

Temporal Proximity Promotes Integration of Overlapping Events

Dagmar Zeithamova¹ and Alison R. Preston²

Abstract

■ Events with overlapping elements can be encoded as two separate representations or linked into an integrated representation, yet we know little about the conditions that promote one form of representation over the other. Here, we tested the hypothesis that the proximity of overlapping events would increase the probability of integration. Participants first established memories for house–object and face–object pairs; half of the pairs were learned 24 hr before an fMRI session, and the other half 30 min before the session. During scanning, participants encoded object–object pairs that overlapped with the initial pairs acquired on the same or prior day. Participants were also scanned as they made inference judgments about the relationships among overlapping pairs learned on the same or different day. Participants were more accurate and faster when inferring relationships among memories learned on the same day relative to those acquired across days,

suggesting that temporal proximity promotes integration. Evidence for reactivation of existing memories—as measured by a visual content classifier—was equivalent during encoding of overlapping pairs from the two temporal conditions. In contrast, evidence for integration—as measured by a mnemonic strategy classifier from an independent study [Richter, F. R., Chanale, A. J. H., & Kuhl, B. A. Predicting the integration of overlapping memories by decoding mnemonic processing states during learning. *Neuroimage*, 124, 323–335, 2016]—was greater for same-day overlapping events, paralleling the behavioral results. During inference itself, activation patterns further differentiated when participants were making inferences about events acquired on the same day versus across days. These findings indicate that temporal proximity of events promotes integration and further influences the neural mechanisms engaged during inference. ■

INTRODUCTION

Events often overlap with one another, sharing content such as people, places, and objects. Overlapping event content may promote interactions among memories triggering one of two mechanisms that determine how new information is represented with respect to existing knowledge. In one case, overlapping content may serve as a retrieval cue to reactivate related memories (Zeithamova, Dominick, & Preston, 2012). The new event and reactivated knowledge may then be represented by overlapping neural populations, leading to the formation of an integrated memory representation that links related events (Schlichting, Zeithamova, & Preston, 2014; Zeithamova, Dominick, et al., 2012; Zeithamova & Preston, 2010; Shohamy & Wagner, 2008; Eichenbaum, 1999). Such integrated memories can support novel decisions, such as inference, by combining knowledge from events experienced at different times (Schlichting, Mumford, & Preston, 2015; Schlichting et al., 2014; Zeithamova & Preston, 2010; Shohamy & Wagner, 2008). In contrast to this integration mechanism, some theoretical models emphasize that overlapping events are represented by distinct or pattern-separated neural populations to minimize interference among similar memories

(Kumaran, Hassabis, & McClelland, 2016; O'Reilly & Rudy, 2000; McClelland, McNaughton, & O'Reilly, 1995). Empirical data provide evidence for both representational strategies; yet, little is known about factors that promote one mechanism versus another. Here, we test the hypothesis that the temporal proximity of events influences whether overlapping experiences are integrated or separated.

Recent data have shown that events close together in time demonstrate more similar representations in the hippocampus, amygdala, and pFC than events separated by longer delays (Cai et al., 2016; Rashid et al., 2016; Ezzyat & Davachi, 2014; Hsieh, Gruber, Jenkins, & Ranganath, 2014). For example, distinct spatial contexts experienced within the same day are represented by a highly overlapping population of neurons within the CA₁ subregion of the hippocampus (Cai et al., 2016). Such integrated representations lead to the generalization of fear responses from one spatial context to another. However, when spatial contexts are experienced 1 week apart, distinct populations of CA₁ cells represent the two contexts, and the generalization of fear across contexts does not occur. These findings suggest that temporal proximity may be a key factor that promotes integration.

Temporal proximity may drive integration through a memory tagging and allocation mechanism, whereby neurons and synapses recruited to represent a recent episode are more readily engaged for new events that occur

¹University of Oregon, ²University of Texas at Austin

within hours of the original episode (Cai et al., 2016; Rashid et al., 2016; Silva, Zhou, Rogerson, Shobe, & Balaji, 2009). Recruitment of the same neural ensembles through tagging and allocation thus results in overlapping population codes for temporally proximal events. At the cognitive level, a similar idea has been expressed by the temporal context model (Howard & Kahana, 2002), which proposes that temporally proximal events are linked by shared context information. Even when events are separated in time, this model suggests that overlapping memory content may serve to reinstate a prior temporal context, leading to the integration of temporally distant events (Howard, Jing, Rao, Probyn, & Datey, 2009). However, it remains an open question how temporal distance constrains the likelihood of memory reinstatement and integration.

Existing research suggests that reactivation of a stored memory may increase its malleability and allow for memory updating (Bridge & Voss, 2014; Diekelmann, Büchel, Born, & Rasch, 2011; Hupbach, Hardt, Gomez, & Nadel, 2008), even when the memory is a day or more old (Rashid et al., 2016; Hupbach et al., 2008). In contrast, the memory tagging and allocation hypothesis (Cai et al., 2016; Rashid et al., 2016; Silva et al., 2009) proposes that recent memories are more likely to be reinstated than temporally distant events, which would result in integration only for temporally proximal events. Furthermore, remote memories that have been stabilized through consolidation are thought to be less susceptible to modification (Frankland & Bontempi, 2005; Squire & Alvarez, 1995). Therefore, even if events encountered on prior days are reactivated during new encoding, a separate memory may be formed for the new information. Here, our goal was to test how the temporal proximity of events impacts reinstatement of prior related memories and integration across events.

Participants underwent fMRI while encoding object pairs that overlapped with previously learned associations. Half of the pairs overlapped with memories (face-object and house-object pairs) acquired on the previous day, whereas the other half overlapped with memories (a different set of face-object and house-object pairs) acquired within the same experimental session. We predicted that unseen, related memory content (face, house) would be more likely to be reactivated and integrated for the same-day than prior-day memories. To test the reactivation hypothesis, participants were also scanned while they viewed individual images of objects, faces, and houses during a functional localizer. Data from this localizer were used to train a multivoxel pattern classifier to differentiate the visual content from individual participants' ventral visual cortex activation patterns; the trained classifier was then applied to the overlapping object pair encoding data to estimate the degree of evidence for reactivation of related face or house information. These reactivation estimates were then compared for events that overlapped with information acquired

on the same day and events that overlapped with memories acquired on the prior day.

We assessed how temporal proximity impacted whether participants integrated overlapping events neurally as well as behaviorally. Behaviorally, we assessed the ease of linking information across related events by having participants perform an inference task, in which they were asked to make judgments about the indirectly related elements of two overlapping pairs. Our prior work has shown that reactivation and integration of prior memories with new content facilitates inference (Schlichting et al., 2014; Zeithamova, Dominick, et al., 2012). Here, we tested the hypothesis that inference would be superior for overlapping events experienced on the same day because of the increased likelihood of integration.

Neurally, we employed a multivariate classification approach to index the degree of integration during encoding of overlapping events, as opposed to simply retrieving old information or encoding the new associations without reference to existing memories. We hypothesized that overlapping events from the same-day condition would be more likely to be associated with an integration strategy. This multivariate approach was based on a neural classifier from an independent study, which demonstrated that the process of memory integration evokes a neutrally distinct signature across the whole brain that can be differentiated from simple retrieval or encoding strategies (Richter, Chanales, & Kuhl, 2016).

