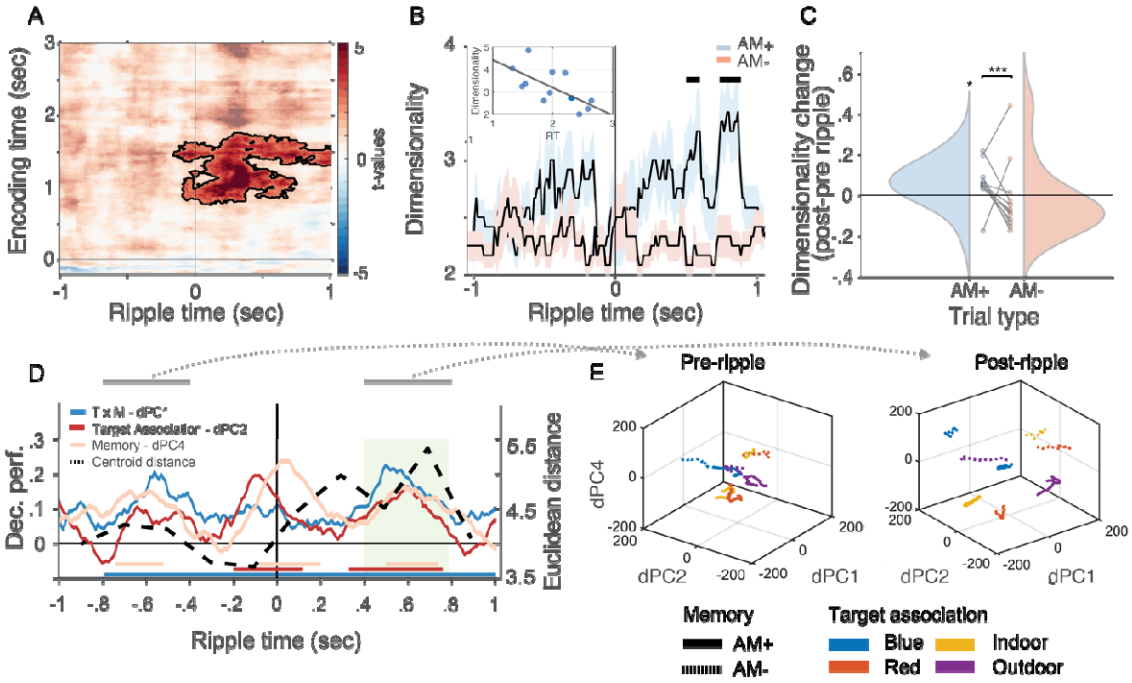

Fig. 3. Target memory decoding and dimensionality transformation are locked to ripple events.

(A) Stimulus-locked encoding x ripple-locked retrieval target decoding accuracies. Results show significant increase of cortical target decodability for successful vs. unsuccessful retrieval following hippocampal ripple events (p < .05, corrected for multiple comparisons). (**B**) Estimating dimensionality for cortical contacts yielded significantly higher dimensionality for AM+ (blue) versus AM- (red) following ripple events (time 0 on x-axis). Inset: The increase in dimensionality was significantly negatively correlated with reaction time. (**C**) Linear mixed-effects modelling (Epoch Half: pre and post ripple, Trial Type: AM+ and AM-) of cortical dimensionality revealed a significant interaction, indicating a greater increase in dimensionality post-ripple for AM+. Asterisks denote significance compared to 0 for AM+ and the interaction of pre and post ripple time-window and AM+ and AM-. **(D)** dPC1, 2, and 4 revealed significant decoding accuracy of target identity in the time window of dimensionality expansion (highlighted in green bar). Significant intervals of decodability are highlighted with horizontal lines. Euclidean distance between clusters increased around ripple onset with a peak 700ms post ripple-events (dashed line). **(E)** Plotting dPCs 1, 2 and 4, showed overlapping representations of experimental variables before ripples (left), whereas their separability increased after ripples (right).

Hippocampal ripple-induced dimensionality expansion increases the separability of cortical representations

These findings suggest that hippocampal ripples initiate a cortical "decompression" of memory representations, allowing latent information to be reconstructed during retrieval (*5*, *10*, *20*, *21*) (Fig. 1C). To formally test this, we first examined changes in representational geometry using principal component analysis (PCA), applied to iEEG data time-locked to ripple events.

We found that cortical dimensionality significantly increased following ripples during AM+ trials, compared to AM− trials, with effects emerging 470–840 ms post-ripple (two clusters; both p<.05). Dimensionality estimates were robust across parameters (fig. S3, A and B) and strongly correlated with the original signal ($r_{spearman}$ .87, p < .001; fig. S3C). Linear mixed-effects modelling confirmed higher dimensionality for AM+ trials overall (Estimate = .09, SE = .03, t(19,988) = 2.82, p = .005), and for post- vs. pre-ripple epochs (Estimate = .14, SE = .03, t(19,988) = 4.39, p < .001), with a significant interaction indicating greater ripple-induced expansion for remembered events (Estimate = -.09, SE = .02, t(19,988) = -4.16, p < .001; Fig. 3C). Paired samples t-tests supported this effect, showing increased dimensionality following ripples in AM+ trials (before: M = 2.43, SE = .05; after: M = 2.50, SE = 0.06; t(11) = 2.22, p = .049) , but no change in AM− trials (before: M = 2.44, SE = 0.06; after: M = 2.41, SE = 0.06;), t(11) = -0.45, p = .66).

Crucially, higher post-ripple dimensionality during AM+ trials was associated with faster response times ($r_{spearman}$ = -.669, p = .02; Inset Fig. 3B), and stronger cortical reinstatement ($r_{spearman}$ = .656, p = .02), linking representational expansion to both behavioural and neural markers of successful retrieval. Control analyses confirmed that these effects were specific to ripple events and not observed in surrogate data (fig. S4), reinforcing the functional role of ripple-triggered dimensionality expansion.

While PCA preserves the geometry of neural activity while maximising variance, and LDA enhances interpretability by maximising class separation, each has limitations. PCA mixes task-related variance and does not isolate the contribution of experimental variables, whereas LDA distorts the original geometry of the neural state space in service of classification. Demixed PCA (dPCA; (*55*)), addresses both limitations by simultaneously reducing population activity and disentangling it into components aligned with specific task parameters. As illustrated in Kobak et al. (*55*), dPCA separates latent structure in a way that retains both class-specific information and the original representational geometry, making it particularly well-suited for characterising the organisation and temporal evolution of memory-related cortical states. Using this approach, we could examine how ripple events affected the organisation of cortical states, not just their dimensionality (*55*).

We reconstructed neural activity using components reflecting target association (blue, red, indoor, outdoor), memory performance (AM+ and AM-), their interaction (target association x memory performance), and a condition-independent component reflecting time. According to our hypothesis, cortical representations of these variables would separate during the period of dimensionality expansion (~400–800 ms post-ripple; highlighted in green in Figure 3D), thereby enhancing decodability in high-dimensional space.

We first confirmed that our reconstruction accurately captured the data. Specifically, the first 50 dPCs accounted for as much variance as our estimated total signal (see Supplementary Materials, and figure S5A). To assess whether these demixed components carried task-relevant information, we applied cross-validated decoding, which tests whether the structure revealed by dPCA generalises to unseen data, providing a direct readout of when and how experimental variables are represented in neural population activity (*55*).

Among all components, the interaction between target association and memory performance (dPC1) explained most variance (fig. S5B), capturing retrieval-related dynamics. This component showed high decoding accuracy across almost the entire epoch and peaked shortly after ripple events. The second component (dPC2), which captured target

association-specific brain dynamics, showed significant decoding accuracy both around ripple events and during dimensionality expansion (Fig. 4A and fig. S5D). Finally, the fourth component (dPC4), linked to memory performance, significantly distinguished AM+ from AM- trials, with effects emerging shortly after ripple events and during dimensionality expansion (Fig. 4A and fig. S5F). For full output, see figure S5.

These findings confirm that dPCA effectively disentangled experimental variables in a lower-dimensional space. To quantify this representational reorganisation, we applied k-means clustering to the dPCA components and assessed the Euclidean distance between clusters. To evaluate clustering quality, we used the Silhouette score, where higher scores indicate stronger clustering. Using 200 ms time bins around ripple events, we observed increased cluster separation beginning at ripple onset and peaking ~700 ms later (Fig. 3D, dashed line). To further quantify this effect, we plotted the state-space organisation of these components during dimensionality expansion (400-800ms post-ripple events) and compared it to corresponding pre-ripple period (400-800ms pre-ripple). Before ripple events, we identified six clusters (silhouette score = .83), whereas eight clusters emerged after ripple events (silhouette score = .90) (Fig. 3E). This was accompanied by a significant increase in Euclidean distance between clusters (pre-ripple: $M = 4.06$, $SE = .044$; post-ripple: $M = 4.78$, $SE = .048$; $t(319) = 5.04$, $p < .01$). A similar pattern emerged when using dPC 1, 3, and 4 (fig. S5H, and Supplementary materials). In both visualisations, we observed a clear separation of all experimental variables following ripple events, a pattern that was less distinct before ripple onset.

Theta-gamma phase-amplitude coupling coordinates information flow between hippocampus and cortex

Our results demonstrate that hippocampal ripples initiate a cortical expansion during episodic memory retrieval in which the separation between experimental variables increases. But what mechanisms enable this transformation of information? A candidate mechanism is the coupling between hippocampal theta and cortical gamma rhythms (*38*, *39*, *44*, *56–58*). Despite theoretical accounts emphasising its importance, there is currently little direct evidence in humans that theta–gamma-phase-amplitude-coupling (TG-PAC) supports long-range communication during episodic memory retrieval.

One likely reason for this gap is the considerable variability in retrieval timing across trials, which poses a challenge for aligning neural events consistently. Our findings suggest that hippocampal ripples can be leveraged to pinpoint the exact moments of reinstatement. Thus, we hypothesised that communication between the hippocampus and cortex would occur

shortly after ripple onset, but prior to the cortical expansion of memory representations, and would be critical for reconstructing the memory trace. To test this, we examined TG-PAC in the time window surrounding ripple events.

