



Published in final edited form as:

*Neuron*. 2017 August 02; 95(3): 709–721.e5. doi:10.1016/j.neuron.2017.06.041.

## Discovering event structure in continuous narrative perception and memory

Christopher Baldassano<sup>\*1</sup>, Janice Chen<sup>2</sup>, Asieh Zadbood<sup>1</sup>, Jonathan W Pillow<sup>1</sup>, Uri Hasson<sup>1</sup>, and Kenneth A Norman<sup>1</sup>

<sup>1</sup>Princeton Neuroscience Institute and Department of Psychology, Princeton University, Princeton, NJ, 08544, USA

<sup>2</sup>Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD, 21218, USA

### Summary

During realistic, continuous perception, humans automatically segment experiences into discrete events. Using a novel model of cortical event dynamics, we investigate how cortical structures generate event representations during narrative perception, and how these events are stored to and retrieved from memory. Our data-driven approach allows us to detect event boundaries as shifts between stable patterns of brain activity without relying on stimulus annotations, and reveals a nested hierarchy from short events in sensory regions to long events in high-order areas (including angular gyrus and posterior medial cortex), which represent abstract, multimodal situation models. High-order event boundaries are coupled to increases in hippocampal activity, which predict pattern reinstatement during later free recall. These areas also show evidence of anticipatory reinstatement as subjects listen to a familiar narrative. Based on these results, we propose that brain activity is naturally structured into nested events, which form the basis of long-term memory representations.

### Keywords

event segmentation; fMRI; perception; memory; hippocampus; hidden markov model; recall; narrative; situation model; event model; reinstatement

---

<sup>\*</sup>Lead Contact: Christopher Baldassano, chrisb@princeton.edu.

**Publisher's Disclaimer:** This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

#### Author Contributions

Conceptualization: C.B., J.C., J.W.P., U.H., and K.A.N.; Methodology: C.B., J.W.P., U.H., and K.A.N.; Software: C.B.; Investigation: C.B., J.C., and A.Z.; Writing – Original Draft: C.B.; Writing – Review and Editing: C.B., J.C., A.Z., J.W.P., U.H., K.A.N.; Supervision: U.H. and K.A.N.

## Introduction

Typically, perception and memory are studied in the context of discrete pictures or words. Real-life experience, however, consists of a continuous stream of perceptual stimuli. The brain therefore needs to structure experience into units that can be understood and remembered: “the meaningful segments of one’s life, the coherent units of one’s personal history” (Beal & Weiss, 2013). Although this question was first investigated decades ago (Newton, Engquist, & Bois, 1977), a general “event segmentation theory” was only proposed recently (Zacks, Speer, Swallow, Braver, & Reynolds, 2007). These and other authors have argued that humans implicitly generate event boundaries when consecutive stimuli have distinct temporal associations (Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013), when the causal structure of the environment changes (Kurby & Zacks, 2008; Radvansky, 2012), or when our goals change (DuBrow & Davachi, 2016).

At what timescale are experiences segmented into events? When reading a story, we could chunk it into discrete units of individual words, sentences, paragraphs, or chapters, and we may need to chunk information on different timescales depending on our goals. Behavioral studies have shown that subjects can segment events into a nested hierarchy from coarse to fine timescales (Kurby & Zacks, 2008; Zacks, Tversky, & Iyer, 2001) and flexibly adjust their units of segmentation depending on their uncertainty about ongoing events (Newton, 1973). The neural basis of this segmentation behavior is unclear; event perception could rely on a single unified system, which segments the continuous perceptual stream at different granularities depending on the current task (Zacks et al., 2007), or may rely on multiple brain areas which segment events at different timescales, as suggested by the selective deficits for coarse segmentations exhibited by some patient populations (Zalla, Pradat-Diehl, & Sirigu, 2003; Zalla, Verlut, Franck, Puzenat, & Sirigu, 2004).

A recent theory of cortical process-memory topography argues that information is integrated at different timescales throughout the cortex. Processing timescales increase from tens of milliseconds in early sensory regions (e.g. for detecting phonemes in early auditory areas), to a few seconds in mid-level sensory areas (e.g. for integrating words into sentences), up to hundreds of seconds in regions including the temporoparietal junction, angular gyrus, and posterior and frontal medial cortex (e.g. for integrating information from an entire paragraphs) (J. Chen et al., 2016; Hasson, Chen, & Honey, 2015). The relationship between the process-memory topography and event segmentation has not yet been investigated. On the one hand, it is possible that cortical representations are accumulated continuously, e.g. using a sliding window approach, at each level of the processing hierarchy (Stephens, Honey, & Hasson, 2013). On the other hand, a strong link between the timescale hierarchy and event segmentation theory would predict that each area chunks experience at its preferred timescale and integrates information within discretized units (e.g. phonemes, words, sentences, paragraphs) before providing its output to the next processing level (Nelson et al., 2017). In this view, “events” in low-level sensory cortex (e.g. a single phoneme, Giraud & Poeppel, 2012) are gradually integrated into minutes-long situation-level events, using a multi-stage nested temporal chunking. This chunking of continuous experience at multiple timescales along the cortical processing hierarchy has not been previously demonstrated in the dynamics of whole-brain neural activity.

A second critical question for understanding event perception is how real-life experiences are encoded into long-term memory. Behavioral experiments and mathematical models have argued that long-term memory reflects event structure during encoding (Ezzyat & Davachi, 2011; Gershman, Radulescu, Norman, & Niv, 2014; Sargent et al., 2013; Zacks, Tversky, et al., 2001), suggesting that the event segments generated during perception may serve as the “episodes” of episodic memory. The hippocampus is thought to bind cortical representations into a memory trace (McClelland, McNaughton, & O’Reilly, 1995; Moscovitch et al., 2005; Norman & O’Reilly, 2003), a process that is typically studied using discrete memoranda (Danker, Tompary, & Davachi, 2016). However, given a continuous stream of information in a real-life context, it is not clear at what timescale memories should be encoded, and whether these memory traces should be continuously updated or encoded only after an event has completed. The hippocampus is connected to long timescale regions including the angular gyrus and posterior medial cortex (Kravitz, Saleem, Baker, & Mishkin, 2011; Ranganath & Ritchey, 2012; Rugg & Vilberg, 2013), suggesting that the main inputs to long-term memory come from these areas, which are thought to represent multimodal, abstract representations of the features of the current event (“situation models”, Johnson-Laird, 1983; Van Dijk & Kintsch, 1983; Zwaan, Langston, & Graesser, 1995; Zwaan & Radvansky, 1998, or more generally, “event models,” Radvansky & Zacks, 2011). Recent work has shown that hippocampal activity peaks at the *offset* of video clips (Ben-Yakov & Dudai, 2011; Ben-Yakov, Eshel, & Dudai, 2013), suggesting that the end of a long-timescale event triggers memory encoding processes that occur *after* the event has ended.

Based on our experiments (described below), we propose that the full life cycle of an event can be described in a unified theory, illustrated in Fig. 1. During perception, each brain region along the processing hierarchy segments information at its preferred timescale, beginning with short events in primary visual and auditory cortex and building to multimodal situation models in long timescale areas, including the angular gyrus and posterior medial cortex. This model of processing requires that: 1) all regions driven by audio-visual stimuli should exhibit event-structured activity, with segmentation into short events in early sensory areas and longer events in high-order areas; 2) events throughout the hierarchy should have a nested structure, with coarse event boundaries annotated by human observers most strongly related to long events at the top of the hierarchy ; 3) event representations in long timescale regions, which build a coarse model of the situation, should be invariant across different descriptions of the same situation (e.g. when the same situation is described visually in a movie or a verbally in a story). We also argue that event structure is reflected in how experiences are stored into episodic memory. At event boundaries in long timescale areas, the situation model is transmitted to the hippocampus, which can later reinstate the situation model in long timescale regions during recall. This implies: 4) the end of an event in long timescale cortical regions should trigger the hippocampus to encode information about the just-concluded event into episodic memory; 5) stored event memories can be reinstated in long timescale cortical regions during recall, with stronger reinstatement for more strongly-encoded events. Finally, this process can come full circle, with prior event memories influencing ongoing processing, such that: 6) prior memory for a narrative should lead to anticipatory reinstatement in long timescale regions.

To test these hypotheses, we need the ability to identify how different brain areas segment events (at different timescales), align events across different datasets with different timings (e.g. to see whether the same situation model is being elicited by a movie vs. a verbal narrative, or a movie vs. recall), and track differences in event segmentations in different observers of the same stimulus (e.g. depending on prior experience). Thus, to search for the neural correlates of event segmentation, we have developed a new data-driven method that allows us to identify events directly from fMRI activity patterns, across multiple timescales and datasets (Fig. 2).

Our analysis approach (described in detail in *STAR Methods*) starts with two simple assumptions: that while processing a narrative stimulus, observers progress through a sequence of discrete event representations (hidden states), and that each event has a distinct (observable) signature (a multi-voxel fMRI pattern) that is present throughout the event. We implement these assumptions using a data-driven event segmentation model, based on a Hidden Markov Model (HMM). Fitting the model to fMRI data (e.g. evoked by viewing a movie) entails estimating the optimal number of events, the mean activity pattern for each event, and when event transitions occur. When applying the model to multiple datasets evoked by the same narrative (e.g. during movie viewing and during later verbal recall), the model is constrained to find the same sequence of patterns (because the events are the same), but the timing of the transitions between the patterns can vary (e.g. since the spoken description of the events might not take as long as the original events).

In prior studies, using human-based segmentation of coarse event structure, we demonstrated that event-related representations generalize across modalities and between encoding and recall (thereby supporting hypotheses 3 and 5, J. Chen et al., 2017; Zadbood, Chen, Leong, Norman, & Hasson, 2016). In the current study, we extend these findings, by testing whether we can use the data driven HMM to detect stable and abstract event boundaries in high order areas without relying on human annotations. This new analysis approach also allows us to test for the first time whether the brain segments information, hierarchically, at multiple timescales (hypotheses 1 and 2), how segmentation of information interacts with the storage and retrieval of this information by the hippocampus (hypothesis 4), and how prior exposure to a sequence of events can later lead to anticipatory reinstatement of those events (hypothesis 6). Taken together, our results provide the first direct evidence that realistic experiences are discretely and hierarchically chunked at multiple time-scales in the brain, with chunks at the top of the processing hierarchy playing a special role in cross-modal situation representation and episodic memory.

## Results

All of our analyses are carried out using our new HMM-based event segmentation model (summarized above, and described in detail in the *Event Segmentation Model* subsection in *Methods*), which can automatically discover the fMRI signatures of each event and its temporal boundaries in a particular dataset. We validated this model using both synthetic data (Fig. S1) and narrative data with clear event breaks between stories (Fig. S2), confirming that we could accurately recover the number of event boundaries and their

locations (see *Materials and Methods*). We then applied the model to test six predictions of our theory of event perception and memory.

### Timescales of cortical event segmentation

We first tested the hypothesis that “all regions driven by audio-visual stimuli should exhibit event-structured activity, with segmentation into short events in early sensory areas and longer events in high-order areas.” We measured the extent to which continuous stimuli evoked the event structure hypothesized by our model (periods with stable event patterns punctuated by shifts between events), and whether the timescales of these events varied along the cortical hierarchy. We tested the model by fitting it to fMRI data collected while subjects watched a 50-minute movie (J. Chen et al., 2017), and then assessing how well the learned event structure explained the activity patterns of a held-out subject. Note that previous analyses of this dataset have shown that the evoked activity is similar across subjects, justifying an across-subjects design (J. Chen et al., 2017). We found that essentially all brain regions that responded consistently to the movie (across subjects) showed evidence for event-like structure, and that the optimal number of events varied across the cortex (Fig. 3). Sensory regions like visual and auditory cortex showed faster transitions between stable activity patterns, while higher-level regions like the posterior medial cortex, angular gyrus, and intraparietal sulcus had activity patterns that often remained constant for over a minute before transitioning to a new stable pattern (see Fig. 3 insets and Fig. S3). This topography of event timescales is broadly consistent with that found in previous work (Hasson et al., 2015) measuring sensitivity to temporal scrambling of a movie stimulus (see Fig. S8).

### Comparison of event boundaries across regions and to human annotations

The second implication of our theory is that “events throughout the hierarchy should have a nested structure, with coarse event boundaries annotated by human observers most strongly related to long events at the top of the hierarchy.” To examine how event boundaries changed throughout the cortical hierarchy, we created four regions of interest, each with 300 voxels, centered on ROIs from prior work (see *STAR Methods*): early visual cortex, late visual cortex, angular gyrus, and posterior medial cortex. We identified the optimal timescale for each region as in the previous analysis, and then fit the event segmentation model at this optimal timescale, as illustrated in Fig. 4a. We found that a significant portion of the boundaries in a given layer were also present in lower layers (Fig. 4b), especially for adjacent layers in the hierarchy. This suggests that event segmentation is at least partially hierarchical, with finer event boundaries nested within coarser boundaries.

We asked four independent raters to divide the movie into “scenes” based on major shifts in the narrative (such as in location, topic, or time). The number of event boundaries identified by the observers varied between 36 and 64, but the boundaries had a significant amount of overlap, with an average pairwise Dice’s coefficient of 0.63 and 20 event boundaries that were labeled by all four raters. We constructed a “consensus” annotation containing boundaries marked by at least two raters, which split the narrative into 54 events, similar to the mean timescale for individual annotators (49.5). We then measured, for each region, what fraction of its fMRI-defined boundaries were close to (within three timepoints of) a consensus event boundary. As shown in Fig. 4c, all regions showed an above-chance match

to human annotations (early visual,  $p=0.0135$ ; late visual,  $p=0.0065$ ; angular gyrus,  $p=0.0011$ ; posterior medial,  $p<0.001$ ), but this match increased across the layers of the hierarchy and was largest in angular gyrus and posterior medial cortex (angular gyrus > early visual,  $p=0.0580$ ; posterior medial > early visual,  $p=0.033$ ).