Specifically, Richter and colleagues presented participants with events that overlapped with previously learned associations, using a paradigm similar to this study. In the first experiment, participants were instructed during overlapping event encoding to use one of three strategies: *retrieve* the prior association without encoding the new one, *encode* the new information without reference to the prior association, or *integrate* the new association with the prior association. A neural classifier trained on normalized, whole-brain activation patterns was able to discriminate the three memory strategies—retrieve, encode, integrate—from one another on a trial-by-trial basis (Richter et al., 2016). Importantly, the mnemonic strategy classifier trained on the group of participants given explicit strategic instructions was then applied to a new set of participants who were not given explicit strategy cues. The classifier successfully predicted spontaneous integration in this independent group of participants. In this study, we used the whole-brain, across-participant mnemonic classifier trained on data from Richter et al. (2016) to test our hypothesis that temporally proximal events, that is, those encoded within the same day, would show greater evidence of memory integration.

A final goal of our study was to determine whether the temporal proximity of events impacts the processes brought to bear as participants make inferences about their relationships. To this end, we also scanned participants as they completed the inference task for which they made novel decisions about the relationships among overlapping

events. Memory integration during encoding has been shown to promote novel inference (Schlichting et al., 2014; Zeithamova, Dominick, et al., 2012; Shohamy & Wagner, 2008). However, successful inference may also be achieved through retrieval and joint consideration of pattern separated memories (Kumaran et al., 2016; Zeithamova & Preston, 2010; Greene, Gross, Elsinger, & Rao, 2006; Preston, Shrager, Dudukovic, & Gabrieli, 2004), albeit at the cost of increased RT and decreased accuracy (Schlichting et al., 2014, 2015). We hypothesized that the inference strategies may differ depending on whether an integrated memory has been already formed during encoding. Behaviorally, we predicted that participants would be faster and more accurate when making inferences involving same-day than prior-day events because of enhanced integration of same-day events during encoding. Neurally, we predicted that activation patterns during novel inferences would differ for the same-day and prior-day conditions, presumably reflecting relative reliance on different inference strategies (retrieval and recombination of separate memories vs. retrieval of integrated memories) across the two conditions. Unlike the hypotheses formulated for reactivation and integration during encoding, we did not have independent data for training of a neural classifier to differentiate between the hypothesized processes during inference. We thus employed a cross-validation approach in which a neural classifier was trained to discriminate whole-brain

activation patterns for inference trials from the two conditions on a within-participant basis.

METHODS

Participants

Thirty-four young adults (age 18–34 years, mean = 23 years; 18 women) participated in the experiment after giving an informed consent in accordance with the University of Texas at Austin institutional review board policy. Two participants did not complete the scanned portion of the experiment, one because of scanner failure and one because of failure to learn the paired associates during the training before scanning. Data from three additional participants were excluded from analysis because of incorrect slice prescription (1), loss of behavioral responses (1), and failure to follow task instructions (1). Data from the remaining 29 participants were used in the final analyses.

Procedures

Stimuli were grayscale images of faces (F), houses (H), and common objects (O). Images were organized into triads consisting of either a face and two objects (60 face–object–object triads) or a house and two objects (60 house–object–object). Stimuli from each triad were presented as two overlapping paired associates (Figure 1). Initial

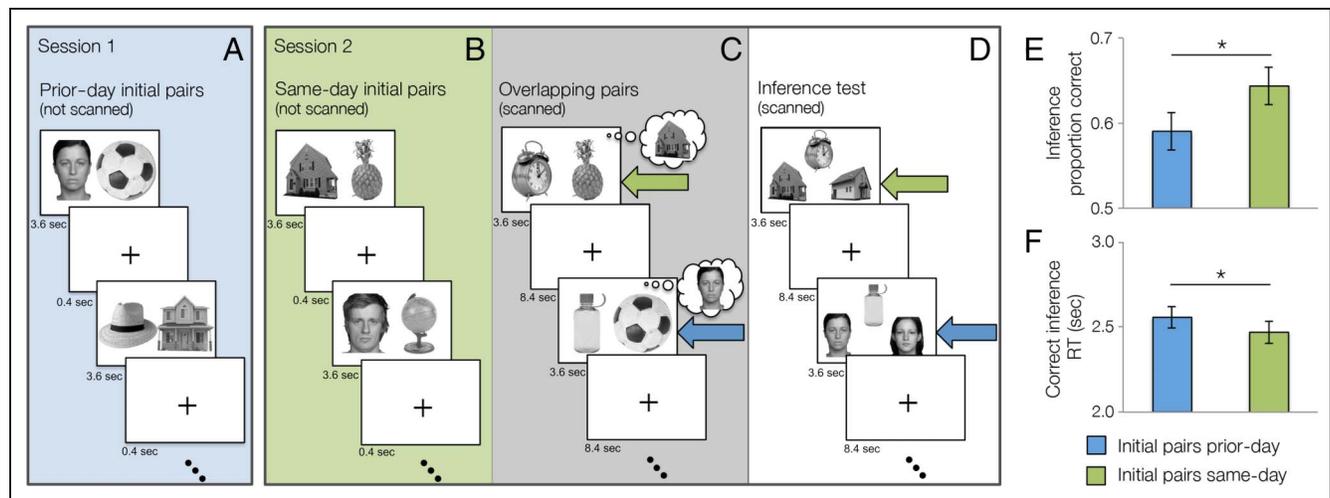


Figure 1. Schematic of study design and behavioral results. (A) Session 1: Training on prior-day initial pairs. One day before scanning, participants were trained on 30 face–object and 30 house–object pairs across three study–test repetitions (test not depicted). (B–D) Session 2 on the following day had three successive phases: (B) Training of same-day initial pairs. Participants were trained outside scanner on a novel set of 30 face–object and 30 house–object pairs across three study–test repetitions (test not depicted). (C) Encoding of overlapping pairs. During slow event-related scanning, participants encoded 120 object–object pairs. One object from each pair was novel and one was encountered previously during prior-day initial training (blue arrow) or same-day initial training (green arrow). (D) Inference test. During slow event-related scanning, participants were given a surprise test of 120 relationships that required linking information across events. Their task was to infer which house or face was indirectly related to objects that shared a common associate. Half of the faces and half of the houses were encountered on prior day (blue arrow) and the other half on the same day (green arrow). No explicit cues distinguished inference trials from the two temporal conditions during the test phase, nor were participants required to retrieve whether the stimuli were experienced on the same or different day. All tests consisted of two-alternative forced-choice judgments with the foil stimuli always of the same type and condition as the correct choice. (E) Inference accuracy split by the two temporal conditions, prior day (blue bar) or same day (green bar). (F) Median RTs for correct inferences for the two conditions. Error bars denote across-participant standard error. In E & F, asterisk (*) denotes a significant difference between conditions (paired *t*-test $p < 0.05$).

associations, referred to as AB pairs, consisted of a face or a house (the “A” item) paired with an object (the “B” item). Overlapping associations, referred to as BC pairs, consisted of the same B objects now paired with novel objects (the “C” items).

The experiment took place over two sessions, approximately 24 hr apart (Figure 1). During Session 1 (not scanned; Figure 1A), participants learned half of the initial pairs (30 face–object and 30 house–object) across three study–test repetitions (hereafter referred to as “prior-day” pairs). During Session 2, participants learned the remaining set of initial pairs (“same-day” pairs: 30 new face–object and 30 new house–object) over three study–test repetitions, 30–60 min before scanning (Figure 1B). For each AB study trial, a pair of stimuli was presented for 3.6 sec followed by a 0.4-sec fixation. The left–right position of the two items was randomized. During each AB test trial, one stimulus of a pair was presented on the top of the screen (e.g., A cue) and two choice stimuli were presented on the bottom of the screen (e.g., a correct B and a foil B from a different AB pair). On each test trial, the foil stimulus was constrained to match the type and condition of the correct stimulus. In other words, if A was a male face from a face–object–object triad, then the B foil was an object paired with a different male face during the same training session. Each AB test trial was equally likely to be cued by the A stimulus or the B stimulus. The correct choice was presented on the left or right position with equal probability. Each test trial lasted for 3.6 sec during which time the participant indicated whether the left or right choice stimulus was paired with the cue stimulus. Each test trial was followed by a 0.4-sec fixation. AB training lasted about a half an hour on each day.

