

Hippocampal theta codes for distances in semantic and temporal spaces

Ethan A. Solomon^a, Bradley C. Lega^b, Michael R. Sperling^c, and Michael J. Kahana^{d,1}

^aDepartment of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104; ^bDepartment of Neurosurgery, University of Texas Southwestern, Dallas, TX 75390; ^cDepartment of Neurology, Thomas Jefferson University Hospital, Philadelphia, PA 19107; and ^dDepartment of Psychology, University of Pennsylvania, Philadelphia, PA 19104

Edited by Gyorgy Buzsáki, New York University Neuroscience Institute, New York, NY, and approved October 16, 2019 (received for review April 19, 2019)

The medial temporal lobe (MTL) is known to support episodic memory and spatial navigation, raising the possibility that its true function is to form “cognitive maps” of any kind of information. Studies in humans and animals support the idea that the hippocampal theta rhythm (4 to 8 Hz) is key to this mapping function, as it has been repeatedly observed during spatial navigation tasks. If episodic memory and spatial navigation are 2 sides of the same coin, we hypothesized that theta oscillations might reflect relations between explicitly nonspatial items, such as words. We asked 189 neurosurgical patients to perform a verbal free-recall task, of which 96 had indwelling electrodes placed in the MTL. Subjects were instructed to remember short lists of sequentially presented nouns. We found that hippocampal theta power and connectivity during item retrieval coded for semantic distances between words, as measured using word2vec-derived subspaces. Additionally, hippocampal theta indexed temporal distances between words after filtering lists on recall performance, to ensure adequate dynamic range in time. Theta effects were noted only for semantic subspaces of 1 dimension, indicating a substantial compression of the possible semantic feature space. These results lend further support to our growing confidence that the MTL forms cognitive maps of arbitrary representational spaces, helping to reconcile longstanding differences between the spatial and episodic memory literatures.

hippocampus | episodic memory | theta | recall | cognitive map

Our ability to form associations rapidly and efficiently has inspired a powerful conceptualization of brain function: It constructs internal “maps” of related information that we mentally navigate when the behavioral context demands it (1–3). These maps need not represent only physical layouts, like an office plan or a route to work—cognitive maps could also reflect associations in domains such as autobiographical memories, perceptual inputs, or social interactions (4). In this way, for instance, we can conceptualize recollection of a life event as taking a mental stroll through a map of our prior experiences, guided by the similarities between items in a multidimensional feature space.

Evidence from studies of memory and spatial navigation provides compelling support for the idea that the brain has a domain-general cognitive mapping ability with a common neural substrate. Decades of lesional, electrophysiological, and imaging studies demonstrate that the human medial temporal lobe (MTL), including the hippocampus and parahippocampal gyrus, supports cognitive operations underlying both spatial navigation and episodic memory (5–8), while studies of rodent navigation have shown that MTL neurons code for spatial locations in a context-dependent manner (9). More recently, microwire recordings in humans have begun to bridge the human–animal divide, showing the presence of place-coding cells in the MTL whose firing rates can be modulated by task contexts and reflect associations between spatial locations and behaviorally relevant content (10–12).

Common to studies of spatial navigation in humans and animals is the emergence of a low-frequency oscillation in the hippocampal local field potential (LFP), called the theta rhythm

(typically 4 to 8 Hz). Elevated theta spectral power is consistently observed in humans during real-world (13) and virtual navigation (14–16) and during tasks that require the encoding or retrieval of spatial information (17, 18). Theta oscillations appear prominently during rodent navigation (19, 20) and relate to neuronal activity, in that the spiking of place-coding cells occurs at drifting phases of the theta rhythm as an animal explores its environment (21, 22). Moreover, these firing patterns may be reinstated during theta states as an animal retrieves information about a previously explored layout (23–25). Although such instances of “phase precession” have not been observed in navigating humans, other studies strongly support the general notion of theta phase coding as humans perform spatial and memory tasks (26–28).

Although it is now clear that theta oscillations correlate with the encoding and retrieval of spatial relations—or maps—there is a more tenuous link between theta and explicitly nonspatial memory. If the MTL truly serves a domain-general mapping function that creates relational networks between information of all types—be it spatial, semantic, autobiographical, or otherwise—we should expect to observe the same theta-mediated phenomena during nonspatial memory processing. Indeed, theta activity has been observed in human MTL and neocortex during episodic memory processing for words and pictures (29–32). However, these studies examine theta through the broad lens of successful-versus-unsuccessful memory, called the

Significance

The medial temporal lobe (MTL) plays a role in both spatial navigation and memory, but how these 2 cognitive processes are linked remains unclear. In particular, theta oscillations (4 to 8 Hz) appear prominently in the hippocampus in rodents and humans as they move through spatial environments, but there is mixed evidence as to whether this signal also emerges while animals search memory for previously acquired information. Using electrodes implanted in human neurosurgical patients, we showed that hippocampal theta oscillations reflect representational distances between word items in memory. This suggests there is a fundamental theta-based mechanism that supports the creation and retrieval of “cognitive maps” in the MTL, for explicitly nonspatial information.

Author contributions: E.A.S., B.C.L., M.R.S., and M.J.K. designed research; E.A.S., B.C.L., and M.R.S. performed research; E.A.S. and M.J.K. contributed new reagents/analytic tools; E.A.S. analyzed data; and E.A.S. and M.J.K. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

Data deposition: Raw electrophysiological data used in this study are available at http://memory.psych.upenn.edu/Electrophysiological_Data. Analysis code and processed data have been deposited in a public interactive Jupyter Notebook, <https://notebooks.azure.com/esolo/projects/theta-semantic>.

¹To whom correspondence may be addressed. Email: kahana@psych.upenn.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1906729116/-DCSupplemental.

First published November 13, 2019.

subsequent memory effect (SME). Although a useful paradigm, the SME is not a direct test that theta specifically supports the formation or retrieval of associations between items in memory, i.e., a cognitive map of recently acquired information. Moreover, many similar studies report the opposite effect: a hippocampal or cortical decrease in theta power during successful encoding or retrieval (33–38). These conflicting findings raise serious questions about whether the neural processing of spatial and episodic information relies on the same theta-based mechanisms.

We set out to directly ask whether the established neural signatures of spatial memory and navigation—MTL theta oscillations—are also present during the retrieval of relational verbal memory. To do this, we tested 189 human neurosurgical patients with indwelling electrodes on a verbal free-recall task. Prior studies indicate that humans principally rely on temporal and semantic relatedness to organize verbal information (39, 40), so our overarching goal was to ascertain whether human behavior and neural activity aligned with measures of temporal and semantic association. We explicitly tested whether theta rhythms in the human MTL code for 1) semantic distances between retrieved nouns and 2) temporal distances based on the serial position of presented items during encoding. In doing so, we sought to establish how a signature of spatial navigation, hippocampal theta, could also play a role in episodic memory

retrieval. Is the encoding of word items akin to guiding subjects through an abstract space, and is retrieval akin to asking them to walk through it themselves?

Using a pretrained word2vec (41) representation of words in a vector space, we found that we could predict a subject’s recall behavior based on interword distances in word2vec subspaces of varying dimensionality. Next, we asked whether interword distances, and temporal distances, correlated with preretrieval theta power in the hippocampus and parahippocampal gyrus (PHG). We discovered that preretrieval hippocampal theta power significantly predicts interword distances in semantic space, but theta was not correlated with temporal distances. However, analytic corrections for the case of temporally restricted recalls revealed a time-associated theta effect. Finally, we assessed interregional theta connectivity, finding that hippocampal–PHG coupling also predicted interword semantic distances. Taken together, we established that retrieval-associated theta activity encodes representational distances in a word space, supporting the idea that theta underlies navigation through domain-general cognitive maps in the MTL.