### Shared event structure across modalities

The third requirement of our theory is that “activity patterns in long timescale regions should be invariant across different descriptions of the same situation.” This hypothesis is based on a prior study (Zadbood et al., 2016), in which one group of subjects watched a 24-minute movie while the other group listened to an 18-minute audio narration describing the events that occurred in the movie. The prior study used a hand-labeled event correspondence between these two stimuli to show that activity patterns evoked by corresponding events in the movie and narration were correlated in a network of regions including angular gyrus, precuneus, retrosplenial cortex, posterior cingulate cortex and mPFC (Zadbood et al., 2016). Here, we use this dataset to ask whether our event segmentation model can replicate these results in a purely unsupervised manner, without using any prior event labeling.

For each cortical searchlight, we first segmented the movie data into events, and then tested whether this same sequence of events from the movie-watching subjects was present in the audio-narration subjects. Regions including the angular gyrus, temporoparietal junction, posterior medial cortex, and inferior frontal cortex showed a strongly significant correspondence between the two modalities (Fig. 5), indicating that a similar sequence of event patterns was evoked by the movie and audio narration irrespective of the modality used to describe the events. In contrast, though low-level auditory cortex was reliably activated by both of these stimuli, there was no above-chance similarity between the series of activity patterns evoked by the two stimuli (movie vs. verbal description), presumably because the low level auditory features of the two stimuli were markedly different.

### Relationship between cortical event boundaries and hippocampal encoding

We tested a fourth hypothesis from our theory, that “the end of an event in long timescale cortical regions should trigger the hippocampus to encode information about the just-concluded event into episodic memory.” Prior work has shown that the end of a video clip is associated with increased hippocampal activity, and the magnitude of the activity predicts later memory (Ben-Yakov & Dudai, 2011; Ben-Yakov et al., 2013). These experiments, however, have used only isolated short video clips with clear transitions between events. Do neurally-defined event boundaries in a continuous movie, evoked by subtler transitions between related scenes, generate the same kind of hippocampal signature? Using a searchlight procedure, we identified event boundaries with the HMM segmentation model for each cortical area across the timescale hierarchy, using the 50-minute movie dataset (J. Chen et al., 2017). We then computed the average *hippocampal* activity around the event boundaries of each cortical area, to determine whether a cortical boundary tended to trigger a hippocampal response. We found that event boundaries in a distributed set of regions including angular gyrus, posterior medial cortex, and parahippocampal cortex all showed a strong relationship to hippocampal activity, with the hippocampal response typically peaking within several timepoints after the event boundary (Fig. 6). This network of regions closely

overlaps with the posterior medial memory system (Ranganath & Ritchey, 2012). Note that both the event boundaries and the hippocampal response are hemodynamic signals, so there is no hemodynamic offset between these two measures. The hippocampal response does start slightly before the event boundary, which could be due to uncertainty in the model estimation of the exact boundary timepoint and/or anticipation that the event is about to end.

### Reinstatement of event patterns during free recall

Our theory further implies that “stored event memories can be reinstated in long timescale cortical regions during recall, with stronger reinstatement for more strongly-encoded events.” After watching the movie, all subjects in this dataset were asked to retell the story they had just watched, without any cues or stimulus (see Chen et al., 2017 for full details). We focused our analyses on the high-level regions that showed a strong relationship with hippocampal activity in the previous analysis (posterior cingulate and angular gyrus), as well as early auditory cortex for comparison.

Using the event segmentation model, we first estimated the (group-average) series of event-specific activity patterns evoked by the movie, and then attempted to segment each subject’s recall data into corresponding events. When fitting the model to the recall data, we assumed that the same event-specific activity patterns seen during the movie-viewing will be reinstated during the spoken recall. Analyzing the spoken recall transcriptions revealed that subjects generally recalled the events in the same order as they appeared in the movie (see table S1 in Chen, Leong, et al., 2016). Therefore, the model was constrained to use the same order of multi-voxel event patterns for recall that it had learned from the movie-watching data. However, crucially, the model was allowed to learn different event timings for the recall data compared to the movie data – this allowed us to accommodate the fact that event durations differed for free recall vs. movie-watching.

For each subject, the model attempted to find a shared sequence of latent event patterns that was shared between the movie and recall, as shown in the example with 25 events in Fig. 7a (see Fig. S6 for examples from all subjects). Green shading indicates the probability that a movie and a recall timepoint belong to the same latent event (with darker shading indicating higher probability), and boxes indicate segments of the movie and recall that were labeled as corresponding to the same event by human annotators. Compared to the null hypothesis that there was no shared event order between the movie and recall, we found significant model fits in both the posterior cingulate ( $p=0.015$ ) and the angular gyrus ( $p=0.002$ ), but not in low-level auditory cortex ( $p=0.277$ ) (Fig. 7b). This result demonstrates that we can identify shared temporal structure between perception and recall without any human annotations. A similar pattern of results can be found regardless of the number of latent events used (see Fig. S6). The identified correspondences for each subject (using both posterior cingulate and angular gyrus) were also significantly similar to the human-labeled correspondences (probability mass inside annotations = 17.8%, significantly greater than null model (11.9%),  $p<0.001$ ) (see Fig. S6).

We then assessed whether the hippocampal response evoked by the end of an event during the encoding of the movie to memory was predictive of the length of time for which the event was strongly reactivated during recall. As shown in Fig. 7c–d, we found that encoding

activity and event reactivation were positively correlated in both angular gyrus ( $r=0.362$ ,  $p=0.002$ ) and the posterior cingulate ( $r=0.312$ ,  $p=0.042$ ), but not early auditory cortex ( $r=0.080$ ,  $p=0.333$ ). Note that there was no relationship between the hippocampal activity at the *starting* boundary of an event and that event's later reinstatement in the angular gyrus ( $r=-0.119$ ,  $p=0.867$ ; difference from ending boundary correlation  $p=0.004$ ) and only a weak, nonsignificant relationship in posterior cingulate ( $r=0.189$ ,  $p=0.113$ ; difference from ending boundary correlation  $p=0.274$ ). The relationship between the *average* hippocampal activity throughout an event and later cortical reinstatement was actually negative (angular gyrus,  $r=-0.247$ ,  $p=0.017$ ; posterior cingulate,  $r=-0.092$ ,  $p=0.175$ ), suggesting that encoding is strongest when hippocampal activity is relatively low during an event and high at its offset.

### Anticipatory reinstatement for a familiar narrative

Finally, we tested a sixth hypothesis, that “prior memory for a narrative should lead to anticipatory reinstatement in long timescale regions.” Our ongoing interpretation of events can be influenced by prior knowledge; specifically, if subjects listening to the audio version of a narrative had already seen the movie version, they may anticipate upcoming events compared to subjects experiencing the narrative for the first time. Detecting this kind of anticipation has not been possible with previous approaches that rely on stimulus annotations, since the difference between the two groups is not in the stimulus (which is identical) but rather in the temporal dynamics of their cognitive processes.

Using data from Zadbood et al. (2016), we simultaneously fit our event segmentation model to three conditions – watching the movie, listening to the narration after watching the movie (“memory”), and listening to the narration without having previously seen the movie (“no memory”) – looking for a shared sequence of event patterns across conditions. By analyzing which timepoints were assigned to the same event, we can generate a timepoint correspondence indicating – for each timepoint during the audio narration datasets – which timepoints of the movie are most strongly evoked (on average) in the mind of the listeners.

We searched for cortical regions along the hierarchy of timescale showing anticipation, in which this correspondence for the memory group was consistently *ahead* of the correspondence for the no-memory group (relative to chance). As shown in Fig. 8, we found anticipatory event reinstatement in the angular gyrus, posterior medial cortex, and medial frontal cortex. Examining the movie-audio correspondences in these regions, the memory group was consistently ahead of the no-memory group, indicating that for a given timepoint of the audio narration the memory group had event representations that corresponded to later timepoints in the movie. A similar result can be obtained by directly aligning the two listening conditions without reference to the movie-watching condition (see Fig. S7).

## Discussion

We found that narrative stimuli evoke event-structured activity throughout the cortex, with pattern dynamics consisting of relatively stable periods punctuated by rapid event transitions. Furthermore, the angular gyrus and posterior medial cortex exhibit a set of overlapping properties associated with high-level situation model representations: long event timescales, event boundaries closely related to human annotations, generalization across



modalities, hippocampal response at event boundaries, reactivation during free recall, and anticipatory coding for familiar narratives. Identifying all of these properties was made possible by using naturalistic stimuli with extended temporal structure, paired with a data-driven model for identifying activity patterns shared across timepoints and across datasets.

### Event segmentation theory

Our results are the first to validate a number of key predictions of event segmentation theory (Zacks et al., 2007) directly from fMRI data of naturalistic narratives, without using specially-constructed stimuli or subjective labeling of where events should start and end. Previous work has shown that hand-labeled event boundaries are associated with univariate activity increases in a network of regions overlapping our high-level areas (Ezzyat & Davachi, 2011; Speer, Zacks, & Reynolds, 2007; Swallow et al., 2011; Whitney et al., 2009; Zacks, Braver, et al., 2001; Zacks, Speer, Swallow, & Maley, 2010), but by modeling fine-scale spatial activity patterns we were able to detect these event changes without an external reference. This allowed us to identify regions with temporal event structure at many different timescales, only some of which matched human-labeled boundaries. Other analyses of these datasets also found reactivation during recall (J. Chen et al., 2017) and shared event structure across modalities (Zadbood et al., 2016); however, because these other analyses defined events based on the narrative rather than brain activity, they were unable to identify differences in event segmentation across brain areas or across groups with different prior knowledge.

### Processing timescales and event segmentation

The topography of event timescales revealed by our analysis provides converging evidence for an emerging view of how information is processed during real-life experience (Hasson et al., 2015). The “process memory framework” argues that perceptual stimuli are integrated across longer and longer timescales along a hierarchy from early sensory regions to regions in the default mode network. Using a variety of experimental approaches, including fMRI, electrocorticography (ECoG), and single-unit recording, this topography has previously been mapped either by temporally scrambling the stimulus at different timescales to see which regions’ responses are disrupted (Hasson, Yang, Vallines, Heeger, & Rubin, 2008; Honey et al., 2012; Lerner, Honey, Silbert, & Hasson, 2011) or by examining the power spectrum of intrinsic dynamics within each region (Honey et al., 2012; Murray et al., 2014; Stephens et al., 2013). Our model and results yield a new perspective on these findings, suggesting that *all* processing regions can exhibit rapid activity shifts, but that these fast changes are much less frequent in long timescale regions. The power spectrum is therefore an incomplete description of voxel dynamics, since the correlation timescale changes dynamically, with faster changes at event boundaries and slower changes within boundaries (see also Fig. S4b). We also found evidence for nested hierarchical structure, suggesting that chunks of information are transmitted from lower to higher levels primarily at event boundaries, as in recent multiscale recurrent neural network models (Chung, Ahn, & Bengio, 2017).

The specific features encoded in the event representations of long timescale regions like the angular gyrus and posterior cingulate cortex during naturalistic perception are still an open question. These areas are involved in high-level, multimodal scene processing tasks

including memory and navigation (Baldassano, Esteva, Beck, & Fei-Fei, 2016; Kumar, Federmeier, Fei-Fei, & Beck, 2017), are part of the “general recollection network” with strong anatomical and functional connectivity to the hippocampus (Rugg & Vilberg, 2013), and are the core components of the posterior medial memory system (Ranganath & Ritchey, 2012), which is thought to represent and update a high-level situation model (Johnson-Laird, 1983; Van Dijk & Kintsch, 1983; Zwaan et al., 1995; Zwaan & Radvansky, 1998). Since event representations in these regions generalized across modalities and between perception and recall, our results provide further evidence that they encode high-level situation descriptions. We also found that event patterns could be partially predicted by key characters and locations from the narrative (see Fig. S4d, and supplementary Fig. 6 in J. Chen et al., 2017), but future work (with richer descriptions of narrative events, Vodrahalli et al., 2016) will be required to understand how event patterns are evoked by semantic content.

### Events in episodic memory

Behavioral experiments have shown that long-term memory reflects event structure during encoding (Ezzyat & Davachi, 2011; Sargent et al., 2013; Zacks, Tversky, et al., 2001). Here, we were able to identify the reinstatement of events that were automatically discovered during perception, extending previous work demonstrating reinstatement of individual items or scenes in angular gyrus and posterior medial cortex (J. Chen et al., 2017; Johnson, McDuff, Rugg, & Norman, 2009; Kuhl & Chun, 2014; Ritchey, Wing, LaBar, & Cabeza, 2013; Wing, Ritchey, & Cabeza, 2015) to continuous perception without any stimulus annotations.

We demonstrated that the hippocampal encoding activity previously shown to be present at the end of movie clips (Ben-Yakov & Dudai, 2011; Ben-Yakov et al., 2013) and at abrupt switches between stimulus category and task (DuBrow & Davachi, 2016) also occurs at the much more subtle transitions between events (defined by pattern shifts in high-level regions), providing evidence that event boundaries trigger the storage of the current situation representation into long-term memory. We have also shown that this post-event hippocampal activity is related to pattern reinstatement during recall, as has been recently demonstrated for the encoding of discrete items (Danker et al., 2016), thereby supporting the view that events are the natural units of episodic memory during everyday life. Changes in cortical activity patterns may drive encoding through a comparator operation in the hippocampus (Lisman & Grace, 2005; Vinogradova, 2001), or the prediction error associated with event boundaries may potentiate the dopamine pathway (Zacks, Kurby, Eisenberg, & Haroutunian, 2011), thereby leading to improved hippocampal encoding (Kempadoo, Mosharov, Choi, Sulzer, & Kandel, 2016; Takeuchi et al., 2016). Notably, the positive relationship between hippocampal activity and subsequent cortical reinstatement was specific to hippocampal activity at the end of an event; there was no significant relationship between hippocampal activity at the start of an event and subsequent reinstatement, and higher hippocampal activity during an event was associated with worse reinstatement (a similar relationship was observed in parietal cortex by Lee, Chun, & Kuhl, 2016). In this respect, our results differ from other results showing that the hippocampal response to novel events drives memory for the novel events *themselves* (for a review, see Ranganath & Rainer, 2003) – here, we show

that the hippocampal response to a new event is linked to subsequent memory for the *previous* event.