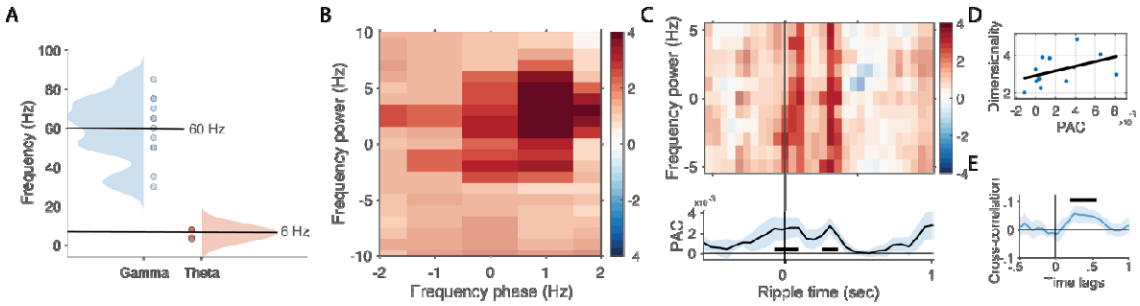

Fig. 4. Phase-amplitude coupling following ripple events are related to dimensionality expansion.

(A) On average gamma oscillations peaked at 60 Hz, whereas theta on average peaked at 6 Hz. **(B)** Significant TG-PAC around peak frequencies for respective frequency (0 on x and y axes). **(C)**. Top: A time-resolved TG-PAC commencing around ripple events and persisted for approximately 400ms after the events. Bottom: Plotting the time-resolved TG-PAC averaged across the y-axis in top panel. **(D)** A positive correlation between TG-PAC and cortical dimensionality expansion, such that stronger TG-PAC was related increased dimensionality expansion. **(F)** A cross-correlation showed that TG-PAC leads dimensionality expansion with a peak at approximately 260ms.

We identified peak theta and gamma frequencies in hippocampal and cortical signals after removing the 1/f background using the FOOOF algorithm (*59*). Time–frequency data were then aligned to these peaks (±2 Hz for theta, ±10 Hz for gamma), and phase–amplitude coupling was computed. Empirical PAC values were compared to a null distribution generated by shuffling trials 250 times per participant and time point. This allowed us to assess TG-PAC in narrow-band activity centred on peak frequencies, time-locked to ripple events (–1 to +1 s). The average peak theta and gamma frequencies were 6.08 Hz (SE = 0.69 Hz) and 59.58 Hz (SE = 4.75 Hz), respectively (Fig. 4A). This revealed significant coupling between hippocampal theta phase and cortical gamma amplitude around the identified peaks ($p_{cluster}$<.05; Fig. 4B).

To examine the timing of TG-PAC relative to ripple events, we conducted a time-resolved analysis using sliding windows. This revealed that coupling emerged shortly after ripple onset and persisted for ~300 ms ($p_{clusters}$ < .05; Fig. 4C). Notably, TG-PAC strength in this post-ripple window correlated with ripple-related increases in cortical dimensionality (r = .59, p = .04; Fig. 4D), linking cross-frequency coupling to representational transformation during retrieval.

We hypothesised that TG-PAC would follow hippocampal ripples but precede cortical dimensionality expansion. To test this, we performed a cross-correlation analysis against a null distribution of 1000 phase-shuffled surrogates. The analysis revealed a significant positive lag, peaking at ~260 ms ($p_{cluster}$ = .02; Fig. 4E), suggesting that TG-PAC may play a mechanistic role in triggering the cortical expansion of memory representations following ripple events.

Discussion

Episodic memory is thought to optimise neural processing by compressing information during encoding and expanding it during retrieval (*3–6*). In particular, the hippocampus plays a central role in this dynamic, matching retrieval cues with stored memory traces and triggering pattern completion to reconstruct the full memory in the cortex (*9–11*). Yet, the neural mechanism through which hippocampal pattern completion gives rise to cortical memory reinstatement remains poorly understood.

We hypothesised that memory retrieval involves a transformation in neural dimensionality. Specifically, that low-dimensional representations in the hippocampus expand into high-dimensional cortical activity patterns. This shift would allow higher order areas to decode rich mnemonic content (*5*, *10*, *20*, *21*). Leveraging the spatial and temporal resolution of human intracranial EEG (iEEG) while participants performed an associative memory task, we demonstrate that that successful memory retrieval was associated with stronger reactivation of encoding-related cortical patterns (Fig. 3A). Notably, this was accompanied by an increase in neural dimensionality (Fig. 3, B–C), which was linked to both faster response times (Fig. 3B) and greater reinstatement strength—suggesting a functional role for dimensional expansion in supporting retrieval.

Our findings position hippocampal ripples as key triggers of cortical reinstatement during episodic memory retrieval. While prior work has linked ripples to hippocampal–cortical communication and memory-related activity in both humans and non-human animals (*29*, *30*, *32–35*, *60–63*), we show that ripples are temporally aligned with both increased cortical dimensionality and reinstatement of encoded information. This extends previous literature by providing direct evidence that ripples mark a transition from compressed to high-dimensional cortical states supporting the reconstruction of specific mnemonic content. Moreover, our results suggest that ripple-related reflect a broader cortical mechanism for memory retrieval.

In our study, ripple-triggered cortical reinstatement occurred reliably during associative retrieval and was tightly time-locked to the onset of ripple events (Fig. 2A, fig. S2). Ripple density was also higher during correct trials, particularly in the pre-response period, suggesting that that ripples functionally drive cortical memory reactivation. Crucially, this reinstatement was accompanied by an increase in neural dimensionality. The finding puts forward a novel account of ripple events that do not simply trigger the reactivation of static memory traces, but instead facilitate an expansion of compressed representations into complex, high-dimensional cortical states. Such expansion may reflect the recruitment of a

broader set of cortical areas and the unfolding of richer, more differentiated representational patterns across the network, which then enables flexible coding of specific mnemonic content.

Our findings are also in line with recent work on the time-scales of hippocampal–cortical interactions during memory retrieval. Michelmann and colleagues (*64*) reported that hippocampal activity precedes cortical reinstatement by ~740 ms during predictive recall under naturalistic conditions, while in another study, the same authors found that hippocampal–cortical information flow precedes cortical state transitions by 500–700 ms during memory search (*65*). In line with these findings, we observed a ~400–800 ms lag from hippocampal ripples to cortical dimensionality expansion and a maximal separability of task-related variables at 700 ms post-ripples. The duration of this effect and its extended temporal window are evidence that hippocampal-driven reactivation unfolds over many synaptic transmission points to realise the expansion of a memory-trace. I.e., hippocampal outputs may initiate distributed cortical reorganisation and enable the dynamic reconstruction of mnemonic content across time and brain regions (*66*, *67*).

To characterise the representational changes that accompany this process, we used demixed principal component analysis (dPCA). Task-relevant variables, such as memory accuracy and target identity, became more separable following ripple events (Fig. 3E), with this separability peaking precisely at the time of dimensionality expansion (Fig. 3D). The neural state space also became more complex, as indicated by an increased number of optimal clusters and greater Euclidean distances between them. These results suggest that ripples support the cortical “decompression” of memory content, allowing for finer-grained separation of mnemonic features.

In our data, the increase in representational structure did not differentiate between correctly remembered and forgotten associations (Fig. 3E, right). This suggests that the ripple-mediated reorganisation of cortical dynamics alone is not diagnostic of retrieval success and other elements can cause memory processes to fail. This mirrors findings in non-human primates, where structured neural representations were observed even in error trials during working memory tasks. For example, Rigotti and colleagues found that although dimensionality collapsed in incorrect trials, task-relevant information could still be decoded—indicating an underlying structure (*68*). Similarly, in our data, incorrect trials showed lower overall dimensionality (Fig. 3B), but retained structured representations of task variables (Fig. 3E). These findings suggest that absolute separability is not the sole determinant of retrieval

success; rather, it is the richness or expressiveness of the neural state—as assessed via PCA (Fig. 3B)—that better predicts retrieval speed and accuracy.

To our knowledge, there is currently no direct evidence in humans for hippocampal theta–cortical gamma phase–amplitude coupling (TG-PAC) during episodic memory retrieval. Here, we show that TG-PAC peaks shortly after ripple events (Fig. 4C) and precedes an increase in cortical dimensionality by ~260 ms (Fig. 4E), suggesting it may act as a bridge between the initiation and reconstruction of retrieved events. This extends previous findings of TG-PAC during encoding and retrieval in the hippocampus (*45*, *47*, *58*, *69*), during spatial memory retrieval (*70*), and across distributed cortical areas during autobiographical recall (*71*).

While prior studies have linked slow and fast gamma to encoding and retrieval, respectively (*44*, *72*, *73*), we found no clear frequency split—participants showed peaks in both bands. These differences may reflect region-specific gamma dynamics outside the medial temporal lobe.

Our findings extend recent work linking neural dimensionality to flexible coding in decision-making and navigation (e.g., (13)), showing that episodic retrieval also involves dynamic, high-dimensional cortical states (5). Events with greater dimensionality during encoding are more likely to be remembered (74), raising the possibility that rich encoding geometry facilitates later reconstruction. Future work should test whether dimensionality at encoding predicts reinstatement fidelity.

We propose a framework in which ripples trigger a shift from compressed hippocampal codes to expanded cortical states. This transition may be mediated by ripple-induced synchrony (75), followed by cortical desynchronisation in alpha/beta bands (76), which has been linked to improved memory fidelity and recall (77, 78). Such desynchronisation may support low-synchrony, high-dimensional representations that promote selective decoding while reducing interference (68, 79–82). Understanding how ripple-induced synchrony interacts with cortical desynchronisation will be an important direction for future research on the temporal dynamics of memory reinstatement.

Acknowledgement

The authors would like to thank Juergen Fell and Bernhard Staresina for sharing their data.