Results

Clustering Effects in Recall. We first sought to examine the degree to which distances between items’ semantic representations predicted both recall behavior and neural activity. Neurosurgical

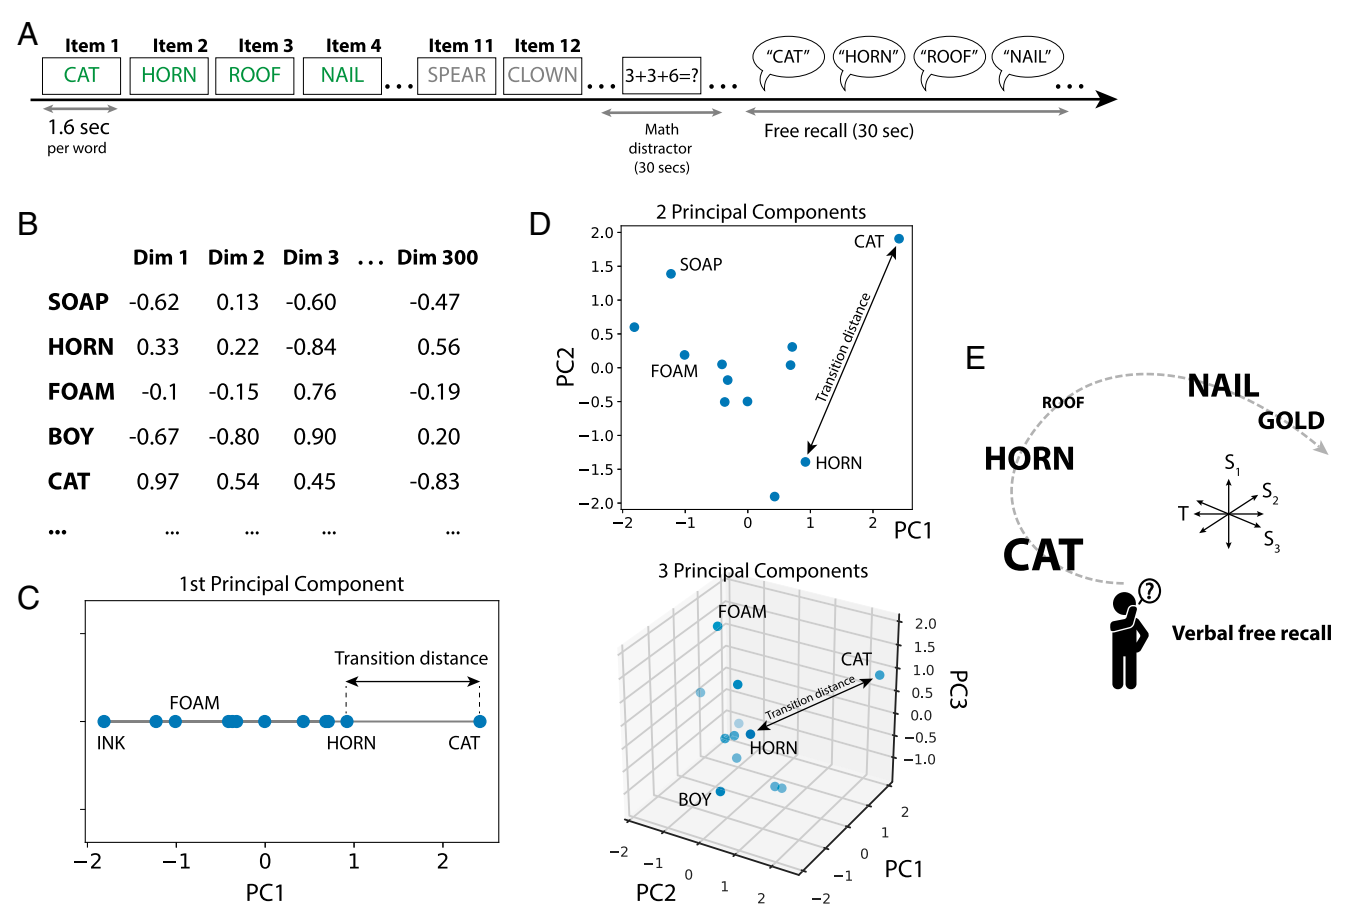


Fig. 1. Embedding of word items in representational semantic spaces. (A) A total of 189 subjects performed a verbal free-recall task, in which they were asked to remember 12-item lists of simple nouns. After a distractor, subjects freely recall as many words as possible from the prior list, which can be conceptualized as a cognitive transition from one word to another. Successful recalls are indicated in green for an example list. (B) Each word occupies a position in the pretrained 300-dimensional word2vec space (SI Appendix, SI Materials and Methods). PCA is used to extract the principal dimensions of variation for each word list. (C) Example interrecall distance in a 1D semantic space, the first principal component. (D) The 2D and 3D embedding of the words from this list. (E) Navigation through a 2D spatial layout is associated with hippocampal theta oscillations, which may also support episodic retrieval processes through a multidimensional semantic–temporal space. Icon designed by Freepik from Flaticon.

patients with indwelling electrodes ($n = 189$) studied and subsequently recalled 12-item word lists. During the study phase, each word appeared for 1.6 s followed by an 800- to 1,200-ms blank interstimulus interval (see Fig. 1*A* for an example list). Following list presentation, subjects first performed a 20-s arithmetic distractor task and then attempted to freely recall items from the preceding list in any order. Subjects completed multiple

sessions of this free-recall task, with each session comprising 12 to 25 study-test lists.

We next used a word2vec semantic representation to create layouts of the words in each list, at varying dimensionalities. Word2vec is a method for creating vector representations of words based on regularities present in text corpora (41). Google has made available a pretrained set of 300-dimensional word