Most extant models of memory consolidation (McClelland et al., 1995) and recall (Polyn, Norman, & Kahana, 2009) have been formulated under the assumption that the input to the memory system is a series of discrete items to be remembered. Although this is true for experimental paradigms that use lists of words or pictures, it was not clear how these models could function for realistic autobiographical memory. Our model connects naturalistic perception with theories about discrete memory traces, by proposing that cortex chunks continuous experiences into discrete events; these events are integrated along the processing hierarchy into meaningful, temporally extended, episodic structures, to be later encoded into memory via interaction with the hippocampus. The fact that the hippocampus is coupled to both angular gyrus and posterior medial cortex, which have slightly different timescales, raises the interesting possibility that these event memories could be stored at multiple temporal resolutions.

### **Our event segmentation model**

Temporal latent variable models have been largely absent from the field of human neuroscience, since the vast majority of experiments have a temporal structure that is defined ahead of time by the experimenter. One notable exception is the recent work of Anderson and colleagues, which has used HMM-based models to discover temporal structure in brain activity responses during mathematical problem solving (Anderson & Fincham, 2014; Anderson, Lee, & Fincham, 2014; Anderson, Pyke, & Fincham, 2016). These models are used to segment problem-solving operations (performed in less than 30 seconds) into a small number of cognitively distinct stages such as encoding, planning, solving and responding. Our work is the first to show that (using a modified HMM and an annealed fitting procedure) this latent-state approach can be extended to much longer experimental paradigms with a much larger number of latent states.

For finding correspondences between continuous datasets, as in our analyses of shared structure between perception and recall or perception under different modalities, several other types of approaches (not based on HMMs) have been proposed in psychology and machine learning. Dynamic time warping (Kang & Wheatley, 2015; Silbert, Honey, Simony, Poeppel, & Hasson, 2014) locally stretches or compresses two timeseries to find the best match, and more complex methods such as conditional random fields (Zhu et al., 2015) allow for parts of the match to be out of order. However, these methods do not explicitly model event boundaries.

Future work will be required to investigate what types of neural correspondences are well modeled by continuous warping versus event-structured models. Logically, a strictly event-structured model (with static event patterns) cannot be a complete description of brain activity during narrative perception, since subjects are actively accumulating information during each event, and extensions of our model could additionally model these within-event dynamics (see Fig. S4b).

## Perception and memory in the wild

Our results provide a bridge between the large literature on long-term encoding of individual items (such as words or pictures) and studies of memory for real-life experience (Nielson, Smith, Sreekumar, Dennis, & Sederberg, 2015; Rissman, Chow, Reggente, & Wagner, 2016). Since our approach does not require an experimental design with rigid timing, it opens the possibility of having subjects be more actively and realistically engaged in a task, allowing for the study of events generated during virtual reality navigation (such as spatial boundaries, Horner, Bisby, Wang, Bogus, & Burgess, 2016) or while holding dialogues with a simultaneously-scanned subject (Hasson, Ghazanfar, Galantucci, Garrod, & Keysers, 2012). The model also is not fMRI-specific, and could be applied to other types of neuroimaging timeseries such as electrocorticography (ECoG), electroencephalography (EEG), or functional near-infrared spectroscopy (fNIRS), including portable systems that could allow experiments to be run outside the lab (Mckendrick et al., 2016).

## Conclusion

Using a novel event segmentation model that can be fit directly to neuroimaging data, we showed that cortical responses to naturalistic stimuli are temporally organized into discrete events at varying timescales. In a network of high-level association regions, we found that these events were related to subjective event annotations by human observers, predicted hippocampal encoding, generalized across modalities and between perception and recall, and showed anticipatory coding of familiar narratives. Our results provide a new framework for understanding how continuous experience is accumulated, stored, and recalled.

## STAR Methods

### Contact for Reagent and Resource Sharing

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Christopher Baldassano ([chrisb@princeton.edu](mailto:chrisb@princeton.edu)).

### Experimental Model and Subject Details

**Interleaved Stories dataset**—22 subjects (all native English speakers) were recruited from the Princeton community (9 male, 13 female, ages 18–26). All subjects provided informed written consent prior to the start of the study in accordance with experimental procedures approved by the Princeton University Institutional Review Board. The study was approximately 2 hours long and subjects received \$20 per hour as compensation for their time. Data from 3 subjects were discarded due to falling asleep during the scan, and 1 due to problems with audio delivery.

### Method Details

**Interleaved Stories dataset**—To test our model in a dataset with clear, unambiguous event boundaries, we used data from subjects who listened to two unrelated audio narratives (J. Chen, Chow, Norman, & Hasson, 2015). We used data from 18 subjects who listened to the two audio narratives in an interleaved fashion, with the audio stimulus switching between the two narratives approximately every 60 seconds at natural paragraph breaks. The

total stimulus length was approximately 29 minutes, during which there were 32 story switches. The audio was delivered via in-ear headphones.

Imaging data were acquired on a 3T full-body scanner (Siemens Skyra) with a 20-channel head coil using a T2\*-weighted echo planar imaging (EPI) pulse sequence (TR 1500 ms, TE 28 ms, flip angle 64, whole-brain coverage 27 slices of 4 mm thickness, in-plane resolution  $3 \times 3$  mm, FOV  $192 \times 192$  mm). Preprocessing was performed in FSL, including slice time correction, motion correction, linear detrending, high-pass filtering (140 s cutoff), and coregistration and affine transformation of the functional volumes to a template brain (MNI). Functional images were resampled to 3 mm isotropic voxels for all analyses.

The analyses in this paper were carried out using data from a posterior cingulate region of interest, the posterior medial cluster in the “dorsal default mode network” defined by whole-brain resting state connectivity clustering (Shirer, Ryali, Rykhlevskaia, Menon, & Greicius, 2012).

**Sherlock Recall dataset**—Our primary dataset consisted of 17 subjects who watched the first 50 minutes of the first episode of BBC’s *Sherlock*, and were then asked to freely recall the episode in the scanner without cues (J. Chen et al., 2017). Subjects varied in the length and richness of their recall, with total recall times ranging from 11 minutes to 46 minutes (and a mean of 22 minutes). Imaging data was acquired using a T2\*-weighted echo planar imaging (EPI) pulse sequence (TR 1500 ms, TE 28 ms, flip angle 64, whole-brain coverage 27 slices of 4 mm thickness, in-plane resolution  $3 \times 3$  mm, FOV  $192 \times 192$  mm). A standard preprocessing pipeline was performed using FSL, including motion correction. Since acoustic output was not correlated across subjects (J. Chen et al., 2017), shared activity patterns at recall are unlikely to be driven by correlated motion artifacts. All subjects were aligned to a common MNI template, and analyses were carried out in this common volume space. We also conducted an alternate version of the segmentation analysis that does not rely on precise MNI alignment (P.-H. (Cameron) Chen et al., 2015) and obtained similar results (see Fig. S4c).

We restricted our searchlight analyses to voxels that were reliably driven by the stimuli, measured using intersubject correlation (Hasson, Nir, Levy, Fuhrmann, & Malach, 2004). Voxels with a correlation less than  $r=0.25$  during movie-watching were removed before running the searchlight analysis.

We defined five regions of interest based on prior work. In addition to the posterior cingulate region defined above, we defined the angular gyrus as area PG (both PGa and PGp) using the maximum probability maps from a cytoarchitectonic atlas (Eickhoff et al., 2005), early auditory cortex as voxels within the Heschl’s gyrus region (Harvard-Oxford cortical atlas) with high intersubject correlation during an audio narrative (“Pieman”, Simony et al., 2016), early visual cortex as voxels near the calcarine sulcus with high intersubject correlation during an audio-visual movie (“The Twilight Zone”, Chen, Honey, et al., 2016), and hV4 based on a group maximum-probability atlas (Wang, Mruczek, Arcaro, & Kastner, 2014).

**Sherlock Narrative dataset**—To investigate cross-modal event representations and the impact of prior memory, we used a separate dataset in which subjects experienced multiple versions of a narrative. One group of 17 subjects watched the first 24 minutes of the first episode of *Sherlock* (a portion of the same episode used in the Sherlock Recall dataset), while another group of 17 subjects (who had never seen the episode before) listened to an 18 minute audio description of the events during this part of the episode (taken from the audio recording of one subject’s recall in the Sherlock Recall dataset). The subjects who watched the episode then listened to the same 18 minute audio description. This yielded three sets of data, all based on the same story: watching a movie of the events, listening to an audio narration of the events *without* prior memory, and listening to an audio narration of the events *with* prior memory. Imaging data was acquired using the same sequence as in Sherlock Recall dataset; see Zadbood et al. (2016) for full acquisition and preprocessing details.

As in the Sherlock Recall experiment, we removed all voxels that were not reliably driven by the stimuli. Only voxels with an intersubject correlation of at least  $r=0.1$  across all three conditions were included in searchlight analyses.

**Event annotations by human observers**—Four human observers were given the video file for the 50-minute *Sherlock* stimulus, and given the following directions: “Write down the times at which you feel like a new scene is starting; these are points in the movie when there is a major change in topic, location, time, etc. Each ‘scene’ should be between 10 seconds and 3 minutes long. Also, give each scene a short title.” The similarity among observers was measured using Dice’s coefficient (number of matching boundaries divided by mean number of boundaries, considering boundaries within three timepoints of one another to match).

**Event Segmentation Model**—Our model is built on two hypotheses: 1) while processing narrative stimuli, observers experience a sequence of discrete events, and 2) each event has a distinct neural signature. Mathematically, the model assumes that a given subject (or averaged group of subjects) starts in event  $s_1=1$  and ends in event  $s_T=K$ , where  $T$  is the total number of timepoints and  $K$  is the total number of events. In each timepoint the subject either remains in the same state or advances to the next state, i.e.  $s_{t+1} \in \{s_t, s_t+1\}$  for all timepoints  $t$ . Each event has a signature mean activity pattern  $m_k$  across all  $V$  voxels in a region of interest, and the observed brain activity  $b_t$  at any timepoint  $t$  is assumed to be highly correlated with  $m_k$ , as illustrated in Fig. 2.

Given the sequence of observed brain activities  $b_t$ , our goal is to infer both the event signatures  $m_k$  and the event structure  $s_t$ . To accomplish this, we cast our model as a variant of a Hidden Markov Model (HMM). The latent states are the events  $s_t$  that evolve according to a simple transition matrix, in which all elements are zero except for the diagonal (corresponding to  $s_{t+1} = s_t$ ) and the adjacent off-diagonal (corresponding to  $s_{t+1} = s_t+1$ ), and

the observation model is an isotropic Gaussian  $p(b_t | s_t=k) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2} \|z(b_t) - z(m_k)\|_2^2}$ , where  $z(x)$  denotes z-scoring an input vector  $x$  to have zero mean and unit variance. Note that, due to this z-scoring, the log probability of observing brain state  $b_t$  in an event with

signature  $m_k$  is simply proportional to the Pearson correlation between  $b_t$  and  $m_k$  plus a constant offset.

The HMM is fit to the fMRI data by using an annealed version of the Baum-Welch algorithm, which iterates between estimating the fMRI signatures  $m_k$  and the latent event structure  $s_t$ . Given the signature estimates  $m_k$ , the distributions over latent events  $p(s_t=k)$  can be computed using the forward-backward algorithm. Given the distributions  $p(s_t=k)$ , the updated estimates for the signatures  $m_k$  can be computed as the weighted average

$$m_k = \frac{\sum_t p(s_t=k) b_t}{\sum_t p(s_t=k)}.$$

To encourage convergence to a high-likelihood solution, we anneal the observation variance  $\sigma^2$  as  $4 \cdot 0.98^i$  where  $i$  is the number of loops of Baum-Welch completed so far. We stop the fitting procedure when the log-likelihood begins to decrease, indicating that the observation variance has begun to drop below the actual event activity variance. We can also fit the model simultaneously to multiple datasets; on each round of Baum-Welch, we run the forward-backward algorithm on each dataset separately, and then average across all datasets to compute a single set of shared signatures  $m_k$ .

The end state requirement of our model - that all states should be visited, and the end state should be symmetrical to all other states - requires extending the traditional HMM by modifying the observation probabilities  $p(b_t | s_t = k)$ . First, we enforce  $s_T = K$  by requiring that, on the final timestep, only the final state  $K$  could have generated the data, by setting  $p(b_T | s_T = k) = 0$  for all  $k \neq K$ . Equivalently, we can view this as a modification of the backwards pass, by initializing the backwards message  $\beta(s_T = k)$  to 1 for  $k = K$  and 0 otherwise. Second, we must modify the transition matrix to ensure that all valid event segmentations (which start at event 1 and end at event  $K$ , and proceed monotonically through all events) have the same prior probability. Formally, we introduce a dummy absorbing state  $K+1$  to which state  $K$  can transition, ensuring that the transition probabilities for state  $K$  are identical to those for previous states, and then set  $p(b_t | s_t = K+1) = 0$  to ensure that this state is never actually used.

Since we do not want to assume that events will have the same relative lengths across different datasets (such as a movie and audio-narration version of the same narrative), we fix all states to have the same probability of staying in the same state ( $s_{t+1} = s_t$ ) versus jumping to the next state ( $s_{t+1} = s_t + 1$ ). Note that the shared probability of jumping to the next state can take any value between 0 and 1 with no effect on the results (up to a normalization constant in the log-likelihood), since every valid event segmentation will contain exactly the same number of jumps ( $K-1$ ).

