

Temporal Context Modulates Encoding and Retrieval of Overlapping Events

Devyn E. Smith, Isabelle L. Moore, and  Nicole M. Long

Department of Psychology, University of Virginia, Charlottesville, Virginia 22904

Overlap between events can lead to interference because of a trade-off between encoding the present event and retrieving the past event. Temporal context information, “when” something occurred, a defining feature of episodic memory, can cue retrieval of a past event. However, the influence of temporal overlap, or proximity in time, on the mechanisms of interference is unclear. Here, by identifying brain states using scalp EEG from male and female human subjects, we show the extent to which temporal overlap promotes interference and induces retrieval. In this experiment, subjects were explicitly directed to either encode the present event or retrieve a past, overlapping event while perceptual input was held constant. We find that the degree of temporal overlap between events leads to selective interference. Specifically, greater temporal overlap between two events leads to impaired memory for the past event selectively when the top-down goal is to encode the present event. Using pattern classification analyses to measure neural evidence for a retrieval state, we find that greater temporal overlap leads to automatic retrieval of a past event, independent of top-down goals. Critically, the retrieval evidence we observe likely reflects a general retrieval mode, rather than retrieval success or effort. Collectively, our findings provide insight into the role of temporal overlap on interference and memory formation.

Key words: context; EEG; encoding; episodic memory; MVPA; retrieval

Significance Statement

When a present event overlaps with an event from the past, this leads to a trade-off between the tendency to encode the present event versus retrieve the past event. Here we show that, when two events are experienced nearby in time, the memory system is biased toward a retrieval state and that subsequent memory for the past event is impaired. These findings suggest an influence of bottom-up temporal factors on both interference and the trade-off between memory states.

Introduction

Overlap between events leads to interference and impairs memory for those events (McGeoch, 1942; Anderson, 2003). For example, at a conference you may talk to a colleague whom you had previously met over Zoom. Later you may have difficulty remembering either the original Zoom meeting or the subsequent conference conversation. The overlap between these events (e.g., the colleague) promotes retrieval of the past event (the meeting on Zoom) while you are trying to encode the present event (your conversation; Kuhl et al., 2010). As retrieval and encoding recruit distinct neural substrates and cannot be

engaged in simultaneously (Hasselmo et al., 2002), retrieving the past comes at the expense of encoding the present (Long and Kuhl, 2019). Although overlap is a critical factor in retrieval-mediated interference, two events may overlap along many dimensions and to varying degrees. Temporal overlap, or proximity in time, has been shown to enhance inference (Zeithamova and Preston, 2017), but it is unclear how temporal overlap contributes to interference. The aim of this study is to investigate the extent to which temporal overlap induces retrieval and, in turn, impacts interference.

Temporal information is a hallmark of episodic memory (Tulving, 1993) and is well known to impact how events are encoded and retrieved. The closer two events are in time and/or space, the more likely they are to be recalled together (Kahana, 1996; Manning et al., 2011) and the greater their neural similarity (Manns et al., 2007; Folkerts et al., 2018). Retrieved context theory (Howard and Kahana, 2002; Sederberg et al., 2008; Polyn et al., 2009; Lohanas and Kahana, 2014) provides an account for these effects whereby spatiotemporal context, an amalgamation of external stimuli and internal states is bound, via the hippocampus, to the present experience (Eichenbaum, 2004; Wang and Diana, 2017; Long and Kahana, 2019; Yonelinas et al., 2019)

Received May 26, 2021; revised Jan. 27, 2022; accepted Feb. 16, 2022.

Author contributions: D.E.S., I.L.M., and N.M.L. designed research; D.E.S. and I.L.M. performed research; D.E.S. and I.L.M. analyzed data; D.E.S. and I.L.M. wrote the first draft of the paper; D.E.S., I.L.M., and N.M.L. edited the paper; D.E.S., I.L.M., and N.M.L. wrote the paper.

N.M.L. is an iTHRIV Scholar. The iTHRIV Scholars Program is supported in part by the National Center for Advancing Translational Sciences of the National Institutes of Health under Awards UL1TR003015 and KL2TR003016. We thank Yujun Hong for assistance with data collection.

The authors declare no competing financial interests.

Correspondence should be addressed to Devyn E. Smith at ds9wk@virginia.edu or Nicole M. Long at nicdong@virginia.edu.

<https://doi.org/10.1523/JNEUROSCI.1091-21.2022>

Copyright © 2022 the authors

and is later used by the hippocampus as a cue to retrieve past experiences (Long et al., 2017). Comparison of activity patterns between study and test items, a recalled word or recognition probe, provides support for retrieved context theory in that the shorter the temporal distance between two items at study, the greater the pattern similarity between the study pattern of one item and the test pattern of the other item (Manning et al., 2011; Howard et al., 2012; El-Kalliny et al., 2019). Although contextually mediated retrieval is typically considered in relation to the test phase of an experiment, in principle, contextually mediated retrieval should occur whenever there is a contextual overlap between items. Such retrieval may occur automatically or independent from top-down demands (Smith et al., 2018). Therefore, we hypothesized that overlap in temporal context between two events produces retrieval during study and in turn promotes interference.

Here, we report a human scalp EEG study in which subjects studied two sets of object images in which the second set categorically overlapped with the first set. During study of the second set of object images, subjects were explicitly instructed to either encode the second (present) object or retrieve the first (past) object. These instructions were intended to bias subjects toward either an encoding or retrieval state. A retrieval state, or mode, is a tonically maintained mental set that is entered when there is need to engage episodic retrieval (Tulving, 1983; Rugg and Wilding, 2000). Our critical manipulation was the temporal distance between the first and second object, whereby the shorter the temporal distance between two objects, the greater their temporal contextual overlap. Following study, subjects completed a recognition task to probe their memory for all previously presented objects. To the extent that temporal contextual overlap influences interference, we should find that temporal distance modulates memory performance for the first and/or second objects. To the extent that temporal contextual overlap promotes retrieval, we should find that subjects are biased toward a retrieval state during second objects that are presented near in time to a categorically overlapping first object.

Materials and Methods

Subjects

Forty (34 female; age range = 18–37 years, mean age = 20.3 years) right-handed, native English speakers from the University of Virginia community participated. This sample size is based on our previous work in which we enrolled 40 participants (Long and Kuhl, 2019). All subjects had normal or corrected-to-normal vision. Informed consent was obtained in accordance with the University of Virginia Institutional Review Board for Social and Behavioral Research, and subjects were compensated for their participation. Three subjects were excluded from the final dataset: one who previously completed a behavioral version of the task, one who had poor task performance (recognition accuracy < 3 SDs of the mean of the full dataset), and one because of technical issues resulting in poor signal quality throughout the majority of the session. Thus, data are reported for the remaining 37 subjects. The raw, deidentified data and the associated experimental and analysis codes used in this study can be accessed via the Long Term Memory lab website (<https://longtermmemorylab.com>).

Mnemonic state task experimental design

Stimuli consisted of 576 object pictures, drawn from an image database with multiple exemplars per object category (Konkle et al., 2010). From this database, we chose 144 unique object categories and four exemplars from each category. For each subject, one exemplar in a set of four served as a List 1 object, one as a List 2 object, and the two remaining exemplars served as lures for the recognition phase. Object condition assignment was randomly generated for each subject.