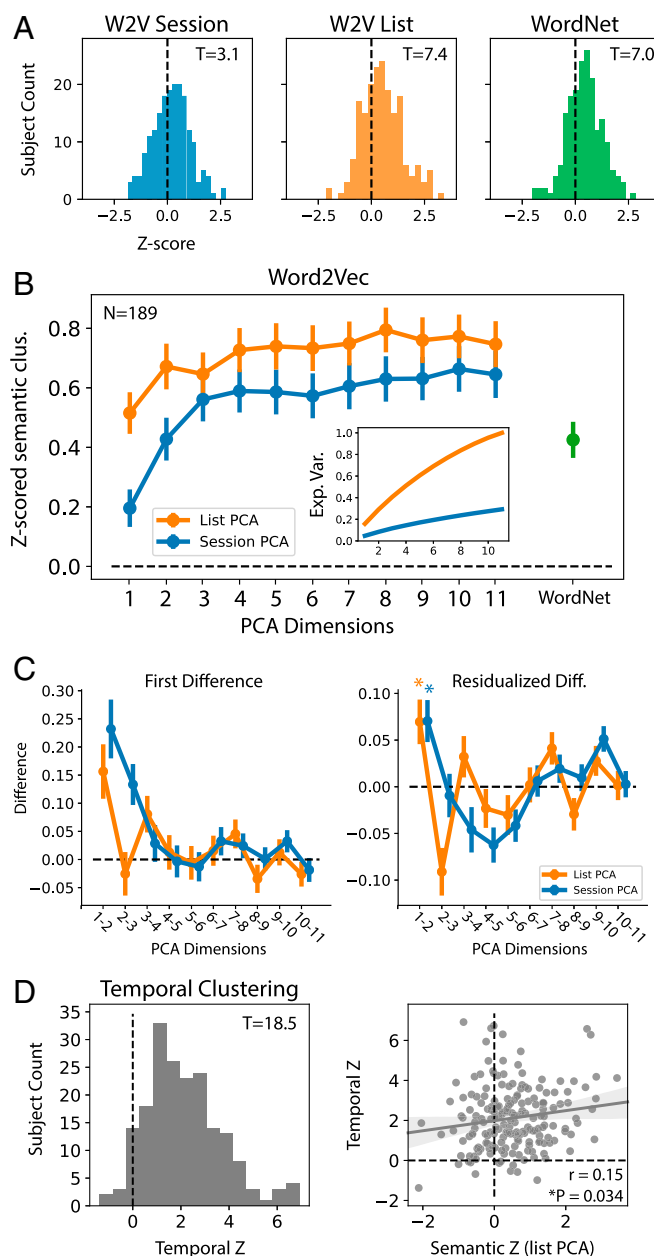


Fig. 2. Recall patterns are predicted by interword distances in semantic spaces. (*A*) Assessing whether interword distances in 1D word2vec-derived spaces predicted the recall of semantically related words. For each subject, a semantic clustering z score was computed, reflecting the degree to which the pattern of recalls correlated with interword distances in PCA-reduced word2vec spaces of n dimension (*SI Appendix, SI Materials and Methods*), here shown for just the 1D space reflecting the first principal component. PCA was applied to the 12-vector group of words from each list individually (list PCA, orange) or applied to the vectors of all words seen in an experimental session together (session PCA, blue). As a benchmark, WordNet lexical similarities were also used to calculate semantic clustering (green). Population-level semantic clustering was significant for all methods (1-sample t test, $P < 0.002$). (*B, Left*) Analysis from *A*, extended to PCA-derived semantic subspaces of higher dimensions. *Inset* shows cumulative explained variance for list- and session-level PCA. (*B, Right*) Mean semantic clustering score using WordNet-derived semantic relations. (*C, Left*) First-order differences of the clustering vs. dimensionality curves in *B*. (*C, Right*) First-order differences residualized on the explained variance, indicating that behavioral clustering increased at a faster rate than the explained variance between the first and second dimensions for both list PCA and session PCA ($P < 0.05$). (*D, Left*) Histogram of subjects' temporal clustering scores, which were significantly above chance ($P < 0.001$). (*D, Right*) Correlation of subjects' semantic clustering and temporal clustering scores, indicating a significant positive relationship ($r = 0.15$, $P = 0.034$).

vectors based on a large Google News text corpus. Word2vec makes explicit the multidimensional nature of word representations, enabling us to ask whether the dimensionality of a word subspace is related to behavioral or neural variables of interest. Accordingly, we used principal component analysis (PCA) to select the top n orthogonal dimensions which explain the greatest variance among list words in their native 300-dimensional (300D) word2vec space (Fig. 1 *B–D* and *SI Appendix, SI Materials and Methods*). The result, for each list, is a set of coordinates that locate each word in an n -dimensional space, where n is a maximum of 11 (since each list consists of 12 words).

Do distances in these semantic representational spaces align with how subjects retrieve words during the free-recall phase? Prior literature has established that other measures of semantic similarity predict the order of word recall, in that semantically related words will tend to be retrieved in clusters (39). Our expectation was that distances in these word2vec spaces would also predict clustered recalls, and they did: Using the semantic clustering score—a metric which indexes whether a subject tends to organize recalls according to semantic relatedness—subjects exhibited recall patterns which were significantly correlated with distances in reduced word2vec spaces of any dimensionality (Fig. 2 *A* and *B*; 1D clustering $t(188) = 7.4$, $P < 0.001$). Semantic clustering scores tended to rise with increasing PCA dimensionality, but beyond 2 dimensions, this rise did not significantly outpace the increase in overall explained variance with additional dimensions; regressing out explained variance from the slope of the curve in Fig. 2*B* showed, across subjects, a significant ($P < 0.05$) rise between 1 and 2 dimensions (Fig. 2 *B*, *Inset* and *C* and *SI Appendix, SI Materials and Methods*).

To ask whether semantically clustered recalls also correlated with interword distances in other spaces, we examined the relation between sequential recalls and reduced word2vec spaces derived from PCA on all of the words presented in an experimental session (typically 144 to 300 words, depending on the number of lists completed). In this way, PCA is identifying the orthogonal dimensions that explain the greatest variance for all of the words seen over the course of the session, as opposed to words seen within a list. Semantic relatedness from session-level PCA also predicted clustering during free recall for all dimensions, although to a lesser degree than list-level PCA (Fig. 2 *A* and *B*; 1D clustering $t(188) = 3.1$, $P = 0.002$). Clustered recalls were better predicted by interword distances at higher dimensions—like those we found in list-level PCA—but as before the rise in clustering scores did not outpace the increase in explained variance beyond 2 dimensions (Fig. 2 *A*, *Inset* and *C*). As a benchmark, we also showed that recalls clustered according to interitem similarities derived from the WordNet lexical database (42), which does not rely on a model-based embedding of words in a high-dimensional vector space (Fig. 2 *A*, *Right* and *B*, *Right*; 1-sample t test, $t(188) = 7.0$, $P < 0.001$).

These data replicate prior findings that verbal recall is governed in part by semantic relations (39), using a multidimensional framework for capturing interword similarities. However, it is well known that humans also organize verbal memory according to the temporal sequence of items during initial encoding (40). The temporal clustering score quantifies this relationship by assessing the likelihood that a subject will recall 2 items in sequence, given their serial position during encoding. As in prior studies, we showed a strong population-level effect of temporal clustering (1-sample t test, $t(188) = 18.5$, $P < 0.001$). To ask whether the tendency to semantically cluster is related to the tendency to temporally cluster, we found that a subject's temporal clustering score was weakly but significantly correlated with the subject's 1D list-PCA semantic clustering score ($r = 0.15$, $P = 0.034$).

Taken together, we demonstrated that humans behaviorally cluster their recalls of verbal items according to interword distances in word2vec semantic spaces of varying dimension. Altogether, this is not surprising—word2vec representations should not be expected to differ dramatically from other methods of capturing semantic similarity, like latent semantic analysis (43). In both list-PCA and session-PCA designs, we found that the relationship between recall structure and semantic dimensionality saturated after 2 dimensions, suggesting that 1 or 2 principle axes of variation explain a substantial fraction of how semantic features are organized in memory. We next sought to find the neural underpinnings of semantically and temporally clustered recalls.

Oscillatory Power during Search through 1D Spaces. Prior studies suggest a fundamental link between hippocampal theta oscillations and the organization of sequential information, such as paths through an environment. In humans, theta power increases as subjects search through a virtual environment (15) or recall trajectories in previously explored spatial layouts (17, 18). Here we ask whether theta also supports the retrieval of nearby items in a semantic representational space.

To assess this, we recorded LFPs from depth electrodes in the hippocampus or PHG of human neurosurgical patients. Our general approach was to use the framework outlined earlier (Fig. 1) to assess whether theta power leading up to a retrieval event predicted the semantic distance between the 2 subsequently recalled words. A positive correlation between theta power and semantic distance would suggest that theta oscillations are serving to reinstantiate the semantic context established during encoding and enable a participant to select the optimal path through a representational space (Fig. 1*E*).