The second phase of Session 2 consisted of scanned encoding of the 120 overlapping object–object (BC) pairs (Figure 1C). Each overlapping pair was studied only once, with half of the BC associates overlapping with prior-day AB pairs and the other half overlapping with same-day AB pairs. On each BC study trial, a pair of objects was presented for 3.6 sec followed by 8.4-sec fixation baseline for a total of 12 sec per trial. BC encoding occurred across three fMRI runs, each lasting 8 min. All triad types were balanced across runs, that is, the BC pairs overlapping with AB pairs from a different temporal condition (prior day, same day) and associated with different stimulus content (faces, houses) were evenly distributed across the runs. Presentation order of the pairs within a run was randomized. Participants were not explicitly informed about the overlap between pairs they learned previously and the new pairs, nor were they asked to retrieve the prior AB associates. Instead, participants were informed that they would be challenged to learn associations in a single trial, as opposed to the three trial learning they had previously experienced.

Following the encoding of overlapping events, participants were given a surprise AC inference test that

required them to relate elements across overlapping pairs (Figure 1D). For the face–object–object triads, participants were told that they had initially seen a person (the A image) owning an object (the B image). Later, they saw the same object (B) paired with another object (C). From these pairings, participants were told that they could infer that the second (C) object must also belong to the same person (A). A similar explanation was given for the house–object–object triads, wherein participants were told they could infer that C items could be found in the same house (A) as their corresponding B associate. We refer to this test as “AC test” or “inference test” as participants needed to infer a relationship between two items (A and C) that were never presented together but were both related to a common associate (the B element). Participants performed the surprise test for 120 inferential AC relationships during fMRI scanning. On each AC inference trial, one stimulus (e.g., a C object) served as a cue presented on the top of the screen, and two choice stimuli (e.g., correct A and a foil A from a different triad) were presented on the bottom of the screen. The foil stimulus was constrained to match the triad type of the correct stimulus. Each AC trial was equally likely to be cued by the A stimulus or the C stimulus. On each trial, the cue and choice stimuli were presented for 3.6 sec followed by 8.4-sec fixation baseline for a total of 12 sec per trial. Inference testing occurred across three fMRI runs, each lasting 8 min, with a balanced distribution of triad types (inferences involving a face or house from the same-day and prior-day conditions). Presentation order within each run was randomized.

After the completion of the inference test, participants underwent testing for their memory of the initial (AB) and overlapping (BC) pairs (not scanned). The test procedure was the same as in previous phases, except that trials lasted 4 sec (a combined 3.6-sec stimulus and response period, 0.4-sec fixation). BC pairs were tested before AB pairs.

Session 2 concluded with the collection of two anatomical scans (see MRI Data Acquisition) followed by a functional localizer task used for the training of a visual content classifier based on multivoxel pattern analysis (MVPA). During the localizer, participants viewed individual objects, faces, and houses in a blocked design, using different stimuli from those presented during the associative inference task. Each block consisted of successive presentation of nine stimuli (2 sec each). Participants performed a 1-back working memory task, in which they pressed a button whenever the current stimulus was a repetition of the one immediately preceding. Within each block, there were one or two stimulus repeats. The order of face, object, and house blocks was counterbalanced within and across participants. In total, there were three localizer runs, each lasting 8 min.

MRI Data Acquisition

Whole-brain imaging data were acquired on a 3.0-T GE Signa MRI system (GE Medical Systems, Milwaukee, WI).

Functional images were acquired using a multiecho GRAPPA parallel EPI sequence with thirty-one 3-mm-thick oblique axial slices (0.6 mm gap), 20° off the AC–PC line (repetition time = 2 sec, echo time = 30 msec, 2 shot, flip angle = 90°, 64 × 64 matrix, 3.75 × 3.75 mm in-plane resolution, interleaved slice acquisition). For each functional scan, the first six EPI volumes were discarded to allow for T1 stabilization. Head movement was minimized using foam padding. Two structural images were collected: a T2-weighted flow-compensated spin-echo pulse sequence with the same slice prescription as the functional images (repetition time = 3 sec, echo time = 68 msec, 256 × 256 matrix, 1 × 1 mm in-plane resolution), and a high-resolution T1-weighted SPGR scan was acquired in the sagittal plane using a 1.3-mm slice thickness with 1-mm² in-plane resolution.

Preprocessing of fMRI Data

Data were preprocessed and analyzed using SPM5 (Wellcome Department of Cognitive Neurology, London, UK) and custom MATLAB (The MathWorks, Natick, MA) routines. Functional images were realigned to the first volume in the time series to correct for motion and coregistered to the T2-weighted structural image. The T2-weighted structural image was coregistered to the high-resolution SPGR, and the coregistration parameters were then applied to the functional images. Finally, the functional images were high-pass filtered with a 128-sec filter and converted to percent signal.

The majority of analyses were performed in the native space of each participant without spatial smoothing. The exception was an analysis for which we assessed memory strategy during overlapping event (BC) encoding by applying the neural classifier from Richter et al. (2016; see below). Because this analysis used a classifier trained on whole-brain, Montreal Neurological Institute (MNI) normalized data from a different group of participants collected in a different lab (Richter et al., 2016), data from each participant in this study were normalized into MNI standard space using ANTs (Advanced Normalization Tools; picsl.upenn.edu/software/ants/). First, ANTs was used to compute coregistration parameters from each participant's reference functional volume (first BOLD image from first run) to their high-resolution T1 anatomy using a rigid affine transformation. Next, each participant's T1 anatomical scan was normalized to a standard 2-mm MNI anatomical template using a diffeomorphic transformation. The transformation parameters were concatenated and applied to the functional data from the overlapping (BC) encoding phase. Finally, the normalized patterns were smoothed with an 8-mm FWHM kernel.

Decoding Reactivation of Prior Related Content during Overlapping Event Encoding

Our prior work (Zeithamova, Dominick, et al., 2012) showed that memories are reactivated during encoding

of overlapping events, with the degree of reactivation predicting participants' ability to subsequently infer novel relationships among related events. Here, we were interested in whether reactivation would be modulated as a function of temporal condition (i.e., prior day, same day). To test this hypothesis, we first verified that patterns of activation within visually selective ventral temporal cortex differentiated between viewing of face, house, and object stimuli during the localizer task.

An anatomical mask of the ventral temporal cortex was used for this analysis, which consisted of the Freesurfer-defined inferior temporal cortex, parahippocampal cortex, and the posterior portion of fusiform. Within each participant, ventral temporal cortex activation patterns for each timepoint from the localizer were labeled according to which stimulus type was being viewed (face, house, object). To account for the hemodynamic lag, condition labels were shifted back by two scans (4 sec) with respect to the functional time series. We then ran MVPA using the Princeton MVPA toolbox (code.google.com/p/princeton-mvpa-toolbox/) and custom code for MATLAB. The visual content classifier was a regularized logistic regression algorithm (penalty = 50), and we first used a leave-one-run-out cross-validation procedure to verify that the classifier could differentiate the visual content. Two runs from the localizer task were used iteratively for classifier training, and the remaining run was used to assess generalization performance. The comparison between the MVPA classifier prediction and the actual condition on each trial provided the cross-validation accuracy of the classifier. Cross-validated classification accuracy was very high for the localizer, averaging 92.2% correct ($SD = 4.9\%$) across participants.