General overview. In each of eight runs, subjects viewed two lists containing 18 object images. For the first list, each object was new (List 1 objects). For the second list (List 2 objects), each object was again new but was categorically related to an object from the first list. For example, if List 1 contained an image of a bench, List 2 would contain an image of a different bench (Fig. 1). During List 1, subjects were instructed to encode each new object. During List 2, however, each trial contained an instruction to either encode the current object (e.g., the new bench) or to retrieve the corresponding object from List 1 (the old bench). The critical manipulation was the distance between the corresponding List 1 and List 2 objects. We divided each list of 18 objects into thirds according to serial position (first [1–6], middle [7–12], and last [13–18]). The objects in the first third of List 1 were “paired” with the objects in the last third of List 2. For example, if List 1 contained an image of a bench in serial position 1, List 2 would contain an image of a different bench in serial position 13–18. The objects in the middle third of List 1 were paired with the objects in the middle third of List 2. The objects in the last third of List 1 were paired with the objects in the first third of List 2. We coded List 1 and List 2 objects as *near* and *far* based on the lag, or difference in serial position, between the two objects in a pair. List 1 and List 2 objects separated by < 18 intervening objects were coded as *near*; List 1 and List 2 objects separated by ≥ 18 intervening objects were coded as *far*. Following eight runs, subjects completed a two-alternative forced-choice recognition test that separately assessed memory for List 1 and List 2 objects.

List 1. On each trial, subjects saw a single object presented for 2000 ms followed by a 1000 ms interstimulus interval. Subjects were instructed to study the presented object in anticipation for a later memory test.

List 2. On each trial, subjects saw a cue word, either “OLD” or “NEW” for 2000 ms. The cue was followed by presentation of an object for 2000 ms, which was followed by a 1000 ms interstimulus interval. All objects in List 2 were nonidentical exemplars drawn from the same category as the objects presented in the immediately preceding List 1. That is, if a subject saw a bench and an apple during List 1, a different bench and a different apple would be presented during List 2. On trials with a NEW instruction (encode trials), subjects were to encode the presented object. On trials with an OLD instruction (retrieve trials), subjects tried to retrieve the categorically related object from the preceding List 1. Importantly, this design prevented subjects from completely ignoring List 2 objects following OLD instructions in that they could only identify the to-be-retrieved object category by processing the List 2 object.

Subjects completed eight runs with two lists in each run (List 1, List 2). Subjects viewed 18 objects per list, yielding a total of 288 object stimuli from 144 unique object categories. Subjects did not make a behavioral response during either List 1 or 2. Following the eight runs, subjects completed a two-alternative forced-choice recognition test.

Recognition phase. Following the eight runs, subjects completed the recognition phase. On each trial, subjects saw two exemplars from the same object category (e.g., two benches; Fig. 1). One object had previously been encountered either during List 1 or List 2. The other object was a lure and had not been presented during the experiment. Because both test probes were from the same object category, subjects could not rely on familiarity or gist-level information to make their response (Brainerd and Reyna, 2002). Trials were self-paced, and subjects selected (via button press) the previously presented object. Trials were separated by a 1000 ms interstimulus interval. There were a total of 288 recognition trials (corresponding to the 288 total List 1 and List 2 objects presented in the experiment). List 1 and List 2 objects never appeared in the same trial together; thus, subjects never had to choose between two previously presented objects. List 1 and List 2 objects were presented randomly throughout the test phase.

EEG data acquisition and preprocessing

EEG recordings were collected using a BrainVision system and an ActiCap equipped with 64 Ag/AgCl active electrodes positioned according to the extended 10–20 system. All electrodes were digitized at a sampling rate of 1000 Hz and were referenced to electrode FCz. Offline, electrodes were later converted to an average reference. Impedances of all electrodes were kept to < 50 kΩ. Electrodes that demonstrated high




Figure 1. Task design. During List 1, subjects studied individual objects (e.g., bench, apple). During List 2, subjects saw novel objects that were from the same categories as the objects shown in List 1 (e.g., a new bench, a new apple). Preceding each List 2 object was an OLD instruction cue or NEW instruction cue. The OLD cue signaled that subjects were to retrieve the corresponding object from List 1 (e.g., the old apple). The NEW cue signaled that subjects were to encode the current object (e.g., the new bench). Each run of the experiment contained a List 1 and List 2; object categories (e.g., bench) were not repeated across runs. List 1 and List 2 objects separated by <18 intervening objects were coded as *near*, and List 1 and List 2 objects separated by ≥ 18 intervening objects were coded as *far* (see Materials and Methods). Lines around the boxes are shown for illustrative purposes and were not present during the actual experiment. After eight runs, subjects completed a two-alternative forced-choice recognition test that tested memory for each List 1 and List 2 object. On each trial, a previously presented object, either from List 1 or List 2, was shown alongside a novel lure from the same category. The subject's task was to choose the previously presented object. List 1 and List 2 objects were never presented together.

impedance or poor contact with the scalp were excluded from the average reference. Bad electrodes were determined by voltage thresholding (see below).

Custom Python codes were used to process the EEG data. We applied a high pass filter at 0.1 Hz, followed by a notch filter at 60 Hz and harmonics of 60 Hz to each subject's raw EEG data. We then performed three preprocessing steps (Nolan et al., 2010) to identify electrodes with severe artifacts. First, we calculated the mean correlation between each electrode and all other electrodes as electrodes should be moderately correlated with other electrodes because of volume conduction. We z-scored these means across electrodes and rejected electrodes with z scores <-3 . Second, we calculated the variance for each electrode as electrodes with very high or low variance across a session are likely dominated by noise or have poor contact with the scalp. We then z-scored variance across electrodes and rejected electrodes with a $|z| > 3$. Finally, we expect many electrical signals to be autocorrelated, but signals generated by the brain versus noise are likely to have different forms of autocorrelation. Therefore, we calculated the Hurst exponent, a measure of long-range autocorrelation, for each electrode and rejected electrodes with a $|z| > 3$. Electrodes marked as bad by this procedure were excluded from the average reference. We then calculated the average voltage across all remaining electrodes at each time sample and rereferenced the data by subtracting the average voltage from the filtered EEG data. We used wavelet-enhanced independent component analysis (Castellanos and Makarov, 2006) to remove artifacts from eyeblinks and saccades.

EEG data analysis

We applied the Morlet wavelet transform (wave number 6) to the entire EEG time series across electrodes, for each of 46 logarithmically spaced frequencies (2–100 Hz; Long and Kahana, 2015). After log-transforming

the power, we downsampled the data by taking a moving average across 100 ms time intervals from either 4000 ms preceding to 4000 ms following object presentation during List 1 and List 2 or 0 ms preceding to 1000 ms following probe presentation for the recognition data. For each phase, we slid the window every 25 ms, resulting in 317 and 37 time intervals, respectively (80 and 10 nonoverlapping). Power values were then z-transformed by subtracting the mean and dividing by the SD power. Mean and SD power were calculated across all List 1 and List 2 objects or all recognition events, across time points for each frequency, which is analogous to performing a prestimulus baseline correction. z-transforming or baseline-correcting spectral power is a necessary step to both reduce the $1/f$ shape of the power spectrum, lower frequencies inherently have more power than higher frequencies, and to perform parametric statistics on the data (Cohen, 2014).

GLM

Trial-specific signals during List 2 were estimated using the GLM implemented via the sklearn linear model module in Python. We ran a separate GLM for each trial in which the trial was modeled as the regressor of interest and all other trials were combined into a single nuisance regressor (Mumford et al., 2012). Serial position (1–36, corresponding to List 1 [1–18] and List 2 [19–36]) was included as a single parametric regressor in each GLM to account for serial position effects. This parametric regressor predicts recognition memory accuracy, such that memory declines as serial position increases (mean = -0.0108 , SD = 0.0196 , $t_{(36)} = -3.3116$, $p = 0.0021$), and is consistent with other approaches for modeling a continuous variable (e.g., Tuladhar et al., 2007; Long et al., 2010; Spitzer et al., 2014). We fit trial-specific GLMs to the z-scored spectral power for each time point and frequency for each electrode to generate trial-level β values. These β values were used in all subsequent analyses.