Fig. 3 demonstrates this procedure in an example subject. In one list, the subject recalled 5 of the original 12 words, recalled in an order that sometimes reflects close semantic relations (e.g., recall of “cloud” was followed by “star”) and other times weak semantic relations (e.g., “star” was followed by “nose”; Fig. 3*A*). For each recall transition, the interword distance is captured by the semantic clustering score, here reported as a percentile rank with 1 indicating the closest-possible semantic transition (*SI Appendix, SI Materials and Methods*). We first computed semantic clustering scores based on the 1D list-PCA space (and later generalized to higher dimensions). We next measured the theta power in 1-s intervals leading up to each recall event, just prior to vocalization (Fig. 3*B*). An ~ 4 -Hz theta oscillation is seen preceding the recall of “nose,” which was followed by the recall of “peach” (0.66 clustering score). Finally, we measured the (z-scored) preretrieval theta power for all words recalled in an experimental session and correlated it with the semantic distance between the 2 subsequently recalled words. This subject shows a significant positive relationship between 1D semantic distance and preretrieval theta power (Fig. 3*C*, Pearson's $r = 0.52$, $P = 0.004$). To assess the actual degree of theta power change relative to average, we also binned semantic distances into “short” (>0.75 clustering score) and “long” (<0.25) and averaged theta power for recalls in both bins (Fig. 3*D*). This demonstrated that theta power increased above average prior to near transitions and fell below average prior to far transitions.

Across all subjects with electrodes placed in the hippocampus (either hemisphere; $n = 70$), we found a significant correlation between preretrieval theta power and 1D transition distance (1-sample t test, $t(69) = 2.19$, $P = 0.032$; Fig. 4*A*, *Top* and *SI Appendix, Fig. S1*), although this effect was primarily driven by activity in the left hippocampus (*SI Appendix, Fig. S2*). Given substantial prior evidence from the literature, we expanded the scope of our analysis to test power in 2 additional frequency bands: gamma (30 to 55 Hz), and high-frequency broadband

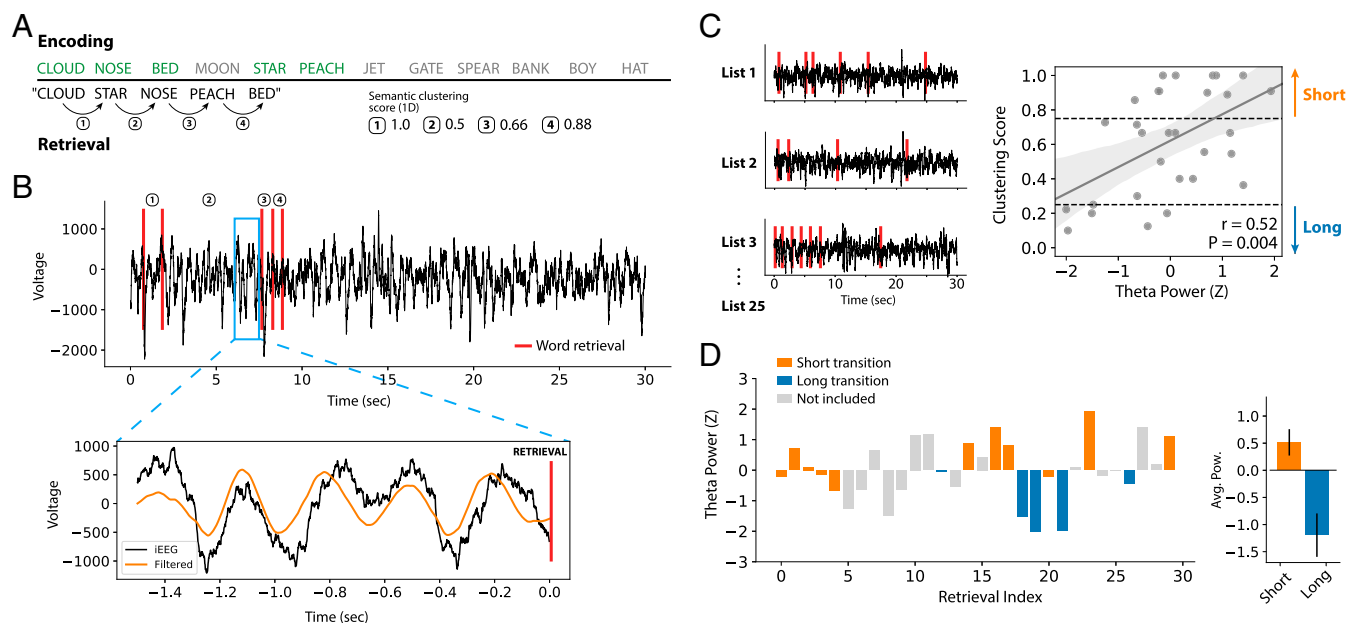


Fig. 3. Correlating preretrieval theta power with semantic transition distance. (A) In an example list, a subject recalls 5 of the original 12 words, making close and far transitions through a 1D semantic space. Transition distances are measured by a percentile rank called the clustering score, with 1 indicating the closest-possible transition through a given semantic space (*SI Appendix, SI Materials and Methods*). (B) Theta power was extracted in 1-s intervals immediately prior to each retrieval event. An ~ 4 -Hz theta oscillation preceding the third recall is highlighted, here measured from an electrode in the CA1 subfield of the hippocampus. (C) Theta power is measured prior to all retrieval events across an experimental session and z scored. Z-scored theta power was correlated with semantic transition distance (i.e., the clustering score), demonstrating a significant positive relationship for this subject ($r = 0.52$, $P = 0.004$). Each point represents a recall event in the example session. (D) To further quantify the magnitude of theta power change, interword distances were binned into "short" (>0.75 clustering score) and "long" (<0.25), and z-scored theta power was averaged for all retrieval events that fell in each bin. In this subject, theta power increased above average for close transitions and fell below average for far transitions.

(HFB) (70 to 150 Hz). Findings in humans and animals implicate MTL gamma oscillations in memory processing (44–46), and HFB has been used as a general marker of neural activation (47). In neither band did we find a significant correlation between preretrieval power and subsequent transition distance (1-sample t test, $P > 0.10$).

A recent study by Herweg et al. (17) noted that theta power in the PHG predicted the recall of items that occurred nearby in a spatial layout. We therefore asked whether PHG theta power demonstrated the same relationship with interword transition distances. Unlike our finding in the hippocampus, PHG theta power was not correlated with ensuing transition distances in a 1D space (1-sample t test, $t(67) = 0.35$, $P = 0.73$), nor did we find a significant relationship in our 2 higher-frequency bands.


Given the strong behavioral tendency for subjects to recall words in a temporally organized manner, we next asked whether distances in a temporal "space" correlated with preretrieval power. To do this, we applied the same analysis procedure as for semantic distances, but instead used a measure of temporal transition distance instead of semantic transition distance; i.e., if a subject makes a recall transition between 2 contiguous words from the original list, the clustering score is 1. We found no significant relationship between preretrieval power and temporal transition distance, in any frequency band or region (Fig. 4B, 1-sample t test, $P > 0.10$).

Our focus on preretrieval measures of spectral power left open the possibility that power more strongly correlated with transition distance during the transition itself. To test this, we correlated 1-s windows of prevocalization spectral power with the transition distances covered during that same time, as schematized in Fig. 4C. Using this measure of power–distance correlation, we found no relationship between power in any band and semantic or temporal transition distance (all $P > 0.10$). This finding suggests that hippocampal theta power rises only prior

to retrieval of semantically related items. Moreover, regression analysis revealed that preretrieval hippocampal theta power predicts subsequent semantic transitions independent of the time between recalls or the output position of a recalled item (*SI Appendix, Fig. S3*). Furthermore, to ensure that overlapping spectral analysis windows between successive recall events did not affect our findings, we reanalyzed the dataset while excluding any transitions for which recalls occurred within 1 s of each other (representing 19.6% of all transitions). We again found that hippocampal theta was correlated with semantic transition distance (Fig. 4D, 1-sample t test, $t(67) = 1.92$, $P = 0.06$), while all other relationships did not meet significance (all $P > 0.1$; *SI Appendix, Fig. S4*).