Next, we turned to estimating the degree of trial-by-trial reactivation of the previously related stimulus (face or house) during encoding of overlapping pairs. To obtain a single activation pattern for every overlapping encoding trial, we averaged data for three time points, corresponding to activation patterns 4 sec, 6 sec, and 8 sec after the stimulus onset. Second, we trained the visual content classifier on all three runs from the functional localizer and then applied the trained classifier to each BC encoding trial pattern. The classifier output for each trial pattern consisted of three numbers representing the similarity to face, house, and object categories. Our critical comparison was between the face and house outputs, which would reflect reactivation of the previously related A item (face or house) during encoding of BC object–object pairs. Specifically, face and house classifier outputs were compared between BC trials related to faces (face–object–object triads) and BC trials related to houses (house–object–object triads). We also computed trial-by-trial reactivation estimates by subtracting *z*-scored face and house classifier outputs on each trial according to the trial type (face minus house output for trials in which the unseen but related A stimulus was a face, house minus face output for trials in which related A stimulus was a house).

To test whether the degree of reactivation differed for prior-day events and same-day events, the trial-by-trial estimates were summed within each temporal condition and compared using a paired *t* test.

Decoding Integration during Encoding

We hypothesized that the memory strategy participants employed while encoding overlapping events might differ for the same-day and prior-day conditions, with participants being more likely to employ an integration strategy as they encoded events that overlapped with associations acquired on the same day. To assess the memory strategy participants employed during encoding of overlapping events from the same-day and prior-day conditions, we employed a multivariate neural classifier derived from an independent study by Richter et al. (2016). We chose to employ an independent neural classifier from a different data set because it allowed us to make forward inferences about the strategies participants spontaneously employed in this study and whether the strategies differed for our two temporal conditions.

In the study by Richter et al., participants were presented with overlapping associations in a similar paradigm to the present work. On each overlapping event trial, they were given one of three explicit instructions for the memory strategy to use to process the overlapping event: (1) retrieve the prior related memory (without encoding the current event), (2) encode the new event (without reference to the prior memory), or (3) integrate the new information with their existing memory. A sparse multinomial logistic regression MVPA classifier (from Princeton MVPA toolbox) was trained on normalized whole-brain activation patterns from all-but-one participant and was then used to predict memory strategy cued on each trial for the remaining participant. The classifier's performance reliably discriminated the three memory strategies from one another, including on a pairwise basis (e.g., integration was separately decoded from both the simple encoding and retrieval strategies). Importantly, the neural classifier was then applied to data from a new set of participants from a second experiment, who were not given explicit task instructions during overlapping event encoding. The neural classifier was also successful at isolating when the new, uninstructed participants spontaneously employed an integration strategy when encoding overlapping events.

Here, we employed the whole-brain mnemonic strategy classifier from Richter et al. (2016) in a similar manner to the application in their second experiment described above. Similar to the second experiment in Richter et al. (2016), our participants were not provided with specific instructions for how to process the overlapping pairs for either temporal condition. We applied the trained whole-brain classifier from Richter et al. (2016) to trial-by-trial

activation patterns (averaged 4–8 sec after stimulus onset) from the encoding phase of our experiment. The output of the mnemonic classifier consisted of three numbers, indicating the degree of evidence for the three mnemonic strategies during overlapping event encoding (retrieve, encode, integrate). To test our hypothesis that overlapping events experienced on the same day are more likely to be integrated, we first averaged classifier evidence for integration separately for same-day condition trials and prior-day condition trials for each participant. We then used a paired *t* test to compare the integration evidence between temporal conditions.

Decoding Temporal Condition during Inference Test

We further predicted that inference judgments about memories acquired on the same day would evoke distinct neural processes from those required for inference about memories acquired across days. This prediction stems from our hypothesis about the strategies brought to bear during encoding of same-day and prior-day overlapping associations. If overlapping pairs from the same-day condition are more likely to be integrated with existing memories during encoding, inference could be supported by retrieval of such an integrated representation during test. In contrast, if overlapping pairs in the prior-day condition are more likely to be encoded separately from existing memories, we would expect inference would rely on retrieval and recombination of two separate representations.

Unlike the two MVPA applications that we used on the encoding data, we did not have independent data for classifier training. We thus employed a leave-one-run cross-validation approach on the inference data to test whether whole-brain activation patterns differentiate between the two temporal conditions. Similar to encoding, we first computed a single pattern for every inference trial by averaging the three time points from the functional time series collected 4 sec, 6 sec, and 8 sec after stimulus onset. The resulting inference trial patterns were labeled based on the day when the corresponding initial memory was first encountered (prior day, same day). We implemented a leave-one-run-out cross-validation analysis within a whole-brain mask using a regularized logistic regression algorithm implemented in the Princeton MVPA toolbox to obtain the day classification accuracy for each participant. A one-sample *t* test across participants was used to assess whether classification was significantly above the chance probability of 0.5 at $p < .05$. To assess whether differential neural recruitment in the two temporal conditions aids inference success, we related the day classification accuracy (i.e., disambiguation of inferences involving prior-day and same-day events) to inference performance across participants using Pearson's correlation.

Effect of Temporal Condition on Neural Mechanisms that Support Inference

Inference across events has been shown to involve hippocampal and prefrontal regions (Zeithamova & Preston, 2010). Importantly, their recruitment during inference may further differ depending on which strategy—retrieval of integrated representation or retrieval and recombination of individual memories—is employed to support performance (Preston & Eichenbaum, 2013). Thus, we further tested whether inference about the relationships for overlapping pairs from the same-day and prior-day associations evoke distinct patterns of activation in pFC and hippocampus. Within pFC, we focused on the inferior frontal gyrus (IFG) given its role in both memory-based inference (Zeithamova & Preston, 2010) and temporal coding (DuBrow & Davachi, 2014; Ezzyat & Davachi, 2014; Jenkins & Ranganath, 2010) and ventromedial PFC (VMPC) based on its hypothesized role in memory integration (Schlichting & Preston, 2015; Kosciak & Tranel, 2012; Zeithamova, Dominick, et al., 2012; van Kesteren, Fernandez, Norris, & Hermans, 2010). Because prior studies have shown functional heterogeneity within IFG in memory tasks (Schlichting et al., 2015; Raposo, Han, & Dobbins, 2009; Badre & Wagner, 2007; Dobbins, 2005), we considered the subregions of the IFG—anterior (pars orbitalis), middle (pars triangularis) and posterior IFG (pars opercularis)—separately. To identify these five ROIs on individual participants, we used FreeSurfer (Martinos Center for Biomedical Imaging, MGH, Charlestown, MA) to create a subcortical segmentation and cortical parcellation based on the high-resolution T1 SPGR. We then employed the same temporal condition classification analysis as described above, assessing cross-validated performance on the inference data within each ROI instead of the whole-brain mask. As with the whole-brain classifier, we related the temporal classifier accuracy to inference performance across participants. To correct for multiple comparisons across the five regions, we adopted a significance threshold of $p < .01$ ($p < .05/5$ ROIs, resulting in Bonferroni-corrected $p < .05$).

Finally, to assess to what degree the multivariate findings may be driven by possible differences in the overall response magnitude for the two temporal conditions, we performed a univariate analysis of the inference data under the assumption of a general linear model. The model consisted of four regressors that separated trials based on the temporal condition (prior day, same day) and inference success (correct, incorrect). This model allowed us to test whether overall activation or the magnitude of success effects (correct–incorrect difference) differ between the temporal conditions. Each trial was modeled as a stick function convolved with a canonical hemodynamic response function as implemented in SPM5. The model also included temporal derivatives for each condition and six motion parameters as regressors of no interest. Parameter estimates (beta weights) were extracted within each ROI for

each participant and condition. A 2×2 repeated-measures ANOVA was used to quantify the main effects of Temporal condition (prior day vs. same day), Inference success (correct vs. incorrect inference), and their interaction.