Our model induces a prior over the locations of the event boundaries. There are a total of  $\binom{T}{K-1}$  equally likely placements of the  $K-1$  event boundaries, and the number of ways to have event boundary  $k$  fall on timepoint  $t$  is the number of ways that  $k-1$  boundaries can be placed in  $t-1$  timepoints times the number of ways that  $(K-1)-(k-1)-1$  boundaries can be

$$p(s_t=k \& s_{t+1}=k+1) = \frac{\binom{t-1}{K-1} \binom{T-t}{K-k-1}}{\binom{T}{K-1}}.$$

placed in  $T-t$  timepoints. Therefore an example of this distribution is shown in Fig. S1. During the annealing process, the distribution over boundary locations starts at this prior, and slowly adjusts to match the event structure of the data.

After fitting the model on one set of data, we can then look for the same sequence of events in another dataset. Using the signatures  $m_k$  learned from the first dataset, we simply perform a single round of the forward-backward algorithm to obtain event estimates  $p(s_t=k)$  on the second dataset. If we expect the datasets to have similar noise properties (e.g. both datasets are group-averaged data from the same number of subjects), we set the observation variance to the final  $\sigma^2$  obtained while fitting the first dataset. When transferring events learned on group-averaged data to individual subjects, we estimate the variance for each event across the individual subjects of the first dataset.

The model implementation was first verified using simulated data. An event-structured dataset was constructed with  $V=10$  voxels,  $K=10$  events, and  $T=500$  timepoints. The event structure was chosen to be either uniform (with 50 timepoints per event), or the length of each event was sampled (from first to last) from  $\mathcal{N}(1,0.25) * (\text{timepoints remaining}) / (\text{events remaining})$ . A mean pattern was drawn for each event from a standard normal distribution, and the simulated data for each timepoint was the sum of the event pattern for that timepoint plus randomly distributed noise with zero mean and varying standard deviation. The noisy data were then input to the event segmentation model, and we measured the fraction of the event boundaries that were exactly recovered from the true underlying event structure. As shown in Fig. S1, we were able to recover a majority of the event boundaries even when the noise level was as large as the signature patterns themselves.

**Finding event structure in narratives**—To validate our event segmentation model on real fMRI data, we first fit the model to group-averaged PCC data from the Interleaved Stories experiment. In this experiment, we expect that an event boundary should be generated every time the stimulus switches stories, giving a ground truth against which to compare the model's segmentations. As shown in Fig. S2, our method was highly effective at identifying events, with the majority of the identified boundaries falling close to a story switch.

The following subsections describe how the model was used to obtain each of the experimental results, with subsection titles corresponding to subsections of the *Results*.

**Timescales of cortical event segmentation**—We applied the model in a searchlight to the whole-brain movie-watching data from the Sherlock Recall study. Cubical searchlights were scanned throughout the volume at a step size of 3 voxels and with a side length of 7 voxels. For each searchlight, the event segmentation model was applied to group-averaged data from all but one subject. We measured the robustness of the identified boundaries by testing whether these boundaries explained the data in the held-out subject.



We measured the spatial correlation between all pairs of timepoints that were separated by four timepoints, and then binned these correlations according to whether the pair of timepoints fell within the same event or crossed over an event boundary. The average difference between the within- versus across-event correlations was used as an index of how well the learned boundaries captured the temporal structure of the held-out subject. The analysis was repeated for every possible held-out subject, and with a varying number of events from  $K=10$  to  $K=120$ . After averaging the results across subjects, the number of events with the best within- versus across-event correlations was chosen as the optimal number of events for this searchlight.

Since the topography of the results was similar to previous work on temporal receptive windows, we compared the map of the optimal number of events with the short and medium/long timescale maps derived by measuring inter-subject correlation for intact versus scrambled movies (J. Chen et al., 2016). The histogram of the optimal number of events for voxels was computed within each of the timescale maps.

#### **Comparison of event boundaries across regions and to human annotations—**

We defined four equally-sized regions along the cortical hierarchy by taking the centers of mass of the early visual cortex, hV4, angular gyrus and PCC ROIs in each hemisphere, and finding the nearest 150 voxels to these centers (yielding 300 bilateral voxels for each region). For each region, we calculated its optimal number of events using the same within-versus across-event correlation procedure described in the previous section, and then fit a final event segmentation model with the optimal number of events (using the Sherlock Recall data). We measured the match between levels of the hierarchy by computing the fraction of boundaries in the “upper” (slower) level that were close to a boundary in the “lower” (faster) level. We defined “close to” as “within three timepoints,” since the typical uncertainty in the model about exactly where an event switch occurred was approximately three timepoints.

To measure similarity to the human annotations, we first constructed a “consensus” annotation from the four observers, consisting of boundary locations that were within three timepoints of boundaries marked by at least two observers. We then measured the match between the consensus boundaries and the boundaries from each region, treating the consensus boundaries as the “upper” level. To ensure that differences between regions were not driven by timescale differences, for this comparison we refit the event segmentation model to each cortical region using the same number of events as in the consensus annotation (rather than using each region’s optimal timescale).

**Shared event structure across modalities—**After fitting the event segmentation model to a searchlight of movie-watching data from the Sherlock Narration experiment, we took the learned event signatures  $m_k$  and used them to run the forward-backward algorithm on the audio narration data, to test whether audio narration of a story elicited the same sequence of events as a movie of that story. Since both the movie and audio data were averaged at the group level, they should have similar levels of noise, and therefore we simply used the fit movie variance  $\sigma^2$  for the observation variance.

### Relationship between cortical event boundaries and hippocampal encoding—

After applying the event segmentation model throughout the cortex to the Sherlock Recall study as described above, we measured whether the data-driven event boundaries were related to activity in the hippocampus. For a given cortical searchlight, we extracted a window of mean hippocampal activity around each of the searchlight's event boundaries. We then averaged these windows together, yielding a profile of boundary-triggered hippocampal response according to this region's boundaries. To assess whether the hippocampus showed a significant increase in activity related to these event boundaries, we measured the mean hippocampal activity for the 10 timepoints following the event boundary minus the mean activity for the 10 timepoints preceding the event boundary.

**Reinstatement of event patterns during free recall—**For each region of interest, we fit the event segmentation model as described above (on the group-averaged Sherlock Recall data). We then took the learned sequence of event signatures  $m_k$  and ran the forward-backward algorithm on each individual subject's recall data. We set the variance of each event's observation model by computing the variance within each event in the movie-watching data of individual subjects, pooling across both timepoints and subjects. The analysis was run for 10 events to 60 events in steps of 5.

We operationalized the overall reinstatement of an event  $k$ , as  $\sum_t p(s_t = k)$ ; that is, the sum across all recall time points of the probability that the subject was recalling perceptual event  $k$  at that time point. We measured whether this per-event re-activation during recall could be predicted during movie-watching, based on the hippocampal response at the end of the event. For each subject, we computed the difference between hippocampal activity after versus before the event boundary as above. We then averaged the event re-activation and hippocampal offset response across subjects, and measured their correlation. For comparison purposes, we also performed the same analysis but with hippocampal differences at the *beginning* of each event, rather than the end, and with the mean hippocampal activity throughout the event.

**Anticipatory reinstatement for a familiar narrative—**To determine whether memory changed the event correspondence between the movie and narration, we then fit the segmentation model simultaneously to group-averaged data from the movie-watching condition, audio narration no-memory condition, and audio narration with memory condition, yielding a sequence of events in each condition with the same activity signatures. We computed the correspondence between the movie states  $s_{m,t}$  and the audio no-memory states  $s_{ann,t}$  as  $p(s_{m,t_1} = s_{ann,t_2}) = \sum_k p(s_{m,t_1} = k) \cdot p(s_{ann,t_2} = k)$ , and similarly for the audio memory states  $s_{am,t}$ . We computed the differences between the group correspondences as  $\sum_{t_1} \sum_{t_2} (p(s_{m,t_1} = s_{ann,t_2}) - p(s_{m,t_1} = s_{am,t_2}))^2$ . For visualization, we also computed how far the memory correspondence was ahead of the no-memory correspondence as the mean over  $t_2$  of the difference in the expected values  $\sum_{t_1} t_1 p(s_{m,t_1} = s_{ann,t_2}) - \sum_{t_1} t_1 p(s_{m,t_1} = s_{am,t_2})$ . We also performed the same analysis but with only the two narration conditions, computing the correspondence between the audio memory and audio no-memory states as  $p(s_{am,t_1} = s_{ann,t_2}) = \sum_k p(s_{am,t_1} = k) \cdot p(s_{ann,t_2} = k)$ . Since deviation from a diagonal correspondence would indicate anticipation in the memory group, we measured the expected deviation from

the diagonal as  $\sum_{t_1} \sum_{t_2} \frac{|t_1 - t_2|}{\sqrt{2}} p(s_{am,t_1} = s_{anm,t_2})$ , and for visualization calculated the amount of anticipation as  $\sum_t t p(s_{am,t_1} = s_{anm,t_2}) - T/2$ .

### Quantification and Statistical Analysis

Permutation or resampling analyses were used to statistically evaluate all of the results. As above, the analyses for each subsection are presented under a corresponding heading.

**Timescales of cortical event segmentation—**To generate a null distribution, the same analysis was performed except that the event boundaries were scrambled before computing the within- versus across-event correlation. This scrambling was performed by reordering the events with their durations held constant, to ensure that the null events had the same distribution of event lengths as the real events. The within versus across difference for the real events compared to 1000 null events was used to compute a z value, which was converted to a p value using the normal distribution. The p values were Bonferroni corrected for the 12 choices of the number of events, and then the false discovery rate q was computed using the same calculation as in AFNI (Cox, 1996).

### Comparison of event boundaries across regions and to human annotations—

For each pairwise comparison between regions or between a region and the human annotations, we scrambled the event boundaries using the same duration-preserving procedure described above to produce 1000 null match values. The true match value was compared to this distribution to compute a z value, which was converted to a p value. To assess whether two cortical regions had significantly *different* matches to human annotations, we scrambled the boundaries from the regions (keeping the human annotations intact), and computed the fraction of scrambles for which the difference in the match to human annotations was larger in the null data than the original data.

**Shared event structure across modalities—**We compared the log-likelihood of the fit to the narration data against a null model in which the movie event signatures  $m_k$  were randomly re-ordered, and computed the z value of the true log-likelihood compared to 100 null shuffles, then converted to a p value. This null hypothesis test therefore assessed whether the narration exhibited *ordered* reactivation of the events identified during movie-watching.

### Relationship between cortical event boundaries and hippocampal encoding—

For each searchlight, we compared the difference in hippocampal activity for the 10 timepoints after an event boundary compared to 10 timepoints before an event boundary, both on the true boundaries and on shuffled event boundaries (using the duration-preserving procedure described above). The z value for this difference was computed to a p value, and then transformed to a false discovery rate q.

**Reinstatement of event patterns during free recall—**As in the movie-to-narration analysis, we compared the log-likelihood of the movie-recall fit to a null model in which the order of the event signatures was shuffled before fitting to the recall data, which yielded a z

value that was converted to a p value. When measuring the match to human annotations, we compared to the same shuffled-event null models.

To assess the robustness of the encoding activity versus reinstatement correlations, we performed a bootstrap test, in which we resampled subjects (with replacement, yielding 17 subjects as in the original dataset) before taking the average and computing the correlation. The p value was defined as the fraction of 1000 resamples that yielded correlations with a different sign from the true correlation.

**Anticipatory reinstatement for a familiar narrative**—To determine if the correspondence with the movie was significantly different between the memory and no-memory conditions, we created null groups by averaging together a random half of the no-memory subjects with a random half of memory subjects, and then averaging together the remaining subjects from each group, yielding two group-averaged timecourses whose correspondences should differ only by chance. We calculated a z value based on the correspondence difference for real versus null groups, which was converted to a p value and then corrected to a false discovery rate q. The analysis using only the narration conditions was performed similarly, computing a z value based on the expected deviation from the diagonal in the real versus null groups.

### Data and Software Availability

All of the primary data used in this study are drawn from other papers (J. Chen et al., 2017; Zadbood et al., 2016), and the “Interleaved Stories” posterior cingulate cortex data are available on request. Implementations of our event segmentation model, along with simulated data examples, are available on GitHub at <https://github.com/intelpni/brainiak> (python) and at <https://github.com/cbaldassano/Event-Segmentation> (Matlab).

### Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

### Acknowledgments

We thank M. Chow for assistance in collecting the Interleaved Stories dataset, M.C. Jordan for help in porting the model implementation to python, and the members of the Hasson and Norman labs for their comments and support. This work was supported by a grant from Intel Labs (CAB), The National Institutes of Health (R01 MH112357-01, UH and KAN, R01-MH094480, UH, and 2T32MH065214-11, KAN), the McKnight Foundation (JWP), NSF CAREER Award IIS-1150186 (JWP), and a grant from the Simons Collaboration on the Global Brain (JWP).