By correlating transition distances in 1D semantic and temporal spaces, we demonstrated that 1) the degree of hippocampal theta power prior to a retrieval event codes for the ensuing semantic transition distance, suggesting that theta serves to reinstantiate semantic context, and 2) short transitions in the time dimension are not significantly predicted by neural activity in our prespecified bands. These findings align with the general idea that theta oscillations support the encoding and retrieval of information organized in a cognitive map. However, it is unclear why the time dimension—shown to powerfully correlate with retrieval behavior—does not have a strong neural correlate (see *Reconciling the Absence of Temporally Related Theta Activity*).

Multidimensional Semantic Spaces. We next asked whether distances in higher-dimensional semantic spaces were also predicted by preretrieval theta power. Asking this question lets us assess how the brain compresses semantic information—for example, Does theta power also correlate with interitem distances in a 2D or 3D semantic layout? Our use of the word2vec semantic representation enabled us to directly test this hypothesis.



Solomon et al.

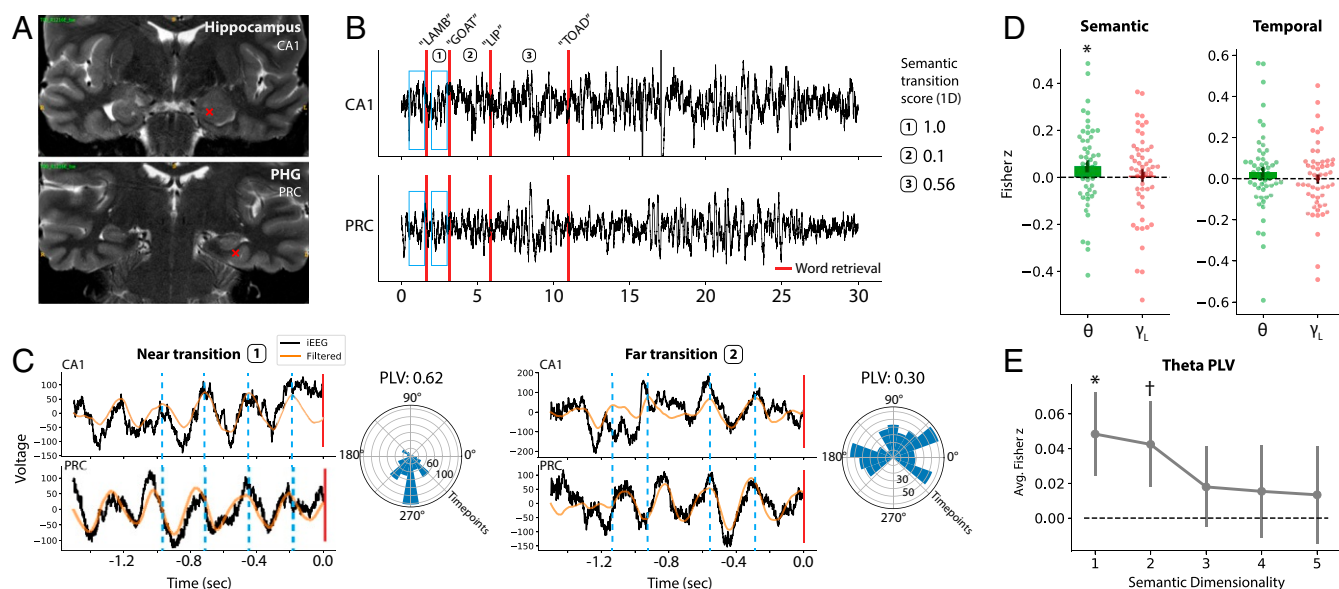


Fig. 5. Assessing preretrieval theta coupling between the PHG and hippocampus. (A) In an example subject, phase locking was assessed between electrodes localized to the hippocampus (CA1) and PHG (perirhinal cortex [PRC]), indicated as red crosses. (B) As in the power analyses, LFP is extracted from each electrode, and theta coupling is assessed in the 1-s intervals leading up to vocalized recalls. A score is assigned to each transition, capturing the closeness of successive recalls in semantic or temporal dimensions. (C) Phase locking is measured by assessing the consistency of theta phase between the electrodes, for each transition. (Left) A close semantic transition is depicted, where the peak of the CA1 oscillation (dashed lines) consistently aligns with the falling phase of the PRC oscillation. (Right) A far semantic transition is shown, with inconsistent phase alignment. The concentration of phase differences over time is quantified as the PLV, visualized as polar histograms. For the example trials shown, the close semantic transition has a higher PLV (0.62) than the far transition (0.30). (D) Hippo-PHG theta coupling mediates semantic search. (Left) Hippocampal-PHG theta coupling significantly correlated with ensuing semantic transition distance in a 1D space (1-sample t test, $t(51) = 2.07$, $P = 0.044$). No effect was observed for phase locking in the gamma band. (Right) Phase locking was not predictive of temporal transition distance in either frequency band ($P > 0.05$). (E) Theta phase locking was marginally significantly correlated with 2D semantic transition ($t(51) = 1.73$, $P = 0.089$), in addition to 1D transitions, as noted in D. * $P < 0.05$, † $P < 0.1$. Error bars show ± 1 SEM.

has been hypothesized to facilitate synaptic plasticity and gate the flow of information between brain regions, by providing a consistent phase reference across regions that process different aspects of an animal's experience (50–52). Having shown that recall behavior and local theta power correlate with distances in 1- to 2-dimensional semantic spaces, we next set out to ask whether interregional theta connectivity also correlated with the retrieval of semantic or temporal context.

Fig. 5 outlines our procedure in an example subject. We used the phase-locking value (PLV) (53), a measure of the consistency of phase differences between 2 electrodes, to assess the degree of theta connectivity between PHG and hippocampus. Here, theta-frequency phases are extracted from 1-s intervals prior to the retrieval of each word, for an electrode in CA1 and an electrode in perirhinal cortex (PRC) (Fig. 5A and B). The PLV is computed on the phase differences between the 2 regions, prior to each retrieval. Fig. 5C illustrates how CA1-PRC theta phases are highly consistent prior to a short semantic transition (1.0 clustering score), since the theta peaks in CA1 always correspond to the falling phase of the oscillation in PRC. Prior to a far semantic transition (0.1 clustering score), the theta peaks in CA1 do not align consistently with the ongoing PRC oscillation. We also note that, to avoid assessing phase locking between bipolar pairs that span the hippocampus and PHG but share a common monopolar contact, our connectivity analyses utilized an MTL-average reference (*SI Appendix, SI Materials and Methods*).

We measured hippocampal-PHG theta phase locking for all subjects with electrodes in both areas ($n = 52$) and, as earlier, correlated those values with the subsequent semantic or temporal transition distance. Given the possibility of intra-MTL gamma connectivity (54, 55), we further included that band in this analysis. Across all subjects, we found a significant corre-

lation between hippocampal-PHG theta phase locking and 1D semantic transition distance (1-sample t test, $t(51) = 2.07$, $P = 0.044$; Fig. 5D) and no effect in the gamma band ($t = 0.38$, $P = 0.71$). Extending the analysis to higher semantic dimensions, we found a marginally significant correlation between theta coupling and semantic distances in a 2D space ($t(51) = 1.73$, $P = 0.089$; Fig. 5E), but no others. Results were qualitatively similar when regressing out power from PLV values, to correct for possible signal-to-noise changes with altered power (*SI Appendix, Fig. S6*). These data suggest that theta coupling between the PHG and hippocampus mediates successful search through a 1D and possibly 2D semantic space.

Does preretrieval hippocampal-PHG theta coupling support the reinstatement of temporal context? Theta coupling was only weakly correlated with temporal transition distance, but did not reach a statistical threshold (1-sample t test, $t(52) = 1.20$, $P = 0.23$; Fig. 5D). This result mirrored our earlier finding that local theta power was not correlated with temporal transition distance, even though time was the most prominent organizing factor in recall behavior.