Summary of Decoding Analyses

To summarize, we used three different MVPA classifiers to test our hypotheses regarding the effect of temporal condition on overlapping event encoding and inference. First, we used a visual content classifier to test the hypothesis that memories acquired within the same day would be more likely to be reactivated during encoding of overlapping events relative to memories acquired on a different day. Second, we used a mnemonic strategy classifier to test the hypothesis that overlapping events would be more likely integrated with existing memories acquired on the same day relative to those experienced on the prior day. Finally, we used a temporal condition classifier on the inference test data to test the hypothesis that temporal condition would impact neural engagement during inference.

RESULTS

Behavioral Performance

To test how the temporal proximity of events impacts the ability to draw inferences about their relationships, we compared inference performance for each temporal condition. As predicted, inference success was greater and RTs on correct trials were faster when initial associations were learned on the same day as compared with the prior day (inference success: $t(28) = 2.25$, $p = .016$; RT for correct inferences: $t(28) = 2.47$, $p = .010$; Figure 1E, F).

However, participants also forgot more prior-day than same-day initial associations (prior day mean = 0.87, same day mean = 0.92, $t(28) = 4.80$, $p < .001$). To ensure that the behavioral benefit for same-day inferences did not solely result from better memory for same-day initial pairs, we performed two control analyses. First, we related the behavioral benefit for same-day inferences to the forgetting cost for prior-day initial memories across participants. If the difference in inference in the two temporal conditions was due to forgetting of prior-day initial pairs, participants who forgot the most prior-day initial pairs compared with same-day initial pairs should also be those that show the greatest difference in inference performance for those conditions. However, the relationship between the same-day inference benefit and the forgetting of prior-day initial memories was not significant ($r = -.132$, $p = .49$). In other words, participants who showed the greatest behavioral benefit for the same-day inferences are not necessarily those who remember same-day initial pairs better than prior-day pairs. We also performed a second analysis of the inference data limited to the trials for which the initial pair was

remembered. Inference performance was still reliably greater and RT faster for same-day associations relative to the prior-day condition (inference success: prior day mean = 0.60, same day mean = 0.65, $t(28) = 2.13$, $p = .021$; RT for correct inference: prior day mean = 2545 msec, same day mean = 2476 msec, $t(28) = 1.76$, $p = .045$).

We further examined memory for the overlapping (BC) events themselves as a function of temporal proximity to the initial associations. Memory for overlapping events was not affected by the temporal condition (recognition: prior day mean = 0.86, same day mean 0.85, $t(28) = 0.73$, $p = .47$; RT: prior day mean = 1701 msec, same day mean = 1705 msec, $t(28) = 0.18$, $p = .86$), indicating that the differences in inference performance did not result from differential memory for overlapping events themselves.

Evidence for Reactivation during Encoding of Overlapping Events Does Not Differ between Temporal Conditions

At the neural level, we first tested whether (1) information about related A items was reactivated during encoding of overlapping BC events and (2) the degree of reactivation differed as a function of temporal condition. We trained an MVPA classifier on independent data from the same participants to differentiate between neural patterns related to face, house, and object categories. We then applied the classifier to index reactivation of the related face or house stimulus from the initial pairs during encoding of each overlapping object–object pair. Replicating our prior study (Zeithamova, Dominick, et al., 2012), we found a significant reactivation of initial memories during encoding of related events (Figure 2A). Contrary to our prediction, evidence for reactivation did not differ for prior-day and same-day memories ($t(28) = -0.87$, $p = .39$; Figure 2B).

To replicate our prior findings, we also tested whether individual differences in reactivation evidence during encoding track individual differences in subsequent inference success (Zeithamova, Dominick, et al., 2012), irrespective of condition. We found a significant positive correlation between reactivation and inference success ($r = .40$, $p = .017$), indicating that participants who show evidence for reactivation during encoding of overlapping events are more successful at inference. When we assessed the correlation between reactivation evidence and inference performance separately for each temporal condition, the across-participant correlation was numerically but not significantly greater for same-day than prior-day conditions (prior day $r = .18$, same day $r = .32$).

Same-day Events Are More Likely to Be Integrated during Encoding

Prior work showed that is possible to decode, from normalized whole-brain activation patterns, whether participants integrate information across events as they learn about new experiences that overlap with existing memories (Richter et al., 2016). Here, we used the classifier from that prior study, which was trained on an independent group of participants, to test our hypothesis that same-day overlapping events would be more likely integrated with existing memories relative to overlapping events from the prior-day condition. Consistent with our prediction, we found greater classifier evidence for integration during same-day overlapping pairs relative to the overlapping pairs in the prior-day condition ($t(28) = 2.04$, $p = .026$; Figure 2C). Thus, although we found comparable evidence for reactivation of related memories across temporal conditions based on the visual content classifier, evidence for integration based on the mnemonic strategy classifier was greater for same-day overlapping events, mirroring the pattern of behavioral data.

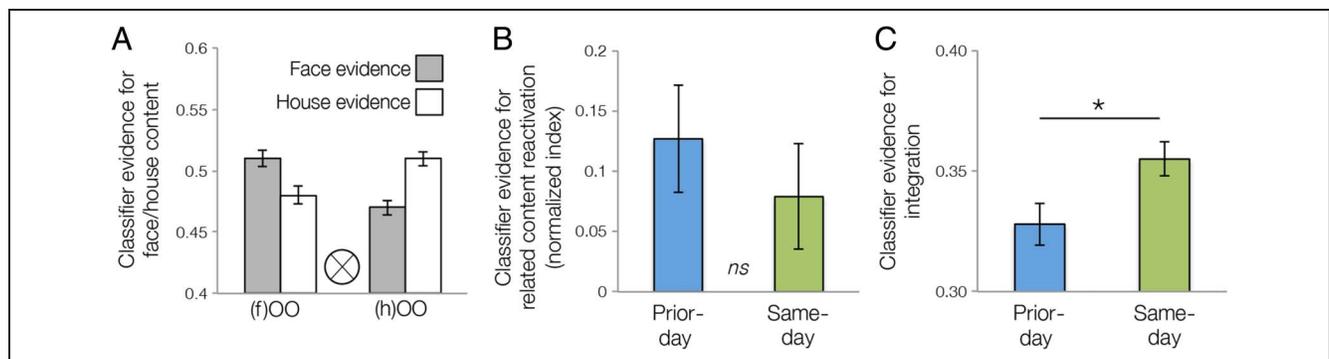


Figure 2. Decoding reactivation and integration during encoding of overlapping events. (A) Evidence for reactivation of faces and houses during encoding of overlapping object–object pairs derived from a visual content classifier. (f)OO: encoding of two objects related to a face. (h)OO: encoding of two objects related to a house. (B) Reactivation evidence (normalized difference between house and face output) during encoding of overlapping pairs from the prior day (blue) and same day (green) conditions. (C) Evidence for use of an integration strategy during encoding of overlapping events for the prior-day (blue) and same-day (green) conditions derived from a mnemonic strategy classifier. Error bars on all panels denote across-participant SEM. Asterisk (*) denoted a significant difference between conditions (paired t -test $p < 0.05$).

