Distinct cortical systems reinstate content and context information during memory search

James E. Kragel,[1] Gregory A. Worrell,[2] Michael R. Sperling,[3] Robert E. Gross,[4] Bradley C. Lega,[5] Barbara C. Jobst,[6] Sameer A. Sheth,[7] Kareem A. Zaghloul,[8] Joel M. Stein,[9] Michael J. Kahana[1]

[1]*Department of Psychology, University of Pennsylvania, Philadelphia PA 19104, USA*
[2]*Department of Neurology, Mayo Clinic, Rochester MN 55905, USA*
[3]*Department of Neurology, Thomas Jefferson University, Philadelphia PA 19107, USA*
[4]*Department of Neurosurgery, Emory School of Medicine, Atlanta GA 30322, USA*
[5]*Department of Neurosurgery, University of Texas Southwestern, Dallas TX 75390, USA*
[6]*Department of Neurology, Dartmouth-Hitchcock Medical Center, Lebanon NH 03756, USA*
[7]*Department of Neurosurgery, Columbia University Medical Center, New York NY 10032, USA*
[8]*Surgical Neurology Branch, National Institutes of Health, Bethesda MD 20814, USA*
[9]*Department of Radiology, Hospital of the University of Pennsylvania, Philadelphia PA 19104, USA*

Running title: Content and context reinstatement

Corresponding Author:
Michael J. Kahana
University of Pennsylvania, Department of Psychology
Suite 263
Stephen A. Levin Building
Philadelphia, PA 19104
e-mail: kahana@psych.upenn.edu
phone: (215) 746–3501
fax: (215) 746–6848

Conflict of Interest: M.J.K. has started a company, Nia Therapeutics, LLC ("Nia"), intended to develop and commercialize brain stimulation therapies for memory restoration. He holds more than 5% equity interest in Nia.

Acknowledgments: We thank Medtronic and Blackrock Microsystems for providing neural recording equipment. This work was supported by the DARPA Restoring Active Memory (RAM) 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 U.S. Government.

Abstract

Episodic recall depends upon the reinstatement of cortical activity present during the formation of a memory. We identified dissociable cortical networks, identified via functional connectivity, that reinstate semantic content and temporal context of previously studied stimuli during free recall. Reinstatement in each of these networks predicted memory organization, demonstrating how specialized neural networks enable the human brain to target specific memories.

The representation of episodic context and semantic content have long been acknowledged as separable components of memory[1]. Based upon converging evidence from fMRI and electrophysiological studies of human memory and parallel investigations in animal models, an emerging consensus has emerged wherein discrete cortical systems support memory for the semantic meaning (content) and the specific time and place (context) of events[2–6]. The posterior medial (PM) network includes multiple cortical sites implicated in representing contextual information, including parahippocampal, retrosplenial, and posterior parietal cortices[7]. The anterior temporal (AT) network contains structures critical for semantic and conceptual memory[8,9], including the ventral temporal pole and perirhinal cortices[10]. During episodic retrieval, neurons in the hippocampus and surrounding medial temporal lobe cortex fire in the moments leading up to the subjective experience of remembering[11]. The firing of these neurons is tuned to different attributes of our memories, with neural activity increasing when memories contain either specific content (e.g., a specific person[11]) or context (e.g., a specific place or time[12]). Functional MRI[13,14] and intracranial EEG[15,16] studies link cortical reinstatement across the brain to these different types of memory. However, because this research typically focuses on either context or content reinstatement in isolation, our understanding of how these types of information are coded across cortical networks is limited. If these systems can independently drive hippocampal-dependent retrieval, they may reflect distinct cortical pathways that guide memory search.

To examine the contributions of cortico-hippocampal networks to memory recall, we utilized a computational modeling technique originally developed to predict patterns of brain activity based on the semantic content individual stimuli[17]. This method takes advantage of the sensitivity of neural signals to semantic attributes of presented items. By learning how activity in the brain is shaped by the semantic properties of items, constructed models can reliably decode the semantic content of stimuli from an observed pattern of brain activity. We extended this technique to develop context-based models of brain activity that were trained to predict patterns of brain activity based on the temporal context (i.e., at which point in time in the experiment) in which stimuli are presented. Here we test whether content- and context-based memory representations are reinstated within distinct cortical networks: the AT and PM networks. We recorded intracranial electroencephalograms (iEEG) from neurosurgical patients ($n = 69$) while they performed a free-recall task (Fig. 1a) with items drawn from 25 distinct categories (Fig. 1b). Recall sequences provided evidence of content and context reinstatement; subject's recalls were ordered based on the semantic content ($t_{68} = 30.66$, $p < 0.0001$) and temporal ordering ($t_{68} = 10.78$, $p < 0.0001$) of stimuli (Fig. 1c). This behavioral evidence suggests

that both the semantic content and temporal context of the studied items was used to guide memory search.