Pattern classification analyses

Pattern classification analyses were performed using penalized (L2) logistic regression (penalty parameter = 1), implemented via the sklearn linear model module in Python. Before pattern classification analyses were performed on the List 2 data, an additional round of z-scoring was performed across features (electrodes and frequencies) to eliminate trial-level differences in spectral power (Kuhl and Chun, 2014; Long and Kuhl, 2018). Therefore, mean univariate activity was matched precisely across all conditions and trial types. Classifier performance was assessed in two ways. “Classification accuracy” represented a binary coding of whether the classifier successfully guessed the instruction condition. We used classification accuracy for general assessment of classifier performance (i.e., whether encode/retrieve instructions could be decoded). “Classifier evidence” was a continuous value reflecting the logit-transformed probability that the classifier assigned the correct instruction for each trial. Classifier evidence was used as a trial-specific, continuous measure of mnemonic state information, which was used to assess the degree of retrieval evidence present on *near* and *far* trials. The logic of using both classifier accuracy and classifier evidence is that, although accuracy indicates how well the classifier can distinguish encode versus retrieve trials, accuracy may obscure differences in conditions on which the classifier was not directly trained, for example, the distance (*near*, *far*) between objects. As an example, the classifier may correctly label both *near* encode and *far* encode trials as “encode”; however, it may have less confidence on the *near* compared with *far* trials, reflecting relatively greater retrieval state evidence on *near* trials.

We trained within-subject classifiers to discriminate List 2 encode versus retrieve trials based on a feature space comprised of all 63 electrodes $\times$ 46 logarithmically spaced frequencies ranging from 2 to 100 Hz. For each subject, we used leave-one-run-out cross-validated classification in which the classifier was trained to discriminate encode from retrieve instructions for seven of the eight runs and tested on the held-out run. For classification analyses in which we assessed classifier accuracy, we averaged β values over the 2000 ms stimulus interval. For analyses measuring classifier evidence, we averaged β values over four separate 500 ms time intervals across the 2000 ms stimulus interval. We assessed classifier evidence as a function of instruction (encode, retrieve), temporal distance (*near*, *far*), and/or retrieval status (success, failure; see below).

To measure the ability of the classifier to generalize across temporal distance, we trained and tested two separate classifiers to distinguish List 2 encode/retrieve trials. One classifier was trained on *near* trials and tested on *far* trials; the other classifier was trained on *far* trials and tested on *near* trials. As there was a slight imbalance in the number of encode and retrieve trials within each distance, we subsampled trials from the condition with the greater number of trials to match the condition with fewer trials. We repeated this procedure for 100 iterations and averaged the resulting classification accuracy values across the 100 iterations.

Retrieval status analysis

Because we did not explicitly measure retrieval success during the List 2 trials, we generated “retrieval success” and “retrieval failure” templates based on the recognition phase data. Specifically, we extracted stimulus-locked z-scored spectral power across 63 electrodes and 46 frequencies separately for hits (trials in which participants selected the target) and misses (trials in which participants selected the categorically related lure). We extracted z-power 500–800 ms following stimulus onset, as this interval has routinely been linked with retrieval success (Friedman and Johnson, 2000; Voss and Paller, 2008; Johnson et al., 2015). We averaged z-power across all subjects to generate a single retrieval success template and a single retrieval failure template.

After having generated the success/failure templates, we applied these templates to the List 2 data. Because we were interested in whether or not the corresponding List 1 object was retrieved at any point within the List 2 trial, we used the trial-level β values averaged across the stimulus interval (2000 ms). We correlated trial-level β values with the success and failure templates using a Pearson correlation. Each trial was assigned a label based on its correlation with the success and failure templates. A trial that was more positively correlated with the success

template was labeled retrieval success or 1, and a trial that was more positively correlated with the failure template was labeled retrieval failure or 0.

We calculated the average label as a function of distance (*near*, *far*) and instruction (encode, retrieve). An average label value of 0.5 means that a given condition was no more likely to be labeled retrieval success than retrieval failure. An average label value >0.5 means that a given condition was more likely to be labeled retrieval success than retrieval failure.

Statistical analyses

We used repeated-measures ANOVAs and paired-sample *t* tests to assess the effect of instruction (encode, retrieve) and temporal distance (*near*, *far*) on behavioral memory performance.

We used paired-sample *t* tests to compare classification accuracy across subjects to chance decoding accuracy, as determined by permutation procedures. Namely, for each subject, we shuffled the condition labels of interest (e.g., encode and retrieve for the List 2 instruction classifier) and then calculated classification accuracy. We repeated this procedure 1000 times for each subject and then averaged the 1000 shuffled accuracy values for each subject. These mean values were used as subject-specific empirically derived measures of chance accuracy.

We used repeated-measures ANOVAs and paired-sample *t* tests to assess the interaction between instruction (encode, retrieve), temporal distance (*near*, *far*), and time interval on retrieval evidence.

Results

Influence of temporal contextual overlap on interference

We first sought to replicate the finding that subjects are able to shift between encoding and retrieval states in a goal-directed manner (Long and Kuhl, 2019), by testing whether instructions influenced performance on the recognition task. Although encode/retrieve instructions only appeared during List 2, we also considered whether memory for List 1 objects was influenced by List 2 instructions (e.g., whether memory for the old bench was influenced by whether the new bench was associated with an encode vs retrieve instruction). A two-way, repeated-measures ANOVA with factors of list (1, 2) and instruction (encode, retrieve) revealed a list by instruction interaction ($F_{(1,36)} = 6.045$, $p = 0.0189$, $\eta_p^2 = 0.14$; Fig. 2A). This interaction was driven by numerically greater recognition for List 2 objects presented with an encode (mean = 82.88%, SD = 8.51%) relative to a retrieve instruction (mean = 80.52%, SD = 7.79%; difference between List 2 encode vs retrieve: $t_{(36)} = 2.1072$, $p = 0.0421$, Bonferroni-corrected $\alpha = 0.025$, Cohen's $d = 0.2938$) and numerically greater recognition for List 1 objects presented with a retrieve (mean = 84.27%, SD = 7.7%) relative to an encode instruction (mean = 83.3%, SD = 7.03%; difference between List 1 encode vs retrieve: $t_{(36)} = -1.7542$, $p = 0.0879$, Bonferroni-corrected $\alpha = 0.025$, Cohen's $d = 0.1324$).

To further demonstrate the impact that encode versus retrieve instructions have on memory behavior, we conducted an analysis of recognition phase reaction times. If subjects are able to shift between encoding and retrieval states, we would expect to find a list by instruction interaction such that memory responses are slowed for List 2 objects associated with a retrieve instruction and List 1 objects associated with an encode instruction. We assessed reaction times from correct trials only. A two-way, repeated-measures ANOVA with factors of list (1, 2) and instruction (encode, retrieve) revealed a significant main effect of list ($F_{(1,36)} = 24.84$, $p < 0.0001$, $\eta_p^2 = 0.41$) driven by faster reaction times for List 1 compared with List 2 objects. There was a main effect of instruction ($F_{(1,36)} = 7.27$, $p = 0.0106$, $\eta_p^2 = 0.17$) driven by faster reaction times for encode compared with retrieve instructions. There was a significant interaction between




Figure 2. Influence of mnemonic instructions on memory behavior. **A**, We assessed recognition accuracy as a function of list (1, 2) and instruction (orange represents encode; teal represents retrieve). We find a significant interaction between list and instruction ($p = 0.0189$) driven by greater accuracy for List 2 objects presented with an encode compared with a retrieve instruction and numerically greater accuracy for List 1 objects presented with a retrieve compared with an encode instruction. **B**, We assessed reaction times as a function of list and instruction. We find a significant interaction between list and instruction ($p = 0.0010$) driven by faster reaction times for List 2 objects presented with an encode compared with a retrieve instruction. **C**, We assessed the relationship between List 1 and List 2 object memory on a pair-by-pair basis for cases where either the List 1 object was remembered and the associated List 2 object was forgotten (L1R-L2F) or the List 1 object was forgotten and the associated List 2 object was remembered (L1F-L2R) separately for encode and retrieve instructions. There was a significant interaction between condition and instruction ($p = 0.0189$) driven by a greater proportion of L1F-L2R items for encode compared with retrieve trials and a numerically greater proportion of L1R-L2F items for retrieve compared with encode trials. * $p < 0.05$, *** $p < 0.001$, uncorrected.

list and instruction ($F_{(1,36)} = 12.9$, $p = 0.0010$, $\eta_p^2 = 0.26$; Fig. 2B). This interaction was driven by faster reaction times for List 2 objects presented with an encode (mean = 1.6456, SD = 0.4405) relative to a retrieve instruction (mean = 1.7904, SD = 0.3954; difference between List 2 encode vs retrieve: $t_{(36)} = -4.248$, $p = 0.0001$, Bonferroni-corrected $\alpha = 0.025$, Cohen's $d = 0.346$).