Funding

Max Planck Society (C.K. and C.F.D.)

Deutsche Forschungsgemeinschaft project 437219953 (S.M.)

Author contributions

Conceptualization: C.K.

Methodology: C.K. and S.M.

Investigation: C.K.

Visualization: C.K.

Funding acquisition: S.M. and C.F.D.

Project administration: C.F.D.

Supervision: S.M. and C.F.D.

Writing – original draft: C.K., S.M. and C.F.D.

Writing – review & editing: C.K., S.M. and C.F.D.

Competing interests

Authors declare that they have no competing interests.

Data and materials availability

The data and code that support the conclusions of this study are available upon reasonable request from C.K. (kerren@cbs.mpg.de) and are also accessible online at GitHub.

Supplementary Materials

Materials and Methods

Figs. S1 to S5

Tables S1

References (*45-46, 55, 59, 83*–94)

Supplementary Materials

Materials and Methods

Participants

A total of 15 patients took part in the study. Three were excluded from further analysis due to clinical monitoring revealing epileptogenic activity in both hippocampi, resulting in a final sample of 12 patients (6 female; mean age = 33 years, SD = 9.3) with pharmacoresistant epilepsy. All participants provided written informed consent, and the study was approved by the Ethics Committee of the Medical Faculty at the University of Bonn. Data were recorded at the Department of Epileptology, University Hospital Bonn.

Experimental procedures

Participants were seated upright in a sound-attenuated room, approximately 50 cm from a laptop screen, and engaged in an associative learning paradigm (Fig. 1A). Each experimental block consisted of an encoding phase, a 1-minute distractor phase, and a retrieval phase.

During encoding, participants were presented with a German noun paired with either a colour or a scene, depending on the run. Colour and scene runs alternated: in colour runs, the noun was paired with either a red or blue square; in scene runs, it was paired with an image depicting either an indoor (e.g. an office) or outdoor (e.g. a nature scene) environment. The task was to form an association between the word and the accompanying stimulus by vividly imagining the object described by the noun in conjunction with the colour or scene (e.g. “a red lemon” or “an elephant in the mountains”), and to rate the plausibility of the imagined scenario. Participants had up to 3 seconds to make their plausibility judgement via button press. Each trial was preceded by a jittered inter-trial interval (ITI) of 700–1300 ms (mean = 1000 ms), during which a fixation cross was displayed at the centre of the screen. Trials ended immediately upon response.

In the retrieval phase, participants were presented with 50 previously seen words randomly intermixed with 25 novel words, along with four response options. Their task was to indicate whether the word was new (‘N’ response), whether it was old but the associated target could not be recalled (‘?’ response), or whether it was old and the associated colour or scene could be correctly retrieved (in which case the appropriate response option was selected). Responses were self-paced with an upper time limit of 5 seconds. As in encoding, trials were

terminated by a button press and were preceded and followed by a jittered ITI (700–1300 ms, mean = 1000 ms) showing a fixation cross. Each run lasted approximately 9 minutes.

Implantation of depth electrodes

Intracranial electroencephalography (iEEG) data were referenced to linked mastoids and recorded from medial temporal lobe regions, including the hippocampus, as well as additional cortical areas, at a sampling rate of 1 kHz (bandpass filter: 0.01 Hz to 300 Hz) (Fig. 1B). Depth electrodes targeting the hippocampus were implanted stereotactically as part of presurgical evaluation, following two different implantation schemes: in eight participants, electrodes were placed along the longitudinal axis of the hippocampus, while in four participants, electrodes were implanted laterally via the temporal lobe. All participants were on anticonvulsive medication, with plasma levels maintained within the therapeutic range. In addition to depth electrodes, scalp electrodes were placed at positions Cz, C3, C4, and Oz according to the international 10–20 system; however, these were excluded from all subsequent analyses.

Electrode selection

Electrode contact localisation was determined using multiple complementary criteria. First, we inspected post-implantation MRI scans and identified electrode contacts located within the hippocampus. Second, pairwise channel coherence in the 4–8 Hz range was calculated during the retrieval phase, based on the assumption that contacts within the same anatomical region would exhibit high coherence (*83*, *84*). Third, event-related potentials (ERPs) were computed for each contact, with the expectation that electrodes in the same region would display similar ERP profiles.

Only hippocampal contacts from the clinically defined healthy hemisphere were included (for one participant, both hemispheres were considered healthy). Following preprocessing (see below) and the application of these localisation criteria, a total of 72 hippocampal contacts across 12 participants were identified as clean and reliably located (mean = 6.0, SD = 2.2 contacts per participant). For extra-hippocampal regions, 647 contacts were retained for analysis (mean = 53.9, SD = 21.6 per participant). See Figure 1B for a summary across participants and figure S1 for individual electrode maps.

Preprocessing

Data processing was carried out using FieldTrip (version 20230422; (*85*)) standard MATLAB functions, and custom-written MATLAB scripts. Line noise was removed using 2-Hz-wide bandstop filters centred at 50, 100, 150, and 200 Hz. Following this, the data were re-

referenced using a common trimmed average approach, implemented via MATLAB's *trimmean* function. Specifically, 20% of the highest and lowest values were trimmed to reduce the influence of outliers before computing the mean, which was then subtracted from each channel.

Next, automatic artefact rejection was performed. The data were z-transformed based on three different metrics: (1) absolute amplitude, (2) gradient (the amplitude difference between two adjacent time points), and (3) high-frequency amplitude, obtained by applying a 250 Hz high-pass filter to the original signal—this final measure was used to detect epileptogenic spikes. A time point was marked as artefactual if it exceeded a z-score of 6 on any single metric, or a z-score of 4 on a conjunction of metrics (i.e. gradient and high-frequency amplitude, or absolute amplitude and high-pass-filtered amplitude). To ensure conservative rejection, an additional 50 ms on either side of each marked segment was also excluded from all subsequent analyses.

Ripple detection

Ripple events were detected in each hippocampal channel using established algorithms (for full details, see: (*86–88*). In brief, the data were bandpass filtered between 80 and 120 Hz, and only events lasting between 38 ms and 500 ms were considered. To qualify as a ripple, the event also had to comprise at least three full cycles in the raw signal. Importantly, we only used data from the non-pathological hemispheres, identified and excluded artifacts using an appropriate algorithm, capable of detecting epileptogenic spikes, excluded false positives, and checked the time-frequency spectra of the ripples. Together, these steps reassured that we found physiological ripples, rather than pathological ones or artifacts.

Ripple density was calculated by dividing the number of detected events by the length of the corresponding trial, up to the participant's reaction time. As ripple detection could not be performed during artefactual segments, any time marked as artefact was subtracted from the trial length prior to this calculation. The resulting value reflects the frequency of ripple occurrence per trial. For all main analyses, we selected a single ripple per trial—the one with the maximum envelope, computed as the sum of the root mean square (RMS) of the ripple signal. Across trials, the median ripple occurred at 1002 ms after cue onset during retrieval.

To examine whether the increase in ripple rate was time-specific, we conducted a time-resolved ripple-rate analysis. Data were aligned to both cue onset and reaction time during retrieval. For each channel, the ripple time series was smoothed with a 400 ms moving average using MATLAB's *smoothdata* function, and z-scored across all conditions. We then

collapsed across channels and conditions and performed a cluster-based permutation test to identify significant changes in ripple rate over time.

Multivariate pattern analysis

The raw iEEG time series were epoched based on two time points: trial onset during encoding and ripple onset during retrieval. This yielded two separate datasets: the encoding-aligned data, used as the training set (from –500 ms to +3000 ms relative to cue onset), and the ripple-aligned data, used as the testing set (from –1000 ms to +1000 ms relative to ripple onset).

Encoding data were downsampled to 100 Hz, smoothed using a 200 ms moving average (using the MATLAB function *smoothdata*) and baseline-corrected using a –200 to 0 ms pre-cue interval. For ripple-aligned retrieval data, the baseline was defined using the pre-cue window from the corresponding encoding trial in which the ripple occurred.

All encoding trials were included in classifier training and split into four stimulus classes: blue vs red (for colour runs) and indoor vs outdoor (for scene runs). These classes were used to train separate classifiers. For retrieval, classification performance was assessed separately for trials with successful associative memory (AM+) and unsuccessful memory (AM–, including both incorrect and 'don't know' responses). After artefact rejection, the average number of trials per participant was similar across conditions (AM+: mean = 58.67, SD = 32.05; AM–: mean = 57, SD = 31.87; $t(11) = 0.11$, $p = .91$). 'Don't know' trials were included for two reasons: (1) to better match the number of trials across conditions, particularly for later dimensionality analyses, and (2) because failure to recall is behaviourally equivalent to an incorrect memory judgement.

Prior to classification, both training and testing datasets were z-scored independently. A linear discriminant analysis (LDA), implemented via the MVPA-Light toolbox (*89*), was used to train and test the classifier at each time point, yielding a time-generalisation matrix (TGM). As the training and testing sets were drawn from independent datasets (encoding vs retrieval), no cross-validation was performed. Statistical comparisons between conditions were assessed using cluster-based permutation tests.

To confirm that reinstatement was specifically linked to the timing of ripple events, we ran two control analyses involving trial-wise ripple time shuffling. First, ripple times were randomly reassigned across AM+ trials 250 times per participant (RandCorrect). Second, ripple times were circularly shifted to the next trial using MATLAB's *circshift* function

(CircCorrect), providing a stricter temporal control. For both controls, the significant temporal cluster identified in the empirical data was applied to the shuffled data, and average reinstatement scores were computed for each participant. Paired-samples t-tests showed that the empirical data exhibited significantly stronger reinstatement than the RandCorrect condition ($t(11) = 2.90$, $p = .02$) and a trend towards stronger reinstatement than CircCorrect ($t(11) = 2.05$, $p = .06$).