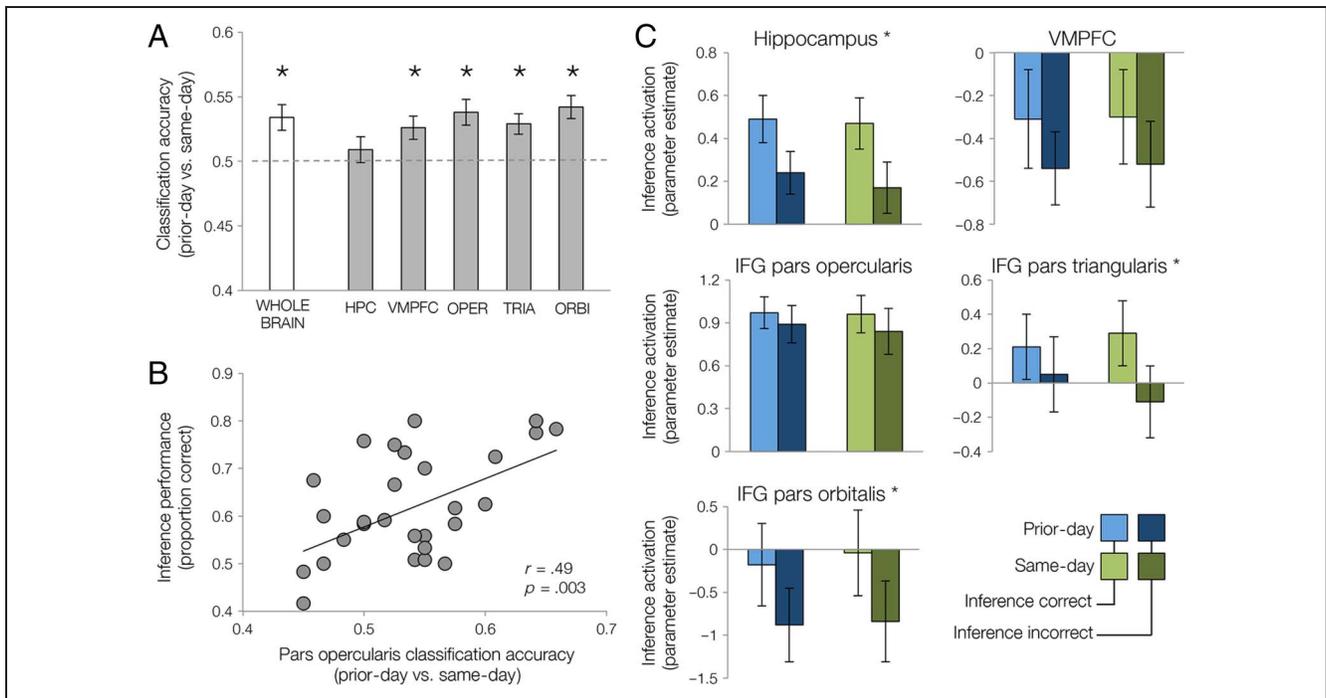


Figure 3. Temporal condition effects during inference. (A, B) Temporal condition decoding during inference. (A) Cross-validated temporal condition classification accuracy in the whole brain (white bar) and each ROI (gray bars). Dashed line represents chance. Error bars denote across-participant SEM. (B) Correlation between IFG pars opercularis day classification and inference performance across participants. (C) Univariate temporal condition and success effects. Mean activation during inference split by temporal condition and inference success. Star (*) next to the region's name denotes significant main effect of Success (corrected for multiple comparisons). No region showed main effect of temporal condition or Temporal condition \times Success interaction. Error bars denote across-participant standard error. Asterisk (*) above individual bars denotes above-chance classification within a given region (one-sample t -test $p < 0.05$, corrected).

We did not have specific a priori predictions about the other mnemonic strategies indexed by the classifier (retrieve old, encode new) for the two temporal conditions. However, because we observed differences in evidence for integration and because the relative evidence is coupled across the three strategies, we tested whether the two strategies showed corresponding differences across temporal conditions. We found no differences in evidence for the retrieve strategy (mean evidence for prior day = 0.327, $SE = 0.006$; same day mean = 0.330, $SE = 0.006$; $t(28) < 1$), but we did observe significantly greater evidence for encoding of new associations for prior-day condition (mean = 0.346, $SE = 0.009$) than same-day condition (mean = 0.315, $SE = 0.006$; $t(28) = 2.24$, $p = .017$).

Temporal Condition Decoding during Inference

Inference across events may rely on retrieval of integrated memory representations or retrieval and recombination of separate memory representations (Schlichting & Preston, 2015; Schlichting et al., 2014; Zeithamova, Schlichting, & Preston, 2012). These two inference strategies are postulated to recruit different neural mechanisms (Preston & Eichenbaum, 2013). Because behavioral and neural evidence suggested greater probability of integration for same-day events, we hypothesized that neural

processes brought to bear to inference would differ for the two temporal conditions. Consistent with this prediction, we found that a whole-brain classifier could distinguish inference trials from the same-day and prior-day conditions (mean = 0.534, $t(28) = 3.59$, $p < .001$; Figure 3A), although nothing on the screen indicated the temporal condition of inference trials. This decoding analysis suggests that inferences involving prior-day memories engage different processes than inferences across same-day events. We further hypothesized that the ability to make successful inferences across events would be influenced by the ability to engage the process best suited to a given temporal condition. To test this hypothesis, we computed the correlation between day classification accuracy and inference accuracy across participant. However, we did not find evidence for such a relationship ($r = .10$, $p > .5$).

Because same-day inferences were overall more successful and faster than prior-day inferences, we performed two control analyses to determine if the temporal condition decoding was driven by differences in difficulty across condition. First, we correlated the behavioral inference benefit for same-day events (both for accuracy and RT) with temporal condition decoding accuracy. If temporal condition decoding is driven by differences in difficulty across conditions, it should be easier to decode temporal condition in those participants who show the

largest difference in accuracy across conditions. We did not find a significant relationship between temporal condition decoding and the same-day inference benefit (accuracy differences: $r = .20, p > .3$; RT differences: $r = .11, p > .5$), suggesting that condition level differences in difficulty were not driving decoding. Second, we performed decoding analysis limited to correct trials only, running 100 Monte Carlo simulations that randomly subsampled trials to equate the number of training trials used for each condition. Cross-validated decoding accuracy remained above chance (mean = 0.545, $t(28) = 3.25, p = .003$).

Prior research suggests that different inference strategies may differentially recruit hippocampus and prefrontal regions (Preston & Eichenbaum, 2013). We tested this hypothesis using the same multivariate approach as we employed with the whole-brain classifier but testing classification performance within hippocampus and each prefrontal ROI separately. Mean day classification accuracy was above chance in VMPFC ($t(28) = 2.87, p = .008$) and all IFG subregions (all $t > 3.47, p < .002$; Figure 3A). Furthermore, IFG pars opercularis classification accuracy positively tracked inference success across participants ($r = .49, p = .006$; Figure 3B), a pattern not present in other regions (all $r < 0$).

Because same-day inferences were overall more successful and faster than prior-day inferences, we performed the same two control analyses as we did with a whole-brain classifier to determine if the temporal condition decoding was driven by differences in difficulty across condition. First, we found no correlation between decoding accuracy and behavioral benefit (in accuracy or RT) for same-day events (accuracy differences: max $r = .07, p > .7$; RT differences: max $r = .18, p > .3$), suggesting that condition level differences in difficulty were not driving decoding. Second, decoding analysis limited to correct trials only (averaging 100 simulations with random subsampling of trials to equate the number of training patterns across conditions) showed that cross-validated accuracy in IFG pars opercularis remained above chance ($t(28) = 2.71, p = .01$), with a similar pattern in IFG pars triangularis ($t(28) = 2.18, p = .04$) that did not reach the corrected threshold.