We next assessed the relationship between List 1 and List 2 object memory on a pair-by-pair basis to investigate the encoding-retrieval trade-off. We isolated cases in which either the List 1 object was remembered and the associated List 2 object was forgotten (L1R-L2F) or the List 1 object was forgotten and the associated List 2 object was remembered (L1F-L2R). To the extent that retrieval of List 1 objects trades off with encoding of List 2 objects, the proportion of L1R-L2F should be greater for retrieve compared with encode instructions and the proportion of L1F-L2R should be greater for encode compared with retrieve instructions. To test this hypothesis, we ran a 2×2 repeated-measures ANOVA with factors of instruction (encode, retrieve) and condition (L1R-L2F, L1F-L2R) and proportion as the dependent variable. There was no main effect of instruction ($F_{(1,36)} = 0.471$, $p = 0.497$, $\eta_p^2 = 0.01$) and the main effect of condition did not reach significance ($F_{(1,36)} = 3.923$, $p = 0.0553$, $\eta_p^2 = 0.10$). There was a significant interaction between condition and instruction ($F_{(1,36)} = 6.045$, $p = 0.0189$, $\eta_p^2 = 0.14$; Fig. 2C). This interaction was driven by a numerically greater proportion of L1F-L2R items when the instruction was to encode (mean = 0.1288, SD = 0.0507) compared with retrieve (mean = 0.1152, SD = 0.0569; difference between L1F-L2R encode vs retrieve: $t_{(36)} = 2.1733$, $p = 0.0364$, Bonferroni-corrected $\alpha = 0.025$, Cohen's $d = 0.2507$) and a numerically greater proportion of L1R-L2F items when the instruction was to retrieve (mean = 0.1528, SD = 0.0597) compared with encode (mean = 0.1329, SD = 0.0599; difference between L1R-L2F encode vs retrieve: $t_{(36)} = -2.0211$, $p = 0.0508$, Bonferroni-corrected $\alpha = 0.025$, Cohen's $d = 0.3325$). Together, these results support the interpretation that encoding and retrieval processes trade-off.

Having replicated our previous finding that instructions to encode and retrieve modulate behavior, we next sought to test the effect of temporal distance on recognition accuracy,




Figure 3. List 1 recognition accuracy by instruction and distance. We assessed recognition accuracy for List 1 objects as a function of instruction (orange represents encode; teal represents retrieve) and distance (near, far). We find a significant interaction between instruction and distance ($p = 0.0435$) driven by greater accuracy for near retrieve trials compared with near encode trials. * $p < 0.05$, uncorrected.

specifically for List 1 objects, as shorter temporal distance may impair List 1 memory specifically for encode trials. The intuition is that automatically retrieved near List 1 objects may be inhibited or suppressed by virtue of being goal-irrelevant during encode trials. This outcome would be analogous to the inhibition that is thought to occur during retrieval-induced forgetting (Anderson et al., 1994; Anderson, 2003).

We assessed whether the distance between objects, as well as the instruction given during List 2, influenced recognition memory of List 1 objects (Fig. 3). A two-way, repeated-measures ANOVA with factors of instruction (encode, retrieve) and distance (near, far), revealed a significant main effect of distance ($F_{(1,36)} = 4.916$, $p = 0.0330$, $\eta_p^2 = 0.12$) driven by greater recognition accuracy for far compared with near objects. The main effect of instruction did not reach significance ($F_{(1,36)} = 3.769$, $p = 0.0601$, $\eta_p^2 = 0.09$). There was a significant interaction between instruction and distance ($F_{(1,36)} = 4.381$, $p = 0.0435$, $\eta_p^2 = 0.11$), driven by greater accuracy for near retrieve trials




Figure 4. Retrieval state evidence. We trained an L2-logistic regression classifier to discriminate encode versus retrieve trials during List 2. The classifier was trained and tested on averaged β values across 63 electrodes and 46 frequencies. **A**, The classifier was trained on average β values across the 2000 ms stimulus interval. Mean classification accuracy across all subjects (solid vertical line) is shown along with a histogram of classification accuracies for individual subjects (gray bars) and mean classification accuracy for permuted data across all subjects (dashed vertical line). Mean classification accuracy for permuted data ranged from 49.7% to 50.27% across individual subjects (1000 permutations per subject). Mean classification accuracy was 53.32%, which differed significantly from chance ($p = 0.0148$). **B–D**, We trained and tested four classifiers on four 500 ms time intervals within the 2000 ms stimulus interval. **B**, When we average retrieval evidence over instruction, we find a significant interaction between distance and time interval ($p = 0.0018$) driven by greater retrieval evidence on *near* compared with *far* trials early in the stimulus interval. **C**, When we average retrieval evidence over distance, we find a significant interaction between instruction and time interval ($p = 0.0026$) driven by greater retrieval evidence on retrieve compared with encode trials late in the stimulus interval. **D**, We do not find a three-way interaction between instruction, distance, and time ($p = 0.869$). Error bars indicate SEM. * $p < 0.05$, *** $p < 0.001$, uncorrected.

(mean = 83.88%, SD = 11.11%) relative to *near* encode trials (mean = 80.9%, SD = 9.17%; difference between *near* encode vs *near* retrieve: $t_{(36)} = -2.6225$, $p = 0.0127$, Bonferroni-corrected $\alpha = 0.0167$, Cohen's $d = 0.2964$). Notably, recognition accuracy on *near* encode trials was significantly worse compared with both *far* encode trials ($t_{(36)} = -3.3417$, $p = 0.0020$, Bonferroni-corrected $\alpha = 0.0167$, Cohen's $d = 0.5561$) and *far* retrieve trials ($t_{(36)} = -3.1204$, $p = 0.0035$, Bonferroni-corrected $\alpha = 0.0167$, Cohen's $d = 0.4653$).

We observed decreased recognition accuracy for List 1 *near* objects when subjects attempted to encode the List 2 object compared with when they attempted to retrieve the *near* List 1 object. Indeed, *near* List 1 objects paired with the encode instruction are remembered worse than all other List 1 objects, strongly suggesting that bottom-up or automatic retrieval of the *near* List 1 object, when coupled with the top-down demand to encode the List 2 object, leads to suppression of the *near* List 1 object.

Influence of temporal contextual overlap on retrieval state

Our first goal was to replicate our previous finding that a pattern classifier trained on spectral signals can distinguish encode and retrieve trials (Long and Kuhl, 2019). We conducted a multivariate pattern classification analysis in which we trained a classifier to discriminate encode versus retrieve List 2 trials based on a feature space comprised of all 63 electrodes and 46 frequencies ranging from 2–100 Hz. For this analysis, we averaged β values over the 2000 ms stimulus interval. Using within-subject, leave-one-run-out classifiers, mean classification accuracy was 53.32% (SD = 7.75%), which was significantly greater than chance, as determined by permutation tests ($t_{(36)} = 2.5595$, $p = 0.0148$, Cohen's $d = 0.6043$; Fig. 4A).