Reconciling the Absence of Temporally Related Theta Activity. Our finding that MTL theta power and connectivity do not relate to temporal clustering came as a surprise—temporal relations are the most prominent organizing principle of verbal free recall, and recent studies support the idea that neurons in the MTL encode aspects of time (56). If time (or ordinal position) is a key dimension along which humans organize episodic memories, and theta rhythms are the physiological substrate of episodic associations, why are temporally clustered recalls not associated with increases in MTL theta activity?

We hypothesized that our temporal theta-correlation analysis was affected by any variables that could limit the dynamic range

of contextual coding during recall. For example, subjects often exclusively retrieve the first few list items as a group, suggesting that these words are associated with a general list context and that such lists do not reflect a dynamic range of temporal contexts. To account for this, we applied a simple modification to the temporal transition analysis presented in Fig. 4, noting that lists with a small number of words recalled (i.e., 4 or less) are more likely to lack temporal dynamic range, while high-performance lists with many recalls are more likely to capture meaningful variability across temporal lags (or “distances”). We therefore computed the correlation between theta power and temporal transition distance at varying list-performance thresholds and compared our findings to the correlations based on the full dataset (we note that, in all analyses in this paper, lists with 3 or fewer recalls were excluded). To ensure that we captured a range of temporal lags, we binned clustering scores into long and short and directly compared the 2 groups, as was demonstrated in prior analyses (Fig. 3*D* and *SI Appendix*, Fig. S5).

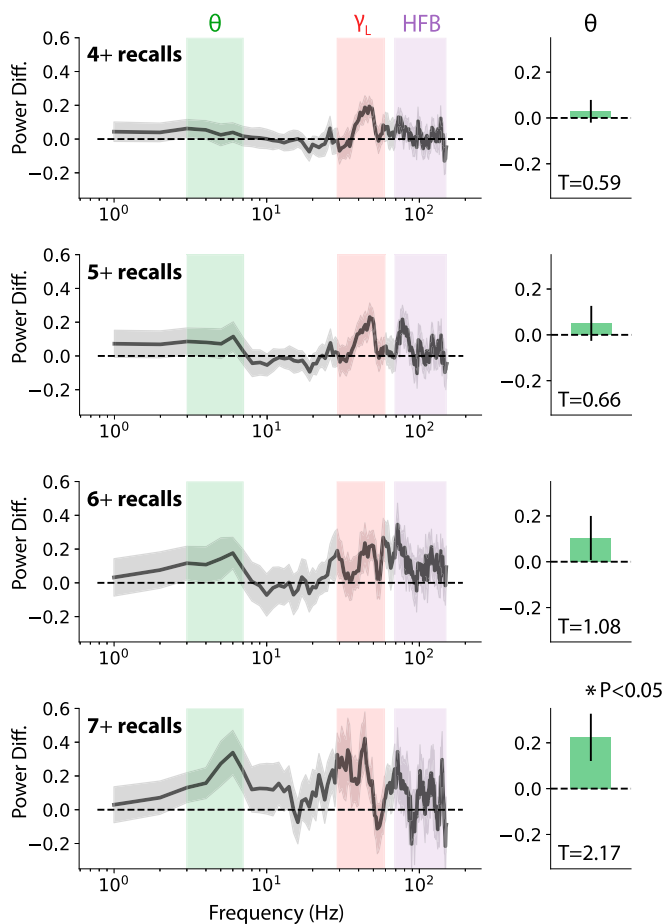


Fig. 6. Reconciling the absence of temporally associated hippocampal theta. To ask whether restricted temporal dynamic range was contaminating measures of time-associated theta power, recall periods were filtered according to the number of recalled items per list. The association between theta power and temporal transition distance is shown for 4 different list-performance thresholds, among the 22 subjects with sufficient data at the highest threshold. The theta power difference between short and long temporal transitions (Fig. 3*D*) is significantly greater than zero for high-performance lists (7-recall threshold; $t(21) = 2.17$, $P = 0.041$) and increases monotonically at each threshold. The power difference spectra, reflecting the degree to which theta power is associated with short vs. long temporal transitions, reveal an emerging spectral peak in the theta band, indicating a significant association between theta power and temporal transition distance. * $P < 0.05$, error bars show ± 1 SEM.

At higher performance thresholds, we noted a shift toward higher correlations between preretrieval theta power and temporal transition distance (Fig. 6). In the 22 subjects with sufficient data in high-performance lists (7 or more recalls per list), higher theta power was significantly associated with shorter temporal transition distances (1-sample t test, $t(21) = 2.17$, $P = 0.041$; Fig. 6), although the association was not apparent at lower performance thresholds (4-recall threshold; $t(21) = 0.77$, $P = 0.44$). On a per-subject basis, filtering for list performance yielded a marginally significant difference between the theta effect in all-list versus high-performance conditions (paired t test, $t(21) = 1.79$, $P = 0.087$). As was demonstrated for semantic transitions, the correlation spectrum revealed a qualitative peak in the theta band (Fig. 6, *Bottom* row).

By focusing on lists with a high number of recalls, and thereby mitigating potential confounds in our temporal correlation analysis, we found evidence for our hypothesized relationship between hippocampal theta and temporal transition distance. This result mirrored our earlier finding of hippocampal theta associated with close semantic transitions, and since word semantics were unrelated to list structure, no corrections were needed to observe the effect. We note that 22 subjects had sufficient high-performance lists to enter this analysis, relative to the 70 subjects included in the semantic analysis. This introduces a potential confound of a subject's general memory performance or task ability and weakens our statistical power. Nonetheless, a performance filter demonstrated the strong possibility that an underlying temporal-theta effect was present, but not directly observable, in our earlier analyses.

Discussion

We set out to determine whether an established neural signature of spatial navigation also applied to navigation of arbitrary representational spaces, in the form of episodic memory search. To do this, we had human neurosurgical patients perform a verbal free-recall task and asked whether MTL theta power and connectivity correlated with pairwise distance in semantic and temporal representational spaces. Across all subjects, we found a reliable correlation between distances in a 1D semantic space and theta activity. Increases in theta power/connectivity were not readily associated with short temporal transitions, although the effect emerged when we considered lists with higher performance, and thereby ensured a dynamic range of temporal context reinstatement.

These results support the notion that MTL theta oscillations support both spatial navigation and episodic memory search, as predicted by the cognitive mapping theory of MTL function. Hippocampal theta has been associated with spatial navigation in rodents (see ref. 57 for a review) and humans (14–16) and is seen during spontaneous retrieval of information encountered in a spatial layout (17). Findings during nonspatial memory tasks have been decidedly more mixed, with many studies showing a decrease in MTL theta power during encoding and retrieval (e.g., refs. 34 and 35). Here, we reconcile these results by using statistical contrasts of spatial or temporal clustering, instead of successful-versus-unsuccessful memory (e.g., the SME). This procedure inherently controls for nonspecific attentional or arousal states, since successful recalls occur in both clustered and nonclustered conditions.

As predicted by models of MTL function (58), a positive association between human episodic memory and hippocampal theta becomes apparent when a more nuanced statistical contrast is employed. Not only does this harmonize results between the spatial and episodic memory literatures, it also suggests the presence of a confound in the popular SME (or retrieval analogues): A nonspecific task contrast between successful and unsuccessful memory operations includes significant brain activity unrelated to contextual memory processing, perhaps reflecting visual

attention, executive control, or general task engagement. Unfortunately, these processes might cause broad decreases in theta power—for reasons still unknown—obscuring the theta rhythms that underlie the formation and retrieval of cognitive maps.