### References

- Anderson JR, Fincham JM. Discovering the sequential structure of thought. *Cognitive Science*. 2014; 38(2):322–352. <http://doi.org/10.1111/cogs.12068>. [PubMed: 23941168]
- Anderson JR, Lee HS, Fincham JM. Discovering the structure of mathematical problem solving. *NeuroImage*. 2014; 97:163–177. <http://doi.org/10.1016/j.neuroimage.2014.04.031>. [PubMed: 24746954]
- Anderson, JR., Pyke, AA., Fincham, JM. Hidden Stages of Cognition Revealed in Patterns of Brain Activation. *Psychological Science*. 2016. <http://doi.org/10.1177/0956797616654912>

- Baldassano C, Esteva A, Beck DM, Fei-Fei L. Two distinct scene processing networks connecting vision and memory. 2016; 3 Retrieved from <http://biorxiv.org/lookup/doi/10.1101/057406>.
- Beal, DJ., Weiss, HM. The Episodic Structure of Life at Work. In: Bakker, AB., Daniels, K., editors. A Day in the Life of a Happy Worker. New York, NY: Psychology Press; 2013. p. 8-24.
- Ben-Yakov A, Dudai Y. Constructing realistic engrams: poststimulus activity of hippocampus and dorsal striatum predicts subsequent episodic memory. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*. 2011; 31(24):9032–9042. <http://doi.org/10.1523/JNEUROSCI.0702-11.2011>. [PubMed: 21677186]
- Ben-Yakov A, Eshel N, Dudai Y. Hippocampal immediate poststimulus activity in the encoding of consecutive naturalistic episodes. *Journal of Experimental Psychology General*. 2013; 142(4):1255–1263. <http://doi.org/10.1037/a0033558>. [PubMed: 23815458]
- Chen, J., Chow, M., Norman, KA., Hasson, U. Society for Neuroscience. 2015. Differentiation of neural representations during processing of multiple information streams.
- Chen J, Honey CJ, Simony E, Arcaro MJ, Norman KA, Hasson U. Accessing Real-Life Episodic Information from Minutes versus Hours Earlier Modulates Hippocampal and High-Order Cortical Dynamics. *Cerebral Cortex*. 2016; 26(8):3428–3441. <http://doi.org/10.1093/cercor/bhv155>. [PubMed: 26240179]
- Chen J, Leong YC, Honey CJ, Yong CH, Norman KA, Hasson U. Shared memories reveal shared structure in neural activity across individuals. *Nature Neuroscience*. 2017; 20(1):115–125. <http://doi.org/10.1038/nn.4450>. [PubMed: 27918531]
- Chen, P-H., (Cameron), Chen, J., Yeshurun, Y., Hasson, U., Haxby, J., Ramadge, PJ. A Reduced-Dimension fMRI Shared Response Model; Neural Information Processing Systems Conference (NIPS). 2015. p. 460-468. Retrieved from <http://papers.nips.cc/paper/5855-a-reduced-dimension-fmri-shared-response-model>
- Chung, J., Ahn, S., Bengio, Y. Hierarchical Multiscale Recurrent Neural Networks. International Conference on Learning Representations. 2017. Retrieved from <https://arxiv.org/pdf/1609.01704v5.pdf>
- Cox RW. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research, an International Journal*. 1996; 29(3):162–173. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8812068>.
- Danker, JF., Tompary, A., Davachi, L. Trial-by-Trial Hippocampal Encoding Activation Predicts the Fidelity of Cortical Reinstatement During Subsequent Retrieval; *Cerebral Cortex*. 2016. p. bhv146 <http://doi.org/10.1093/cercor/bhv146>
- DuBrow S, Davachi L. Temporal binding within and across events. *Neurobiology of Learning and Memory*. 2016; 134:107–114. <http://doi.org/10.1016/j.nlm.2016.07.011>. [PubMed: 27422018]
- Eickhoff SB, Stephan KE, Mohlberg H, Grefkes C, Fink GR, Amunts K, Zilles K. A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. *NeuroImage*. 2005; 25(4):1325–1335. <http://doi.org/10.1016/j.neuroimage.2004.12.034>. [PubMed: 15850749]
- Ezzyat Y, Davachi L. What constitutes an episode in episodic memory? *Psychological Science*. 2011; 22(2):243–252. <http://doi.org/10.1177/0956797610393742>. [PubMed: 21178116]
- Gershman SJ, Radulescu A, Norman KA, Niv Y. Statistical Computations Underlying the Dynamics of Memory Updating. *PLoS Computational Biology*. 2014; 10(11):e1003939. <http://doi.org/10.1371/journal.pcbi.1003939>. [PubMed: 25375816]
- Giraud A-L, Poeppel D. Cortical oscillations and speech processing: emerging computational principles and operations. *Nature Neuroscience*. 2012; 15(4):511–517. <http://doi.org/10.1038/nn.3063>. [PubMed: 22426255]
- Hasson U, Chen J, Honey CJ. Hierarchical process memory: Memory as an integral component of information processing. *Trends in Cognitive Sciences*. 2015; 19(6):304–313. <http://doi.org/10.1016/j.tics.2015.04.006>. [PubMed: 25980649]
- Hasson U, Ghazanfar AA, Galantucci B, Garrod S, Keysers C. Brain-to-brain coupling: A mechanism for creating and sharing a social world. *Trends in Cognitive Sciences*. 2012; 16(2):114–121. <http://doi.org/10.1016/j.tics.2011.12.007>. [PubMed: 22221820]

- Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R. Intersubject synchronization of cortical activity during natural vision. *Science (New York, N.Y.)*. 2004; 303(5664):1634–1640. <http://doi.org/10.1126/science.1089506>.
- Hasson U, Yang E, Vallines I, Heeger DJ, Rubin N. A Hierarchy of Temporal Receptive Windows in Human Cortex. *Journal of Neuroscience*. 2008; 28(10):2539–2550. <http://doi.org/10.1523/JNEUROSCI.5487-07.2008>. [PubMed: 18322098]
- Honey CJ, Thesen T, Donner TH, Silbert LJ, Carlson CE, Devinsky O, Hasson U. Slow Cortical Dynamics and the Accumulation of Information over Long Timescales. *Neuron*. 2012; 76(2):423–434. <http://doi.org/10.1016/j.neuron.2012.08.011>. [PubMed: 23083743]
- Horner AJ, Bisby JA, Wang A, Bogus K, Burgess N. The role of spatial boundaries in shaping long-term event representations. *Cognition*. 2016; 30:1–30. <http://doi.org/10.1016/j.cognition.2016.05.013>.
- Johnson JD, McDuff SGR, Rugg MD, Norman KA. Recollection, Familiarity, and Cortical Reinstatement: A Multivoxel Pattern Analysis. *Neuron*. 2009; 63(5):697–708. <http://doi.org/10.1016/j.neuron.2009.08.011>. [PubMed: 19755111]
- Johnson-Laird, PN. *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press; 1983.
- Kang O, Wheatley T. Pupil dilation patterns reflect the contents of consciousness. *Consciousness and Cognition*. 2015; 35(August):128–135. <http://doi.org/10.1016/j.concog.2015.05.001>. [PubMed: 26002764]
- Kempadoo KA, Mosharov EV, Choi SJ, Sulzer D, Kandel ER. Dopamine release from the locus coeruleus to the dorsal hippocampus promotes spatial learning and memory. *Proceedings of the National Academy of Sciences*. 2016; 113(51):14835–14840. <http://doi.org/10.1073/pnas.1616515114>.
- Kravitz DJ, Saleem KS, Baker CI, Mishkin M. A new neural framework for visuospatial processing. *Nature Reviews Neuroscience*. 2011; 12(4):217–230. <http://doi.org/10.1038/nrn3008>. [PubMed: 21415848]
- Kuhl, Ba, Chun, MM. Successful remembering elicits event-specific activity patterns in lateral parietal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*. 2014; 34(23):8051–8060. <http://doi.org/10.1523/JNEUROSCI.4328-13.2014>. [PubMed: 24899726]
- Kumar, M., Federmeier, KD., Fei-Fei, L., Beck, DM. Evidence For Similar Patterns of Neural Activity Elicited by Picture- and Word-based Representations of Natural Scenes. *NeuroImage*. 2017. <http://doi.org/10.1016/j.neuroimage.2017.03.037>
- Kurby CA, Zacks JM. Segmentation in the perception and memory of events. *Trends in Cognitive Sciences*. 2008; 12(2):72–79. <http://doi.org/10.1016/j.tics.2007.11.004>. [PubMed: 18178125]
- Lee, H., Chun, MM., Kuhl, BA. Lower Parietal Encoding Activation Is Associated with Sharper Information and Better Memory; *Cerebral Cortex (New York, N.Y. : 1991)*. 2016. p. bhw097 <http://doi.org/10.1093/cercor/bhw097>
- Lerner Y, Honey CJ, Silbert LJ, Hasson U. Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*. 2011; 31(8):2906–2915. <http://doi.org/10.1523/JNEUROSCI.3684-10.2011>. [PubMed: 21414912]
- Lisman JE, Grace AA. The hippocampal-VTA loop: Controlling the entry of information into long-term memory. *Neuron*. 2005; 46(5):703–713. <http://doi.org/10.1016/j.neuron.2005.05.002>. [PubMed: 15924857]
- McClelland, JL., McNaughton, BL., O'Reilly, RC. 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*. 1995. <http://doi.org/10.1037/0033-295X.102.3.419>
- Mckendrick R, Parasuraman R, Murtza R, Formwalt A, Baccus W, Paczynski M, Ayaz H. Into The Wild: Neuroergonomic Differentiation of Hand-Held and Augmented Reality Wearable Displays During Outdoor Navigation with Functional Near Infrared Spectroscopy. *Frontiers in Human Neuroscience*. 2016; 10(216):1–15. <http://doi.org/10.3389/fnhum.2016.00216>. [PubMed: 26858619]

- Moscovitch M, Rosenbaum RS, Gilboa A, Addis DR, Westmacott R, Grady C, Nadel L. Functional neuroanatomy of remote episodic, semantic and spatial memory: A unified account based on multiple trace theory. *Journal of Anatomy*. 2005; 207(1):35–66. <http://doi.org/10.1111/j.1469-7580.2005.00421.x>. [PubMed: 16011544]
- Murray JD, Bernacchia A, Freedman DJ, Romo R, Wallis JD, Cai X, Wang X-J. A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*. 2014; 17(12):1661–1663. <http://doi.org/10.1038/nn.3862>. [PubMed: 25383900]
- Nelson, M., Karoui, IEL, Giber, K., Yang, X., Cohen, L., Koopman, H., Angeles, L. Neurophysiological dynamics of phrase-structure building during sentence processing. 2017. p. 1-21. <http://doi.org/10.1073/pnas.1701590114>
- Newton D. Attribution and the unit of perception of ongoing behavior. *Journal of Personality and Social Psychology*. 1973; 28(1):28–38. <http://doi.org/10.1037/h0035584>.
- Newton D, Engquist GA, Bois J. The objective basis of behavior units. *Journal of Personality and Social Psychology*. 1977; 35(12):847–862. <http://doi.org/10.1037/0022-3514.35.12.847>.
- Nielson DM, Smith TA, Sreekumar V, Dennis S, Sederberg PB. Human hippocampus represents space and time during retrieval of real-world memories. *Proceedings of the National Academy of Sciences of the United States of America*. 2015; 112(35):11078–11083. <http://doi.org/10.1073/pnas.1507104112>. [PubMed: 26283350]
- Norman KA, O'Reilly RC. Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychological Review*. 2003; 110(4):611–646. <http://doi.org/10.1037/0033-295X.110.4.611>. [PubMed: 14599236]
- Polyn SM, Norman Ka, Kahana MJ. A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*. 2009; 116(1):129–156. <http://doi.org/10.1037/a0014420>. [PubMed: 19159151]
- Radvansky GA. Across the Event Horizon. *Current Directions in Psychological Science*. 2012; 21(4):269–272. <http://doi.org/10.1177/0963721412451274>.
- Radvansky GA, Zacks JM. Event perception. *Wiley Interdisciplinary Reviews: Cognitive Science*. 2011; 2(6):608–620. <http://doi.org/10.1002/wcs.133>. [PubMed: 23082236]
- Ranganath C, Rainer G. Neural mechanisms for detecting and remembering novel events. *Nature Reviews Neuroscience*. 2003; 4(3):193–202. <http://doi.org/10.1038/nrn1052>. [PubMed: 12612632]
- Ranganath C, Ritchey M. Two cortical systems for memory-guided behaviour. *Nature Reviews Neuroscience*. 2012; 13(10):713–726. <http://doi.org/10.1038/nrn3338>. [PubMed: 22992647]
- Rissman J, Chow TE, Reggente N, Wagner AD. Decoding fMRI Signatures of Real-world Autobiographical Memory Retrieval. *Journal of Cognitive Neuroscience*. 2016; 28(4):604–620. [http://doi.org/10.1162/jocn\\_a\\_00920](http://doi.org/10.1162/jocn_a_00920). [PubMed: 26741803]
- Ritchey M, Wing EA, LaBar KS, Cabeza R. Neural Similarity Between Encoding and Retrieval is Related to Memory Via Hippocampal Interactions. *Cerebral Cortex*. 2013; 23(12):2818–2828. <http://doi.org/10.1093/cercor/bhs258>. [PubMed: 22967731]
- Rugg MD, Vilberg KL. Brain networks underlying episodic memory retrieval. *Current Opinion in Neurobiology*. 2013; 23(2):255–260. <http://doi.org/10.1016/j.conb.2012.11.005>. [PubMed: 23206590]
- Sargent JQ, Zacks JM, Hambrick DZ, Zacks RT, Kurby CA, Bailey HR, Beck TM. Event segmentation ability uniquely predicts event memory. *Cognition*. 2013; 129(2):241–255. <http://doi.org/10.1016/j.cognition.2013.07.002>. [PubMed: 23942350]
- Schapiro AC, Rogers TT, Cordova NI, Turk-Browne NB, Botvinick MM. Neural representations of events arise from temporal community structure. *Nature Neuroscience*. 2013; 16(4):486–492. <http://doi.org/10.1038/nn.3331>. [PubMed: 23416451]
- Shirer WR, Ryali S, Rykhlevskaia E, Menon V, Greicius MD. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral Cortex*. 2012; 22(1):158–165. <http://doi.org/10.1093/cercor/bhr099>. [PubMed: 21616982]
- Silbert LJ, Honey CJ, Simony E, Poeppel D, Hasson U. Coupled neural systems underlie the production and comprehension of naturalistic narrative speech. *Proceedings of the National Academy of Sciences*. 2014; 111(43):E4687–E4696. <http://doi.org/10.1073/pnas.1323812111>.