We next sought to investigate the effect of temporal overlap on retrieval. If greater temporal contextual overlap between two events promotes retrieval, we would expect to find greater evidence for a retrieval state on *near* compared with *far* trials. Moreover, to the extent that this retrieval occurs automatically, we would expect to find greater evidence for a retrieval state early in the stimulus interval. Although temporal distance could interact with instruction, evidence for a retrieval state may be particularly strong for *near* retrieve trials, given that

temporal distance did not enhance memory for *near* List 1 objects on retrieve trials or impact memory for List 2 objects, we do not anticipate an interaction between temporal distance and instruction.

To investigate the effect of temporal distance on retrieval state evidence over time, we trained classifiers to discriminate encode versus retrieve trials using the average betas from four 500 ms time intervals across the 2000 ms stimulus interval. We conducted a repeated-measures ANOVA in which true (nonpermuted) retrieval evidence was the dependent variable and with factors of instruction (encode, retrieve), distance (*near*, *far*), and time interval (four 500 ms time intervals). We find a significant two-way interaction between distance and time interval ($F_{(3,108)} = 5.355$, $p = 0.0018$, $\eta_p^2 = 0.13$) whereby retrieval evidence is greater for *near* compared with *far* trials during the first two 500 ms time intervals (*near* vs *far*: 0–500, $t_{(36)} = 2.4899$, $p = 0.0175$, Cohen's $d = 0.598$; 500–1000, $t_{(36)} = 4.159$, $p = 0.0002$, Cohen's $d = 0.9056$; Bonferroni-corrected $\alpha = 0.0125$). Retrieval evidence does not differ during the second two 500 ms time intervals (*near* vs *far*: 1000–1500, $t_{(36)} = -1.2887$, $p = 0.2057$, Cohen's $d = 0.2772$; 1500–2000, $t_{(36)} = 0.9867$, $p = 0.3304$, Cohen's $d = 0.2492$; Bonferroni-corrected $\alpha = 0.0125$). We also find a main effect of distance ($F_{(1,36)} = 8.649$, $p = 0.0057$, $\eta_p^2 = 0.19$; Fig. 4B), with greater retrieval evidence for *near* compared with *far* trials. We find a significant two-way interaction between instruction and time interval ($F_{(3,108)} = 5.041$, $p = 0.0026$, $\eta_p^2 = 0.12$), whereby the largest differences in retrieval evidence between retrieve and encode trials occur during the last two 500 ms time intervals (encode vs retrieve: 0–500, $t_{(36)} = -1.4215$, $p = 0.1638$, Cohen's $d = 0.3759$; 500–1000, $t_{(36)} = -1.9205$, $p = 0.0628$, Cohen's $d = 0.4996$; 1000–1500, $t_{(36)} = -4.2349$, $p = 0.0002$, Cohen's $d = 1.2395$; 1500–2000, $t_{(36)} = -4.4573$, $p = 0.0001$, Cohen's $d = 1.2841$; Bonferroni-corrected $\alpha = 0.0125$). We find a significant main effect of instruction ($F_{(1,36)} = 22.31$, $p < 0.0001$, $\eta_p^2 = 0.38$; Fig. 4C), consistent with the results of the classifier trained on the full 2000 ms interval above. The two-way interaction between instruction and distance was not significant ($F_{(1,36)} = 1.932$, $p = 0.173$, $\eta_p^2 = 0.05$) nor was the three-way interaction between instruction, distance, and time interval ($F_{(3,108)} = 0.239$, $p = 0.869$, $\eta_p^2 = 0.0066$; Fig. 4D). Together, these results suggest that greater temporal contextual




Figure 5. Cross distance mnemonic state decoding. We trained two L2-logistic regression classifiers to discriminate encode versus retrieve based on average β values over the 2000 ms stimulus interval with 63 electrodes and 46 frequencies used as features. For each classifier, we show mean classification accuracy across all subjects (solid vertical line) along with a histogram of classification accuracies for individual subjects (gray bars) and mean classification accuracy for permuted data across all subjects (dashed vertical line). **A**, We trained the classifier on only List 2 *near* trials and tested the classifier on List 2 *far* trials. Mean classification accuracy for permuted data ranged from 49.73% to 50.40% across individual subjects (1000 permutations per subject). Mean classification accuracy was 52.98%, which was significantly greater than chance performance ($p = 0.0042$). **B**, We trained the classifier on only List 2 *far* trials and tested the classifier on List 2 *near* trials. Mean classification accuracy for permuted data ranged from 49.27% to 50.46% across individual subjects (1000 permutations per subject). Mean classification accuracy was 52.85%, which was significantly greater than chance performance ($p = 0.0038$).

overlap induces automatic retrieval independent of the actual instruction to either encode or retrieve.

Retrieval state mechanisms

We have found an increase in retrieval state evidence when objects appear closer together in time. Although our hypothesis is that this dissociation reflects greater instantiation of a retrieval state, the classifier may be indexing retrieval success or retrieval effort as opposed to a general retrieval state or mode (Rugg and Wilding, 2000). Specifically, by virtue of the shorter temporal distance, retrieval success might be greater for *near* compared with *far* objects. Likewise, by virtue of the longer temporal distance, retrieval might be more effortful for *far* compared with *near* objects. In our previous classification analysis, the classifier was trained using data from both *near* and *far* trials, meaning that the dissociation between encode/retrieve trials, and consequently, *near/far* trials, could be based on information exclusively from either *near* or *far* trials. Put another way, the classifier may have learned to distinguish either encode and retrieval success (i.e., *near* retrieve) trials or encode and retrieval effort (i.e., *far* retrieve) trials. Therefore, to demonstrate that a general retrieval state or mode underlies the dissociation between *near* and *far* trials, we trained two separate classifiers to distinguish encode/retrieve using only *near* or only *far* trials, and tested the classifiers on the other held-out distance (*far* or *near*) trials. The logic is that to the extent that the dissociation between encode/retrieve is supported by the same mechanism on both *near* and *far* trials, classifiers trained on one distance should generalize, reflected by above chance (50%) performance, to the other distance. To the extent that the dissociation between encode/retrieve is driven either by retrieval success or retrieval effort, the classifiers should fail to generalize to the other distance.

We conducted a multivariate pattern classification analysis in which we trained a classifier on only *near* or *far* trials to discriminate encode versus retrieve trials. We averaged β values across the stimulus interval (2000 ms) and used leave-one-run-out cross-validated classification. First, we trained a classifier to distinguish encode versus retrieve List 2 *near* trials and tested the classifier on the List 2 *far* trials (Fig. 5A). Mean classification accuracy was 52.98% (SD = 5.77%), which was significantly greater than chance performance ($t_{(36)} = 3.0602$, $p = 0.0042$, Cohen's

$d = 0.7225$; Fig. 5A). Next, we trained a classifier to distinguish encode versus retrieve List 2 *far* trials and tested the classifier on List 2 *near* trials (Fig. 5B). Mean classification accuracy was 52.85% (SD = 5.48%), which was significantly above chance ($t_{(36)} = 3.0933$, $p = 0.0038$, Cohen's $d = 0.7377$; Fig. 5B). The ability of these classifiers to generalize across distance suggests that neural signals during encode and retrieve trials are similar across temporal distance.

The cross-distance decoding analysis suggests that a general retrieval mode is present during both *near* and *far* trials. However, it is possible that the dissociation we observe between *near* and *far* trials in our prior analysis of retrieval state evidence is still driven in some part by retrieval success. Namely, greater retrieval state evidence may specifically be tracking *near* success trials.

To adjudicate between the possibilities that elevated retrieval evidence on *near* trials is because of a retrieval mode versus retrieval success, it is necessary to account for retrieval success during each List 2 trial. By design, there are no behavioral responses made during List 2 trials to equate the behavioral output across instructions. Therefore, we do not have a direct measure of retrieval success. However, we can generate a proxy of retrieval success by leveraging the recognition phase data. Specifically, we created a retrieval success and a retrieval failure template (Fig. 6A) across all subjects and assigned a retrieval "status" label to each List 2 trial of either retrieval success (1) or retrieval failure (0) based on how well a given trial correlated with each template (see Materials and Methods).