Dimensionality transformation, reconstruction of original data and decoding

The same procedure as described above was used to epoch and preprocess the data. All subsequent analyses were conducted separately for AM+ and AM– trials. Ripple-aligned data were segmented into temporal windows using a 60 ms sliding window with 90% overlap, spanning from –1 to +1 second around ripple onset. This windowing approach was chosen to increase the signal-to-noise ratio for dimensionality estimation and to produce a smooth temporal profile of representational complexity. Similar results were observed when using 100 ms or 200 ms sliding windows (see fig. S3A and B).

For each time window, we estimated the embedding dimensionality using principal component analysis (PCA). The eigenvalues of the covariance matrix were extracted, and the number of retained components was determined via the second derivative method—i.e., identifying the "elbow" point at which the explained variance sharply declined across subsequent components. This data-driven approach avoids the arbitrary selection of a fixed variance threshold (e.g., 85–90%), which is known to be problematic in PCA analyses (*90*). The number of retained components for each time window served as our estimate of dimensionality (variance explained within significant time window, AM+: $M = 53.85\%$, $SD = 10.06\%$; AM-: $M = 46.87\%$, $SD = 13.17\%$; paired samples t-test between trial types: $t(1,11) = 2.80$, $p = .02$). Almost identical results were observed when calculating the distance to a fitted power-law distribution across time.

To statistically assess the temporal difference in dimensionality between AM+ and AM– trials, we applied cluster-based permutation testing. Additionally, to confirm that dimensionality expansion was specifically linked to ripple timing, we repeated the same control analyses used in the reinstatement analyses (see fig. S4). Both control comparisons revealed significantly greater dimensionality in the empirical data compared to the shuffled conditions (RandCorrect: $t(11) = 3.36$, $p < .01$; CircCorrect: $t(11) = 3.36$, $p < .01$).

In a follow-up analysis, we correlated the significant time points of dimensionality estimates with (1) participants' reaction times (correct trials only), and (2) decoding accuracy (i.e.,

classifier performance for correct minus incorrect trials) using two-sided Spearman's rank correlations.

To further assess and to gain more specificity of the dimensionality transformation, we split the data into blocks (3-6 blocks per participant) and fitted a linear mixed effect model with the fixed effects being conditions (AM+ and AM-), dimensionality and time around ripple (pre and post [-1 to 0 and 0 to 1]) (table S1). Random effects were participant and blocks. We used the MATLAB function *fitlme* for the model fitting. Comparable results were obtained when using effective dimensionality (ED), an alternative metric that estimates the intrinsic dimensionality of neural population activity (*91*):

$$\mathrm{ED} = \frac{(\sum \lambda_i)^{\,2}}{\sum \lambda_i^2}$$

where λ are the eigenvalues (variance explained by each principal component).

An important step in dimensionality reduction is to validate the information retained in the reduced components. To this end, we performed a reconstruction analysis. First, we computed the within- and between-class covariance matrices and retained the eigenvectors corresponding to the number of components previously extracted via PCA. We then conducted an inverse transformation by multiplying the class centroids with the encoding and retrieval eigenvectors separately, effectively reconstructing the data in its original feature space.

To evaluate performance, we subtracted the centroids of the colour and scene classes from each test projection and assigned a class label (1 or 2) depending on which centroid the projection was closest to. This label was then compared to the actual test label. This yielded a time generalisation matrix equivalent in structure to that obtained with the original LDA analysis.

We averaged across encoding time for both the original and reconstructed data, separately for each participant, and computed the Spearman correlation between the resulting matrices. This yielded a two-dimensional correlation matrix showing, for each ripple-aligned retrieval time point, its similarity to all encoding time points. The diagonal of this matrix reflects the correspondence between original and reconstructed decoding performance at each matched time point (see fig. S3C).

To further summarise this relationship, we averaged across retrieval time as well, producing a single decoding accuracy value for each participant and trial type. These values were

strongly correlated across participants (Spearman's $\rho = .87$, $p < .001$), indicating a high degree of correspondence between original and reconstructed data. While we used decoding accuracy as our index of retained information, any metric distinguishing AM+ and AM– trials could have been used in principle.

De-mixed Principal component analysis

The same procedures described previously were used for epoching, preprocessing, and ripple detection. We then divided the ripple-aligned data into the four target association conditions (red, blue, indoor, outdoor) and memory outcome (AM+ and AM–). For each participant, all cortical channels were included. These data were then combined, resulting in a four-dimensional array (channels × target association × memory × time).

Data organisation followed the procedure outlined in (*55*) and made use of open-source code provided in (*92*). For analysis, default algorithm parameters were used. The marginalisation variables of interest were: (1) target association, (2) memory, (3) their interaction, and (4) an independent component. The lambda parameter was optimised by running 100 iterations using *dpca_optimizeLamda*, which yielded the decoder (w) and encoder (v) matrices for each component. In addition, the noise covariance matrix was estimated using *dpca_getNoiseCovariance*, and explained variance was quantified *dpca_explainedVariance*. This explained variance estimate, based on the noise covariance, enabled us to distinguish structured neural signals from random noise. Since all estimated signal was captured by the first 50 demixed principal components (dPCs), any remaining variance was likely attributable to noise (see fig. S5A) (*55*).

For decoding, we included the top 3 components for each marginalisation and repeated the classification 100 times using *dpca_classificationAccuracy*. To obtain a baseline, label-shuffled data were used, with 500 shuffles repeated 100 times using *dpca_classificationShuffled*. Significant time windows were determined via 500 iterations of stratified Monte Carlo leave-group-out cross-validation. Trial labels were shuffled 500 times using a stratified approach to maintain an equal number of trials per condition. See (*55*) for full implementation details. The number of repetitions and shuffles was selected to provide stable estimates of the decoding distribution.

To assess statistically significant differences from the shuffled baseline, we applied a cluster-based permutation test using z-scored decoding accuracy values (note: this procedure is not part of the original dPCA toolbox but was added to control for multiple comparisons). Z-scores were computed by subtracting the mean of the shuffled distribution from the empirical

decoding values and dividing by the standard deviation. Consecutive time points exceeding a z-threshold of ±1.96 (two-sided α = .05) were grouped into clusters using binary connected component labelling (*bwlabel* in MATLAB). For each cluster, the sum of z-values was computed as the test statistic. A null distribution was generated via 1000 random permutations of the accuracy time series, and the maximum cluster-level statistic was recorded for each permutation. Observed cluster statistics were compared to the 97.5th percentile of the null distribution, and clusters exceeding this threshold were deemed statistically significant. This procedure was conducted separately for each dPC. Because each marginalisation involved a different number of conditions (e.g., interaction: 8 classes, chance = 12.5%; memory: 2 classes, chance = 50%), decoding results in Fig. 3D are plotted as the empirical accuracy minus the shuffled baseline. In figure S5C–G, both empirical and baseline decoding values are shown.

To visualise the components, we projected the original data onto the weight matrix and plotted the first dPCs corresponding to the target association × memory interaction, target association, and memory effects (dPCs 1, 2, and 4). To visualise neural trajectories in state space, we multiplied each component by the original data and separated trials into pre- and post-ripple windows. For post-ripple, we included data from 400 to 800 ms after ripple onset, corresponding to the peak of dimensionality expansion. To ensure symmetry in comparison, we used –800 to –400 ms for the pre-ripple window.

K-means clustering was performed for cluster counts ranging from 1 to 10 using MATLAB's built-in *kmeans* function, with each k run 10 times using random initial centroids. K-means clustering identifies clusters by minimising the distance between data points and their respective cluster centres. The silhouette score, which measures how well-separated data points are within their assigned clusters with values ranging from zero to one, where higher scores indicate stronger clustering, was calculated using MATLAB's *silhouette* function, and we plotted the mean silhouette value across cluster sizes. The optimal cluster number was identified for both pre- and post-ripple windows. Statistical results for dPCs 1, 2, and 4 are reported in the main text. Additionally, when plotting dPCs 1, 3, and 4 (fig. S5H), we observed that the optimal cluster count increased from four clusters pre-ripple (silhouette score = 0.85) to eight clusters post-ripple (silhouette score = 0.91). The Euclidean distance between cluster centroids also significantly increased (pre-ripple: M = 3.51, SE = 0.043; post-ripple: M = 4.66, SE = 0.052; $t(319) = 7.12$, $p < .01$).

Theta-gamma phase-amplitude coupling

The same procedure as for decoding and dimensionality estimation was used for epoching, processing and detecting ripples in the data, except that we now used a much broader time window of -4 to 6 seconds around ripple peak to account for later time-frequency transformation. We defined the peak frequency in the theta and gamma range, separately. To isolate oscillatory contributions (*45*) and to find the peak frequency of low (1-30Hz, in steps of 1Hz) and high (30-150Hz, in steps of 5Hz) frequencies, 1/f activity was attenuated in the time-frequency domain using the FOOOF algorithm (*59*) as implemented in the Fieldtrip toolbox (*85*). We then defined the peak theta frequency as the frequency with the highest power between 3 and 8Hz in hippocampal channels and for gamma frequency between 40 and 140Hz for cortical channels (to allow for ±2 and 10Hz for theta and gamma, respectively; see next). This procedure was done to ensure that the phase-amplitude coupling was performed on a narrow-band oscillation rather than broadband. Once defined we centred the data on these frequencies with a span for theta being ±2Hz and for gamma frequency being ±10Hz, with the peak frequency in the middle (frequency 0). Once having defined the peak frequency, the original data was subjected to another decomposition. Again, we divided high and low frequencies, using different methods to estimate phase and power. For low frequencies we convolved the data using a wavelet transformation with a hanning taper, with the number of cycles being roughly 500ms for each frequency, but never less than 5 cycles. For high frequencies, we estimated power using a multitaper method based on Slepian sequences as tapers. Frequency smoothing was set to one quarter of the frequency of interest and temporal smoothing was set to 200ms (*45*). The data were then baseline corrected between -500 to -100 pre-cue onset. The phase-amplitude coupling was performed per channel between -1 and 1 second, to comply with all other analyses. We performed the phase-amplitude coupling as in (*46*, *93*). For each phase-amplitude sample, we also ran a permuted baseline, where we shuffled the trials 250 times. The resulting data show the contrast between empirical and shuffled phase-amplitude coupling with the peak frequency for theta and gamma as 0 on the x and y-axis, respectively. The contrast was statistically tested by running a two-sided cluster-based permutation.

