Reinstatement of distributed cortical oscillations occurs with precise spatiotemporal dynamics during successful memory retrieval

Robert B. Yaffe[a], Matthew S. D. Kerr[a,1], Srikanth Damera[b,1], Sridevi V. Sarma[a], Sara K. Inati[c], and Kareem A. Zaghloul[b,2]

[a]Institute for Computational Medicine, Department of Biomedical Engineering, The Johns Hopkins University, Baltimore, MD 21218; and [b]Surgical Neurology Branch and [c]Office of the Clinical Director, National Institute of Neurological Disorders and Stroke, National Institutes of Health, Bethesda, MD 20892

Edited by James L. McClelland, Stanford University, Stanford, CA, and approved November 18, 2014 (received for review September 4, 2014)

Reinstatement of neural activity is hypothesized to underlie our ability to mentally travel back in time to recover the context of a previous experience. We used intracranial recordings to directly examine the precise spatiotemporal extent of neural reinstatement as 32 participants with electrodes placed for seizure monitoring performed a paired-associates episodic verbal memory task. By cueing recall, we were able to compare reinstatement during correct and incorrect trials, and found that successful retrieval occurs with reinstatement of a gradually changing neural signal present during encoding. We examined reinstatement in individual frequency bands and individual electrodes and found that neural reinstatement was largely mediated by temporal lobe theta and high-gamma frequencies. Leveraging the high temporal precision afforded by intracranial recordings, our data demonstrate that high-gamma activity associated with reinstatement preceded theta activity during encoding, but during retrieval this difference in timing between frequency bands was absent. Our results build upon previous studies to provide direct evidence that successful retrieval involves the reinstatement of a temporal context, and that such reinstatement occurs with precise spatiotemporal dynamics.

memory | intracranial EEG | oscillations | reinstatement

Reinstatement of neural activity is hypothesized to underlie our ability to recover the internal representation of a previous experience, a process described as mental time travel (1–4). These internal representations, which may reflect the external environment or internal mental states, form the context in which an episodic memory is embedded. Central to the hypothesis of mental time travel is that context representations in the brain gradually change over time, and that successful retrieval of an episodic memory involves mentally jumping back in time to reexperience a particular context (5–8). Consistent with this paradigm, when an episode is successfully retrieved from memory, the memory for neighboring episodes that occurred close in time is enhanced, an effect known as contiguity (9). However, despite substantial behavioral data supporting this hypothesis, a number of important yet unanswered questions remain regarding its underlying neural mechanisms.

Empiric support for neural reinstatement has largely emerged from functional MRI (fMRI) studies that have used multivoxel pattern analysis (MVPA) (10–12). MVPA relies on classifying neural activity during retrieval to dissociate broad manipulations such as category or task that are present during encoding (13–16). MVPA, however, is unable to directly examine whether successful retrieval reinstates the neural representations of individual items. Representational similarity analysis supports neural reinstatement of individual stimuli (17–19), but this alternative fMRI approach does not examine to what extent retrieval reinstates a changing neural representation of context. In addition, the limited temporal resolution of fMRI studies makes them unable to identify the precise spatiotemporal dynamics of neural activity that distinguish successful and unsuccessful retrieval of individual events.

Intracranial EEG (iEEG) recordings offer an opportunity to explore the neural mechanisms of reinstatement with high temporal and spatial precision. Oscillatory activity during retrieval has recently been shown to resemble activity present during encoding of neighboring items, providing neural evidence for the contiguity effect particular to mental time travel (20). However, because retrieval in this study was unconstrained, one possibility is that the observed reinstatement was a consequence, rather than a cause, of retrieval contiguity. Spiking activity in the medial temporal lobe has also been shown to reinstate activity present during encoding (21, 22). However, it is unknown how the temporal dynamics of neural activity distinguish reinstatement between correct and incorrect memory retrieval across broader cortical regions.

Here, we directly investigate whether reinstatement of a gradually changing neural representation of context occurs during successful retrieval and explore the precise spatiotemporal dynamics that underlie this process. We investigate the reinstatement of spectral power across multiple frequency bands and multiple electrode locations during encoding and retrieval as patients with subdural electrodes participated in a verbal paired-associates memory task. Paired-associates memory tasks offer an opportunity to directly explore these questions as explicit associations are formed, and subsequently recalled, between pairs of items (23, 24). Hence, rather than rely on a participant's own retrieval strategy, we use direct experimental control over retrieval to examine the relationship between retrieval cues and the reinstatement of a drifting context representation.

Significance

Our results represent significant contributions to understanding the neural mechanisms and spatiotemporal dynamics governing neural reinstatement in two important ways. First, by using a cued recall memory task, our paradigm offers experimental control over retrieval. We compare reinstatement during correct and incorrect retrieval, and provide evidence that retrieval recovers a gradually changing representation of temporal context. These data provide support for mental time travel hypothesized to underlie episodic memory. Second, leveraging the high temporal precision afforded by intracranial recordings, we investigate the precise timing of reinstatement and demonstrate that retrieval may reactivate cortical representations of a memory on a faster timescale than during encoding. Our data complement previous studies demonstrating faster replay of patterns associated with a prior episode.

Author contributions: R.B.Y., M.S.D.K., S.V.S., and K.A.Z. designed research; R.B.Y., S.D., S.K.I., and K.A.Z. performed research; R.B.Y. and K.A.Z. analyzed data; and R.B.Y. and K.A.Z. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]M.S.D.K. and S.D. contributed equally to this work.

[2]To whom correspondence should be addressed. Email: kareem.zaghloul@nih.gov.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1417017112/-/DCSupplemental.