To validate our proxy of retrieval success, we first assessed whether temporal overlap impacts retrieval success. Given that retrieval success should be more likely for *near* compared with *far* objects, we predicted that *near* trials should be labeled "retrieval success" more often than *far* trials, reflected by an average label value closer to 1. We conducted a two-way, repeated-measures ANOVA with factors of instruction (encode, retrieve) and distance (*near*, *far*) and the average retrieval status label as the dependent variable (Fig. 6B). We find a significant main effect of distance ($F_{(1,36)} = 11.32$, $p = 0.0018$, $\eta_p^2 = 0.24$) driven by greater assignment of retrieval success for *near* compared with *far* trials. We find no main effect of instruction ($F_{(1,36)} = 0.104$, $p = 0.749$, $\eta_p^2 = 0.0029$) and no interaction




Figure 6. Impact of retrieval success on retrieval state evidence. **A**, We derived retrieval success and retrieval failure templates from hit and miss trials during the recognition phase. Each panel shows an across-subject electrode-frequency spectrogram of z-power during retrieval success (hits; left) and retrieval failure (misses; right) in which red represents z-power increases and blue represents z-power decreases. **B**, We assessed average retrieval status label as a function of instruction (orange represents encode; teal represents retrieve) and distance (*near*, *far*). We find a significant main effect of distance ($p = 0.0018$) driven by greater assignment of retrieval success for *near* compared with *far* trials. **C**, We assessed retrieval state evidence as a function of distance (solid line represents *near*; dashed line represents *far*) and retrieval status (red represents success; blue represents failure). We find a significant interaction between distance and time interval ($p = 0.0010$) driven by greater retrieval evidence on *near* compared with *far* trials early in the stimulus interval. Error bars indicate SEM. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, uncorrected.

between instruction and distance ($F_{(1,36)} = 0.351$, $p = 0.557$, $\eta_p^2 = 0.0097$).

Having established that our proxy for retrieval success matches our predictions, we next sought to test whether retrieval state evidence differs as a function of retrieval success. If the output of a classifier trained on all List 2 trials purely reflects a retrieval mode, *near* trials should show greater retrieval state evidence than *far* trials regardless of retrieval success. If the classifier purely reflects retrieval success, retrieval success trials should show greater retrieval state evidence than retrieval failure trials regardless of distance. We conducted a repeated-measures ANOVA in which true (nonpermuted) retrieval evidence was the dependent variable with factors of retrieval status (success, failure), distance (*near*, *far*), and time interval (four 500 ms time intervals). We find a significant main effect of distance ($F_{(1,36)} = 7.564$, $p = 0.0093$, $\eta_p^2 = 0.17$; Fig. 6C) driven by greater retrieval evidence on *near* compared with *far* trials. We find a significant two-way interaction between distance and time interval ($F_{(3,108)} = 5.853$, $p = 0.0010$, $\eta_p^2 = 0.14$) whereby retrieval evidence is greater for *near* compared with *far* trials for the first two 500 ms time intervals (*near* vs *far*: 0–500, $t_{(36)} = 2.579$, $p = 0.0141$, Cohen's $d = 0.6058$; 500–1000, $t_{(36)} = 3.973$, $p = 0.0003$, Cohen's $d = 0.8835$; Bonferroni-corrected $\alpha = 0.0125$). Retrieval evidence does not differ during the second two 500 ms time intervals (*near* vs *far*: 1000–1500, $t_{(36)} = -1.5492$, $p = 0.1301$, Cohen's $d = 0.3336$; 1500–2000, $t_{(36)} = 0.8985$, $p = 0.3749$, Cohen's $d = 0.2254$; Bonferroni-corrected $\alpha = 0.0125$). The three-way interaction between retrieval status, distance, and time interval was not significant ($F_{(3,108)} = 0.703$, $p = 0.552$, $\eta_p^2 = 0.02$). Bayes factor analysis revealed that a model without the three-way interaction term is preferred to a model with the three-way interaction by a factor of 13.1333. Together, these results suggest that, although retrieval may be more successful on *near* compared with *far* trials, retrieval success does not influence the dissociation in retrieval evidence between *near* and *far* trials.

Discussion

Here we show that temporal contextual overlap between events selectively increases interference and induces automatic retrieval. We used scalp EEG to measure memory brain states in a task

during which subjects were explicitly instructed to either encode the present event or retrieve a past, overlapping event. We find behavioral evidence that temporal overlap selectively leads to interference for past events when the top-down goal is to encode the present event. We find neural evidence that temporal overlap induces automatic retrieval independent from top-down demands to encode or retrieve. Critically, our neural results suggest that the retrieval state we observe is likely the result of a general retrieval mode (Rugg and Wilding, 2000), rather than a reflection of retrieval success or effort. Collectively, these findings demonstrate a link between temporal context, interference, and memory brain states.

We find that greater temporal overlap between events leads to a selective memory deficit for a past event when the top-down demand is to encode the present event. Overlap between events can lead to both proactive interference, in which learning about a past event impairs memory for the present, and retroactive interference, in which learning about a present event impairs memory for the past (Underwood, 1948; Crowder, 1976). Here we find that greater temporal overlap between two events leads to an increase in retroactive interference; however, this increase is selective for conditions in which subjects' top-down goal is to encode the currently presented stimulus. This result has striking similarity with retrieval-induced forgetting (Anderson et al., 1994; Anderson and Spellman, 1995). In paradigms that produce retrieval-induced forgetting, subjects retrieve a target (e.g., strawberry) based on a word stem (e.g., s____) and a cue (e.g., food) that is associated with other nontargets (e.g., tomato). Researchers theorize that cue-driven retrieval of the nontarget leads to suppression or inhibition which impairs subsequent memory for the nontarget (compare Perfect et al., 2004). As the strength, typically framed in terms of semantic overlap, between nontarget and cue increases, there is an increase in memory impairment, putatively because of stronger inhibition (Anderson et al., 1994). We extend these findings by showing that temporal overlap can likewise impair memory for nontargets, suggesting that greater temporal overlap may lead to inhibition of automatically retrieved items that are not goal-relevant.

Although in our study we find that temporal overlap is detrimental to later memory, there is evidence that temporal overlap

between events can facilitate behavior. Participants are better at associative inference tasks when associated events are studied close together in time (Zeithamova and Preston, 2017). Events presented close together in time are often recalled together (temporally clustered; Kahana, 1996; Long and Kahana, 2015) and overall recall performance increases as more events are temporally clustered (Sederberg et al., 2010; Healey et al., 2014). Temporal overlap may promote the integration of two separate events (Schlichting and Preston, 2015; Richter et al., 2016), which enables those events to cue one another during a memory test. It is possible that in our study the explicit instruction to encode interrupts or prevents integration leading to worse memory for the past event. Follow-up studies investigating the influence of temporal overlap in the absence of explicit instructions to encode or retrieve are needed to test this possibility.

We find induction of a retrieval state early in the stimulus interval when objects are closer together in time. We anticipated that greater temporal overlap would lead to increased retrieval on the basis of retrieved context theory. According to retrieved context theory (Howard and Kahana, 2002; Sederberg et al., 2008; Polyn et al., 2009; Lohanas and Kahana, 2014), spatiotemporal context is bound to items during study and used as a retrieval cue during test (Long et al., 2017), enabling items with overlapping spatiotemporal contexts to cue retrieval of one another (Manning et al., 2011). Consistent with retrieved context theory, we find more retrieval state evidence for objects with greater temporal overlap (*near* compared with *far* objects). Our observation of elevated retrieval state evidence on *near* trials, even when the instruction is to encode the present (or, conversely, when the instruction is to not retrieve the past), suggests that the retrieval we observe is the result of a bottom-up or stimulus-driven property of the object (e.g., its temporal contextual overlap with a past object) rather than the result of top-down or goal-driven demands. The dissociation in retrieval evidence as a function of temporal overlap may reflect the engagement of an automatic retrieval process, given that automatic retrieval is thought to be a fast, bottom-up process that can occur without top-down control (Moscovitch, 1994). That the largest retrieval state evidence dissociations between *near* and *far* trials occur within the first 1000 ms following stimulus onset is consistent with this interpretation. Collectively, these findings indicate that memory brain states can be impacted by both bottom-up and top-down influences.