To assess when in time around the ripple events the TG-PAC was strongest, we ran the same analysis, but now binning the data into 500ms bins, with 90% overlap and only including peak gamma ±5Hz. We selected the length of the time window to allow for a minimum of one full cycle for each frequency before and after the ripple event (-1 to 1), constraining the binning to at least 500ms (1Hz frequency needs 1 second for a full cycle). Again, we ran a two-sided cluster-based permutation to test for significance between empirical data and a shuffled baseline (now between -.5 to 1 sec as we were mainly

interested in the post-ripple effect). In a subsequent analysis, we correlated the dimensionality transformation with the TG-PAC using a two-sided Spearman's correlation.

Lastly, to understand the temporal directionality, we performed a cross-correlation between TG-PAC and dimensionality expansion. Due to difference number of sample points for the two vectors, we linearly interpolated them using MATLAB's *interp1* function. For each participant, we then generated 1000 null-distributions. To keep the temporal autocorrelation in the surrogate data, we used the Iterative Amplitude Adjusted Fourier Transform (*94*), which instead of randomly shuffling time points, shuffles phase-values. We then z-scored the observed group-average cross-correlation against the distribution of permuted cross-correlations at each time lag. To correct for multiple comparisons across lags, we applied a cluster-based approach: z-score values exceeding a threshold of 1.96 (alpha .05 for two-sided test) were binarised, and temporally contiguous clusters of supra-threshold points were identified using connected component labeling (*bwlabel* in MATLAB). The sum of z-scores within each cluster was computed as a cluster-level statistic, and the maximum cluster sum across the entire lag window was retained for the real data. This procedure was repeated for each of the 500 permutations to generate a null distribution of maximum cluster statistics. A p-value was computed by comparing the real maximum cluster statistic to this null distribution, quantifying the probability that a cluster of equal or greater strength would be observed under the null hypothesis.

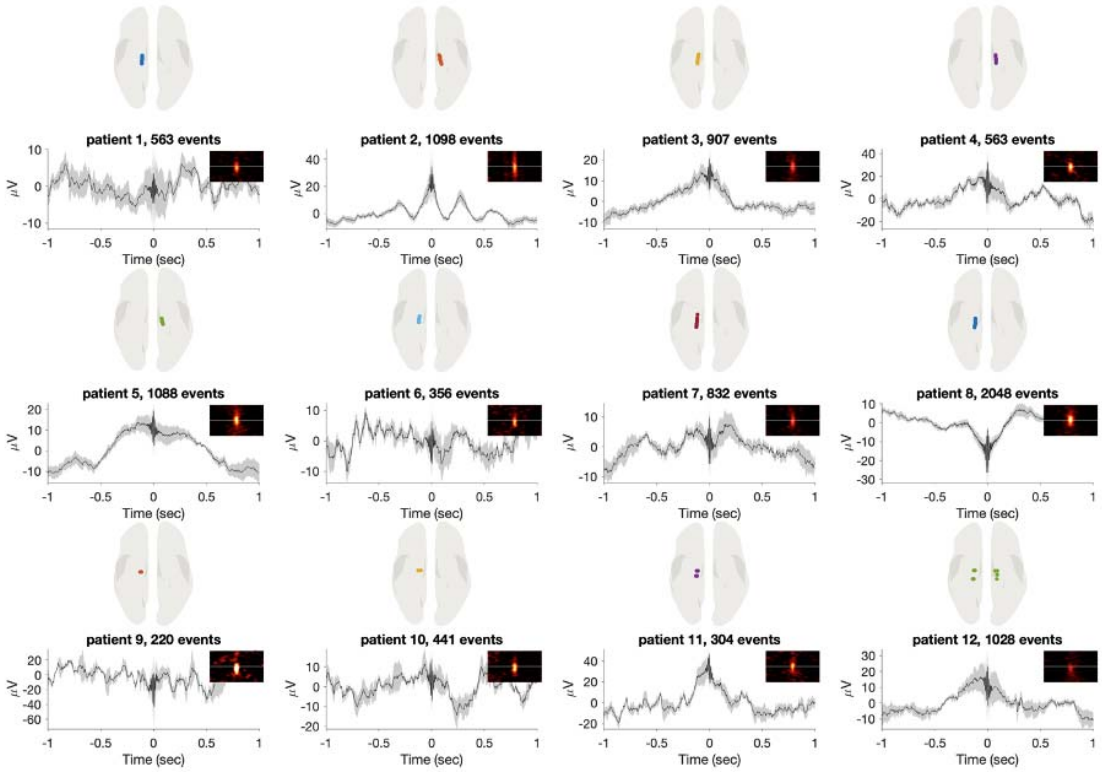

Fig. S1. Electrode implantation, ripple counts, and spectral characteristics per participant. Each panel displays data from a single participant. The top row illustrates electrode implantation sites within the hippocampus. The middle row shows the number of detected ripple events. The bottom row presents the average ripple waveform alongside the corresponding spectral power profile.

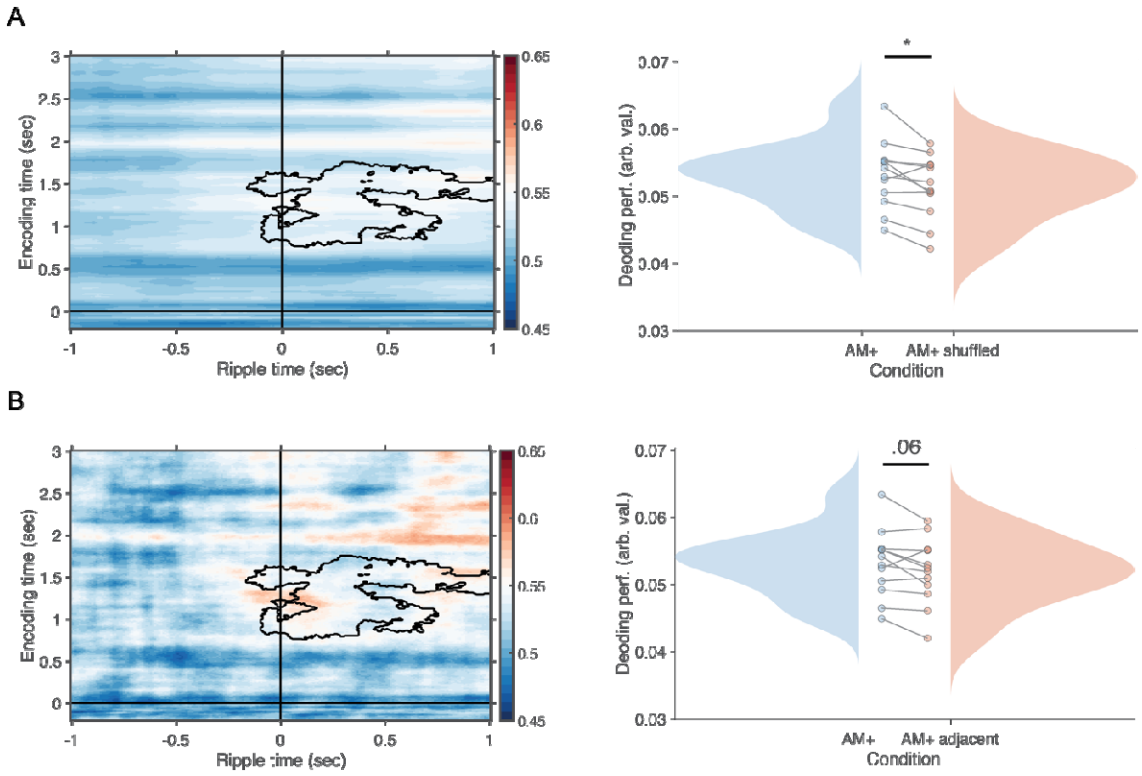

Fig. S2. Control Analyses Confirming Ripple-Triggered Reinstatement Specificity. **(A)** When AM+ trials were shuffled 1000 times, decoding accuracy within the significant cluster identified in Figure 2A (highlighted in black, left) was significantly higher for the empirical AM+ data compared to the shuffled data. Individual participant values are shown on the right. **(B)** A similar analysis was performed using a stricter control in which ripple times were swapped with those of adjacent trials (i.e., neighbouring trial shuffling). This revealed a trend towards higher decoding in the empirical AM+ condition compared to the adjacent-trial control.

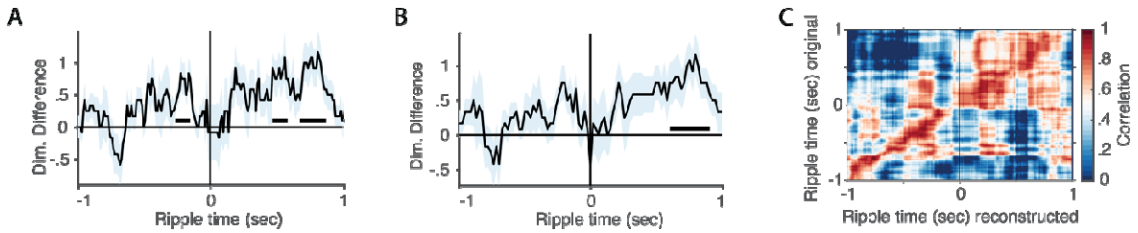

Fig. S3. Robustness of Dimensionality and Reconstruction Analysis

(A, B) Using a 100 ms **(A)** or 200 ms **(B)** sliding window for dimensionality estimation produced results highly similar to those obtained with a 60 ms window, demonstrating the robustness of the approach. **(C)** Correlation between original and reconstructed data, using the same number of components as in the PCA analysis, revealed a strong diagonal pattern—indicating a high-fidelity reconstruction of the original data structure.