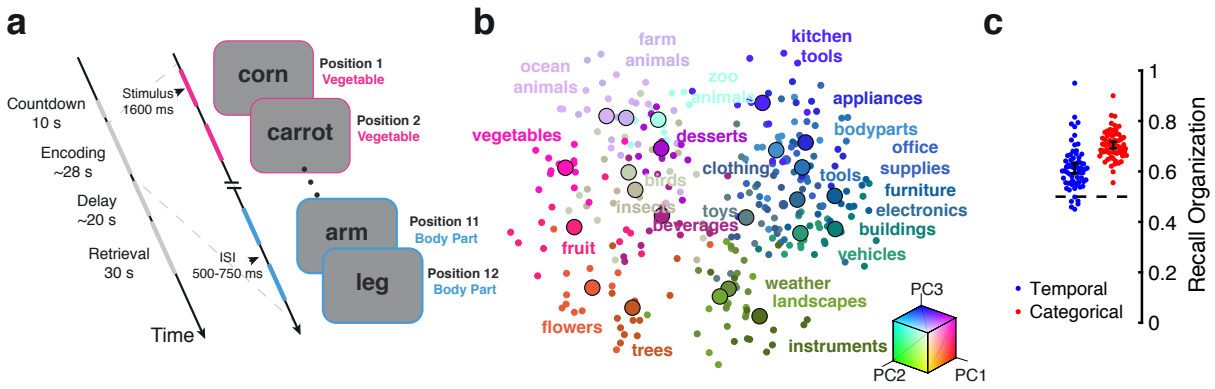

Figure 1 Task and behavioral results. **a**, Schematic of the free-recall task. Left: Task periods within a trial. Right: List structure and timing during the encoding period. Subjects completed 25 lists per experimental session. **b**, Category representations of stimuli. Each point denotes a word's semantic representation along the first three principal components of the word2vec[18] space. Larger points denote the midpoint of all stimuli within a category. Words with similar meaning are spatially proximal. **c**, Behavioral performance. Measures of temporal and categorical recall organization are compared to chance (dashed line). Error bars denote $\pm$ 1 s.e.m.

To examine how PM and AT networks contribute to these distinct types of reinstatement, we assigned cortical parcels to either network based on their intrinsic connectivity in an independent study utilizing resting-state fMRI (Fig. 2a; see Methods for details)[19]. For the AT network, we identified regions with connectivity to temporal polar cortex, revealing a distributed network including inferior prefrontal cortex, anterior temporal cortex, and inferior angular gyrus (Fig. 2a, green). For the PM network, we included regions with connectivity to the posterior angular gyrus (Fig. 2a, purple). Cortical regions included within the network included the posterior cingulate cortex, precuneus, parahippocampal cortex, and posterior parietal cortex.

To validate our assignment of recording sites to either network, we asked whether intrinsic connectivity was stronger within as opposed to between the two networks. We measured intrinsic connectivity by correlating spectral power across individual trials of the free-recall task. Such trial-by-trial variability in neural activity is thought to reflect endogenous fluctuations across neural networks, accounting for the common activation of intrinsic networks during cognitive tasks[20–22]. Because these estimates of functional connectivity can be spuriously elevated due to proximity of

recording sites[23], we used a bootstrap matching procedure to control for differences in the distance between contacts located within the same network or across both networks (see Methods). When comparing the strength of functional connections within the AT network to those between networks, we found significantly increased ($p < 0.05$, permutation FWER) connectivity in the theta band ($3 - 10$ Hz; Fig. 2a). Connections within the PM network exhibited significantly greater strength than connections between the two networks within the low beta band. These results demonstrate that fluctuations in neural activity are correlated across performance of the recall task, consistent with their description as intrinsic cortical networks.