We interpret the retrieval state effects that we observe as reflecting a general retrieval mode rather than serial position effects, retrieval success, or retrieval effort. By design, *near* and *far* objects occurred in systematically distinct serial positions (primacy and recency, respectively). To address this inherent confound, we fit a trial-level GLM to the *z*-transformed spectral power and included serial position as a parametric regressor based on a logistic-regression model fit of the behavioral data. We used this approach to limit the potential contribution of serial position to the observed retrieval state effects as distinct neural signals are recruited across primacy and nonprimacy positions (Sederberg et al., 2006). Given that the GLM cannot completely eliminate serial position effects, lingering primacy-related signals could contribute to the observed dissociation in retrieval evidence between *near* and *far* trials. However, we note that, as the pattern classifier is trained on data across all serial positions, we expect such a contribution to be limited.

The dissociation between *near* and *far* trials could be the result of other retrieval processes rather than a more general retrieval mode (Tulving, 1985; Rugg and Wilding, 2000). “Retrieval” as it

stands is a broad concept and can encompass multiple different subprocesses. We consider a retrieval state or mode as a content-independent process. Although typically retrieval mode has been considered within the framework of goal-directed or intentional remembering, we expect that a retrieval mode can also be engaged automatically based on bottom-up inputs (as demonstrated in the current study) and may align or be synonymous with the internal axis of attention (Chun et al., 2011). A retrieval mode is thought to be distinct from retrieval “orientation” in which specific cues or features are used to guide memory (Herron and Wilding, 2004; Hornberger et al., 2006a, b). Finally, both retrieval mode and orientation are separate from retrieval success and retrieval effort. After directing attention internally and orienting to particular cues to guide retrieval, an individual will either bring to mind the target item (retrieval success) or fail to bring to mind the target item, leading to effortful retrieval.

The retrieval process that we observe in the current study likely reflects a retrieval state given that a pattern classifier can distinguish encoding and retrieval across both *near* and *far* trials and that retrieval state evidence does not differ as a function of retrieval success. If the processes underlying *near* and *far* trials were entirely the product of retrieval success and retrieval effort, respectively, the cross-distance pattern classifier would be unable to distinguish encoding and retrieval across these trials. This is not to say that there are not potential differences in terms of retrieval success or effort between *near* and *far* trials, only that there exist shared mechanisms which enable the pattern classifier to generalize across these trials. Although we cannot rule out the potential influence of retrieval effort, the interpretation that elevated retrieval evidence on *near* compared with *far* trials reflects decreased retrieval effort would be inconsistent with our findings of greater retrieval evidence on retrieve compared with encode trials, given that one would expect more retrieval effort for retrieve trials. By leveraging the recognition phase data, we indexed retrieval success across *near* and *far* trials and found that retrieval state evidence is modulated by distance, but not retrieval success. It is important to note that our index of retrieval success is more likely to capture recollection-based as opposed to familiarity-based retrieval processes, although given the strong categorical overlap between the object pairs, we would anticipate high levels of familiarity for all objects regardless of temporal overlap. As the content of retrieval varies on every trial, it is unlikely that retrieval orientation differs systematically across *near* and *far* trials. Thus, the account best supported by these findings is that the dissociation in retrieval state evidence reflects a general retrieval mode. These results present an exciting avenue for future work to further dissociate these different retrieval subprocesses via multivariate methods and to more generally relate memory retrieval to internal attention.

Our results add to a growing body of work demonstrating the presence of neurally dissociable mnemonic states (Hasselmo et al., 2002; Hasselmo, 2005). Like other brain states (e.g., Kay and Frank, 2019), mnemonic states likely reflect sustained brain activity configurations. The shift between encoding and retrieval can occur on the order of milliseconds via theta oscillations which drive rapid shifts in entorhinal-hippocampal connectivity (Hasselmo et al., 2002). However, these states may operate along slower timescales and be mediated by acetylcholine (Hasselmo and McGaughy, 2004; Meeter et al., 2004). Mnemonic states predict subsequent memory (Long and Kuhl, 2019), impact the cortical location of stimulus representations (Long and Kuhl, 2021), and can influence behavior and decision-making

(Duncan et al., 2012; Duncan and Shohamy, 2016; Patil and Duncan, 2018). Memory encoding and retrieval may reflect two states along a continuum within the broader framework of external and internal attention, respectively (Chun et al., 2011). Here we show that mnemonic states in the cortex persist for several hundred milliseconds and are influenced by bottom-up stimulus properties, in addition to explicit top-down demands. We expect that mnemonic states fluctuate based on both stimuli and goals; to the extent that events overlap, there is the potential for automatic retrieval and a shift into a retrieval state. Tracking mnemonic state fluctuations will be critical for understanding both how these states are induced and how these states in turn impact behavior.

In conclusion, we show that temporal overlap between events induces retrieval and selectively impairs memory performance. These findings are consistent with theoretical models which propose that temporal information can cue retrieval (Howard and Kahana, 2002) and behavioral findings that retrieving non-goal-relevant information can lead to memory impairments (Anderson et al., 1994). More broadly, these findings point to a role for bottom-up stimulus features in driving mnemonic brain states.