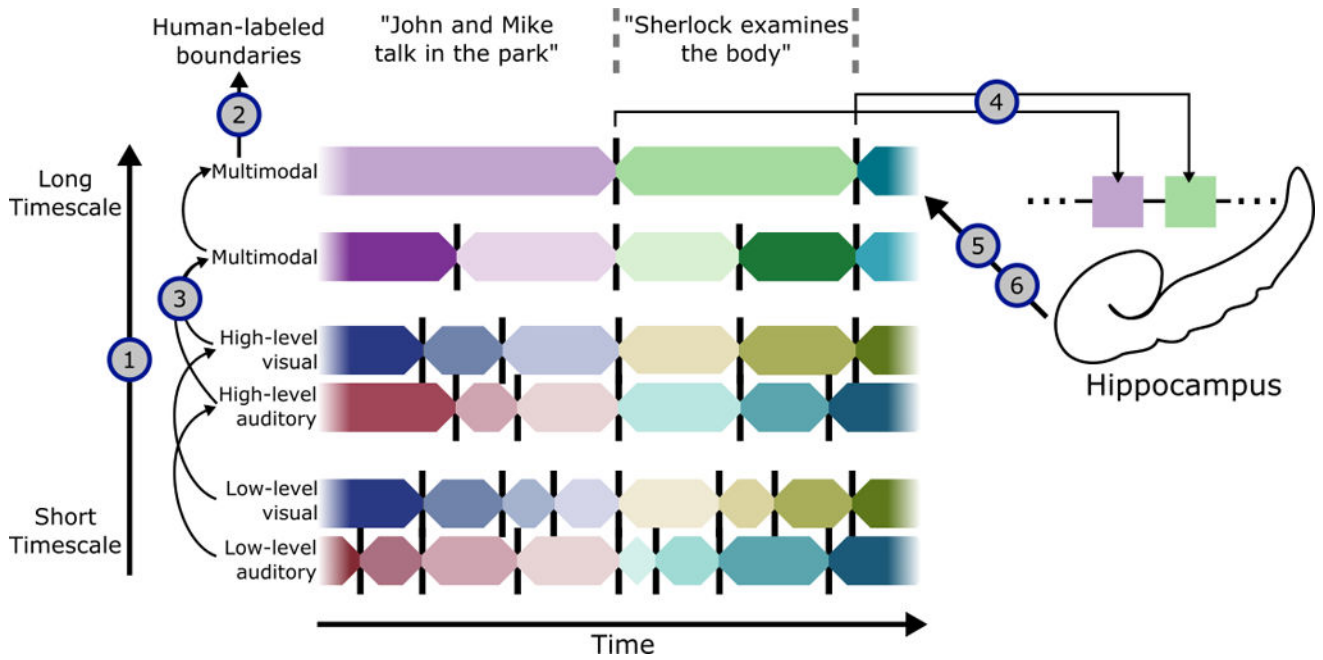
- Simony E, Honey CJ, Chen J, Lositsky O, Yeshurun Y, Wiesel A, Hasson U. Dynamical reconfiguration of the default mode network during narrative comprehension. *Nature Communications*. 2016; 7(May 2015):1–13. <http://doi.org/10.1038/ncomms12141>.
- Speer NK, Zacks JM, Reynolds JR. Human Brain Activity Time- Locked to Narrative Event Boundaries. *Psychological Science*. 2007; 18(5):449–455. <http://doi.org/10.1111/j.1467-9280.2007.01920.x>. [PubMed: 17576286]
- Stephens GJ, Honey CJ, Hasson U. A place for time: the spatiotemporal structure of neural dynamics during natural audition. *The Journal of Neuroscience*. 2013; 110(9):2019–2026. <http://doi.org/10.1152/jn.00268.2013>.
- Swallow KM, Barch DM, Head D, Maley CJ, Holder D, Zacks JM. Changes in Events Alter How People Remember Recent Information. *Journal of Cognitive Neuroscience*. 2011; 23(5):1052–1064. <http://doi.org/10.1162/jocn.2010.21524>. [PubMed: 20521850]
- Takeuchi T, Duzkiewicz AJ, Sonneborn A, Spooner PA, Yamasaki M, Watanabe M, Morris RGM. Locus coeruleus and dopaminergic consolidation of everyday memory. *Nature*. 2016; 537(7620):357–362. <http://doi.org/10.1038/nature19325>. [PubMed: 27602521]
- Van Dijk, TA., Kintsch, W. *Strategies of discourse comprehension*. New York: Academic Press; 1983.
- Vinogradova OS. Hippocampus as comparator: Role of the two input and two output systems of the hippocampus in selection and registration of information. *Hippocampus*. 2001; 11(5):578–598. <http://doi.org/10.1002/hipo.1073>. [PubMed: 11732710]
- Vodrahalli, K., Chen, P-H., Liang, Y., Baldassano, C., Chen, J., Yong, E., Arora, S. Mapping Between fMRI Responses to Movies and their Natural Language Annotations. arXiv. 2016. Retrieved from <http://arxiv.org/abs/1610.03914>
- Wang, L., Mruczek, REB., Arcaro, MJ., Kastner, S. Probabilistic Maps of Visual Topography in Human Cortex; *Cerebral Cortex*. 2014. p. bhu277 <http://doi.org/10.1093/cercor/bhu277>
- Whitney C, Huber W, Klann J, Weis S, Krach S, Kircher T. Neural correlates of narrative shifts during auditory story comprehension. *NeuroImage*. 2009; 47(1):360–366. <http://doi.org/10.1016/j.neuroimage.2009.04.037>. [PubMed: 19376237]
- Wing EA, Ritchey M, Cabeza R. Reinstatement of Individual Past Events Revealed by the Similarity of Distributed Activation Patterns during Encoding and Retrieval. *Journal of Cognitive Neuroscience*. 2015; 27(4):679–691. [http://doi.org/10.1162/jocn\\_a\\_00740](http://doi.org/10.1162/jocn_a_00740). [PubMed: 25313659]
- Zacks JM, Braver TS, Sheridan MA, Donaldson DI, Snyder AZ, Ollinger JM, Raichle ME. Human brain activity time-locked to perceptual event boundaries. *Nature Neuroscience*. 2001; 4(6):651–655. <http://doi.org/10.1038/88486>. [PubMed: 11369948]
- Zacks JM, Kurby CA, Eisenberg ML, Haroutunian N. Prediction Error Associated with the Perceptual Segmentation of Naturalistic Events. *Journal of Cognitive Neuroscience*. 2011; 23(12):4057–4066. [http://doi.org/10.1162/jocn\\_a\\_00078](http://doi.org/10.1162/jocn_a_00078). [PubMed: 21671745]
- Zacks JM, Speer NK, Swallow KM, Braver TS, Reynolds JR. Event perception: a mind-brain perspective. *Psychological Bulletin*. 2007; 133(2):273–293. <http://doi.org/10.1037/0033-2909.133.2.273>. [PubMed: 17338600]
- Zacks JM, Speer NK, Swallow KM, Maley CJ. The Brain's Cutting-Room Floor: Segmentation of Narrative Cinema. *Frontiers in Human Neuroscience*. 2010; 4(October):1–15. <http://doi.org/10.3389/fnhum.2010.00168>. [PubMed: 20204154]
- Zacks JM, Tversky B, Iyer G. Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology General*. 2001; 130(1):29–58. <http://doi.org/10.1037/0096-3445.130.1.29>. [PubMed: 11293458]
- Zadbood, A., Chen, J., Leong, YC., Norman, KA., Hasson, U. How we transmit memories to other brains: constructing shared neural representations via communication. bioRxiv. 2016. <http://doi.org/10.1101/081208>
- Zalla T, Pradat-Diehl P, Sirigu A. Perception of action boundaries in patients with frontal lobe damage. *Neuropsychologia*. 2003; 41(12):1619–1627. [http://doi.org/10.1016/S0028-3932\(0300098-8\)](http://doi.org/10.1016/S0028-3932(0300098-8)). [PubMed: 12887987]
- Zalla T, Verlut I, Franck N, Puzenat D, Sirigu A. Perception of dynamic action in patients with schizophrenia. *Psychiatry Research*. 2004; 128(1):39–51. <http://doi.org/10.1016/j.psychres.2003.12.026>. [PubMed: 15450913]



- Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., Fidler, S. Aligning Books and Movies: Towards Story-like Visual Explanations by Watching Movies and Reading Books; International Conference on Computer Vision; 2015. p. 19-27.
- Zwaan R, Langston M, Graesser A. The Construction of Situation Models in Narrative Comprehension: An Event-Indexing Model. *Psychological Science*. 1995; 6(5):292–297. <http://doi.org/10.1111/j.1467-9280.1995.tb00513.x>.
- Zwaan R, Radvansky G. Situation models in language comprehension and memory. *Psychological Bulletin*. 1998; 123(2):162–185. <http://doi.org/10.1037/0033-2909.123.2.162>. [PubMed: 9522683]

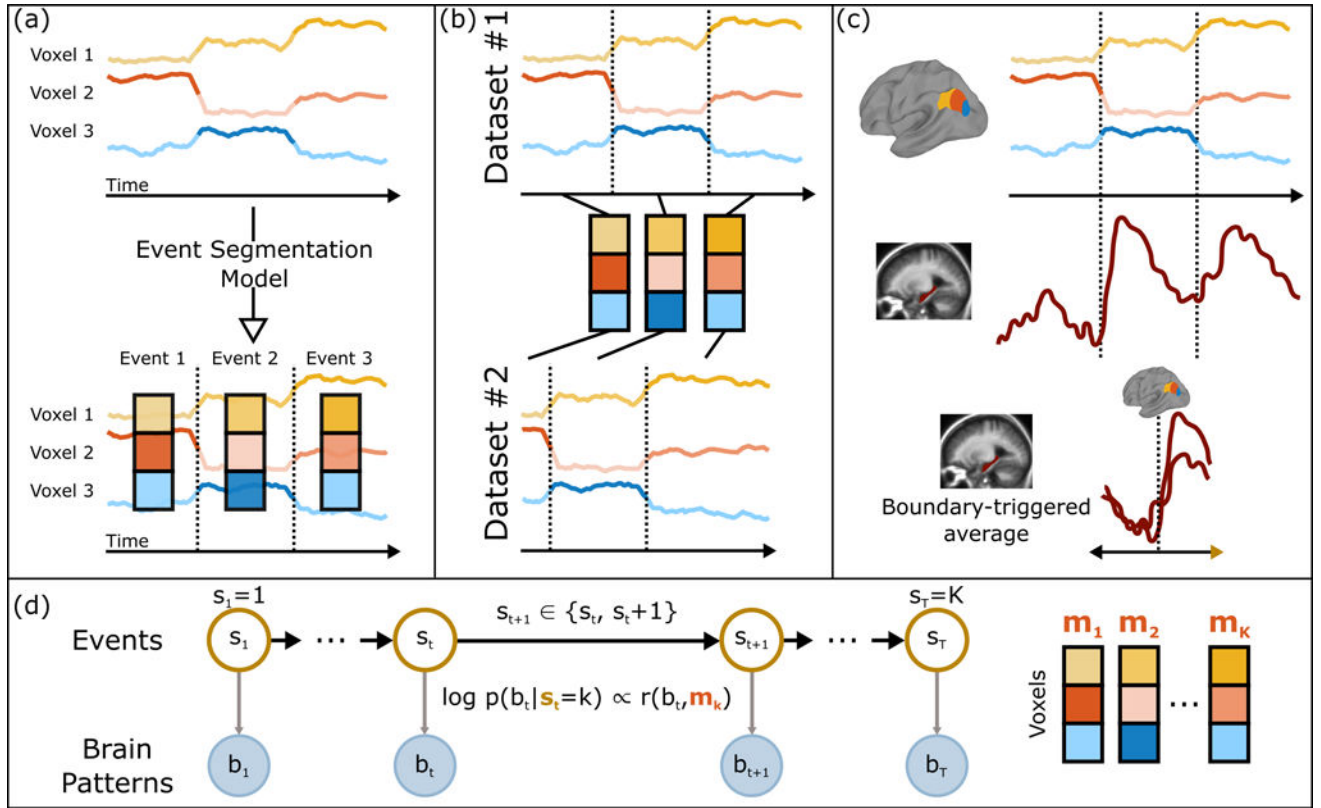
**Highlights**

- Event boundaries during perception can be identified from cortical activity patterns
- Event timescales vary from seconds to minutes across the cortical hierarchy
- Hippocampal activity following an event predicts reactivation during recall
- Prior knowledge of a narrative enables anticipatory reinstatement of event patterns