Overall Activation and Inference Success Effects in Hippocampus and IFG Are Equivalent across Temporal Conditions

To test whether temporal condition decoding in these regions was driven by overall activation differences between conditions and to test whether they are engaged during inference in the current task, we performed a univariate analysis with Temporal condition and Inference success as factors. Consistent with prior reports of hippocampal and IFG role in mnemonic inference (Zeithamova & Preston, 2010), we observed a significant main effect of Inference success in the hippocampus, IFG pars triangu-

laris, and IFG pars orbitalis (all $F(1, 28) > 15.9, p < .001$), with greater activation during correct than incorrect inferences (Figure 3C). Success effects in IFG pars opercularis ($F(1, 28) = 2.59$, uncorrected $p = .12$) and VMPFC ($F(1, 28) = 4.05$, uncorrected $p = .054$) did not reach significance threshold. No main effects of Temporal condition (all $F < 1$) nor Temporal condition \times Success interactions (all $F(1, 28) < 1.33, p > .26$; Figure 3C) were found in any region, suggesting that mean activation differences were not driving decoding of temporal condition in these regions.

DISCUSSION

In this study, we tested the hypothesis that temporally proximal events would be more likely to be integrated, resulting in enhanced inference for the relationships among those episodes. Consistent with this prediction, participants were faster and more accurate when making novel inferences—our behavioral marker of integration—that involved events experienced within the same day relative to those experienced on different days. Using a mnemonic strategy classifier from an independent study (Richter et al., 2016), we were also able to quantify the neural evidence for integration during new encoding. Similar to the behavioral findings, we found greater evidence for integration when participants encoded events overlapping with memories acquired on the same day versus those overlapping with prior-day events. In contrast, neural evidence for reactivation during encoding was equivalent for prior-day and same-day related content, suggesting that reactivation does not always lead to integration.

During inference itself, whole-brain activation patterns distinguished when participants were making decisions about the relationships among events experienced on the same day from those about events experienced on different days. Differentiation of the temporal conditions during inference was particularly apparent in IFG pars opercularis, where discrimination accuracy tracked individual differences in inference performance. These findings may indicate that different strategies are brought to bear when inferring relationships among integrated versus separated memories and that deployment of optimal inference strategies aids performance. Collectively, these data indicate that temporal proximity of events promotes memory integration and modulates neural mechanisms brought to bear during inference.

The temporal context model (Howard & Kahana, 2002) proposes that events experienced close in time are linked together via a shared temporal context. Consistent with this notion, recent neuroimaging studies have shown greater neural overlap when events are experienced in immediate succession (Ezzyat & Davachi, 2014; Hsieh et al., 2014). However, the temporal context model further predicts that events do not have to immediately follow one another to become linked in memory.

Events that share overlapping content with prior experience may serve to reactivate a previous temporal context, providing a mechanism for integration of even temporally distant memories (Howard et al., 2009). In this study, however, we demonstrate that there are boundary conditions that may limit the ability to bridge memories across time. We show evidence for reactivation and integration of events experienced minutes apart within the same experimental session. In contrast, events experienced on different days were less likely to become integrated despite comparable evidence for reactivation. Our prior work showed that integration during encoding aids subsequent inference performance (Schlichting et al., 2014; Zeithamova, Dominick, et al., 2012). In cases when integration does not occur, inference can still be performed by retrieving and recombining two individual associations at test, but at a cost of reduced accuracy and increased RT (Schlichting et al., 2014, 2015). The current study demonstrates that integration during encoding is less likely for temporally distant events. Consequently, inference performance for the prior-day condition is both slower and less accurate than the same-day condition because fewer integrated representations are formed for events experienced across days.

At a cellular level, recent activation of a synapse is thought to temporarily increase its excitability and tag it for subsequent long-term potentiation to support memory formation (Frey & Morris, 1997). This synaptic tagging mechanism has also been proposed to guide allocation of neurons during new encoding (Silva et al., 2009). Recently tagged neurons would be more likely to be recruited during new event encoding, which would promote memory integration. Such enhanced recruitment of tagged neurons, however, is thought to be temporally limited, occurring only when events are experienced within minutes or hours of one another. Two recent rodent studies evince this mechanism showing that overlapping neural populations represent events experienced on the same day, whereas those experienced on different days are representationally distinct (Cai et al., 2016; Rashid et al., 2016).

In one of these studies (Rashid et al., 2016), mice were exposed to two events that were separated by a variable time interval. Two events separated by up to 6 hr recruited overlapping ensembles of amygdala neurons, and memories of the two events became integrated, as evidenced by generalization of fear responses across events. In contrast, events separated by 18 or 24 hr recruited distinct populations of amygdala neurons and did not show behavioral evidence of integration. Similar findings have been observed within the CA₁ subfield of the hippocampus, for which overlapping neural populations represented two events separated by 5 hr but not two events separated by 7 days (Cai et al., 2016). These findings parallel the behavioral and fMRI results observed in the current study and suggest a potential cellular mechanism underlying preferential integration of temporally

proximal events. Our study extends these findings to demonstrate how different representational schemes for temporally proximal and distant events impacts individuals' ability to extract new knowledge and reason about the relationships among memories.

Although these converging animal (Cai et al., 2016; Rashid et al., 2016) and human findings demonstrate that temporal proximity is a critical factor determining the likelihood of integration, other factors may play an important role in determining whether memories are integrated or separated. For instance, memory strength may be another boundary condition of integration, as suggested by two sets of findings (Schlichting et al., 2015; Tse et al., 2007, 2011). In rodents, new flavor–location associations are acquired rapidly when experienced within a well-learned spatial configuration acquired across many sessions relative to a spatial configuration that was frequently changed (Tse et al., 2007, 2011). Presumably, such facilitation occurs because well-learned events are more readily reactivated when learning new content that overlaps with existing knowledge. In humans, stronger memories also facilitate integration (Schlichting et al., 2015). When initial events (AB associations) are repeated many times before the introduction of an overlapping event (BC associations), hippocampus and pFC form integrated representations of the overlapping events. However, when overlapping events were presented in alternation, hippocampal and pFC form pattern separated memories. In the latter condition, the initial event is experienced only once before introduction of the overlapping event; thus, the initial memory is likely weaker and less likely to be reactivated during the overlapping event.

In one aspect, these findings may seem somewhat contradictory to the present data, as they show enhanced integration for events in which the first experience of the initial memory was more temporally distant from the first presentation of the overlapping event. However, temporal distance between the last presentation of the initial memory and the first presentation of the overlapping memory was equated across conditions in both of these studies. The critical events in the strong and weak conditions were both 2 days apart in the studies by Tse and colleagues. In Schlichting et al. (2015), the last presentation of the initial memory and the first presentation of the overlapping event were a few trials apart in both strong and weak conditions. Thus, when temporal distance from the last experience is similar, stronger memories are preferentially integrated. As time and memory strength are often intertwined, it is not possible to completely rule out memory strength differences between conditions in this study. Nevertheless, our results do indicate that a benefit for same-day inferences was observed in this study even after forgetting of initial events was taken into account, both within and across participants. In light of this finding, the present results suggest that temporal proximity provides an additional boundary

condition on integration across events that may not be explained by strength alone.

Our study further extends prior research on the boundary conditions of memory integration to show how temporal proximity impacts processes brought to bear during inference itself. Specifically, multivariate activation patterns across the brain and in IFG specifically discriminated when participants inferred relationships about events experienced within the same day from inferences about events that occurred on different days. Furthermore, discrimination accuracy in IFG (pars opercularis) was related to superior inference. One interpretation of these findings is that distinct mechanisms may support inference when an integrated representation has been formed during encoding (as was more likely in the same-day condition) relative to when separate representations are formed. Integrated representations would support direct retrieval of the indirect relationship between items that share a common associate; inference from separated representations would require successful retrieval of two memories that would then be jointly considered to support inference (Zeithamova, Schlichting, et al., 2012). Prior work has shown that distinct subregions of IFG and VMPFC support integrated and separated representations (Schlichting et al., 2015). Here, prefrontal activation patterns that discriminate between temporal conditions during inference may reflect retrieval of integrated representations for same-day events and retrieval and recombination of separate memory traces for events encoded across days. The relationship between prefrontal discrimination accuracy and performance further suggests that the ability to access the appropriate representations for a given condition may contribute to individual differences in inference success.