References

- Anderson MC (2003) Rethinking interference theory: executive control and the mechanisms of forgetting. *J Mem Lang* 49:415–445.
- Anderson MC, Spellman BA (1995) On the status of inhibitory mechanisms in cognition: memory retrieval as a model case. *Psychol Rev* 102:68–100.
- Anderson MC, Bjork RA, Bjork EL (1994) Remembering can cause forgetting: retrieval dynamics in long-term memory. *J Exp Psychol Learn Mem Cogn* 20:1063–1087.
- Brainerd C, Reyna V (2002) Fuzzy-trace theory and false memory. *Curr Dir Psychol Sci* 11:164–169.
- Castellanos NP, Makarov VA (2006) Recovering EEG brain signals: artifact suppression with wavelet enhanced independent component analysis. *J Neurosci Methods* 158:300–312.
- Chun MM, Golomb JD, Turk-Browne NB (2011) A taxonomy of external and internal attention. *Annu Rev Psychol* 62:73–101.
- Cohen M (2014) Analyzing neural time series data: theory and practice. Cambridge, MA: MIT Press.
- Crowder R (1976) Principles of learning and memory. Hillsdale, NJ: Lawrence Erlbaum.
- Duncan KD, Sadanand A, Davachi L (2012) Memory's penumbra: episodic memory decisions induce lingering mnemonic biases. *Science* 337:485–487.
- Duncan KD, Shohamy D (2016) Memory states influence value-based decisions. *J Exp Psychol Gen* 145:1420–1426.
- Eichenbaum H (2004) Hippocampus: cognitive processes and neural representations that underlie declarative memory. *Neuron* 44:109–120.
- El-Kalliny MM, Wittig JH, Sheehan TC, Sreekumar V, Inati SK, Zaghoul KA (2019) Changing temporal context in human temporal lobe promotes memory of distinct episodes. *Nat Commun* 10:203.
- Folkerts S, Rutishauser U, Howard MW (2018) Human episodic memory retrieval is accompanied by a neural contiguity effect. *J Neurosci* 38:4200–4211.
- Friedman D, Johnson R Jr (2000) Event-related potential (ERP) studies of memory encoding and retrieval: a selective review. *Microsc Res Tech* 51:6–28.
- Hasselmo ME (2005) What is the function of hippocampal theta rhythm? Linking behavioral data to phasic properties of field potential and unit recording data. *Hippocampus* 15:936–949.
- Hasselmo ME, McGaughy J (2004) High acetylcholine levels set circuit dynamics for attention and encoding and low acetylcholine levels set dynamics for consolidation. *Prog Brain Res* 145:207–231.
- Hasselmo ME, Bodelon C, Wyble BP (2002) A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput* 14:793–817.
- Healey MK, Crutchley P, Kahana MJ (2014) Individual differences in memory search and their relation to intelligence. *J Exp Psychol Gen* 143:1553–1569.
- Herron J, Wilding E (2004) An electrophysiological dissociation of retrieval mode and retrieval orientation. *Neuroimage* 22:1554–1562.
- Hornberger M, Rugg MD, Henson RN (2006a) ERP correlates of retrieval orientation: direct versus indirect memory tasks. *Brain Res* 1071:124–136.
- Hornberger M, Rugg MD, Henson RN (2006b) fMRI correlates of retrieval orientation. *Neuropsychologia* 44:1425–1436.
- Howard MW, Kahana MJ (2002) A distributed representation of temporal context. *J Math Psychol* 46:269–299.
- 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* 22:1833–1847.
- Johnson JD, Price MH, Leiker EK (2015) Episodic retrieval involves early and sustained effects of reactivating information from encoding. *Neuroimage* 106:300–310.
- Kahana MJ (1996) Associative retrieval processes in free recall. *Mem Cognit* 24:103–109.
- Kay K, Frank LM (2019) Three brain states in the hippocampus and cortex. *Hippocampus* 29:184–238.
- Konkle T, Brady TF, Alvarez GA, Oliva A (2010) Conceptual distinctiveness supports detailed visual long-term memory for real-world objects. *J Exp Psychol Gen* 139:558–578.
- Kuhl BA, Chun MM (2014) Successful remembering elicits event-specific activity patterns in lateral parietal cortex. *J Neurosci* 34:8051–8060.
- Kuhl BA, Shah AT, DuBrow S, Wagner AD (2010) Resistance to forgetting associated with hippocampus-mediated reactivation during new learning. *Nat Neurosci* 13:501–508.
- Lohnas LJ, Kahana MJ (2014) Compound cuing in free recall. *J Exp Psychol Learn Mem Cogn* 40:12–24.
- Long NM, Kahana MJ (2015) Successful memory formation is driven by contextual encoding in the core memory network. *Neuroimage* 119:332–337.
- Long NM, Kahana MJ (2019) Hippocampal contributions to serial-order memory. *Hippocampus* 29:252–259.
- Long NM, Kuhl BA (2018) Bottom-up and top-down factors differentially influence stimulus representations across large-scale attentional networks. *J Neurosci* 38:2495–2504.
- Long NM, Kuhl BA (2019) Decoding the tradeoff between encoding and retrieval to predict memory for overlapping events. *Neuroimage* 201:116001.
- Long NM, Kuhl BA (2021) Cortical representations of visual stimuli shift locations with changes in memory states. *Curr Biol* 31:1119–1126.
- Long NM, Öztekin I, Badre D (2010) Separable prefrontal cortex contributions to free recall. *J Neurosci* 30:10967–10976.
- Long NM, Sperling MR, Worrell GA, Davis KA, Gross RE, Lega BC, Jobst BC, Sheth SA, Zaghoul K, Stein JM, Kahana MJ (2017) Contextually mediated spontaneous retrieval is specific to the hippocampus. *Curr Biol* 27:1074–1076.
- 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* 108:12893–12897.
- Manns JR, Howard MW, Eichenbaum H (2007) Gradual changes in hippocampal activity support remembering the order of events. *Neuron* 56:530–540.
- McGeoch JA (1942) The psychology of human learning: an introduction. New York: Longmans, Green and Co.
- Meeter M, Murre J, Talamini L (2004) Mode shifting between storage and recall based on novelty detection in oscillating hippocampal circuits. *Hippocampus* 14:722–741.
- Moscovitch M (1994) Cognitive resources and dual-task interference effects at retrieval in normal people: the role of the frontal lobes and medial temporal cortex. *Neuropsychology* 8:524–534.
- Mumford JA, Turner BO, Ashby FG, Poldrack RA (2012) Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *Neuroimage* 59:2636–2643.
- Nolan H, Whelan R, Reilly R (2010) FASTER: fully automated statistical thresholding for EEG artifact rejection. *J Neurosci Methods* 192:152–162.
- Patil A, Duncan KD (2018) Lingering cognitive states shape fundamental mnemonic abilities. *Psychol Sci* 29:45–55.

- Perfect TJ, Stark LJ, Tree JJ, Moulin CJ, Ahmed L, Hutter R (2004) Transfer appropriate forgetting: the cue-dependent nature of retrieval-induced forgetting. *J Mem Lang* 51:399–417.
- Polyn SM, Norman KA, Kahana MJ (2009) A context maintenance and retrieval model of organizational processes in free recall. *Psychol Rev* 116:129–156.
- Richter FR, Chanales AJ, Kuhl BA (2016) Predicting the integration of overlapping memories by decoding mnemonic processing states during learning. *Neuroimage* 124:323–335.
- Rugg MD, Wilding EL (2000) Retrieval processing and episodic memory. *Trends Cogn Sci* 4:108–115.
- Schlichting ML, Preston AR (2015) Memory integration: neural mechanisms and implications for behavior. *Curr Opin Behav Sci* 1:1–8.
- Sederberg PB, Gauthier LV, Terushkin V, Miller JF, Barnathan JA, Kahana MJ (2006) Oscillatory correlates of the primacy effect in episodic memory. *Neuroimage* 32:1422–1431.
- Sederberg PB, Howard MW, Kahana MJ (2008) A context-based theory of recency and contiguity in free recall. *Psychol Rev* 115:893–912.
- Sederberg PB, Miller JF, Howard MW, Kahana MJ (2010) The temporal contiguity effect predicts episodic memory performance. *Mem Cognit* 38:689–699.
- Smith SM, Handy JD, Hernandez A, Jacoby LL (2018) Context specificity of automatic influences of memory. *J Exp Psychol Learn Mem Cogn* 44:1501–1513.
- Spitzer B, Gloel M, Schmidt TT, Blankenburg F (2014) Working memory coding of analog stimulus properties in the human prefrontal cortex. *Cereb Cortex* 24:2229–2236.
- Tuladhar AM, Huurne N, Schoffelen JM, Maris E, Oostenveld R, Jensen O (2007) Parieto-occipital sources account for the increase in alpha activity with working memory load. *Hum Brain Mapp* 28:785–792.
- Tulving E (1983) *Elements of episodic memory*. Oxford: Oxford UP.
- Tulving E (1985) How many memory systems are there? *Am Psychologist* 40:385–398.
- Tulving E (1993) What is episodic memory? *Curr Dir Psychol Sci* 2:67–70.
- Underwood B (1948) Retroactive and proactive inhibition after five and forty-eight hours. *J Exp Psychol* 38:29–38.
- Voss J, Paller K (2008) Neural substrates of remembering: electroencephalographic studies; learning and memory: a comprehensive reference. In: *Memory systems* (Byrne JH, ed). Oxford: Elsevier.
- Wang F, Diana RA (2017) Temporal context in human fMRI. *Curr Opin Behav Sci* 17:57–64.
- Yonelinas A, Ranganath C, Ekstrom A, Wiltgen B (2019) A contextual binding theory of episodic memory: systems consolidation reconsidered. *Nat Rev Neurosci* 20:364–375.
- Zeithamova D, Preston AR (2017) Temporal proximity promotes integration of overlapping events. *J Cogn Neurosci* 29:1311–1323.