Results

Thirty-two participants (18 males; age, 33.5 ± 2.2 y) with medication-resistant epilepsy who underwent a surgical procedure for placement of intracranial electrodes for seizure monitoring participated in a verbal paired-associates task (Fig. 1*A*). For each trial, participants responded with the correct word, with a wrong word (intrusions), or with no word. Responses were designated as passes when no response was made or when the participant vocalized the word "pass." We considered intrusion and pass trials as incorrect trials for subsequent analyses. Participants studied 223 ± 23 word pairs and successfully recalled 41.0 ± 8.0% words with a mean response time of 1,831 ± 80 ms (Fig. S1*B*). On 11.8 ± 1.6% of trials, participants responded with an incorrect word (intrusions). The mean response time for intrusions was 2,736 ± 116 ms. For the remaining study word pairs, participants either made no response to the cue word, or vocalized the word "pass." Participants vocalized the word "pass" with a mean response time of 3,334 ± 207 ms. A one-way ANOVA revealed no significant effect of study-test lag on recall probability [$F_{(4,155)} = 1.06$, $P = 0.38$] (Fig. S1*A*).

To assess the reinstatement of neural activity between encoding and retrieval, we compared multidimensional representations of oscillatory power at every time point during encoding and retrieval. Briefly, for every 500-ms temporal epoch, spaced every 100 ms (80% overlap), during the encoding and retrieval periods, we constructed a feature vector containing oscillatory power information from five frequency bands (theta, 3.5–8 Hz; alpha, 8–12 Hz; beta, 13–25 Hz; low gamma, 30–58 Hz; high gamma, 62–100 Hz) and from all electrode locations (*SI Materials and Methods*). We quantified reinstatement between every temporal epoch during encoding and retrieval by calculating the cosine similarity between feature vectors. Thus, for a single trial, we generate a precise temporal map of neural reinstatement between the encoding and retrieval periods (Fig. 1*B*).

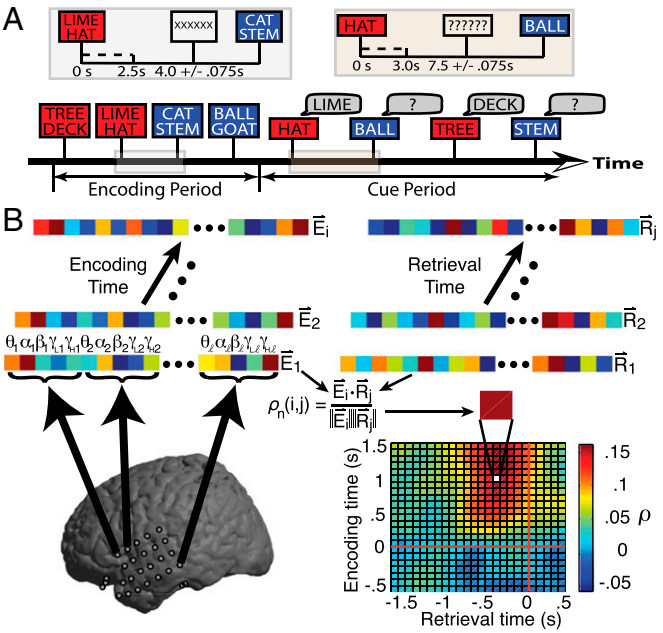

Fig. 1. Paired-associates task and neural reinstatement. (*A*) Task presentation. Timing of word and cue presentation is shown in the *Inset*. Red and blue boxes indicate correct and incorrect trials, respectively. (*B*) Reinstatement of distributed oscillatory power. For every temporal epoch during encoding and retrieval, the *z*-scored power from every electrode and every frequency band is combined to create a single feature vector for that epoch. Reinstatement (cosine similarity) for all encoding–retrieval pairs is shown for a single trial. The red lines on the reinstatement map correspond to the time of word presentation during encoding and the time of vocalization during retrieval.

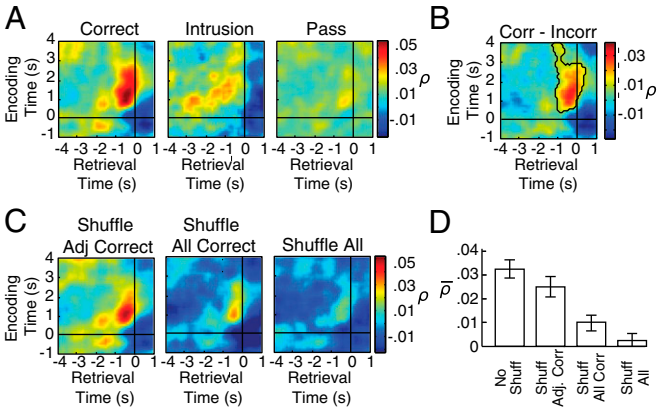

Fig. 2. Correct memory retrieval exhibits trial-specific reinstatement of distributed oscillatory power. (*A*) Mean reinstatement across all participants during correct, intrusion, and pass trials. Black lines indicate the time of word presentation during encoding and the time of vocalization during retrieval. (*B*) The difference in mean reinstatement between correct and incorrect trials across all participants. We defined a temporal region of interest (tROI) for subsequent analyses as all encoding–retrieval time pairs that exhibited significant differences between the two trial types, outlined in black. (*C*) Mean reinstatement during correct trials across all participants when shuffling neighboring correct retrieval periods (*Left*), all correct retrieval periods (*Middle*), or all retrieval periods from all trial types (*Right*). (*D*) Mean reinstatement across all participants, averaged over the tROI for each participant, during each shuffled permutation shown in *C*. Error bars represent SEM across participants.