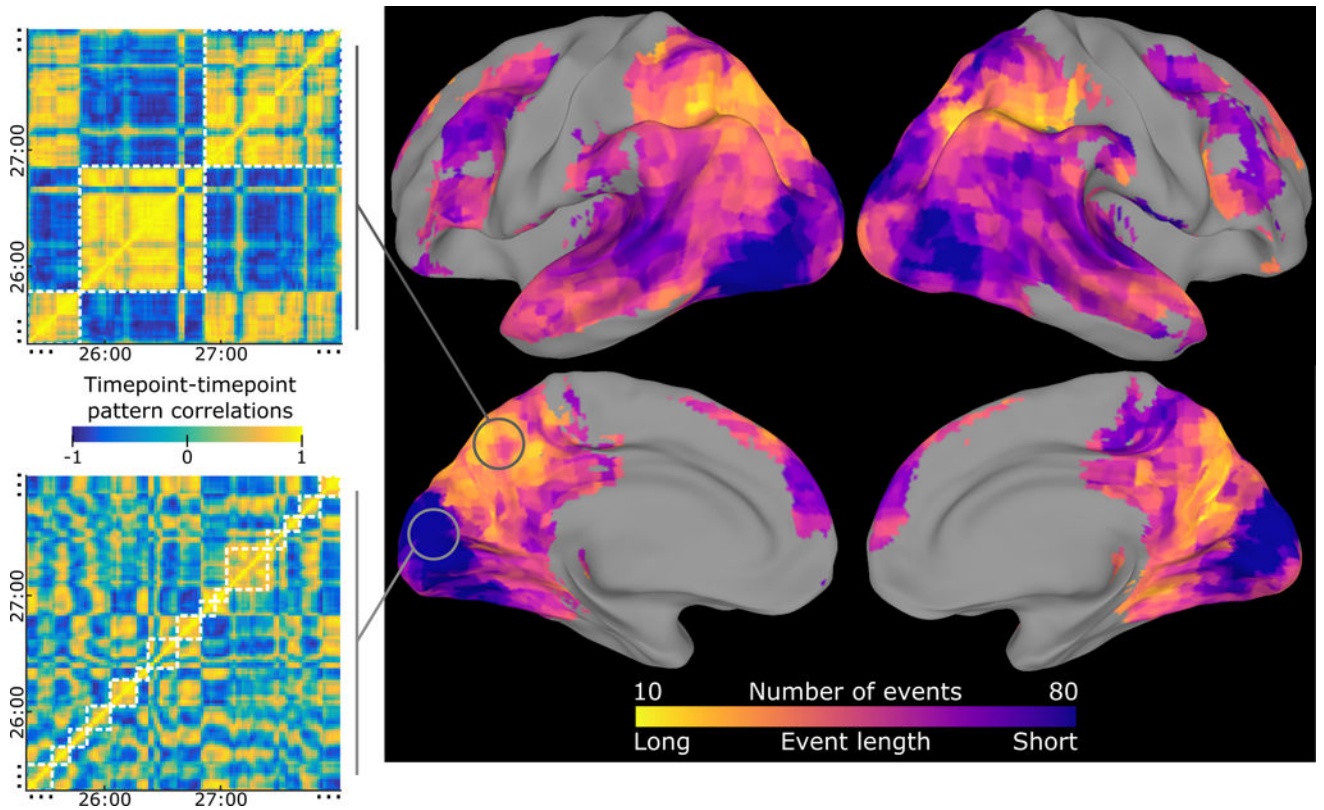
**Figure 1. Theory of event segmentation and memory**

During perception, events are constructed at a hierarchy of timescales (1), with short events in primary sensory regions and long events in regions including the angular gyrus and posterior medial cortex. These high-level regions have event boundaries that correspond most closely to putative boundaries identified by human observers (2), and represent abstract narrative content that can be drawn from multiple input modalities (3). At the end of a high-level event, the situation model is stored into long-term memory (4) (resulting in post-boundary encoding activity in the hippocampus), and can be reinstated during recall back into these cortical regions (5). Prior event memories can also influence ongoing processing (6), facilitating prediction of upcoming events in related narratives. We test each of these hypotheses using a data-driven event segmentation model, which can automatically identify transitions in brain activity patterns and detect correspondences in activity patterns across datasets.



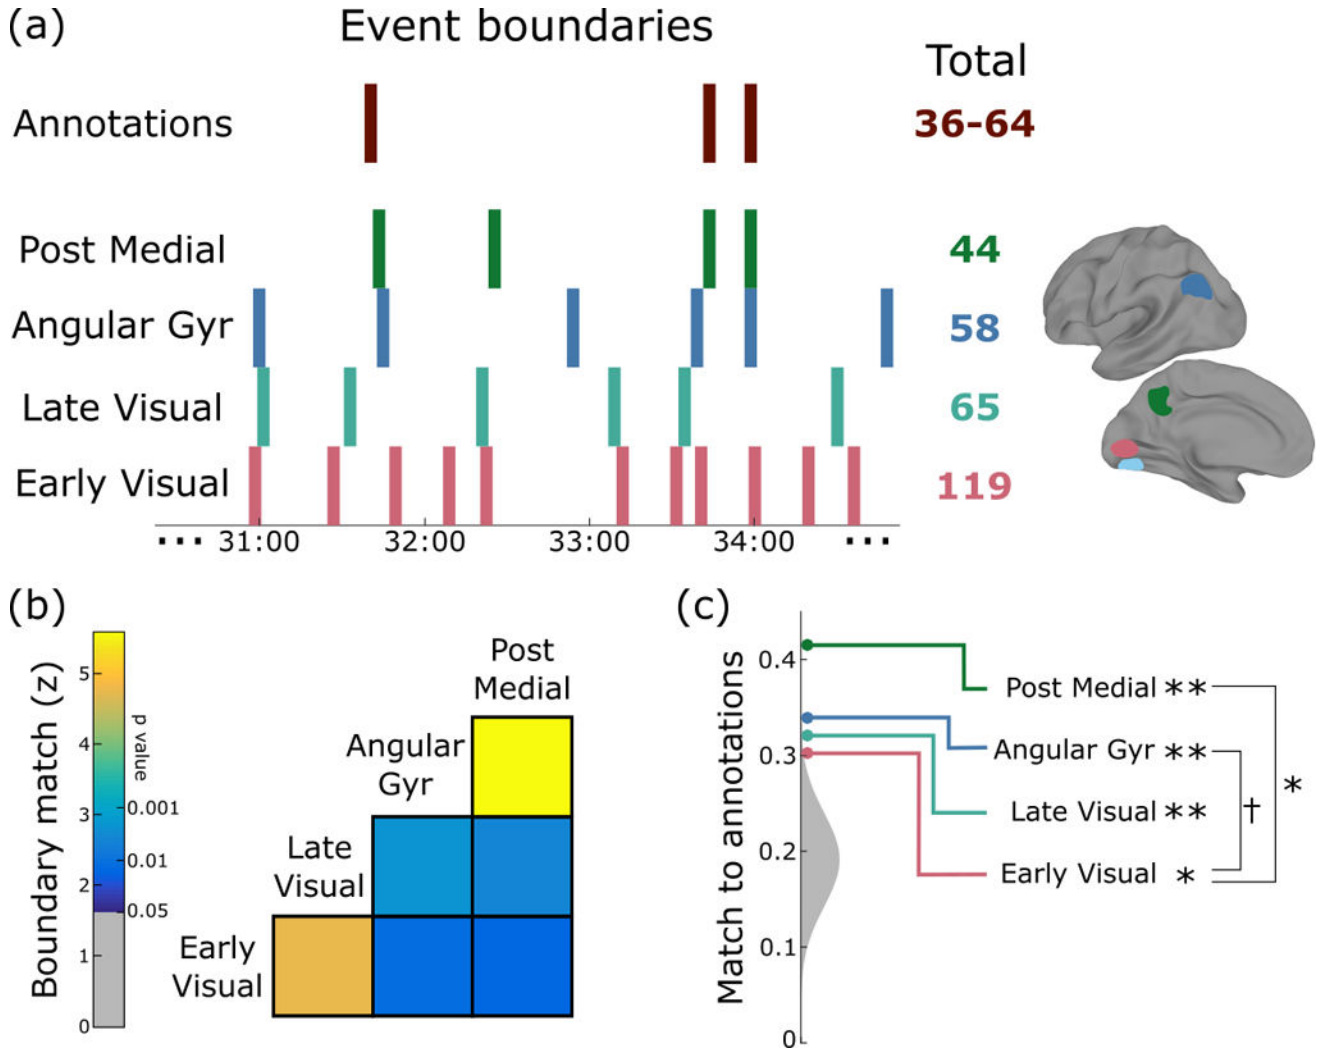
**Figure 2. Event segmentation model**

(a) Given a set of (unlabeled) timecourses from a region of interest, the goal of the event segmentation model is to temporally divide the data into “events” with stable activity patterns, punctuated by “event boundaries” at which activity patterns rapidly transition to a new stable pattern. The number and locations of these event boundaries can then be compared across brain regions or to stimulus annotations. (b) The model can identify event correspondences between datasets (e.g. responses to movie and audio versions of the same narrative) that share the same sequence of event activity patterns, even if the duration of the events is different. (c) The model-identified boundaries can also be used to study processing evoked by event transitions, such as changes in hippocampal activity coupled to event transitions in the cortex. (d) The event segmentation model is implemented as a modified Hidden Markov Model (HMM) in which the latent state  $s_t$  for each timepoint denotes the event to which that timepoint belongs, starting in event 1 and ending in event K. All datapoints during event k are assumed to exhibit high similarity with an event-specific pattern  $m_k$ . See also Figs. S1, S2, S3.



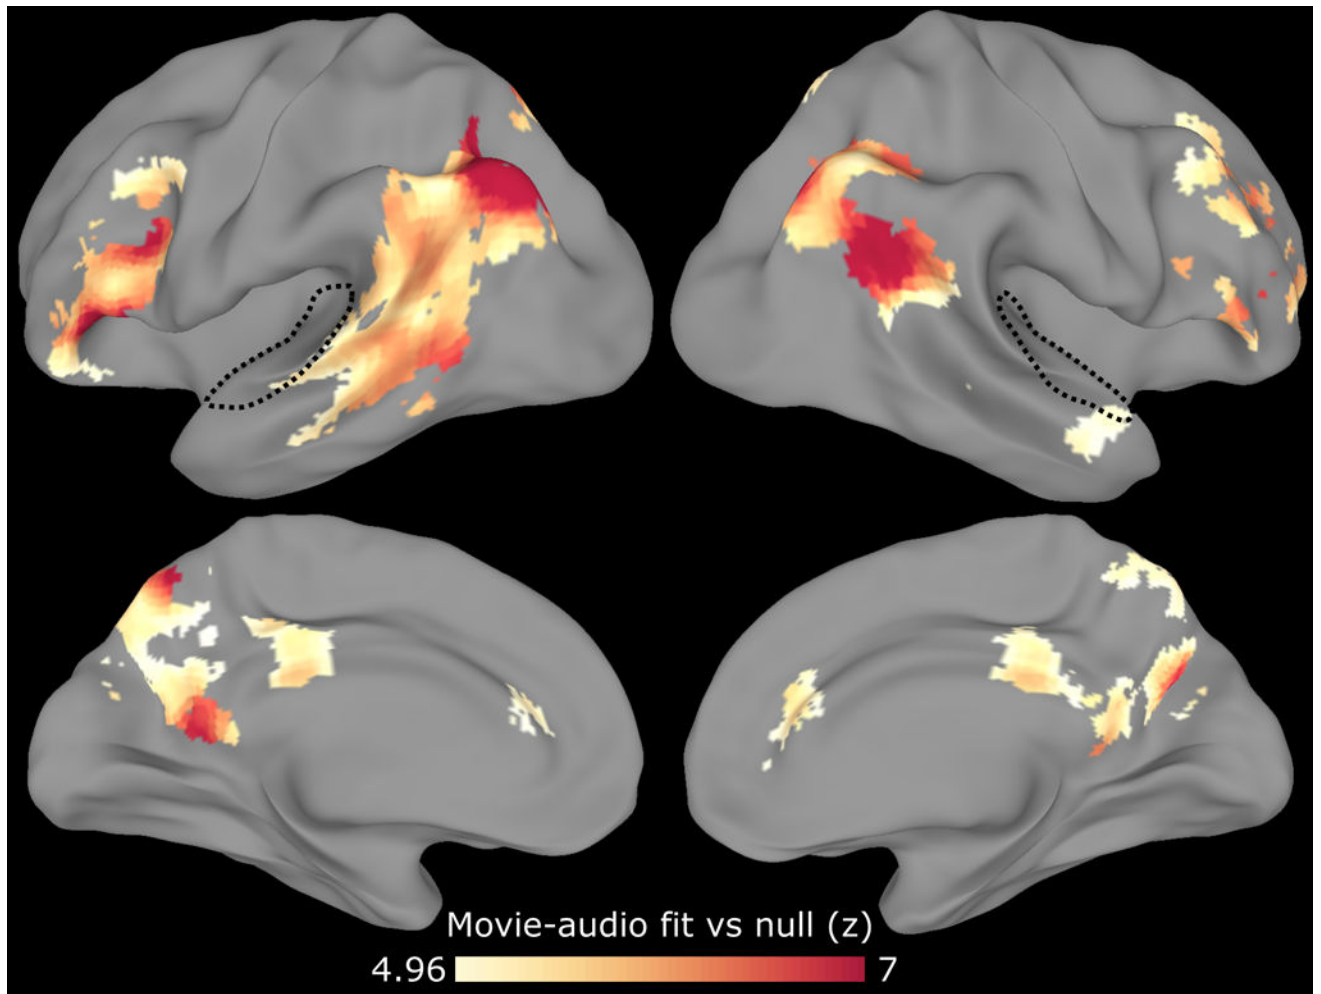
**Figure 3. Event segmentation model for movie-watching data reveals event timescales**

The event segmentation model identifies temporally-clustered structure in movie-watching data throughout all regions of cortex with high intersubject correlation. The optimal number of events varied by an order of magnitude across different regions, with a large number of short events in sensory cortex and a small number of long events in high-level cortex. For example, the timepoint correlation matrix for a region in the precuneus exhibited coarse blocks of correlated patterns, leading to model fits with a small number of events (white squares), while a region in visual cortex was best modeled with a larger number of short events (note that only ~3 minutes of the 50 minute stimulus are shown, and that the highlighted searchlights were selected post-hoc for illustration). The searchlight was masked to include only regions with intersubject correlation  $> 0.25$ , and voxelwise thresholded for greater within- than across-event similarity,  $q < 0.001$ . See also Figs. S3, S4, S8.