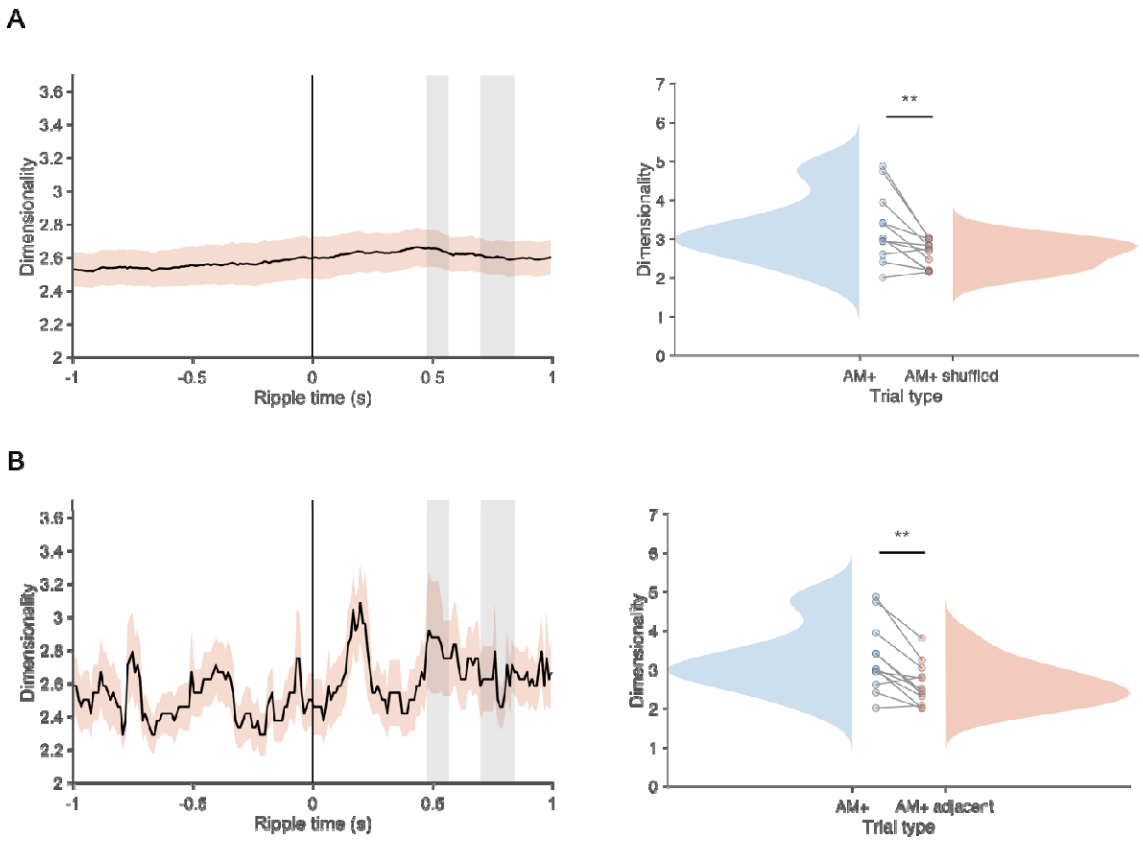

Fig. S4. Control Analyses Confirming Specificity of Dimensionality Expansion **(A)** Shuffling AM+ trials 1000 times revealed significantly greater dimensionality for empirical AM+ trials compared to the shuffled distribution, when averaging across the significant time points identified in Figure 2B (highlighted as grey bars, left). **(B)** A similar contrast between empirical AM+ and adjacent trials (i.e. temporally shifted controls) also yielded a significant difference, further supporting the specificity of ripple-linked dimensionality expansion. Each participant shown in right panels.

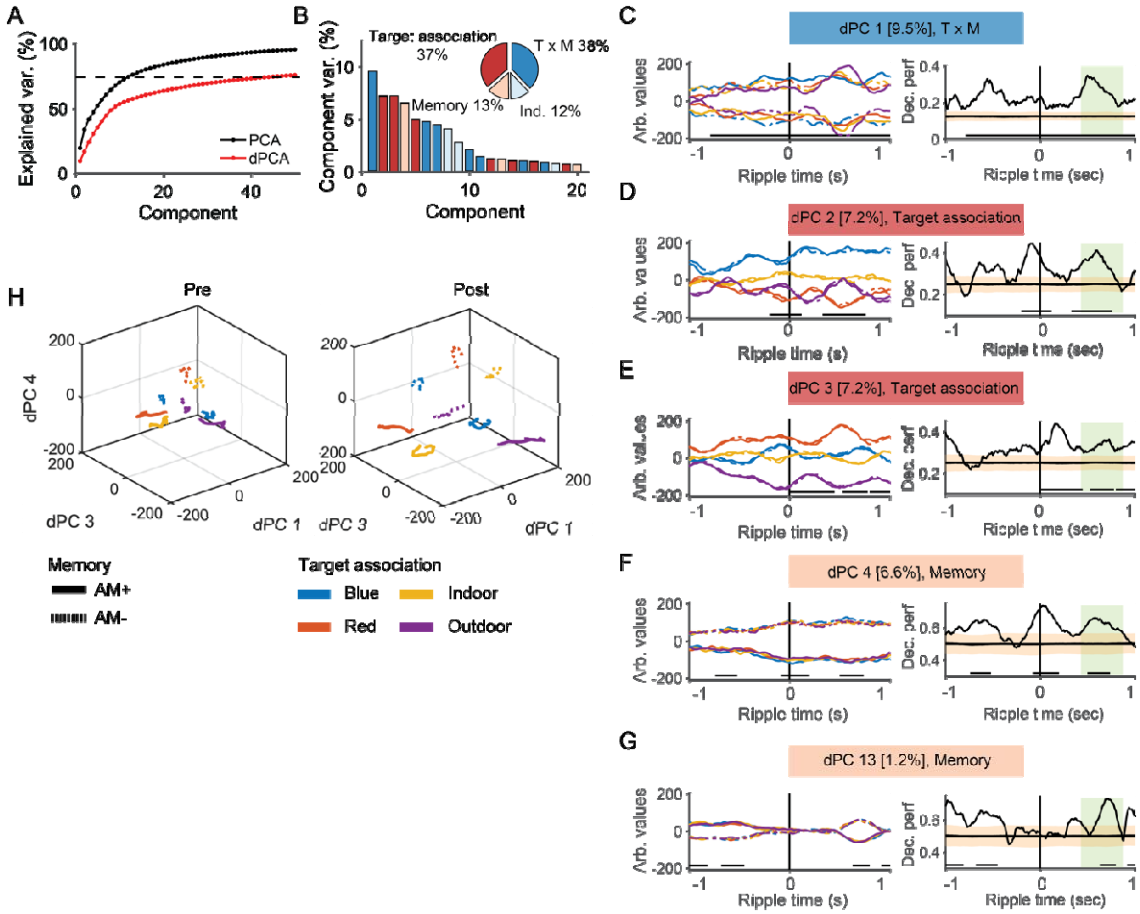

Fig. S5. Demixed Principal Component Analysis (dPCA). ***(A***) Cumulative variance explained by the first 50 components exceeded the estimated signal variance (dashed line). Standard PCA (black) accounted for more total variance than dPCA. **(B)** The first four dPCs primarily reflected a mixture of "target association" and "memory" (dark blue), "target association only" (red), and "memory only" (beige). **(C-G)** Decoding accuracy for individual components, shown only for those with significant decoding between 400–800 ms post-ripple peak (highlighted by green bars). Left panels display component activation; right panels show decoding performance. Significant decoding intervals are marked with black horizontal lines. Orange shaded areas represent the shuffled distributions, with their mean indicated in black. **(H)** Neural state space plots of dPCs 1, 3, and 4 show that pre-ripple representations were more overlapping (left), whereas post-ripple activity exhibited greater separation between experimental conditions (right), indicating increased representational differentiation.

Table S1. Linear mixed-effects model of dimensionality transformation.

Name	Estimate	SE	t	df	p	95% CI (lower)	95% CI (upper)
Intercept	2.28	.07	32.47	19,988	<.001	2.14	2.42
Trial Type	.09	.03	2.82	19,988	.005	.03	.16
Half	.14	.03	4.39	19,988	<.001	.08	.21
Trial Type x Half	-.09	.02	-4.16	19,988	<.001	-.13	-.05

References

1. E. Tulving, *Elements of Episodic Memory* (Oxford University Press, 1983).
2. L. de Almeida, M. Idiart, J. E. Lisman, Memory retrieval time and memory capacity of the CA3 network: Role of gamma frequency oscillations. *Learn. Mem.* 14, 795–806 (2007).
3. C. J. Bates, R. A. Jacobs, Efficient data compression in perception and perceptual memory. *Psychological Review* 127, 891–917 (2020).
4. H. B. Barlow, "Possible Principles Underlying the Transformations of Sensory Messages" in *Sensory Communication*, W. A. Rosenblith, Ed. (The MIT Press, 2012; https://academic.oup.com/mit-press-scholarship-online/book/20714/chapter/180090664), pp. 216–234.
5. C. Kerrén, D. Reznik, C. F. Doeller, B. J. Griffiths, Exploring the role of dimensionality transformation in episodic memory. *Trends in Cognitive Sciences*, doi: 10.1016/j.tics.2025.01.007 (2025).
6. D. Reznik, P. Majka, M. G. P. Rosa, M. P. Witter, C. F. Doeller, Phylogeny of neocortical-hippocampal projections provides insight in the nature of human memory. bioRxiv [Preprint] (2024). https://doi.org/10.1101/2024.05.09.593130.
7. H. Eichenbaum, T. Otto, N. J. Cohen, Two functional components of the hippocampal memory system. *Behavioral and Brain Sciences* 17, 449–517 (1994).
8. D. O. Hebb, *The Organization of Behavior; a Neuropsychological Theory* (Wiley, Oxford, England, 1949) *The organization of behavior; a neuropsychological theory*.
9. K. Duncan, N. Ketz, S. J. Inati, L. Davachi, Evidence for area CA1 as a match/mismatch detector: A high□resolution fMRI study of the human hippocampus. *Hippocampus* 22, 389–398 (2012).
10. 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, 419–457 (1995).
11. H. Eichenbaum, N. J. Cohen, *From Conditioning to Conscious Recollection: Memory Systems of the Brain* (Oxford University Press, 2004; https://doi.org/10.1093/acprof:oso/9780195178043.001.0001).
12. C. K. Machens, R. Romo, C. D. Brody, Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. *J Neurosci* 30, 350–360 (2010).

