Repeated Viewing of a Narrative Movie Changes Event Timescales in The Brain

Narjes Al-Zahli[1], Mariam Aly[1]*, Christopher Baldassano[2]*

[1]Department of Psychology, University of California, Berkeley; Berkeley, CA, USA
[2]Department of Psychology, Columbia University; New York, NY, USA

*These authors contributed equally to this work

Please address correspondence to:
Narjes Al-Zahli <na2852@columbia.edu>

ORCiD:
Narjes Al-Zahli: https://orcid.org/0009-0008-0997-3251
Mariam Aly: https://orcid.org/0000-0003-4033-6134
Christopher Baldassano: https://orcid.org/0000-0003-3540-5019

Acknowledgements:
This work was funded by a National Institutes of Health Research Project Grant (R01EY034436-01) to M.A. and C.B. We would like to thank the Alyssano Group for helpful advice on this project.

Abstract

Many experiences occur repeatedly throughout our lives: we might watch the same movie more than once and listen to the same song on repeat. How does the brain modify its representations of events when experiences are repeated? We hypothesized that, with repeated viewing of a narrative movie, brain regions would adapt their event representations by becoming either finer (more detailed) or coarser (more generalized). To test this hypothesis, we analyzed data from 30 human participants who underwent functional magnetic resonance imaging (fMRI) while watching three 90-second clips from "The Grand Budapest Hotel" six times each. We used hidden Markov models and pattern similarity analysis applied to searchlights across the brain to quantify the strength of event structure at different timescales for each clip presentation. We then tested how event structure strength changed at both slow and fast timescales with repeated viewings. Most brain regions exhibited stability in the strength of event structure at both slow and fast timescales. Other regions, however, showed flexible event representations that became more or less granular across repeated clip presentations. Notably, several brain regions exhibited consistent changes in the strength of event structure at a slow timescale across different movie clips. Furthermore, in lateral occipital cortex and middle temporal gyrus, greater loss of event structure at a slow timescale predicted more detailed memory recall. These results highlight that event dynamics in the brain are not fixed, but can change flexibly with experience.

Significance Statement

Many experiences are not unique—they repeat multiple times. We asked how increasing familiarity with an experience changes the brain's response to it. Individuals repeatedly watched narrative movie clips while undergoing fMRI. This allowed us to examine how the brain's temporal representations of events in the movie changed when the movie was viewed multiple times. As the movie clip became familiar, some brain regions fine-tuned their event representations, dividing the movie clip into smaller events. Other brain regions showed coarser event representations with increasing experience by grouping previously distinct events. Finer event representations in the brain predicted more detailed memory. These results show that the brain flexibly changes how it represents events, and this flexibility may help memory.

Introduction

Experience unfolds over multiple timescales—from the millisecond flux of sensory signals to extended narratives that span minutes. Converging neuroimaging, electrophysiological, and computational work shows that the brain mirrors this temporal structure by processing information hierarchically across multiple timescales. Lower-level sensory cortex operates on short timescales, rapidly tracking moment-to-moment changes in sensory input, whereas higher-order association areas accumulate information over progressively longer periods. This cortical gradient of "temporal receptive windows" (Hasson et al., 2008; Honey et al., 2012; Murray et al., 2014; Geerligs et al., 2022) supports the segmentation of continuous experience into discrete events at multiple timescales, which in turn scaffolds situation models and episodic memories that weave discrete moments into meaningful narratives (Zacks et al., 2001; Baldassano et al., 2017).

While research has established that different brain regions integrate information over different timescales, it remains an open question how fixed or flexible these timescales are. Some research suggests that a given region's temporal integration window may be determined by local circuit properties and the region's position in the cortical hierarchy (Chaudhuri et al., 2015; Hasson et al., 2008), indicating that a brain region's timescale of processing may be relatively fixed. However, emerging evidence points to the malleability of timescales: the brain can adapt its temporal processing to stimulus properties, context, goals, and expectations (Lerner et al., 2014; Baumgarten et al., 2021; Shin & DuBrow, 2021; De Soares et al., 2024; Çatal et al., 2024; Lee et al., 2021). Further, studies of memory have shown that as content becomes familiar, the brain uses prior knowledge to anticipate upcoming information (Günseli & Aly, 2020; Tarder Stoll et al., 2024), potentially shifting the timing of event encoding (Shin & DuBrow, 2021; Lee et al., 2021; De Soares et al., 2024). Inspired by these findings, we hypothesized that repeated exposure to an event may change neural representations of event structure at multiple timescales. Repeated experience may allow the brain to group sub-events that were previously perceived to be separate, or fine-tune larger event structures into subcomponents. Such a change in neural event structure may be detectable as speeding or slowing of the brain's event representations with repeated experience, even if the external stimulus remains unchanged.