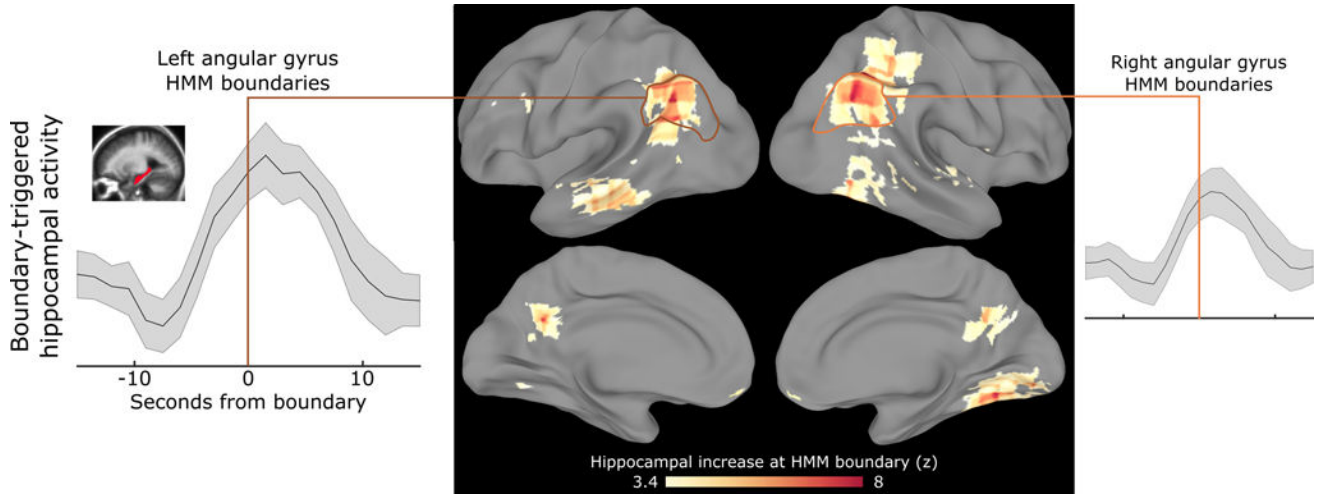
**Figure 4. Cortical event boundaries are hierarchically structured and are related to human-labeled event boundaries, especially in posterior medial cortex**

(a) An example of boundaries evoked by the movie over a four-minute period shows how the number of boundaries decreases as we proceed up the hierarchy, with boundaries in posterior medial cortex most closely related to human annotations of event transitions. (b) Event boundaries in higher regions are present in lower regions at above-chance levels (especially pairs of regions that are close in the hierarchy), suggesting that event segmentation is in part hierarchical, with lower regions subdividing events in higher regions. (c) All four levels of the hierarchy show an above-chance match to human annotations (the null distribution is shown in gray), but the match increases significantly from lower to higher levels ( $\dagger$   $p=0.058$ , \*  $p<0.05$ , \*\*  $p<0.01$ ).



**Figure 5. Movie-watching model generalizes to audio narration in high-level cortex**

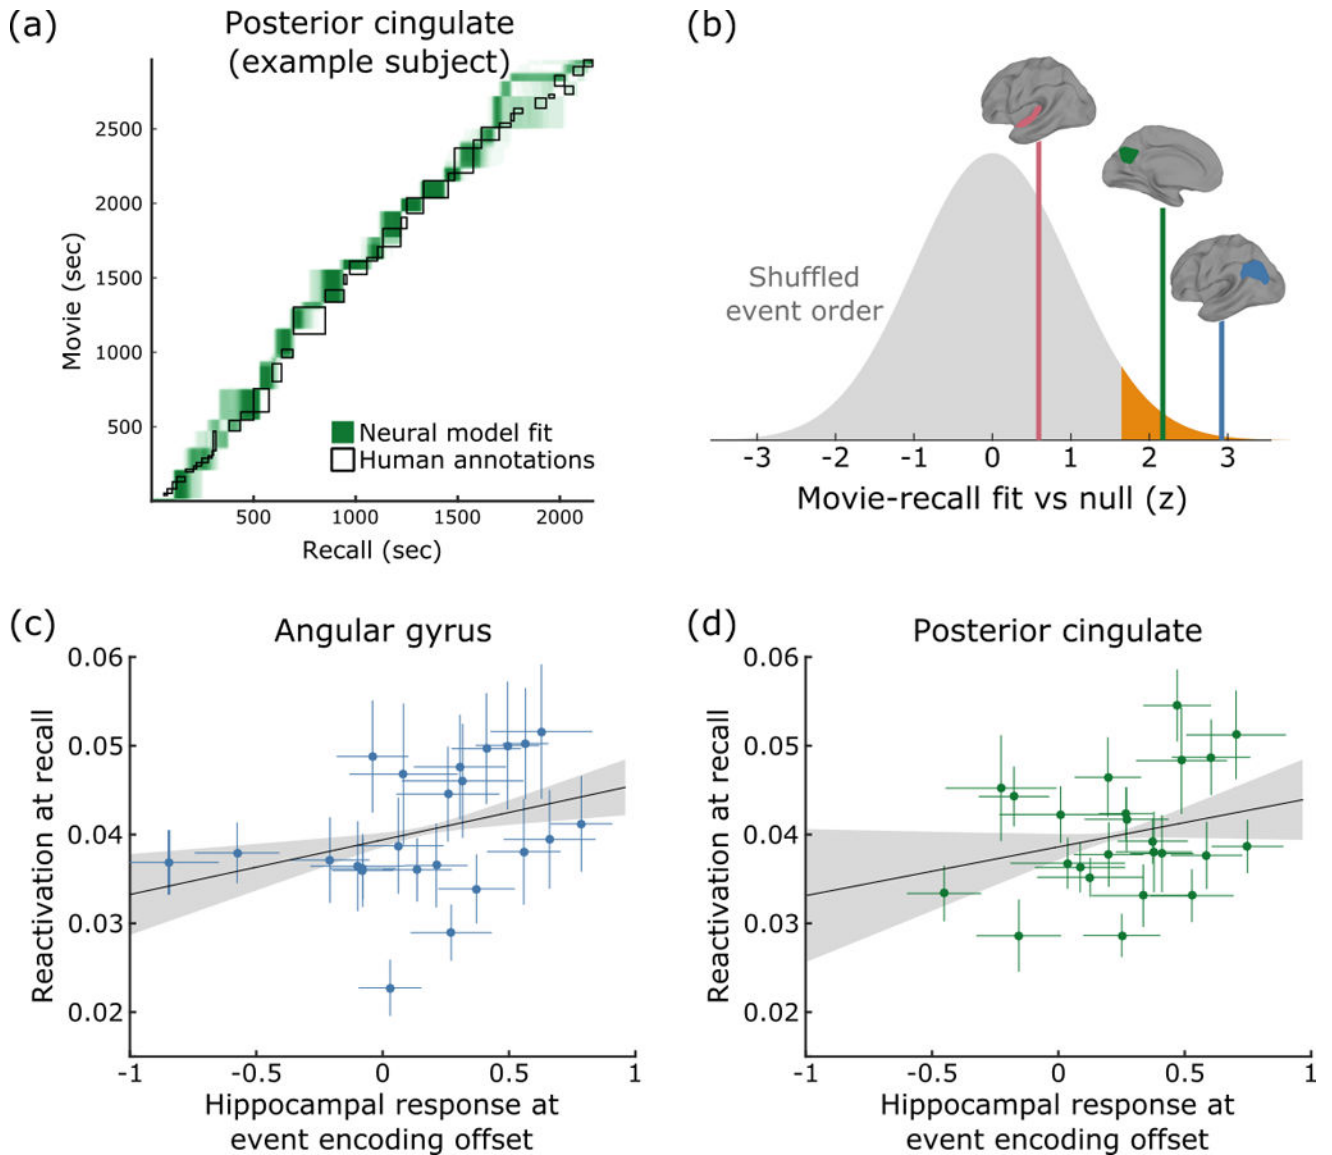
After identifying a series of event patterns in a group of subjects who watched a movie, we tested whether this same series of events occurred in a separate group of subjects who heard an audio narration of the same story. The movie and audio stimuli were not synchronized and differed in their duration. We restricted our searchlight to voxels that responded to both the movie and audio stimuli (having high intersubject correlation within each group). Movie-watching event patterns in early auditory cortex (dotted line) did not generalize to the activity evoked by audio narration, while regions including the angular gyrus, temporoparietal junction, posterior medial cortex, and inferior frontal cortex exhibited shared event structure across the two stimulus modalities. This analysis, conducted using our data driven model, replicates and extends the previous analysis of this dataset (Zadbood et al., 2016) in which the event correspondence between the movie and audio narration was specified by hand. The searchlight is masked to include only regions with intersubject correlation  $> 0.1$  in all conditions, and voxelwise thresholded for above-chance movie-audio fit,  $q < 10^{-5}$ . See also Fig. S8.



**Figure 6. Hippocampal activity increases at cortically-defined event boundaries**

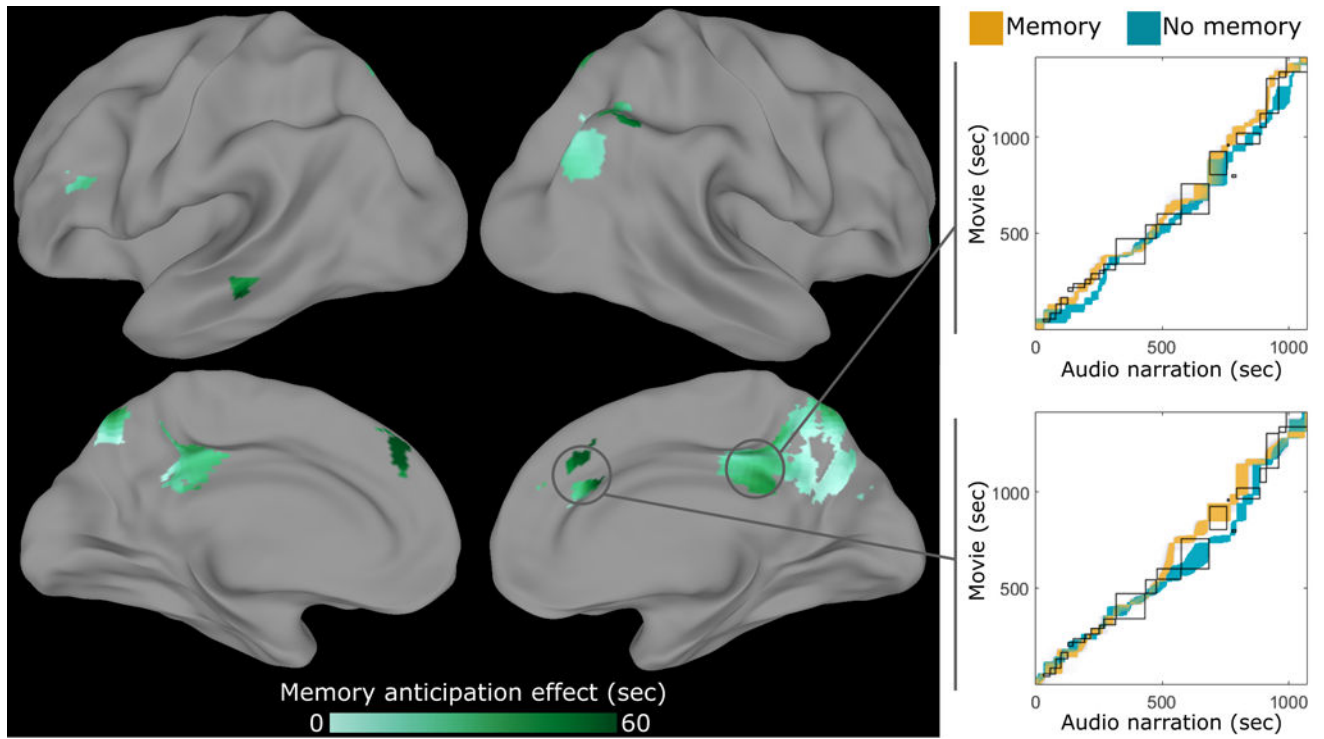
To determine whether event boundaries may be related to long-term memory encoding, we identify event boundaries based on a *cortical* region and then measure *hippocampal* activity around those boundaries. In a set of regions including angular gyrus, posterior medial cortex, and parahippocampal cortex, we find that event boundaries robustly predict increases in hippocampal activity, which tends to peak just after the event boundary. The searchlight is masked to include only regions with intersubject correlation  $> 0.25$ , and voxelwise thresholded for post-boundary hippocampal activity greater than pre-boundary activity,  $q < 0.001$ . See also Figs. S5, S8.





**Figure 7. Movie-watching events are reactivated during individual free recall, and reactivation is related to hippocampal activation at encoding event boundaries**

(a) We can obtain an estimated correspondence between movie-watching data and free-recall data in individual subjects by identifying a shared sequence of event patterns, shown here for an example subject using data from posterior cingulate cortex. (b) For each region of interest, we tested whether the movie and recall data shared an ordered sequence of latent events (relative to a null model in which the order of events was shuffled between movie and recall). We found that both angular gyrus (blue) and posterior cingulate cortex (green) showed significant reactivation of event patterns, while early auditory cortex (red) did not. (c–d) Events whose offset drove a strong hippocampal response during encoding (movie-watching) were strongly reactivated for longer fractions of the recall period, both in the angular gyrus and the posterior cingulate. Error bars for event points denote s.e.m. across subjects, and error bars on the best-fit line indicate 95% confidence intervals from bootstrapped best-fit lines. See also Fig. S6.



**Figure 8. Prior memory shifts movie-audio correspondence**

The event segmentation model was fit simultaneously to a data from a group watching the movie, the same group listening to the audio narration after having seen the movie (“memory”), and a separate group listening to the audio narration for the first time (“no memory”). By examining which timepoints were estimated to fall within the same latent event, we obtained a correspondence between timepoints in the audio data (for both groups) and timepoints in the movie data. We found regions in which the correspondence in both groups was close to the human-labeled correspondence between the movie and audio stimuli (black boxes), but the memory correspondence (orange) significantly led the non-memory correspondence (blue) (indicated by an upward shift on the correspondence plots; note that the highlighted searchlights were selected post-hoc for illustration). This suggests that cortical regions of the memory group were anticipating events in the narration based on knowledge of the movie. The searchlight is masked to include only regions with intersubject correlation  $> 0.1$  in all conditions, and voxelwise thresholded for above-chance differences between memory and no memory groups,  $q < 0.05$ . See also Fig. S7, S8.