13. A. Kikumoto, A. Bhandari, K. Shibata, D. Badre, A transient high-dimensional geometry affords stable conjunctive subspaces for efficient action selection. *Nat Commun* 15, 8513 (2024).
14. R. Bartolo, R. C. Saunders, A. R. Mitz, B. B. Averbeck, Dimensionality, information and learning in prefrontal cortex. *PLOS Computational Biology* 16, e1007514 (2020).
15. V. Mante, D. Sussillo, K. V. Shenoy, W. T. Newsome, Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84 (2013).
16. C. D. Harvey, P. Coen, D. W. Tank, Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484, 62–68 (2012).
17. S. Hedayati, R. E. O'Donnell, B. Wyble, A model of working memory for latent representations. *Nat Hum Behav* 6, 709–719 (2022).
18. J. D. Murray, J. Jaramillo, X.-J. Wang, Working Memory and Decision-Making in a Frontoparietal Circuit Model. *J. Neurosci*. 37, 12167–12186 (2017).
19. S. L. Brincat, M. Siegel, C. von Nicolai, E. K. Miller, Gradual progression from sensory to task-related processing in cerebral cortex. *Proceedings of the National Academy of Sciences* 115, E7202–E7211 (2018).
20. E. Spens, N. Burgess, A generative model of memory construction and consolidation. *Nat Hum Behav*, doi: 10.1038/s41562-023-01799-z (2024).
21. M. Jazayeri, S. Ostojic, Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity. *Current Opinion in Neurobiology* 70, 113–120 (2021).
22. J. Linde-Domingo, M. S. Treder, C. Kerrén, M. Wimber, Evidence that neural information flow is reversed between object perception and object reconstruction from memory. *Nat Commun* 10, 179 (2019).
23. J. Lifanov, J. Linde-Domingo, M. Wimber, Feature-specific reaction times reveal a semanticisation of memories over time and with repeated remembering. *Nat Commun* 12, 3177 (2021).
24. S. Mirjalili, P. Powell, J. Strunk, T. James, A. Duarte, Context Memory Encoding and Retrieval Temporal Dynamics are Modulated by Attention across the Adult Lifespan. *eNeuro* 8, ENEURO.0387-20.2020 (2021).
25. C. Kerrén, Y. Zhao, B. J. Griffiths, A reduction in self-reported confidence accompanies the recall of memories distorted by prototypes. *Commun Psychol* 2, 1–11 (2024).

26. C. Kerrén, J. Linde-Domingo, B. Spitzer, Prioritization of semantic over visuo-perceptual aspects in multi-item working memory. bioRxiv [Preprint] (2022). https://doi.org/10.1101/2022.06.29.498168.

27. D. Khodagholy, J. N. Gelinas, G. Buzsáki, Learning-enhanced coupling between ripple oscillations in association cortices and hippocampus. *Science* 358, 369–372 (2017).

28. S. P. Jadhav, G. Rothschild, D. K. Roumis, L. M. Frank, Coordinated Excitation and Inhibition of Prefrontal Ensembles during Awake Hippocampal Sharp-Wave Ripple Events. *Neuron* 90, 113–127 (2016).

29. J. Karimi Abadchi, M. Nazari-Ahangarkolaee, S. Gattas, E. Bermudez-Contreras, A. Luczak, B. L. McNaughton, M. H. Mohajerani, Spatiotemporal patterns of neocortical activity around hippocampal sharp-wave ripples. *Elife* 9, e51972 (2020).

30. N. K. Logothetis, O. Eschenko, Y. Murayama, M. Augath, T. Steudel, H. C. Evrard, M. Besserve, A. Oeltermann, Hippocampal-cortical interaction during periods of subcortical silence. *Nature* 491, 547–553 (2012).

31. H. R. Joo, L. M. Frank, The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. *Nat Rev Neurosci* 19, 744–757 (2018).

32. Y. Norman, E. M. Yeagle, S. Khuvis, M. Harel, A. D. Mehta, R. Malach, Hippocampal sharp-wave ripples linked to visual episodic recollection in humans. *Science* 365, eaax1030 (2019).

33. Y. Norman, O. Raccah, S. Liu, J. Parvizi, R. Malach, Hippocampal ripples and their coordinated dialogue with the default mode network during recent and remote recollection. *Neuron* 109, 2767-2780.e5 (2021).

34. A. P. Vaz, S. K. Inati, N. Brunel, K. A. Zaghloul, Coupled ripple oscillations between the medial temporal lobe and neocortex retrieve human memory. *Science* 363, 975–978 (2019).

35. A. P. Vaz, J. H. Wittig, S. K. Inati, K. A. Zaghloul, Replay of cortical spiking sequences during human memory retrieval. *Science* 367, 1131–1134 (2020).

36. S. Henin, A. Shankar, H. Borges, A. Flinker, W. Doyle, D. Friedman, O. Devinsky, G. Buzsáki, A. Liu, Spatiotemporal dynamics between interictal epileptiform discharges and ripples during associative memory processing. *Brain* 144, 1590–1602 (2021).

37. J. J. Sakon, M. J. Kahana, Hippocampal ripples signal contextually mediated episodic recall. *Proc Natl Acad Sci U S A* 119, e2201657119 (2022).

38. E. Nyhus, T. Curran, Functional role of gamma and theta oscillations in episodic memory. *Neurosci Biobehav Rev* 34, 1023–1035 (2010).

39. P. Fries, Rhythms for Cognition: Communication through Coherence. *Neuron* 88, 220–235 (2015).

40. R. T. Canolty, R. T. Knight, The functional role of cross-frequency coupling. *Trends Cogn Sci* 14, 506–515 (2010).

41. A. Sirota, S. Montgomery, S. Fujisawa, Y. Isomura, M. Zugaro, G. Buzsáki, Entrainment of Neocortical Neurons and Gamma Oscillations by the Hippocampal Theta Rhythm. *Neuron* 60, 683–697 (2008).

42. A. Hyafil, A.-L. Giraud, L. Fontolan, B. Gutkin, Neural Cross-Frequency Coupling: Connecting Architectures, Mechanisms, and Functions. *Trends in Neurosciences* 38, 725–740 (2015).

43. J. Fell, N. Axmacher, The role of phase synchronization in memory processes. *Nat Rev Neurosci* 12, 105–118 (2011).

44. L. L. Colgin, Theta-gamma coupling in the entorhinal-hippocampal system. *Curr Opin Neurobiol* 31, 45–50 (2015).

45. B. J. Griffiths, M. C. Martín-Buro, B. P. Staresina, S. Hanslmayr, Disentangling neocortical alpha/beta and hippocampal theta/gamma oscillations in human episodic memory formation. *NeuroImage* 242, 118454 (2021).

46. R. T. Canolty, E. Edwards, S. S. Dalal, M. Soltani, S. S. Nagarajan, H. E. Kirsch, M. S. Berger, N. M. Barbaro, R. T. Knight, High Gamma Power Is Phase-Locked to Theta Oscillations in Human Neocortex. *Science* 313, 1626–1628 (2006).

47. L. Saint Amour di Chanaz, A. Pérez-Bellido, X. Wu, D. Lozano-Soldevilla, D. Pacheco-Estefan, K. Lehongre, E. Conde-Blanco, P. Roldan, C. Adam, V. Lambrecq, V. Frazzini, A. Donaire, M. Carreño, V. Navarro, A. Valero-Cabré, L. Fuentemilla, Gamma amplitude is coupled to opposed hippocampal theta-phase states during the encoding and retrieval of episodic memories in humans. *Current Biology* 33, 1836-1843.e6 (2023).

48. E. L. Newman, S. N. Gillet, J. R. Climer, M. E. Hasselmo, Cholinergic Blockade Reduces Theta-Gamma Phase Amplitude Coupling and Speed Modulation of Theta Frequency Consistent with Behavioral Effects on Encoding. J Neurosci 33, 19635–19646 (2013).

49. J. Daume, J. Kamiński, A. G. P. Schjetnan, Y. Salimpour, U. Khan, M. Kyzar, C. M. Reed, W. S. Anderson, T. A. Valiante, A. N. Mamelak, U. Rutishauser, Control of working memory by phase–amplitude coupling of human hippocampal neurons. *Nature* 629, 393–401 (2024).

50. F. Mormann, J. Fell, N. Axmacher, B. Weber, K. Lehnertz, C. E. Elger, G. Fernández, Phase/amplitude reset and theta–gamma interaction in the human medial temporal lobe during a continuous word recognition memory task. *Hippocampus* 15, 890–900 (2005).

51. A. Fernández-Ruiz, A. Oliva, G. A. Nagy, A. P. Maurer, A. Berényi, G. Buzsáki, Entorhinal-CA3 Dual-Input Control of Spike Timing in the Hippocampus by Theta-Gamma Coupling. *Neuron* 93, 1213-1226.e5 (2017).

52. A. B. L. Tort, R. W. Komorowski, J. R. Manns, N. J. Kopell, H. Eichenbaum, Theta-gamma coupling increases during the learning of item-context associations. *Proc Natl Acad Sci U S A* 106, 20942–20947 (2009).

53. N. Axmacher, M. M. Henseler, O. Jensen, I. Weinreich, C. E. Elger, J. Fell, Cross-frequency coupling supports multi-item working memory in the human hippocampus. *Proc Natl Acad Sci U S A* 107, 3228–3233 (2010).