To compare the extent of reinstatement during correct, intrusion, and pass trials, we averaged reinstatement maps across all trial types for each participant (see Fig. S2 for single-participant examples). When we examined reinstatement during the retrieval period time locked to vocalization across participants, we found significantly greater reinstatement during correct trials compared with both pass trials and intrusion trials (Fig. 2*A* and Fig. S3). We identified temporal encoding and retrieval epochs where reinstatement during correct trials was significantly greater than during incorrect trials, and defined this as our temporal region of interest (tROI) for further analyses ($P < 0.001$, permutation test; *SI Materials and Methods*; black outline, Fig. 2*B*). In a post hoc analysis, we found that the mean reinstatement over the tROI, averaged across participants, was significantly greater during correct compared with both intrusion and pass trials [$t_{(29)} = 5.19$, $P = 0.00001$ for correct vs. pass, paired, two-tailed; $t_{(24)} = 4.47$, $P = 0.0002$ for correct vs. intrusion]. We identified the encoding–retrieval time pair within the tROI for each participant that exhibited the greatest reinstatement during correct trials and found that the peak encoding time occurred 1.68 ± 0.17 s after the presentation of the study pair, and the peak retrieval time occurred 0.44 ± 0.064 s before vocalization.

Reinstatement during memory encoding and retrieval may reflect trial-specific reinstatement of neural activity or general encoding and retrieval mechanisms. To investigate these possibilities, we compared the reinstatement observed during correct trials to reinstatement computed when we shuffled trial labels. We first shuffled the labels of all retrieval periods and calculated reinstatement between the true correct encoding periods and the shuffled retrieval periods. We found that the average reinstatement in the tROI originally observed with our data were significantly greater than reinstatement in this shuffled distribution [$t_{(29)} = 8.50$, $P < 10^{-8}$, paired, two-tailed; Fig. 2 *C* and *D*, Shuffle All]. We next shuffled the labels of only the correct retrieval periods and calculated reinstatement between the true correct encoding periods and the shuffled correct retrieval periods.

If reinstatement reflects a general encoding and retrieval process, then reinstatement calculated using shuffled correct retrieval periods should be identical to that observed using the original correct trials. Instead, we found that mean reinstatement in the tROI was significantly less when we used the shuffled compared with the true correct trials [$t_{(29)} = 6.58$, $P < 10^{-6}$, paired, two-tailed; Fig. 2 *C* and *D*, Shuffle All Correct; see also Fig. S4]. When we restricted shuffling to only swap retrieval periods from adjacent correct trials, we also found that mean reinstatement in the tROI during correct trials was significantly greater than that calculated using the shuffled adjacent correct trials [$t_{(29)} = 4.81$, $P = 0.00004$, paired, two-tailed; Fig. 2 *C* and *D*, Shuffle Adjacent Correct].

The decreases in mean reinstatement observed with greater shuffling may be related to the amount of elapsed time between the true encoding and shuffled retrieval period. We therefore examined whether neural activity during an experimental session exhibited a slow temporal drift. We chose an identical temporal epoch from every encoding event (0–500 ms before stimulus presentation) and calculated the cosine similarity of neural activity between all pairs of the selected time epochs. We averaged the resulting values over all pairs separated by the same amount of time within and then across participants and binned the data separately into 5-s and 60-s bins (Fig. 3 *A* and *B*). Because 5-s time bins approximately correspond to the timing of individual item presentation, the resulting data reflect the neural similarity between items separated by different lags within the same list. We examined the slopes of the resulting regressions and found that neural similarity was significantly related to the time between experimental events on both small and large timescales [5-s bins, slope $= -0.0141 \pm 0.0021$, $t_{(29)} = -6.71$, $P < 10^{-6}$, one-sample, two-tailed; 60-s bins, slope $= -0.0053 \pm 0.0011$, $t_{(29)} = -4.67$, $P = 0.00006$].

We hypothesized that reinstatement during correct trials principally arose because neural activity jumped back in time to reinstate context present during encoding. To directly examine this, we identified the serial position during encoding of each correct trial and paired the associated retrieval period with all encoding periods from all serial positions within the same study list. In this manner, a true pairing would have a lag of 0, whereas a lag of 1 (−1) would correspond to a pairing between the true retrieval period and the encoding period of the subsequent (previous) item in the study list. We computed the average tROI reinstatement for each pairing across participants. If the reinstatement of temporal context governs the observed similarity, then neural activity during retrieval should be most similar to activity during the encoding period of the same pair, and should fall off with both positive and negative lag. Conversely, if temporal drift accounts for the observed similarity, then neural activity during retrieval should be most similar to activity during the most recent encoding period (i.e., most positive lag).

For correct trials, we found that each retrieval period exhibited maximum reinstatement with the true encoding period, and that when paired with other encoding periods from the same list, mean reinstatement decreased with both positive and negative lag (Fig. 3*C*). We regressed mean reinstatement with lag and found an average slope that was positive across participants when retrieval periods were paired with previous encoding periods (negative lag, slope $= 0.0068 \pm 0.0018$), and an average slope that was negative when retrieval periods were paired with subsequent encoding periods (positive lag, slope $= -0.0037 \pm 0.0014$). The decrease in mean reinstatement observed with positive and negative lags during correct retrievals suggests that successful reinstatement was shaped by the reinstatement of temporal context during encoding. We performed the same analysis using incorrect trials and found that mean reinstatement across participants exhibited a progressive increase with encoding lag (slope $= 0.0023 \pm 0.0006$). The increase with encoding lag suggests that measured reinstatement during incorrect trials is governed most by temporal proximity to the item at test. To directly assess whether the reinstatement of temporal context during encoding was particular to correct retrieval, we compared the relation between positive lag (0–2) and mean reinstatement during correct trials to the relation between positive lag and mean reinstatement during incorrect trials. We found that the relation between reinstatement and positive lag was significantly different between correct and incorrect trials [$t_{(29)} = -3.41$, $P = 0.0019$, paired, two-tailed].