Several prior studies have addressed similar questions, but did not typically report strong hippocampal theta effects. Long and Kahana (59) examined neural activity during word encoding by comparing subsequently temporally clustered vs. not-clustered recalls, finding no increase in hippocampal theta power. However, theta power was not significantly decreased between the 2 conditions either, unlike the strong theta decrease seen in the classic SME reported in that same study. This suggests a possible underlying increase in hippocampal theta that was obscured by residual noncontextual processes, which are especially strong during initial encoding. A related scalp electroencephalography (EEG) study by Long and Kahana (60) examined encoding-related activity for semantically clustered retrievals and reported a trend toward increases in low-theta power—a finding that may have been significant with direct recordings of hippocampal LFPs. Future work should thoroughly examine whether intracranial measures of hippocampal activity show evidence of theta power increases during encoding, for words that are subsequently clustered in temporal or semantic space.

In this paper, we found that theta power and connectivity were only significantly associated with the first principal word2vec dimension for each word list. Additional semantic dimensions were found to explain recall clustering at the behavioral level, but beyond the second dimension, the marginal explained variance fully accounts for this increase. In this sense, our behavioral and neurophysiological results align: One semantic dimension accounts for the majority of 1) recall clustering behavior and 2) MTL theta activity. However, we suspect that these low-dimensional findings are partially a function of the list structure in our free-recall task. Since PCA was performed on the 12 words in each list, high dimensions are likely to contain substantial noise and therefore not correlate with hippocampal theta. Taken together, semantic memories are surely represented in the brain by more than 1 dimension, but theta activity reliably correlates with the most salient organizing feature of word lists.

The findings here show that the same theta signature of spatial navigation occurs during spontaneous episodic recall of word items. We therefore reconceptualize free recall as navigation of a representational space with time and semantics as the most salient dimensions. The idea of a 2D semantic–temporal space suggests that other signatures of 2D spatial navigation, such as hexadirectional modulation of neural activity, may also be detectable during episodic recall. Two groups recently showed that theta power in the human entorhinal cortex is hexadirectionally modulated during virtual navigation (61, 62). In the fMRI literature, it has been shown that hippocampal activity encodes for spatial (63) and abstract representational distances (64) and that hexadirectional modulation of the BOLD signal is observed during mental travel through abstract or imagined

spaces (65, 66). Several recent studies have suggested that theta power may itself correlate with spatial distances (14, 16, 17), just as we showed a correlation between theta and representational distances in this paper. If free recall truly constitutes a form of navigation through a semantic–temporal space, might we observe hexadirectionally modulated theta during episodic retrieval? And might future studies find further evidence that theta power in humans correlates with traversed navigational distances?

Neuroscientists have long sought to understand how the medial temporal lobe simultaneously supports spatial navigation and episodic memory. This paper is another link in a growing chain of evidence that says memory and navigation are fundamentally supported by the same underlying cognitive process: the formation of associations between items, reflected as distances in arbitrary representational spaces. Until recently, spatial layouts have been the most obvious way to study these spaces. Here we demonstrated that temporal and semantic distances between word items were also captured by a classic hippocampal signature of navigation, lending strong support to the idea that the MTL performs domain-general cognitive mapping.

Materials and Methods

We recorded intracranial EEG from 189 adult patients undergoing seizure monitoring, using depth electrodes placed in the medial temporal lobe. We asked subjects to perform a verbal free-recall task in which lists of 12 nouns were visually presented, and after a brief arithmetic distractor task, subjects vocalized as many words as possible from the prior list. Using spectral methods, we correlated power and connectivity in subregions of the MTL with the semantic or temporal distance between items vocalized during recall. Across subjects, we tested whether power–distance or power–connectivity correlations differed significantly from zero in prespecified frequency bands, including theta (4 to 8 Hz), gamma (30 to 55 Hz), and high-frequency broadband (70 to 150 Hz). Prior to data collection, our research protocol was approved by the Institutional Review Board at the University of Pennsylvania, Thomas Jefferson University Hospital, and the University of Texas Southwestern, and informed consent was obtained from each participant. See *SI Appendix* for further details.

Data and Code Availability. Raw electrophysiological data used in this study are freely available at <http://memory.psych.upenn.edu/Electrophysiological>. Data. Analysis code and processed data are available on a public interactive Jupyter Notebook (<https://notebooks.azure.com/esolo/projects/theta-semantics>). Other custom processing scripts are available by request to E.A.S.

ACKNOWLEDGMENTS. We thank Blackrock Microsystems and Medtronic for providing neural recording equipment. This work was supported by National Institutes of Health Grant MH55687, as well as the Defense Advanced Research Projects Agency (DARPA) Restoring Active Memory program (Cooperative Agreement N66001-14-2-4032). We are indebted to all patients who have selflessly volunteered their time to participate in our study. The views, opinions, and/or findings contained in this material are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the US Government. We also thank Dr. Nora Herweg for providing valuable feedback on this work.

1. N. Burgess, E. A. Maguire, J. O'Keefe, The human hippocampus and spatial and episodic memory. *Neuron* **35**, 625–641 (2002).
2. J. L. S. Bellmund, P. Gärdenfors, E. I. Moser, C. F. Doeller, Navigating cognition: Spatial codes for human thinking. *Science* **362**, eaat6766 (2018).
3. J. O'Keefe, J. Dostrovsky, The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175 (1971).
4. D. Schiller et al., Memory and space: Towards an understanding of the cognitive map. *J. Neurosci. Off. J. Soc. Neurosci.* **35**, 13904–13911 (2015).
5. H. Eichenbaum, A cortical–hippocampal system for declarative memory. *Nat. Rev. Neurosci.* **1**, 41–50 (2000).
6. L. R. Squire, S. Zola-Morgan, The medial temporal lobe memory system. *Science* **253**, 1380–1386 (1991).
7. L. R. Squire, J. T. Wixted, The cognitive neuroscience of human memory since H.M. *Annu. Rev. Neurosci.* **34**, 259–288 (2011).
8. N. A. Herweg, M. J. Kahana, Spatial representations in the human brain. *Front. Hum. Neurosci.* **12**, 297 (2018).
9. E. I. Moser, E. Kropff, M.-B. Moser, Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.* **31**, 69–89 (2008).
10. A. D. Ekstrom et al., Cellular networks underlying human spatial navigation. *Nature* **425**, 184–188 (2003).
11. J. F. Miller et al., Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. *Science* **342**, 1111–1114 (2013).
12. J. Jacobs et al., Direct recordings of grid-like neuronal activity in human spatial navigation. *Nat. Neurosci.* **16**, 1188–1190 (2013).
13. V. D. Bohbot, M. S. Copara, J. Gotman, A. D. Ekstrom, Low-frequency theta oscillations in the human hippocampus during real-world and virtual navigation. *Nat. Commun.* **8**, 14415 (2017).
14. L. K. Vass et al., Oscillations go the distance: Low-frequency human hippocampal oscillations code spatial distance in the absence of sensory cues during teleportation. *Neuron* **89**, 1180–1186 (2016).
15. A. D. Ekstrom et al., Human hippocampal theta activity during virtual navigation. *Hippocampus* **15**, 881–889 (2005).