54. B. P. Staresina, M. Wimber, A Neural Chronometry of Memory Recall. *Trends Cogn Sci* 23, 1071–1085 (2019).

55. D. Kobak, W. Brendel, C. Constantinidis, C. E. Feierstein, A. Kepecs, Z. F. Mainen, X.-L. Qi, R. Romo, N. Uchida, C. K. Machens, Demixed principal component analysis of neural population data. *Elife* 5, e10989 (2016).

56. G. Buzsáki, A. Draguhn, Neuronal oscillations in cortical networks. *Science* 304, 1926–1929 (2004).

57. M. Tamura, T. J. Spellman, A. M. Rosen, J. A. Gogos, J. A. Gordon, Hippocampal-prefrontal theta-gamma coupling during performance of a spatial working memory task. *Nat Commun* 8, 2182 (2017).

58. B. Lega, J. Burke, J. Jacobs, M. J. Kahana, Slow-Theta-to-Gamma Phase–Amplitude Coupling in Human Hippocampus Supports the Formation of New Episodic Memories. *Cereb. Cortex* 26, 268–278 (2016).

59. T. Donoghue, M. Haller, E. J. Peterson, P. Varma, P. Sebastian, R. Gao, T. Noto, A. H. Lara, J. D. Wallis, R. T. Knight, A. Shestyuk, B. Voytek, Parameterizing neural power spectra into periodic and aperiodic components. *Nat Neurosci* 23, 1655–1665 (2020).

60. R. F. Helfrich, J. D. Lendner, B. A. Mander, H. Guillen, M. Paff, L. Mnatsakanyan, S. Vadera, M. P. Walker, J. J. Lin, R. T. Knight, Bidirectional prefrontal-hippocampal dynamics organize information transfer during sleep in humans. *Nat Commun* 10, 3572 (2019).

61. B. P. Staresina, J. Niediek, V. Borger, R. Surges, F. Mormann, How coupled slow oscillations, spindles and ripples coordinate neuronal processing and communication during human sleep. *Nat Neurosci* 26, 1429–1437 (2023).
62. Z. Xiao, X. Wang, J. Zhang, J. Ou, L. He, Y. Qu, X. Hu, T. Behrens, Y. Liu, Human hippocampal ripples predict the alignment of experience to a grid-like schema. bioRxiv [Preprint] (2025). https://doi.org/10.1101/2025.01.08.632069.
63. R. Kaplan, M. H. Adhikari, R. Hindriks, D. Mantini, Y. Murayama, N. K. Logothetis, G. Deco, Hippocampal Sharp-Wave Ripples Influence Selective Activation of the Default Mode Network. *Curr Biol* 26, 686–691 (2016).
64. S. Michelmann, A. R. Price, B. Aubrey, C. K. Strauss, W. K. Doyle, D. Friedman, P. C. Dugan, O. Devinsky, S. Devore, A. Flinker, U. Hasson, K. A. Norman, Moment-by-moment tracking of naturalistic learning and its underlying hippocampo-cortical interactions. *Nat Commun* 12, 5394 (2021).
65. S. Michelmann, P. Dugan, W. Doyle, D. Friedman, L. Melloni, C. K. Strauss, S. Devore, A. Flinker, O. Devinsky, U. Hasson, K. A. Norman, Fast-timescale hippocampal processes bridge between slowly unfurling neocortical states during memory search. [Preprint] (2025). https://doi.org/10.1101/2025.02.11.637471.
66. D. Reznik, R. Trampel, N. Weiskopf, M. P. Witter, C. F. Doeller, Dissociating distinct cortical networks associated with subregions of the human medial temporal lobe using precision neuroimaging. *Neuron* 111, 2756-2772.e7 (2023).
67. D. Reznik, D. S. Margulies, M. P. Witter, C. F. Doeller, Evidence for convergence of distributed cortical processing in band-like functional zones in human entorhinal cortex. *Current Biology* 34, 5457-5469.e2 (2024).
68. M. Rigotti, O. Barak, M. R. Warden, X.-J. Wang, N. D. Daw, E. K. Miller, S. Fusi, The importance of mixed selectivity in complex cognitive tasks. *Nature* 497, 585–590 (2013).
69. A. C. Heusser, D. Poeppel, Y. Ezzyat, L. Davachi, Episodic sequence memory is supported by a theta-gamma phase code. Nat Neurosci 19, 1374–1380 (2016).
70. U. Vivekananda, D. Bush, J. A. Bisby, S. Baxendale, R. Rodionov, B. Diehl, F. A. Chowdhury, A. W. McEvoy, A. Miserocchi, M. C. Walker, N. Burgess, Theta power and theta-gamma coupling support long-term spatial memory retrieval. *Hippocampus* 31, 213–220 (2021).
71. N. Roehri, L. Bréchet, M. Seeber, A. Pascual-Leone, C. M. Michel, Phase-Amplitude Coupling and Phase Synchronization Between Medial Temporal, Frontal and

Posterior Brain Regions Support Episodic Autobiographical Memory Recall. *Brain Topogr* 35, 191–206 (2022).

72. L. L. Colgin, T. Denninger, M. Fyhn, T. Hafting, T. Bonnevie, O. Jensen, M.-B. Moser, E. I. Moser, Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* 462, 353–357 (2009).

73. B. J. Griffiths, G. Parish, F. Roux, S. Michelmann, M. van der Plas, L. D. Kolibius, R. Chelvarajah, D. T. Rollings, V. Sawlani, H. Hamer, S. Gollwitzer, G. Kreiselmeyer, B. Staresina, M. Wimber, S. Hanslmayr, Directional coupling of slow and fast hippocampal gamma with neocortical alpha/beta oscillations in human episodic memory. *Proc Natl Acad Sci U S A* 116, 21834–21842 (2019).

74. J. Sheng, L. Zhang, C. Liu, J. Liu, J. Feng, Y. Zhou, H. Hu, G. Xue, Higher-dimensional neural representations predict better episodic memory. Sci Adv 8, eabm3829 (2022).

75. G. Buzsáki, L. W. Leung, C. H. Vanderwolf, Cellular bases of hippocampal EEG in the behaving rat. *Brain Res* 287, 139–171 (1983).

76. G. Parish, S. Hanslmayr, H. Bowman, “The Sync/deSync model: How a synchronized hippocampus and a de-synchronized neocortex code memories” (preprint, Neuroscience, 2017); https://doi.org/10.1101/185231.

77. B. J. Griffiths, S. D. Mayhew, K. J. Mullinger, J. Jorge, I. Charest, M. Wimber, S. Hanslmayr, Alpha/beta power decreases track the fidelity of stimulus-specific information. *eLife* 8, e49562 (2019).

78. M. C. Martín-Buro, M. Wimber, R. N. Henson, B. P. Staresina, Alpha Rhythms Reveal When and Where Item and Associative Memories Are Retrieved. *J. Neurosci*. 40, 2510–2518 (2020).

79. N. A. Cayco-Gajic, C. Clopath, R. A. Silver, Sparse synaptic connectivity is required for decorrelation and pattern separation in feedforward networks. *Nat Commun* 8, 1116 (2017).

80. F. Lanore, N. A. Cayco-Gajic, H. Gurnani, D. Coyle, R. A. Silver, Cerebellar granule cell axons support high dimensional representations. *Nature neuroscience* 24, 1142 (2021).

81. C. Higgins, Y. Liu, D. Vidaurre, Z. Kurth-Nelson, R. Dolan, T. Behrens, M. Woolrich, Replay bursts in humans coincide with activation of the default mode and parietal alpha networks. *Neuron* 109, 882-893.e7 (2021).

82. S. Michelmann, B. P. Staresina, H. Bowman, S. Hanslmayr, Speed of time-compressed forward replay flexibly changes in human episodic memory. *Nat Hum Behav* 3, 143–154 (2019).

83. F. Mormann, S. Kornblith, R. Q. Quiroga, A. Kraskov, M. Cerf, I. Fried, C. Koch, Latency and Selectivity of Single Neurons Indicate Hierarchical Processing in the Human Medial Temporal Lobe. J. Neurosci. 28, 8865–8872 (2008).

84. B. P. Staresina, R. N. A. Henson, N. Kriegeskorte, A. Alink, Episodic reinstatement in the medial temporal lobe. J Neurosci 32, 18150–18156 (2012).

85. R. Oostenveld, P. Fries, E. Maris, J.-M. Schoffelen, FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Comput Intell Neurosci 2011, 156869 (2011).

86. M. Mölle, T. O. Bergmann, L. Marshall, J. Born, Fast and slow spindles during the sleep slow oscillation: disparate coalescence and engagement in memory processing. *Sleep* **34**, 1411–1421 (2011).

87. M. Mölle, L. Marshall, S. Gais, J. Born, Grouping of spindle activity during slow oscillations in human non-rapid eye movement sleep. *J Neurosci* **22**, 10941–10947 (2002).

88. B. P. Staresina, T. O. Bergmann, M. Bonnefond, R. van der Meij, O. Jensen, L. Deuker, C. E. Elger, N. Axmacher, J. Fell, Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. *Nat Neurosci* **18**, 1679–1686 (2015).

89. M. S. Treder, MVPA-Light: A Classification and Regression Toolbox for Multi-Dimensional Data. *Front. Neurosci*. **14** (2020).

90. J. Dien, The ERP PCA Toolkit: An open source program for advanced statistical analysis of event-related potential data. *Journal of Neuroscience Methods* **187**, 138–145 (2010).

91. M. Hörnquist, J. Hertz, M. Wahde, Effective dimensionality for principal component analysis of time series expression data. *Biosystems* **71**, 311–317 (2003).

92. D. Fetterhoff, M. Costa, R. Hellerstedt, R. Johannessen, L. Imbach, J. Sarnthein, B. A. Strange, Neuronal population representation of human emotional memory. *Cell Reports* **43** (2024).

93. A. Bragin, G. Jandó, Z. Nádasdy, J. Hetke, K. Wise, G. Buzsáki, Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. *J Neurosci* **15**, 47–60 (1995).

94. T. Schreiber, A. Schmitz, Improved Surrogate Data for Nonlinearity Tests. *Phys Rev Lett* **77**, 635–638 (1996).