To assess the contributions of different frequency bands to reinstatement, we calculated the mean reinstatement during correct trials over our tROI using only features from each frequency band and found a significant effect of frequency on mean reinstatement [$F_{(4,145)} = 3.34$, $P = 0.012$, one-way ANOVA; Fig. 4*A*]. Pairwise *t* tests showed that reinstatement was significantly greater in both high gamma and theta compared with alpha and beta [$t_{(29)} > 2.34$, $P < 0.03$, paired, two-tailed], but there was no difference between theta and high gamma [$t_{(29)} = -1.26$, $P = 0.22$]. We next examined the temporal dynamics underlying the contributions of different frequencies to reinstatement during correct trials. First, we calculated the mean reinstatement for all encoding–retrieval time pairs across participants using only theta or high-gamma features and observed that the time periods exhibiting significant differences in reinstatement between correct and incorrect trials varied between frequency bands (Fig. 4*B*). We then averaged mean reinstatement, calculated using only features drawn separately from each frequency band, over all encoding temporal epochs defined in the tROI. High-gamma and theta features exhibited the largest rises in reinstatement during correct trials, peaking immediately before vocalization (Fig. 4*C*, *Top*). We then averaged mean reinstatement over all retrieval epochs defined in the tROI separately for each frequency band. First, high gamma, and then, theta exhibited peaks in reinstatement around 1 s following word presentation during encoding (Fig. 4*C*, *Bottom*).

As reinstatement seemed to be principally mediated by theta and high-gamma frequencies, we visualized the anatomic regions, separately for each frequency band, that contained features exhibiting significant differences in mean reinstatement over the tROI between correct and incorrect trials across participants. Only regions that had electrodes from five or more participants were included in this analysis (Fig. 4*D*). We rendered all spatial regions exhibiting significant differences across participants ($P < 0.004$, permutation procedure, FDR corrected; *SI Materials and Methods*), where color intensity represents mean reinstatement during correct trials for that spatial region and frequency (Fig. 4*E* and Fig. S5; *SI Materials and Methods*). We found significantly greater reinstatement for correct versus incorrect trials in the inferior temporal lobes bilaterally in both theta and

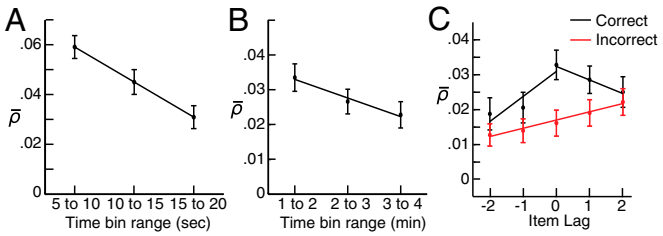

Fig. 3. Neural reinstatement is shaped by temporal context. (*A*) Temporal autocorrelation of distributed oscillatory power during encoding. Plot shows the cosine similarity of neural activity between all pairs of time epochs from 0 to 500 ms before stimulus presentation during encoding. Data are shown for events separated by different lengths of time, separated into 5-s bins. (*B*) Same as *A*, with data separated into 60-s bins. (*C*) Mean tROI reinstatement across all participants between each retrieval period and each encoding period from the associated study list (black, correct; red, incorrect). For *A*–*C*, solid lines indicate the best-fit regression line averaged across participants, and error bars represent SEM across participants.

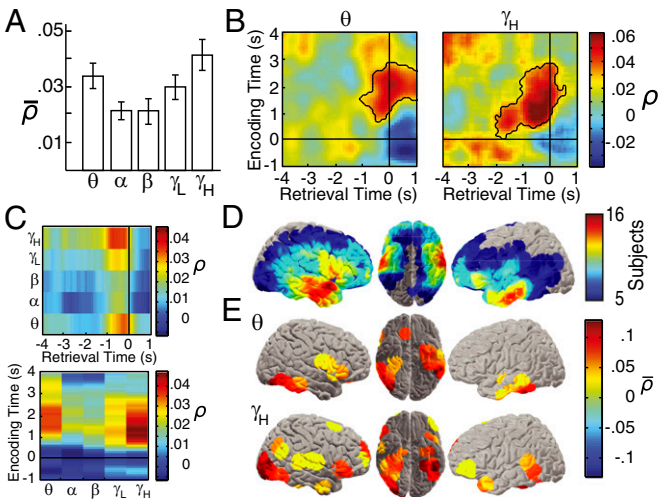

Fig. 4. Neural reinstatement is mediated by temporally precise theta and high-gamma frequency activity in specific anatomic locations. (*A*) Mean reinstatement for all features in each frequency band averaged over the tROI for each participant. Error bars represent SEM across participants. (*B*) Mean reinstatement of correct trials across all participants for theta (*Left*) and high-gamma (*Right*) frequency bands. Regions of significant difference between correct and incorrect trials are outlined in black. Black lines indicate time of vocalization during retrieval and time of pair presentation during encoding. (*C*) Mean reinstatement for each frequency, averaged over all encoding (*Top*) or retrieval (*Bottom*) time periods defined by the tROI, for all retrieval (*Top*) or encoding (*Bottom*) times. Color represents the average mean reinstatement for each frequency across participants. The black lines indicate the point of vocalization during retrieval (*Top*) and the point of stimulus presentation during encoding (*Bottom*). (*D*) Electrode coverage across all 32 participants. (*E*) Anatomic regions exhibiting significant differences in reinstatement between correct and incorrect trials for all theta and high-gamma features across all participants.

high-gamma frequency bands. We also found significant reinstatement in the right ventrolateral prefrontal cortex in theta, and bilaterally in high gamma. In high gamma, we also found areas of significant reinstatement that included the prefrontal cortex, the right supplementary motor area, the right temporo-parietal junction, the left posterior lateral temporal lobe, and the right occipital lobe.