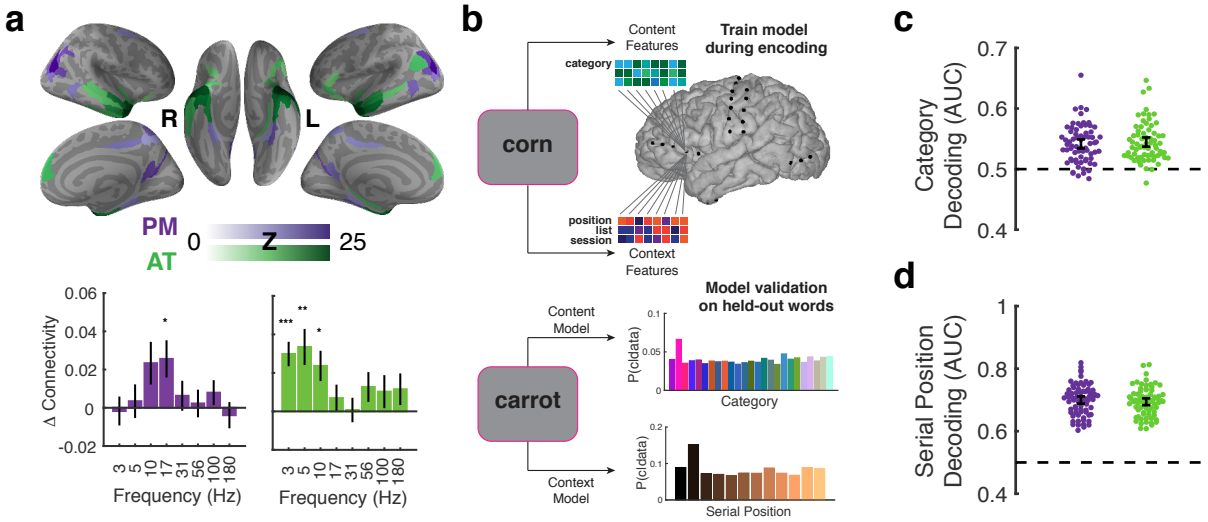

Figure 2 Network-based decoding content and context. **a**, rsfMRI connectivity identifies distinct cortical networks. Top: Functional connectivity (Fisher Z) to cortical targets of the PM and AT networks. Bottom: difference in connectivity of recording sites within vs. between the two networks. Frequency specificity was observed within each network ($p < 0.05$, FWER corrected). **b**, Modeling schematic. Ridge regression modeled evoked neural activity during stimulus encoding. Models were validated by decoding the category and serial position of held out stimuli during encoding. **c**, Reliable category decoding in the PM (one tailed sign-test, $p < 0.0001$) and AT (one tailed sign-test, $p < 0.0001$) networks during encoding. **d**, Serial position decoding for the PM (one tailed sign-test, $p < 0.0001$) and AT (one tailed sign-test, $p < 0.0001$) networks during encoding. $N = 69$ subjects. All tests are two-tailed unless otherwise noted. Error bars denote $\pm$ 1 s.e.m. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{**}p < 0.001$

We next characterized the presence of context and content information within the PM and AT networks based on

patterns of activity observed during the encoding phase of the free-recall task (Fig. 2b). We examined content information by modeling neural responses from the semantic features of each category (estimated using word2vec [18], see Methods). In the context model, neural activity was modeled based on temporal features (i.e., the serial position, list, and session) of each stimulus. This model identifies population-level neural activity that may code specific moments in time or serial positions, in a similar fashion to recently identified 'time cells' [24,25].

Prior to examining reinstatement, we first tested the ability of each model to decode these two aspects of memory as patients studied items on each list. We used a bootstrap procedure to estimate the amount of information contained within five randomly sampled recording sites from each network. Because electrodes typically target deep-brain structures in the MTL (as determined by clinical needs), coverage was biased to record structures in the AT network. This sampling procedure ensured that superior decoding performance did not result from differences in the amount of information available. While this approach provides a conservative estimate of the information contained within each network, it allowed us to assess the presence of memory representations distributed across each network. Applying this technique, we could reliably decode both the stimulus category (Fig. 1c) and serial position (Fig. 1d) of presented stimuli in both the AT and PM networks (median AUCs > 0.54, ps < 0.001). These results provide evidence that the connectivity of distinct cortical networks code for semantic and temporal information during encoding.

We leveraged the high temporal resolution of iEEG to identify when content or context reinstatement occurred within 20 msec intervals starting 900 msec prior to overt recall responses (Fig. 3a). Following our encoding analysis, content reinstatement was measured by the ability to decode the semantic category of the recalled item; context reinstatement was measured by the ability to decode when in the study list a recalled item was presented. In the moment preceding recall, we observed reinstatement of temporal information in the PM network and semantic information in the AT network. Moreover, these two forms of reinstatement were specific to each network (Fig. 3b; $p < 0.05$, FWER corrected). These results attribute reinstatement of content [16] and context [15] to distinct cortical networks and characterize the precise timing of reinstatement within each cortical system. Further, they illustrate that information that codes for specific moments in time (independent of content) are reinstated within the PM prior to recall.

Cortical reinstatement has been shown to predict how memory search unfolds [26,27], with cortical representations providing a top-down cue for the memory system to retrieve specific information. As such, it is possible that the