16. D. Bush *et al.*, Human hippocampal theta power indicates movement onset and distance travelled. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 12297–12302 (2017).
17. N. A. Herweg *et al.*, Reactivated spatial context guides episodic recall. *bioRxiv*: 399972v2 (13 September 2018).
18. J. Miller *et al.*, Lateralized hippocampal oscillations underlie distinct aspects of human spatial memory and navigation. *Nat. Commun.* **9**, 2423 (2018).
19. G. Buzsáki, E. I. Moser, Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat. Neurosci.* **16**, 130–138 (2013).
20. G. Buzsáki, S. L.-W. Leung, C. H. Vanderwolf, Cellular bases of hippocampal EEG in the behaving rat. *Brain Res. Rev.* **6**, 139–171 (1983).
21. J. Huxter, N. Burgess, J. O'Keefe, Independent rate and temporal coding in hippocampal pyramidal cells. *Nature* **425**, 828–832 (2003).
22. J. O'Keefe, M. L. Recce, Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**, 317–330 (1993).
23. A. Johnson, A. D. Redish, Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* **27**, 12176–12189 (2007).
24. A. M. Wikenheiser, A. D. Redish, Hippocampal theta sequences reflect current goals. *Nat. Neurosci.* **18**, 289–294 (2015).
25. D. J. Foster, M. A. Wilson, Hippocampal theta sequences. *Hippocampus* **17**, 1093–1099 (2007).
26. U. Rutishauser, I. B. Ross, A. N. Mamelak, E. M. Schuman, Human memory strength is predicted by theta-frequency phase-locking of single neurons. *Nature* **464**, 903–907 (2010).
27. A. J. Watrous, J. Miller, S. E. Qasim, I. Fried, J. Jacobs, Phase-tuned neuronal firing encodes human contextual representations for navigational goals. *eLife* **7**, e32554 (2018).
28. J. Jacobs, M. J. Kahana, A. D. Ekstrom, I. Fried, Brain oscillations control timing of single-neuron activity in humans. *J. Neurosci.* **27**, 3839–3844 (2007).
29. J.-J. Lin *et al.*, Theta band power increases in the posterior hippocampus predict successful episodic memory encoding in humans. *Hippocampus* **27**, 1040–1053 (2017).
30. B. C. Lega, J. Jacobs, M. Kahana, Human hippocampal theta oscillations and the formation of episodic memories. *Hippocampus* **22**, 748–761 (2012).
31. N. A. Herweg *et al.*, Theta-alpha oscillations bind the hippocampus, prefrontal cortex, and striatum during recollection: Evidence from simultaneous EEG-fMRI. *J. Neurosci.* **36**, 3579–3587 (2016).
32. N. W. Morton *et al.*, Category-specific neural oscillations predict recall organization during memory search. *Cerebr. Cortex* **23**, 2407–2422 (2013).
33. J. A. Greenberg, J. F. Burke, R. Haque, M. J. Kahana, K. A. Zaghloul, Decreases in theta and increases in high frequency activity underlie associative memory encoding. *NeuroImage* **114**, 257–263 (2015).
34. J. F. Burke *et al.*, Synchronous and asynchronous theta and gamma activity during episodic memory formation. *J. Neurosci.* **33**, 292–304 (2013).
35. E. A. Solomon *et al.*, Widespread theta synchrony and high-frequency desynchronization underlies enhanced cognition. *Nat. Commun.* **8**, 1704 (2017).
36. E. A. Solomon *et al.*, Dynamic theta networks in the human medial temporal lobe support episodic memory. *Curr. Biol.* **29**, 1100–1111.e4 (2019).
37. Y. Ezzyat *et al.*, Direct brain stimulation modulates encoding states and memory performance in humans. *Curr. Biol.* **27**, 1251–1258 (2017).
38. M.-C. Fellner *et al.*, Spectral fingerprints or spectral tilt? Evidence for distinct oscillatory signatures of memory formation. *PLoS Biol.* **17**, e3000403 (2019).
39. M. W. Howard, M. J. Kahana, When does semantic similarity help episodic retrieval? *J. Mem. Lang.* **46**, 85–98 (2002).
40. M. J. Kahana, Associative retrieval processes in free recall. *Mem. Cogn.* **24**, 103–109 (1996).
41. T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space. *arXiv:1301.3781* (16 January 2013).
42. G. A. Miller, WordNet: A lexical database for English. *Commun. ACM* **38**, 39–41 (1995).
43. T. K. Landauer, S. T. Dumais, A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychol. Rev.* **104**, 211–240 (1997).
44. L. L. Colgin, E. I. Moser, Gamma oscillations in the hippocampus. *Physiology* **25**, 319–329 (2010).
45. 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. *Cerebr. Cortex* **26**, 268–278 (2016).
46. P. B. Sederberg *et al.*, Hippocampal and neocortical gamma oscillations predict memory formation in humans. *Cerebr. Cortex* **17**, 1190–1196 (2006).
47. K. J. Miller *et al.*, Cortical activity during motor execution, motor imagery, and imagery-based online feedback. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 4430–4435 (2010).
48. A. J. Watrous, N. Tandon, C. R. Conner, T. Pieters, A. D. Ekstrom, Frequency-specific network connectivity increases underlie accurate spatiotemporal memory retrieval. *Nat. Neurosci.* **16**, 349–356 (2013).
49. J. Fell *et al.*, Rhinal-hippocampal theta coherence during declarative memory formation: Interaction with gamma synchronization? *Eur. J. Neurosci.* **17**, 1082–1088 (2003).
50. J. Fell, N. Axmacher, The role of phase synchronization in memory processes. *Nat. Rev. Neurosci.* **12**, 105–118 (2011).
51. J. E. Lisman, O. Jensen, The theta-gamma neural code. *Neuron* **77**, 1002–1016 (2013).
52. A. Sirota *et al.*, Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron* **60**, 683–697 (2008).
53. J. P. Lachaux, E. Rodriguez, J. Martinerie, F. J. Varela, Measuring phase synchrony in brain signals. *Hum. Brain Mapp.* **8**, 194–208 (1999).
54. J. Fell *et al.*, Human memory formation is accompanied by rhinal-hippocampal coupling and decoupling. *Nat. Neurosci.* **4**, 1259–1264 (2001).
55. L. L. Colgin *et al.*, Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* **462**, 353–357 (2009).
56. H. Eichenbaum, Time cells in the hippocampus: A new dimension for mapping memories. *Nat. Rev. Neurosci.* **15**, 732–744 (2014).
57. G. Buzsáki, Theta oscillations in the hippocampus. *Neuron* **33**, 325–340 (2002).
58. G. Buzsáki, Theta rhythm of navigation: Link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus* **15**, 827–840 (2005).
59. N. M. Long, M. J. Kahana, Successful memory formation is driven by contextual encoding in the core memory network. *NeuroImage* **119**, 332–337 (2015).
60. N. M. Long, M. J. Kahana, Modulation of task demands suggests that semantic processing interferes with the formation of episodic associations. *J. Exp. Psychol. Learn. Mem. Cogn.* **43**, 167–176 (2017).
61. S. Maidenbaum, J. Miller, J. M. Stein, J. Jacobs, Grid-like hexadirectional modulation of human entorhinal theta oscillations. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 10798–10803 (2018).
62. D. Chen *et al.*, Hexadirectional modulation of theta power in human entorhinal cortex during spatial navigation. *Curr. Biol.* **28**, 3310–3315.e4 (2018).
63. L. K. Morgan, S. P. Macevoy, G. K. Aguirre, R. A. Epstein, Distances between real-world locations are represented in the human hippocampus. *J. Neurosci.* **31**, 1238–1245 (2011).
64. S. Theves, G. Fernandez, C. F. Doeller, The hippocampus encodes distances in multidimensional feature space. *Curr. Biol.* **29**, 1226–1231.e3 (2019).
65. A. J. Horner, J. A. Bisby, E. Zotow, D. Bush, N. Burgess, Grid-like processing of imagined navigation. *Curr. Biol.* **26**, 842–847 (2016).
66. A. O. Constantinescu, J. X. O'Reilly, T. E. J. Behrens, Organizing conceptual knowledge in humans with a gridlike code. *Science* **352**, 1464–1468 (2016).