Because we observed that the specific time periods of reinstatement varied between participants, between frequency bands, and between lobes, we hypothesized that the time course of reinstatement for individual neural features would also vary. We examined differences in reinstatement between correct and incorrect trials for individual features in individual participants (*SI Materials and Methods*). For each participant, we visualized the anatomic regions and frequency bands that contained features exhibiting significant differences between correct and incorrect trials (Fig. 5*A*; *SI Materials and Methods*). On average, 11.3 ± 0.9% of features demonstrated significant differences in reinstatement for a given participant ($P < 0.05$, permutation test). The encoding–retrieval time regions exhibiting significant differences in reinstatement varied among features (Fig. 5*A*). For each feature, we identified the encoding–retrieval time pair that exhibited the maximum reinstatement during correct trials within these identified time regions. When we examined the timing distribution of all significant theta and high-gamma features, we found that peak encoding times were significantly greater for theta features than for high-gamma features [$t_{(414)} = 2.90$, $P = 0.004$, unpaired, two-tailed; Fig. 5*B*]. Peak retrieval times did not significantly differ between theta and high-gamma features [$t_{(414)} = 0.37$, $P = 0.71$].

To further investigate the differences in the timing of reinstatement, for each participant, we calculated the difference between the peak encoding and peak retrieval times for all significant theta and high-gamma features. We first only included individual electrodes that demonstrated both significant theta and high-gamma reinstatement, and found that peak high-gamma activity significantly preceded peak theta activity during encoding across participants [Fig. 5*C*; $t_{(20)} = 3.46$, $P = 0.003$, one-sample, two-tailed]. Notably, during retrieval, we found no difference in peak theta and high-gamma times [$t_{(20)} = 1.49$, $P = 0.15$]. The difference we observed in encoding times between high gamma and theta was significantly greater than the difference in retrieval times in these electrodes ($P = 0.039$, permutation test), which were mostly located in the left posterior temporal lobe, the right fronto-temporal region, and the bilateral temporo-occipital areas (Fig. 5*D*). We next included all electrodes that exhibited either significant high-gamma or theta reinstatement, and in each participant computed the average peak encoding and retrieval times for each frequency for different anatomic regions. Across participants, the left temporal lobe exhibited a significant difference in peak time of high-gamma and theta reinstatement during encoding [$t_{(10)} = 4.18$, $P = 0.002$, one-sample, two-tailed; Fig. 5*C*; right medial temporal lobe, $t_{(5)} = 2.74$, $P = 0.041$, not shown; other regions, $P > 0.05$]. Conversely, during retrieval, we found no anatomic region that exhibited significant differences in the peak time of theta and high-gamma reinstatement ($P > 0.05$).

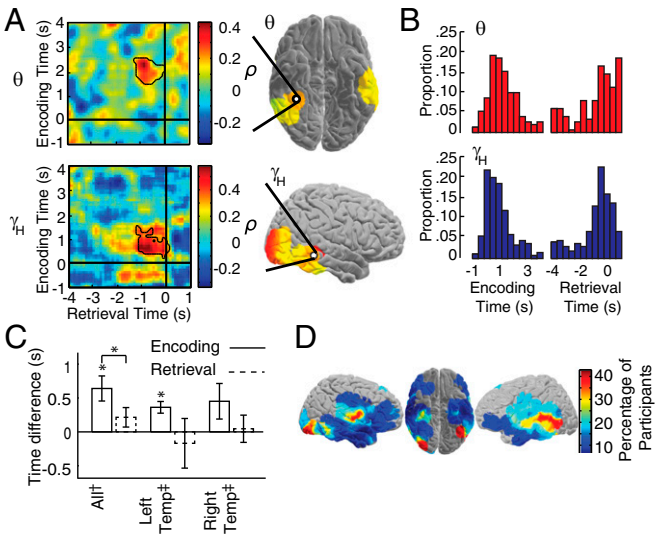

Fig. 5. High-gamma reinstatement precedes theta during encoding but not during retrieval. (*A*) Mean reinstatement during correct trials of one theta (*Top*) and one high-gamma (*Bottom*) feature in a single participant. Encoding–retrieval time pairs exhibiting significant differences between correct and incorrect trials are outlined in black. The anatomic locations of the individual features are represented as white circles on the topographic plots. Topographic plots represent anatomic regions exhibiting significant differences in reinstatement between correct and incorrect trials for all theta and high-gamma features for this participant. (*B*) Distributions of peak encoding and retrieval times across all participants for theta (*Top*) and high-gamma (*Bottom*) features exhibiting significant differences in reinstatement between correct and incorrect trials. (*C*) Differences in peak reinstatement times between theta and high-gamma features during encoding and retrieval. Differences in peak time are shown for all electrodes that exhibit significant reinstatement in both theta and high gamma on the *Left* (single dagger). Peak time differences are shown for all electrodes in the left and right temporal lobes that exhibit either theta or high-gamma reinstatement are shown in the *Middle* and on the *Right*, respectively (double dagger). Asterisks indicate a difference between peak theta and high-gamma time that is significantly greater than zero ($P < 0.05$) or a significant difference between encoding time differences and retrieval time differences ($P < 0.05$). Error bars represent SEM across participants. (*D*) Percentage of participants with electrodes in a given spatial ROI that exhibited significant reinstatement in both theta and high gamma.