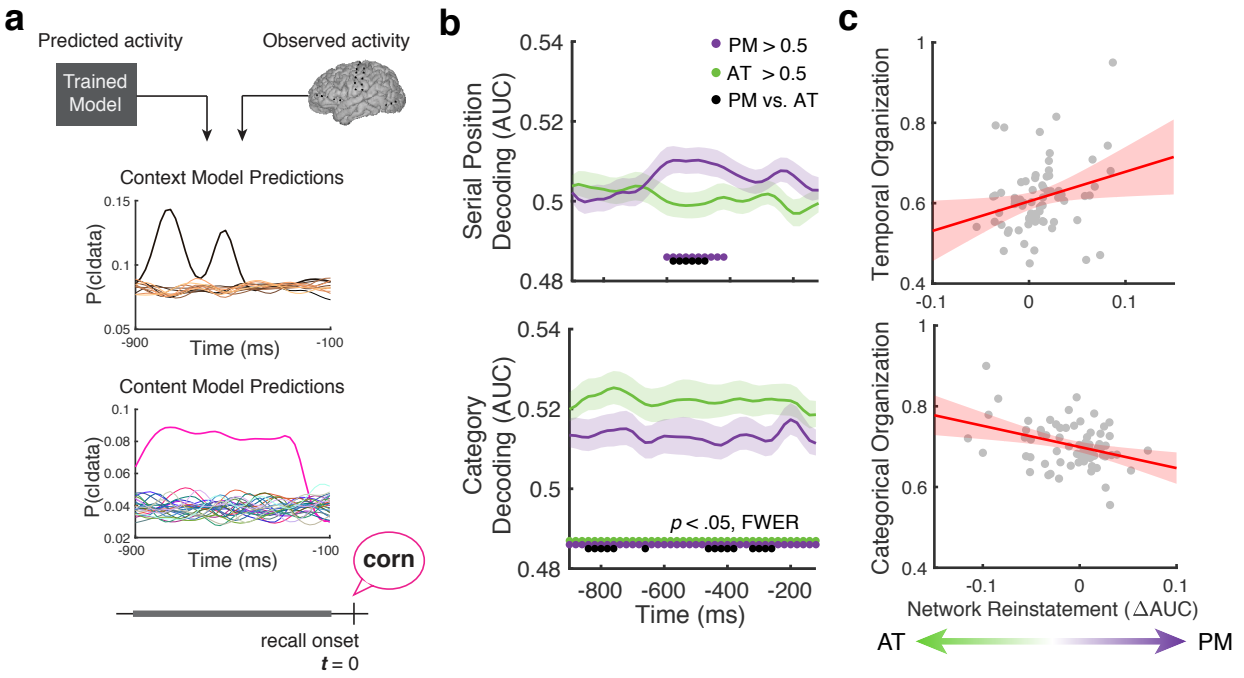

Figure 3 Reinstatement of content and context information during memory search. **a**, Reinstatement analysis schematic. Models trained during encoding predict the probably that each epoch precedes recall of an item from one of 12 serial positions (context model, top) or 25 categories (content model, bottom). **b**, Decoding the serial position (top) and semantic category (bottom) from reinstated patterns of neural activity within each network. **c**, Relationship between network reinstatement and organization of recall sequences along temporal and categorical dimensions. Shaded regions denote s.e.m. across ($N = 69$) subjects.

PM and AT networks are responsible for targeting stored memories with specific content- or context-based features. To test this hypothesis, we examined whether differences in reinstatement across the two networks predicted the tendency of subjects to organize their recall sequences along temporal or categorical dimensions. Consistent with prior work linking cortical reinstatement to memory organization [15,16], variability in the organization of subject's memory was predicted by cortical reinstatement during memory search (Fig. 3c). Greater reinstatement of context- based representations within the PM network tracked the tendency to consecutively recall items from nearby serial positions ($r_{67} = 0.26$, $p = 0.028$). Increased reinstatement within the AT network was associated with greater organization based on the categorical structure of the list ($r_{67} = -0.35$, $p = 0.003$). These results demonstrate that representations encoded in distinct cortical systems can differentially guide memory search, biasing how and what we remember.

We report evidence that the reinstatement of content- and context- based information occurs within distinct cortical networks. Furthermore, the quality of reinstatement within a given network was predictive of a subject's tendency to organize their memories along either temporal or categorical dimensions. These results help resolve conflicting viewpoints on how distinct neural representations contribute to memory search. It is widely held that episodic recall occurs when a retrieval cue converges on the hippocampal formation, prompting associative recall of individual memories[28]. The temporal context model[29] proposes that a slowly drifting context-based representation provides a mechanism with which we can target memories from a specific time. It has been proposed that specific cortical structures are responsible for such a representation, including lateral prefrontal cortex[30] and its projections to the hippocampus within the parahippocampal gyrus[31]. Despite evidence suggesting these structures are involved in representing temporal information[13,15,32], it has been unclear whether these cortical structures are the primary drivers of episodic recall. Our findings provide support for an alternative hypothesis regarding how neural representations guide memory search. One possibility is that semantic representations within the AT system can cue memories in a context-independent manner, prompting recall of memories with similar semantic attributes. If content-based representations can guide memory search in this fashion, one would expect individuals to organize their memories based on semantic content, rather than the temporal order in which it was learned. This is precisely the pattern of results that we observed. This hypothesis is further supported by evidence that category-specific activity in the inferior temporal cortex can be used as top-down signal to bias retrieval to targeted information[33]. Along these lines, the contextual information represented across the PM network allows us to focus memory to a specific time or place. However, this process is not ubiquitous to episodic recall. Reinstatement of activity within specialized neural networks enables the human brain to target memories with specific content or from a specific time.

Methods

Participants

69 patients (40 male) with medication-resistant epilepsy underwent neurosurgical procedures to implant intracranial electrodes (subdural, depth, or both) to determine epileptogenic regions. Data were collected at Dartmouth Medical Center (Hanover, NH), Emory University Hospital (Atlanta, Georgia), Hospital of the University of Pennsylvania (Philadelphia, PA), Mayo Clinic (Rochester, MN), Thomas Jefferson University Hospital (Philadelphia, PA), Columbia

University Medial Center (New York, NY), and University of Texas Southwestern Medical Center (Dallas, TX). Prior to data collection, the research protocol was approved by the institutional review board at each hospital. Informed consent was obtained from either the participant or their guardians. De-identified data and analysis code is available at `http://memory.psych.upenn.edu/Electrophysiological_Data`.