Conclusions

Our findings demonstrate that temporal proximity of events promotes the formation of integrated memory representations for overlapping events that facilitate novel inferences about their relationships. These data shed light on an important boundary condition—time—that influences when overlapping events are represented by integrated or separated neural codes, thus converging with recent cellular work in animals. Furthermore, our findings indicate that the nature of the representations formed during encoding influence the processes brought to bear during inference. More broadly, these findings contribute to our growing understanding that temporal structure plays an important role in organizing memories.

Acknowledgments

This work was supported by a National Science Foundation CAREER Award (A. R. P.), NIH-NIMH R01 MH100121-01 (A. R. P.), and NIH-NIMH National Research Service Award F32MH094085

(D. Z.). We thank Brice A. Kuhl and Avi J. H. Chanals for providing us with the neural classifier from Richter et al. (2016). We thank Amelia R. Wattenberger for help with data collection.

Reprint requests should be sent to Alison R. Preston, 1 University Stop A8000, Austin, TX 78712, or via e-mail: apreston@utexas.edu or Dagmar Zeithamova, 1227 University of Oregon, Eugene, OR 97403, or via e-mail: dasa@uoregon.edu.

REFERENCES

- Badre, D., & Wagner, A. D. (2007). Left ventrolateral prefrontal cortex and the cognitive control of memory. *Neuropsychologia*, *45*, 2883–2901.
- Bridge, D. J., & Voss, J. L. (2014). Hippocampal binding of novel information with dominant memory traces can support both memory stability and change. *Journal of Neuroscience*, *34*, 2203–2213.
- Cai, D. J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., et al. (2016). A shared neural ensemble links distinct contextual memories encoded close in time. *Nature*, *534*, 115–118.
- Diekelmann, S., Büchel, C., Born, J., & Rasch, B. (2011). Labile or stable: Opposing consequences for memory when reactivated during waking and sleep. *Nature Publishing Group*, *14*, 381–386.
- Dobbins, I. G. (2005). Domain-general and domain-sensitive prefrontal mechanisms for recollecting events and detecting novelty. *Cerebral Cortex*, *15*, 1768–1778.
- DuBrow, S., & Davachi, L. (2014). Temporal memory is shaped by encoding stability and intervening item reactivation. *Journal of Neuroscience*, *34*, 13998–14005.
- Eichenbaum, H. (1999). The hippocampus and mechanisms of declarative memory. *Behavioural Brain Research*, *103*, 123–133.
- Ezzyat, Y., & Davachi, L. (2014). Similarity breeds proximity: Pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron*, *81*, 1179–1189.
- Frankland, P. W., & Bontempi, B. (2005). The organization of recent and remote memories. *Nature Reviews Neuroscience*, *6*, 119–130.
- Frey, U., & Morris, R. G. (1997). Synaptic tagging and long-term potentiation. *Nature*, *385*, 533–536.
- Greene, A. J., Gross, W. L., Elsinger, C. L., & Rao, S. M. (2006). An fMRI analysis of the human hippocampus: Inference, context, and task awareness. *Journal of Cognitive Neuroscience*, *18*, 1156–1173.
- Howard, M. W., Jing, B., Rao, V. A., Provyn, J. P., & Datey, A. V. (2009). Bridging the gap: Transitive associations between items presented in similar temporal contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 391–407.
- Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*, 269–299.
- Hsieh, L.-T., Gruber, M. J., Jenkins, L. J., & Ranganath, C. (2014). Hippocampal activity patterns carry information about objects in temporal context. *Neuron*, *81*, 1165–1178.
- Hupbach, A., Hardt, O., Gomez, R., & Nadel, L. (2008). The dynamics of memory: Context-dependent updating. *Learning & Memory*, *15*, 574–579.
- Jenkins, L. J., & Ranganath, C. (2010). Prefrontal and medial temporal lobe activity at encoding predicts temporal context memory. *Journal of Neuroscience*, *30*, 15558–15565.
- Koscik, T. R., & Tranel, D. (2012). The human ventromedial prefrontal cortex is critical for transitive inference. *Journal of Cognitive Neuroscience*, *24*, 1191–1204.

- Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in Cognitive Sciences*, *20*, 512–534.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (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. *Psychological Review*, *102*, 419–457.
- O'Reilly, R. C., & Rudy, J. W. (2000). Computational principles of learning in the neocortex and hippocampus. *Hippocampus*, *10*, 389–397.
- Preston, A. R., & Eichenbaum, H. (2013). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*, *23*, R764–R773.
- Preston, A. R., Shrager, Y., Dudukovic, N. M., & Gabrieli, J. D. E. (2004). Hippocampal contribution to the novel use of relational information in declarative memory. *Hippocampus*, *14*, 148–152.
- Raposo, A., Han, S., & Dobbins, I. G. (2009). Ventrolateral prefrontal cortex and self-initiated semantic elaboration during memory retrieval. *Neuropsychologia*, *47*, 2261–2271.
- Rashid, A. J., Yan, C., Mercaldo, V., Hsiang, H.-L. L., Park, S., Cole, C. J., et al. (2016). Competition between engrams influences fear memory formation and recall. *Science*, *353*, 383–387.
- Richter, F. R., Chanales, A. J. H., & Kuhl, B. A. (2016). Predicting the integration of overlapping memories by decoding mnemonic processing states during learning. *Neuroimage*, *124*, 323–335.
- Schlichting, M. L., Mumford, J. A., & Preston, A. R. (2015). Learning-related representational changes reveal dissociable integration and separation signatures in the hippocampus and prefrontal cortex. *Nature Communications*, *6*, 1–10.
- Schlichting, M. L., & Preston, A. R. (2015). Memory integration: Neural mechanisms and implications for behavior. *Current Opinion in Behavioral Sciences*, *1*, 1–8.
- Schlichting, M. L., Zeithamova, D., & Preston, A. R. (2014). CA₁ subfield contributions to memory integration and inference. *Hippocampus*, *24*, 1248–1260.
- Shohamy, D., & Wagner, A. D. (2008). Integrating memories in the human brain: Hippocampal-midbrain encoding of overlapping events. *Neuron*, *60*, 378–389.
- Silva, A. J., Zhou, Y., Rogerson, T., Shobe, J., & Balaji, J. (2009). Molecular and cellular approaches to memory allocation in neural circuits. *Science*, *326*, 391–395.
- Squire, L. R., & Alvarez, P. (1995). Retrograde amnesia and memory consolidation: A neurobiological perspective. *Current Opinion in Neurobiology*, *5*, 169–177.
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., et al. (2007). Schemas and memory consolidation. *Science*, *316*, 76–82.
- Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., et al. (2011). Schema-dependent gene activation and memory encoding in neocortex. *Science*, *333*, 891–895.
- van Kesteren, M. T. R., Fernandez, G., Norris, D. G., & Hermans, E. J. (2010). Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proceedings of the National Academy of Sciences, U.S.A.*, *107*, 7550–7555.
- Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron*, *75*, 168–179.
- Zeithamova, D., & Preston, A. R. (2010). Flexible memories: Differential roles for medial temporal lobe and prefrontal cortex in cross-episode binding. *Journal of Neuroscience*, *30*, 14676–14684.
- Zeithamova, D., Schlichting, M. L., & Preston, A. R. (2012). The hippocampus and inferential reasoning: Building memories to navigate future decisions. *Frontiers in Human Neuroscience*, *6*, 70.