Discussion

Our data demonstrate that reinstatement of a drifting neural representation of temporal context occurs during successful memory retrieval with precise spatiotemporal dynamics. Our approach capitalizes on the sensitivity of examining multivariate activity (25) and extends previous work in several ways. We demonstrate that patterns of spectral power distributed across multiple spatial locations and multiple frequency bands gradually change over multiple timescales. By examining trial-specific reinstatement during correct and incorrect memory retrieval, we show that during successful retrieval this neural signal is recovered, providing direct evidence of mental time travel hypothesized to underlie episodic memory formation (1, 2, 4, 6, 8). We subsequently identify the precise contributions of separate anatomic locations and frequency bands to neural reinstatement both in individuals and across our population of participants. We found that successful neural reinstatement is largely mediated by cortical high-gamma activity that precedes theta oscillatory activity in the temporal lobe during encoding, but this difference in timing is absent during retrieval.

One requirement of the hypothesis that successful retrieval involves mental time travel is that the neural representation of context gradually changes with time (6, 8, 26). Consistent with this, we found that distributed patterns of oscillatory activity exhibited a slow temporal drift across multiple timescales. Temporal autocorrelations have traditionally been attributed to the statistical properties of complex neural signals, including resting state functional imaging (27), scalp and intracranial EEG (28). However, one possibility is that these correlations actually reflect a slowly changing representation of temporal context shaped by both external inputs and a continuously changing internal state (1, 26). This representation should be accessed by the brain if mental time travel governs memory formation and retrieval (8, 29).

Hence, the second requirement of the mental time travel hypothesis is that retrieval recovers this gradually changing signal (1, 6, 8). By comparing reinstatement during correct and incorrect retrieval, our findings build upon recent evidence that trial-specific reinstatement of this neural representation of temporal context occurs during successful retrieval (20, 21). If successful retrieval is associated with the recovery of a temporal context representation, then we would expect to find a graded decrease in reinstatement as retrieval periods were paired with encoding periods separated by longer time intervals in both the forward and backward direction (6, 8, 21). Consistent with this prediction, our data demonstrate a contiguity effect for correct trials within lists. Because the list length used here was only four items, drawing significant conclusions using only correct trial contiguity is limited and raises the possibility that the observed effects may just reflect the reinstatement of content rather than context (8). However, we also found that incorrect trials, used here as a control condition, exhibited similarity that was principally shaped by the drifting temporal context representation itself and not by the recovery of this signal. Hence, by directly comparing reinstatement contiguity between correct and incorrect trials, our data are consistent with the hypothesis that the reinstatement of temporal context occurs during correct retrievals, and suggest that an underlying component characterized by temporal drift is evident and unmasked during incorrect trials.

Although we demonstrate a significant difference in reinstatement between correct and incorrect trials, one concern is that the observed differences in reinstatement arise as a result rather than a cause of successful retrieval (30). Although our experimental paradigm offers some control over when reinstatement and retrieval occur, our data are unable to explicitly distinguish causality between one and the other. A second possible confound in these data are that neural features responsive to visual presentation, activated both during encoding and retrieval, underlie reinstatement, and that the observed differences between correct and incorrect trials arise solely from differences in response time distributions (*SI Text*). We note, however, that when locked to cue presentation when response time differences do not affect reinstatement, we also observe significant differences in reinstatement (Fig. S6). We corrected for the possible confounds associated with visual responses and confirmed that a significant component of the observed reinstatement is separately mediated by memory encoding and retrieval mechanisms (Fig. S7).

We found significant differences in reinstatement when we examined distributed patterns of neural oscillations, but our analysis also enables us to investigate how individual spatial locations and frequencies contribute to reinstatement. We found that the largest frequency contributions to cortical reinstatement arose from theta and high-gamma frequencies, consistent with the role these frequency bands play during memory encoding and retrieval (24, 31). Within these frequency bands, we found that the greatest anatomic contributions arise from the temporal and occipital lobes, although the ventrolateral prefrontal cortex and the right temporo-parietal junction were also involved. These contributions are also not unexpected, as temporal lobe regions have been implicated in verbal memory processes, and temporal and prefrontal regions may be involved in maintaining temporal context (29, 32). Whether reinstatement of episodic memory is specifically limited to these anatomic regions remains to be determined, as interpreting the spatial extent of reinstatement from our data is difficult given the limits imposed by our electrode coverage.

Previous studies that provide empiric support for neural reinstatement have not investigated the temporal dynamics mediating this process (11, 12, 17, 19). We build upon these studies, as the intracranial recordings we use here enable us to examine neural reinstatement with high temporal precision across multiple frequency bands. Hence, the timing of reinstatement can inform us about the timing of encoding and retrieval processes that are particularly relevant for successful memory formation. Notably, we found that reinstatement associated with each frequency band emerges in a temporally precise manner. When both theta and high-gamma frequencies were involved in reinstatement in a single location, we found that the difference in timing between encoding theta and high-gamma activity was significantly greater than the difference during retrieval. These data provide evidence that memory-relevant oscillatory activity within individual cortical locations during encoding may be compressed in time during retrieval. One possibility may be that theta and high-gamma activity coordinate separate processes relevant for memory, such as attention or semantic processing, and that these processes occur separately in time during encoding and simultaneously during retrieval. However, another possibility is that the observed temporal compression complements previous studies suggesting that retrieval reactivates internal representations of an experience on a faster timescale than the original experience (33, 34).