Free-recall task

Each subject performed a categorized free-recall task in which they studied a list of words with the intention to commit the items to memory. The task was performed at the bedside on a laptop, using PyEPL software[34]. Analog pulses were sent to available recording channels to enable alignment of experimental events with the recorded iEEG signal. Word presentation lasted for a duration of 1600 ms, followed by a blank inter-stimulus interval (ISI) of 750 to 1000 ms (see Fig. 1a). Each list contained items from three distinct categories (four items per category), with two same-category items presented consecutively. The total word pool consisted of 25 distinct categories, with individual items selected as prototypical items within each category[35]. Presentation of word lists was followed by a 20 s post-encoding delay. Subjects performed an arithmetic task during the delay in order to disrupt memory for end-of-list items. Math problems of the form A+B+C=?? were presented to the participant, with values of A, B, and C set to random single digit integers. After the delay, a row of asterisks, accompanied by an 800 Hz auditory tone, was presented for a duration of 300 ms to signal the start of the recall period. Subjects were instructed to recall as many words as possible from the most recent list, in any order during the 30 s recall period. Vocal responses were digitally recorded and parsed offline using Penn TotalRecall (`http://memory.psych.upenn.edu/TotalRecall`). Subjects performed up to 25 lists in a single recall session.

Behavioral analysis

To compute behavioral measures of categorical and temporal clustering, we used the temporal factor[36]. Temporal factor measures the percentile rank of the absolute lag between successive recalls from the full distribution of available lags for items that have yet to be recalled. To measure category clustering, we computed a category factor which assumes that all items within the same category have a distance of zero and items of different categories have a distance of one. These metrics measure the degree to which recall sequences exhibit organization along temporal or

categorical dimensions, with a random recall sequences falling at the median of the distribution (i.e., 0.5).

Electrophysiological recordings and data processing

iEEG signal was recorded using subdural grids and strips (contacts spaced 10 mm apart) or depth electrodes (contacts spaced 5-10 mm apart) using recording systems at each clinical site. iEEG systems included DeltaMed & XlTek (Natus), Grass Telefactor, and Nihon-Kohden EEG systems. Signals were sampled at 500, 512, 1000, 1024, or 2000 Hz, depending on clinical site. Signals recorded at individual contacts were converted to a bipolar montage by computing the difference in signal between adjacent electrode pairs on each strip, grid, and depth electrode[37]. Bipolar signal was notch filtered at 60 Hz with a fourth order 2 Hz stop-band butterworth notch filter in order to remove the effects of line noise on the iEEG signal.

Anatomical localization

Anatomical localization of electrode placement was accomplished using independent processing pipelines for depth and surface electrodes. Post-implant CT images were coregistered with presurgical T1 and T2 weighted structural scans using Advanced Normalization Tools[38]. For patients with MTL depth electrodes, hippocampal subfields and MTL cortices were automatically labeled in a pre-implant, T2-weighted MRI using the automatic segmentation of hippocampal subfields multi-atlas segmentation method[39]. Subdural electrodes were localized by reconstructing whole-brain cortical surfaces from pre-implant T1-weighted MRIs using Freesurfer[40], and snapping electrode centroids to the cortical surface using an energy minimization algorithm[41]. Reconstructed surfaces were additionally mapped to a population-average surface[42] that we used to assign network membership based on resting state connectivity of cortical regions defined by a multi-modal cortical parcellation[19].

Network assignment and analysis

We assignment recording sites to networks of interest was based on resting state functional connectivity. From the Glasser et al. parcellation[19], we assigned parcels with high connectivity to the posterior angular gyrus (area PGp) to the PM network and parcels with high partial correlations to the ventral temporal pole (area TGv) to the AT network. Bipolar pairs were assigned to the nearest network if the bipolar centroid was within 8mm of a parcel within either

network, with the exclusion criteria that they could not be within 4mm of the other network.

To evaluate the properties of these two cortical networks, we used a bootstrap sampling procedure to randomly sample from bipolar pairs within each network. In our assessment of functional connectivity, we controlled for the effect of distance between recording sites on functional connectivity. As the distance between recording sites is greater between rather than within functional networks (leading to a biased estimate of connectivity within each network), we controlled the distances within each distribution of connections (i.e., within AT, within PM, or between networks). For each subject, we randomly sampled connections between the two networks. Connections within a given network were sampled without replacement to match the distribution of between network connections by minimizing the differences in connection length. The functional connectivity within and between each network was computed as the Pearson product-moment correlation across encoding events in the experiment. This procedure was repeated 1000 times, and the intrinsic connectivity at a given frequency was estimated from the bootstrap distribution.

A similar bootstrap procedure was used to estimate the decoding accuracy of models trained from a given particular network. Networks with greater electrode coverage are likely to have higher decoding accuracy due to the number of features alone. As a result, we randomly sampled five electrodes from each network prior to estimating the ability to decode content and context information from patterns of brain activity. This sampling procedure was repeated 1000 times per model evaluation, and the average performance across bootstrap distributions was used to indicate model performance for a given subject.

Spectral power

To compute spectral power during word encoding, we applied the Morlet wavelet transform (wave number 5) to all bipolar electrode EEG signals from the onset to the offset of stimulus presentation, across 8 logarithmically spaced frequencies from 3 to 180 Hz. Spectral power during recall was estimated from 900 ms to 100 ms preceding the onset of response vocalization for correct recalls. Recall events were required to be free of vocalization onsets in the preceding 1500 ms. Power estimates were log transformed and down sampled to 50 Hz. To avoid edge artifacts, we included buffers of 1000 ms surrounding events of interest during the computation of spectral power; mirrored buffering was applied to all retrieval-period data. Prior to modeling, all power estimates were standardized using the

mean and standard deviation of each session.

Model fitting and testing

We constructed two models to predict stimulus-related patterns of neural activity as a function of either 1) the taxonomic category of presented stimuli or 2) the serial position of each presented item within the experiment. For the content model, we modeled the neural response as a function of 300 intermediate category features computed using the word2vec model[18] trained from Google News corpora. Category features were computed by averaging semantic representations across all words presented from a given category. This results in the construction of a high dimensional space that respects the semantic relationships between all of the categories. For the context model, we modeled the neural response to stimuli as a function of serial position within the list, across lists, and across sessions. For subjects in which only a single session was run, the regressors predicting the effect of session on neural activity was excluded.

We fit both the content and context models separately to each neural feature (i.e., spectral power at a given frequency and bipolar pair). Ridge regression was used to identify model parameters that minimized prediction error on the training data, which was randomly assigned using a 5-fold cross-validation procedure, holding out individual words. Within each training fold, we performed an additional 10-fold cross-validation procedure (i.e., nested cross-validation[43]) to optimize the regularization coefficient used in each fold. Across the ten folds, we selected the regularization coefficient (from 50 potential parameters log-spaced between 10^{-2} to 10^{10}) that minimized the mean squared error across of the model predictions across the training set. The resulting set of model weights across neural features was used to decode either the serial position or category membership for unseen patterns of brain activity. We computed the probability that the observed brain data belonged to each of the 25 categories (or 12 serial positions) by computing the Pearson product-moment correlation between the observed data and the predicted pattern of brain activity for each model. A softmax function was applied to the resultant evidence for each class, and decoding accuracy was measured using area under the receiver operating characteristic curve (AUC) metric across all held out data.

To quantify reinstatement effects across each network, we applied models trained to predict patterns of brain activity observed during encoding of words to epochs just prior (from 900 ms to 100 ms before vocalization onset). For each sample within this window, we computed the ability of each model to decode either the category or serial position of

recalled words in the validation data. The resulting AUC timeseries were smoothed with with a 7 ms FWHM Gaussian kernel prior to statistical analysis for noise reduction.

Statistics

Significant differences were assessed using unpaired, two-tailed *t* tests between x and y, from a sampling distribution of subjects. One-tailed tests were used to assess differences versus chance performance in evaluating decoding accuracy and recall organization. Theoretical chance levels (e.g., an AUC of 0.5 for binary classification) were further tested by constructing null distributions via permutation of class labels (i.e., category or serial position).

We corrected for multiple comparisons (across time and frequencies) using a nonparametric permutation procedure. In our analysis of intrinsic connectivity, we performed a nonparametric one sample *t* test by constructing a null distribution of the maximum *t* statistic across frequencies, assuming no difference in intrinsic connectivity for connections within and between networks. This assumption was satisfied by random sign flipping of observed values at the subject level. Significance of observed differences in connectivity strength within versus between networks were compared to a distribution constructed from 2000 random permutations, yielding two tailed significance with $p_{FWE} < 0.05$.

Multiple comparison correction for network reinstatement followed a similar procedure. We identified significant ($p_{FWE} < 0.05$) clusters in the moments leading up to recall using threshold-free cluster enhancement (TFCE)[44]. The TCFE statistic was computed by taking the original test statistics over the pre-recall period and adjusting by weighting by the height (h) and cluster extent (e):

$$\sum_k h^H e(h_k)^E dh, \tag{1}$$

where h_k is one of k cluster forming thresholds ($h_k = h_0, h + dh, ..., h_{max}$). Height ($H$) and extent ($E$) exponents were set to 2 and 0.5 respectively, with the step size in the cluster threshold (dh) set to 0.01. After computing TCFE statistics, we used previously described nonparametric *t* tests to identify significant clusters based on null distributions ($n = 1000$).

Given the limited availability of iEEG data, no statistical methods were used to predetermine sample sizes (number of

subjects). The number of trials included in the experiment was determined by practical constraints of patient testing at epilepsy monitoring units. The subjects in this study were a subset of those included in a multi-site collaboration to investigate modulation of human memory via direct electrical stimulation. Subjects with at least 5 electrodes located within both the PM and AT networks who performed categorized free recall were included in the present study.

References

1. E. Tulving. Episodic and semantic memory. In E. Tulving and W. Donaldson, editors, *Organization of Memory.*, pages 381–403. Academic Press, New York, 1972.