Together, our data demonstrate that successful retrieval is associated with reinstatement of a gradually changing temporal context representation, providing direct evidence of mental time travel. Furthermore, individual cortical oscillations contribute to overall reinstatement and exhibit their own time course of activation during encoding and retrieval. It is notable that, given the observed temporal variability of individual neural oscillations and spatial locations, averaging across all neural features still demonstrated significant differences in reinstatement between correct and incorrect trials across participants. Neural reinstatement may actually be much stronger than our analysis would suggest, and it is possible that if all oscillations and spatial locations involved in encoding and retrieval were selectively chosen, neural reinstatement would be observed with higher fidelity (Fig. S8). Taken as a whole, our data suggest that these distributed

patterns of activity are coordinated to create a representation of temporal context, which is then recovered during the retrieval of a coherent memory trace.

Materials and Methods

Participants. Thirty-two participants with medication-resistant epilepsy underwent a surgical procedure in which platinum recording contacts were implanted subdurally on the cortical surface as well as deep within the brain parenchyma. Each patient participated in a paired-associates task (Fig. 1*A*; *SI Materials and Methods*).

For each participant, the clinical team determined the placement of the contacts so as to best localize epileptogenic regions. Data were collected at three different hospitals: Clinical Center at the National Institutes of Health (NIH) (Bethesda, MD), the Hospital of the University of Pennsylvania (UP) (Philadelphia, PA), and Thomas Jefferson University Hospital (TJ) (Philadelphia, PA). The research protocol was approved by the institutional review board at each hospital, and informed consent was obtained from the participants and their guardians.

Metrics of Reinstatement. We binned continuous time wavelet transforms into 500-ms epochs spaced every 100 ms (80% overlap) and averaged the instantaneous power over each epoch. To account for changes in power across experimental sessions, we *z*-transformed power values separately for each frequency and for each session using the mean and SD of all 500-ms epochs for that session (*SI Materials and Methods*). We *z*-transformed power values separately for each frequency and for each session using the mean and SD of all 500-ms epochs for that session. For each temporal epoch, we subsequently averaged the *z*-transformed power across five frequency bands: theta (3.5–8 Hz), alpha (8–12 Hz), beta (13–25 Hz), low gamma (30–58 Hz), and high gamma (62–100 Hz).

For every temporal epoch in each trial, we constructed a feature vector composed of the average *z*-scored power for every electrode and for every frequency band. For each encoding temporal epoch, *i*, and for each retrieval temporal epoch, *j*, we define feature vectors as follows:

$$\vec{E}_i = \left[z_{1,1}(i)\ldots z_{1,F}(i)\ldots z_{L,F}(i)\right],$$

$$\vec{R}_j = \left[z_{1,1}(j)\ldots z_{1,F}(j)\ldots z_{L,F}(j)\right],$$

where $z_{l,f}(i)$ is the *z*-transformed power of electrode $l = 1\ldots L$ at frequency band $f = 1\ldots F$ in temporal epoch *i*. For *L* electrodes and *F* frequency bands, we thus create a feature vector at each temporal epoch that contains $K = L*F$ features.

To quantify reinstatement during trial *n*, we calculated the cosine similarity between all encoding and retrieval feature vectors $\vec{E}_i$ and $\vec{R}_j$ for all pairs of encoding and retrieval temporal epochs during that trial (Fig. 1*D*). Cosine similarity gives a measure of how close the angles of two vectors are in a multidimensional space. We chose cosine similarity over Pearson's correlation to measure reinstatement because if all of the elements of two feature vectors show increases in power from baseline, with small additional random noise, then these two vectors should have high measured reinstatement. Pearson's correlation, a centered version of cosine similarity, would give a low correlation in this case because of the noise fluctuations, whereas the cosine similarity would be high, consistent with our interpretation of reinstatement. Thus, for each trial, *n*, we generate a temporal map of reinstatement values:

$$C_n(i,j) = \frac{\vec{E}_i \cdot \vec{R}_j}{\|\vec{E}_i\|\|\vec{R}_j\|}, \qquad \text{[S1]}$$

where $C_n(i,j)$ corresponds to the reinstatement of neural activity across all electrodes and all frequencies between encoding epoch *i* and retrieval epoch *j* during trial *n*. For every patient, we compute the reinstatement maps separately for all correct, intrusion, and pass trials.

ACKNOWLEDGMENTS. We thank Marc W. Howard, John F. Burke, John H. Wittig, Jr., and Michael J. Kahana for helpful and insightful comments on the manuscript. We thank the clinical teams at the University of Pennsylvania Hospital and Thomas Jefferson Hospital for assistance in data collection. We are indebted to all patients who have selflessly volunteered their time to participate in this study. This work was supported by the intramural research program of the National Institute of Neurological Disorders and Stroke. Data collection at the University of Pennsylvania and at Thomas Jefferson University was supported by National Institutes of Health Grants MH055687, MH061975, NS067316, and MH017168. S.V.S. was supported by the National Science Foundation (NSF) Career Award 1055560 and Burroughs Wellcome Fund Career Award at the Scientific Interface 1007274. M.S.D.K. was supported by NSF Graduate Research Fellowship Program (GRFP) DGE-1232825 and the Achievement Rewards for College Scientists (ARCS) Foundation. R.B.Y. was supported by the Epilepsy Foundation Pre-doctoral Research Training Fellowship.