2. LR Squire and S. Zola-Morgan. The medial temporal lobe memory system. *Science*, 253(5026):1380–1386, 1991.

3. H. Eichenbaum, AP Yonelinas, and C. Ranganath. The medial temporal lobe and recognition memory. *Annual Review of Neuroscience*, 30:123–152, 2007.

4. L Davachi. Item, context and relational episodic encoding in humans. *Current Opinion in Neurobiology*, 16(6): 693—700, 2006. doi: 10.1016/j.conb.2006.10.012.

5. C.M. Bird and N. Burgess. The hippocampus and memory: insights from spatial processing. *Nature Reviews Neuroscience*, 9(3):182–194, 2008.

6. C Ranganath and M Ritchey. Two cortical systems for memory-guided behavior. *Nature Reviews Neuroscience*, 13:713 – 726, 2012.

7. Kaia L Vilberg and Michael D Rugg. Memory retrieval and the parietal cortex: a review of evidence from a dual-process perspective. *Neuropsychologia*, 46(7):1787–1799, 2008.

8. Jeffrey R. Binder, Rutvik H. Desai, William W. Graves, and Lisa L. Conant. Where Is the Semantic System? A Critical Review and Meta-Analysis of 120 Functional Neuroimaging Studies. *Cerebral Cortex*, 19(12): 2767–2796, 2009. ISSN 1047-3211. doi: 10.1093/cercor/bhp055.

9. Karalyn Patterson, Peter J Nestor, and Timothy T Rogers. Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, 8(12):976–987, 2007.

10. Alex Clarke and Lorraine K. Tyler. Object-Specific Semantic Coding in Human Perirhinal Cortex. *Journal of Neuroscience*, 34(14):4766–4775, 2014. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.2828-13.2014.

11. H. Gelbard-Sagiv, R. Mukamel, M. Harel, R. Malach, and I. Fried. Internally generated reactivation of single neurons in human hippocampus during free recall. *Science*, 3:96–101, 2008.

12. J. F. Miller, M. Neufang, A. Solway, A. Brandt, M. Trippel, I. Mader, S. Hefft, M. Merkow, S. M. Polyn, J. Jacobs, M. J. Kahana, and A. Schulze-Bonhage. Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. *Science*, 342(6162):1111–1114, 2013. doi: 10.1126/science.1244056.

13. L.-T. Hsieh, M.J. Gruber, L.J. Jenkins, and C. Ranganath. Hippocampal activity patterns carry information about objects in temporal cortex. *Neuron*, 81:1165–1178, 2014.

14. Liang-Tien Hsieh and Charan Ranganath. Cortical and subcortical contributions to sequence retrieval: schematic coding of temporal context in the neocortical recollection network. *NeuroImage*, 121:78–90, 2015. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2015.07.040.

15. J. R. Manning, S. M. Polyn, G. Baltuch, B. Litt, and M. J. Kahana. Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Sciences, USA*, 108(31): 12893–12897, 2011. doi: 10.1073/pnas.1015174108.

16. J. R. Manning, M. R. Sperling, A. Sharan, E. A. Rosenberg, and M. J. Kahana. Spontaneously reactivated patterns in frontal and temporal lobe predict semantic clustering during memory search. *Journal of Neuroscience*, 32(26): 8871–8878, 2012. doi: 10.1523/JNEUROSCI.5321-11.2012.

17. T.M. Mitchell, S.V. Shinkareva, A. Carlson, K.M. Chang, V.L. Malave, R.A. Mason, and M.A. Just. Predicting human brain activity associated with the meanings of nouns. *Science*, 320(5880):1191–1195, 2008.

18. T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781v3*, 2013.

19. Matthew F Glasser, Timothy S Coalson, Emma C Robinson, Carl D Hacker, John Harwell, Essa Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F Beckmann, Mark Jenkinson, Stephen M Smith, and David C Van Essen. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171–178, 2016. doi: 10.1038/nature18933.

20. Brett L. Foster, Vinitha Rangarajan, William R. Shirer, and Josef Parvizi. Intrinsic and Task-Dependent Coupling of Neuronal Population Activity in Human Parietal Cortex. *Neuron*, 86(2):578–590, April 2015. ISSN 0896-6273. doi: 10.1016/j.neuron.2015.03.018.

21. Christopher M. Lewis, Conrado A. Bosman, Thilo Womelsdorf, and Pascal Fries. Stimulus-induced visual cortical networks are recapitulated by spontaneous local and interareal synchronization. *Proceedings of the National Academy of Sciences*, 113(5):E606–E615, 2016. doi: 10.1073/pnas.1513773113.

22. Amy L Daitch, Brett L Foster, Jessica Schrouff, Vinitha Rangarajan, Itır Kaşikçi, Sandra Gattas, and Josef Parvizi. Mapping human temporal and parietal neuronal population activity and functional coupling during mathematical cognition. *Proceedings of the National Academy of Sciences*, pages E7277–E7286, 2016.

23. Carl D. Hacker, Abraham Z. Snyder, Mrinal Pahwa, Maurizio Corbetta, and Eric C. Leuthardt. Frequency-specific electrophysiologic correlates of resting state fMRI networks. *NeuroImage*, 149:446–457, 2017. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2017.01.054.

24. C.J. MacDonald, K.Q. Lepage, U.T. Eden, and H. Eichenbaum. Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, 71(4):737–749, 2011.

25. Howard Eichenbaum. Time cells in the hippocampus: a new dimension for mapping memories. *Nature Reviews Neuroscience*, 15(11):732–744, 2014.

26. S. M. Polyn, V. S. Natu, J. D. Cohen, and K. A. Norman. Category-specific cortical activity precedes retrieval during memory search. *Science*, 310:1963–1966, 2005.

27. N. W. Morton, M. J. Kahana, E. A. Rosenberg, M. R. Sperling, A. D. Sharan, and S. M. Polyn. Category-specific neural oscillations predict recall organization during memory search. *Cerebral Cortex*, 23(10):2407–2402, 2013. doi: 10.1093/cercor/bhs229.

28. M. W. Brown and J. P. Aggleton. Recognition memory: what are the roles of the perirhinal cortex and hippocampus. *Nature Reviews Neuroscience*, 2:51–61, 2001.

29. M. W. Howard and M. J. Kahana. A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46(3):269–299, 2002. doi: 10.1006/jmps.2001.1388.

30. S. M. Polyn and M. J. Kahana. Memory search and the neural representation of context. *Trends in Cognitive Sciences*, 12(1):24–30, 2008. doi: 10.1016/j.tics.2007.10.010.

31. M. W. Howard, M. S. Fotedar, A. V. Datey, and M. E. Hasselmo. The temporal context model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. *Psychological Review*, 112(1):75–116, 2005.

32. Lucas J Jenkins and Charan Ranganath. Prefrontal and medial temporal lobe activity at encoding predicts temporal context memory. *Journal of Neuroscience*, 30(46):15558–15565, 2010.

33. Yitzhak Norman, Erin M. Yeagle, Michal Harel, Ashesh D. Mehta, and Rafael Malach. Neuronal baseline shifts underlying boundary setting during free recall. *Nature Communications*, 8(1):1301, 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-01184-1.

34. A. S. Geller, I. K. Schleifer, P. B. Sederberg, J. Jacobs, and M. J. Kahana. PyEPL: A cross-platform experiment-programming library. *Behavior Research Methods*, 39(4):950–58, 2007.

35. C. T. Weidemann, J. E. Kragel, B. C. Lega, G. A. Worrell, M. R. Sperling, A. D. Sharan, B. C. Jobst, F. Khadjevand, K. A. Davis, P. A. Wanda, A. Kadel, D. S. Rizzuto, and M. J. Kahana. Neural activity reveals interactions between episodic and semantic memory systems during retrieval. *Journal of Experimental Psychology: General*, in press.

36. S. M. Polyn, K. A. Norman, and M. J. Kahana. A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, 116(1):129–156, 2009. doi: 10.1037/a0014420.

37. J. F. Burke, K. A. Zaghloul, J. Jacobs, R. B. Williams, M. R. Sperling, A. D. Sharan, and M. J. Kahana. Synchronous and asynchronous theta and gamma activity during episodic memory formation. *Journal of Neuroscience*, 33(1):292–304, 2013.

38. Brian B Avants, Charles L Epstein, Murray Grossman, and James C Gee. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, 12(1):26–41, 2008.

39. Paul A Yushkevich, John B Pluta, Hongzhi Wang, Long Xie, Song-Lin Ding, Eske C Gertje, Lauren Mancuso, Daria Kliot, Sandhitsu R Das, and David A Wolk. Automated volumetry and regional thickness analysis of

hippocampal subfields and medial temporal cortical structures in mild cognitive impairment. *Human Brain Mapping*, 36(1):258–287, 2015.

40. Bruce Fischl, André van der Kouwe, Christophe Destrieux, Eric Halgren, Florent Ségonne, David H Salat, Evelina Busa, Larry J Seidman, Jill Goldstein, David Kennedy, et al. Automatically parcellating the human cerebral cortex. *Cerebral Cortex*, 14(1):11–22, 2004.

41. Andrew R Dykstra, Alexander M Chan, Brian T Quinn, Rodrigo Zepeda, Corey J Keller, Justine Cormier, Joseph R Madsen, Emad N Eskandar, and Sydney S Cash. Individualized localization and cortical surface-based registration of intracranial electrodes. *NeuroImage*, 59(4):3563–3570, 2012.

42. David C. Van Essen, Matthew F. Glasser, Donna L. Dierker, John Harwell, and Timothy Coalson. Parcellations and Hemispheric Asymmetries of Human Cerebral Cortex Analyzed on Surface-Based Atlases. *Cerebral Cortex*, 22(10):2241–2262, 2012. ISSN 1047-3211. doi: 10.1093/cercor/bhr291.

43. Sudhir Varma and Richard Simon. Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics*, 7:91, 2006. doi: 10.1186/1471-2105-7-91.

44. Stephen M. Smith and Thomas E. Nichols. Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference. *NeuroImage*, 44(1):83–98, January 2009.