1. Tulving E (1972) *Organization of Memory*, eds Tulving E, Donaldson W (Academic, New York), pp 381–403.
2. Alvarez P, Squire LR (1994) Memory consolidation and the medial temporal lobe: A simple network model. *Proc Natl Acad Sci USA* 91(15):7041–7045.
3. McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* 102(3):419–457.
4. Norman KA, O'Reilly RC (2003) Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychol Rev* 110(4):611–646.
5. Howard MW, Kahana MJ (2002) A distributed representation of temporal context. *J Math Psychol* 46(3):269–299.
6. Howard MW, Fotedar MS, Datey AV, Hasselmo ME (2005) The temporal context model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. *Psychol Rev* 112(1):75–116.
7. Sederberg PB, Howard MW, Kahana MJ (2008) A context-based theory of recency and contiguity in free recall. *Psychol Rev* 115(4):893–912.
8. Polyn SM, Norman KA, Kahana MJ (2009) A context maintenance and retrieval model of organizational processes in free recall. *Psychol Rev* 116(1):129–156.
9. Kahana MJ, Howard MW, Polyn SM (2008) *Cognitive Psychology of Memory* (Vol 2 of *Learning and Memory: A Comprehensive Reference*, 4 Vols, ed Byrne J), ed Roediger HL, III (Elsevier, Oxford).
10. Johnson JD, Rugg MD (2007) Recollection and the reinstatement of encoding-related cortical activity. *Cereb Cortex* 17(11):2507–2515.
11. Danker JF, Anderson JR (2010) The ghosts of brain states past: Remembering reactivates the brain regions engaged during encoding. *Psychol Bull* 136(1):87–102.
12. Rissman J, Wagner AD (2012) Distributed representations in memory: Insights from functional brain imaging. *Annu Rev Psychol* 63:101–128.
13. Polyn SM, Natu VS, Cohen JD, Norman KA (2005) Category-specific cortical activity precedes retrieval during memory search. *Science* 310(5756):1963–1966.
14. Johnson JD, McDuff SG, Rugg MD, Norman KA (2009) Recollection, familiarity, and cortical reinstatement: A multivoxel pattern analysis. *Neuron* 63(5):697–708.
15. McDuff SG, Frankel HC, Norman KA (2009) Multivoxel pattern analysis reveals increased memory targeting and reduced use of retrieved details during single-agenda source monitoring. *J Neurosci* 29(2):508–516.
16. Kuhl B, Rissman J, Wagner A (2012) Multi-voxel patterns of visual category representation during episodic encoding are predictive of subsequent memory. *Neuropsychologia* 50(4):458–469.
17. Staresina BP, Henson RN, Kriegeskorte N, Alink A (2012) Episodic reinstatement in the medial temporal lobe. *J Neurosci* 32(50):18150–18156.
18. Ritchey M, Wing EA, LaBar KS, Cabeza R (2013) Neural similarity between encoding and retrieval is related to memory via hippocampal interactions. *Cereb Cortex* 23(12):2818–2828.
19. Deuker L, et al. (2013) Memory consolidation by replay of stimulus-specific neural activity. *J Neurosci* 33(49):19373–19383.
20. Manning JR, Polyn SM, Baltuch GH, Litt B, Kahana MJ (2011) Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proc Natl Acad Sci USA* 108(31):12893–12897.
21. Howard MW, Viskontas IV, Shankar KH, Fried I (2012) Ensembles of human MTL neurons "jump back in time" in response to a repeated stimulus. *Hippocampus* 22(9):1833–1847.
22. Miller JF, et al. (2013) Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. *Science* 342(6162):1111–1114.
23. Kahana MJ (2012) *Foundations of Human Memory* (Oxford Univ Press, New York).
24. Hanslmayr S, Staudigl T (2014) How brain oscillations form memories—a processing based perspective on oscillatory subsequent memory effects. *Neuroimage* 85(Pt 2):648–655.
25. Norman KA, Newman E, Detre G, Polyn S (2006) How inhibitory oscillations can train neural networks and punish competitors. *Neural Comput* 18(7):1577–1610.
26. Estes WK (1955) Statistical theory of spontaneous recovery and regression. *Psychol Rev* 62(3):145–154.
27. Zarahn E, Aguirre GK, D'Esposito M (1997) Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *Neuroimage* 5(3):179–197.
28. Linkenkaer-Hansen K, Nikouline VV, Palva JM, Ilmoniemi RJ (2001) Long-range temporal correlations and scaling behavior in human brain oscillations. *J Neurosci* 21(4):1370–1377.
29. Jenkins LJ, Ranganath C (2010) Prefrontal and medial temporal lobe activity at encoding predicts temporal context memory. *J Neurosci* 30(46):15558–15565.
30. Rugg MD, Johnson JD, Park H, Uncapher MR (2008) Encoding-retrieval overlap in human episodic memory: A functional neuroimaging perspective. *Prog Brain Res* 169:339–352.
31. Sederberg PB, et al. (2007) Hippocampal and neocortical gamma oscillations predict memory formation in humans. *Cereb Cortex* 17(5):1190–1196.
32. Davachi L, Mitchell JP, Wagner AD (2003) Multiple routes to memory: Distinct medial temporal lobe processes build item and source memories. *Proc Natl Acad Sci USA* 100(4):2157–2162.
33. Davidson TJ, Kloosterman F, Wilson MA (2009) Hippocampal replay of extended experience. *Neuron* 63(4):497–507.
34. Carr MF, Jadhav SP, Frank LM (2011) Hippocampal replay in the awake state: A potential substrate for memory consolidation and retrieval. *Nat Neurosci* 14(2):147–153.