Here, we examined such experience-driven shifts in event timescales in the brain. Characterizing event timescales can be complex because even within a single brain region, multiple timescales of activity can be hierarchically nested (shorter events embedded within longer events) (Stephens et al., 2013; Kanter et al., 2025). To detect event structure at multiple concurrent timescales within individual brain regions, we developed a method that first segments continuous neural activity into discrete events at multiple timescales, and then quantifies the strength of event structure at each timescale. We applied this method to fMRI data from Aly et al. (2018), in which participants repeatedly viewed the same movie clips. The clips varied in narrative coherence (one with a continuous storyline; two with scrambled scene order), allowing us to examine how neural event structure changes across different story structures. We used a Hidden Markov Model-based

event segmentation approach ([Baldassano et al., 2017](#)) to identify event boundaries in each brain region across multiple timescales, and used pattern similarity analyses to examine the strength of event representations within and across event boundaries. This approach, applied to each movie viewing separately, allowed us to: (i) determine how event structure at slow and fast timescales changed with experience; (ii) assess which changes in event structure were specific to a particular movie clip vs. generalized across clips with different content and temporal structure; and (iii) relate these neural changes to subsequent memory. Together, this allowed us to test flexibility vs. stability in the timescales of event processing in the brain, and the relevance of these changes for memory.

Figure 1. Stimuli and Experimental Design. Participants watched three 90-second clips from *The Grand Budapest Hotel*. Each clip was watched six times. Each clip was segmented into smaller continuous chunks, 2.4-5.5s long. The numbers in each clip refer to a segment within each clip. Each clip was assigned to one of three conditions. In the *Intact* condition, the clip segments were shown in

their original order. In the *Scrambled-Fixed* condition, the clip segments were randomly shuffled, but shown in the same scrambled sequence for all clip presentations. In the *Scrambled-Random* condition, the clip segments were randomly shuffled and shown in a different order for each clip presentation.

Materials and Methods

Experimental Design

Grand Budapest Hotel dataset

We used data collected by Aly et al. (2018). Data and stimuli are publicly available on OpenNeuro: https://openneuro.org/datasets/ds001545/versions/1.1.1.

The dataset included thirty right-handed participants (12 men, average age = 23.0 years, SD = 4.2; average education = 15.3 years, SD = 3.2). During functional magnetic resonance imaging (fMRI), these participants watched three 90-second clips from *The Grand Budapest Hotel* (**Fig. 1**). Each clip was shown six times. The clips were pseudo-randomly interleaved, such that all three were shown one time each before the next set of presentations began. None of the participants had seen the movie before. After viewing the three clips six times each, participants were asked to recall the clips. They were given a blank text document and instructed to type their memories for the three clips in any order they chose.

The study included three different types of clips. Each 90-second clip was divided into smaller segments, 2.4-5.5 seconds in length. The first clip, *Intact*, was shown in its original narratively coherent order, i.e., the segments were viewed in the same order as the original movie. The second clip, *Scrambled-Fixed*, featured randomly ordered segments, but the order was consistent across viewings. This resulted in a stable but incoherent structure. The last clip, *Scrambled-Random*, featured segments that were reordered differently each time, creating an incoherent and unstable structure.

Three 90-second movie scenes were selected: an interview scene, a painting theft scene, and a chase scene. There were two counterbalancing groups, in which the assignment of movie scene to condition (*Intact, Scrambled-Fixed, Scrambled-Random*) was partially counterbalanced. For both groups, the *Intact* clip was the interview scene. The *Scrambled-Fixed* clip was the painting theft scene for one group and the chase scene for the other group, and vice versa for the *Scrambled-Random* clip.

fMRI Methods

fMRI Acquisition

Full details of the MRI acquisition and preprocessing procedures are reported in Aly et al. (2018). Functional MRI data were acquired using a 3-Tesla Siemens Prisma scanner equipped with a 64-channel head/neck coil. Functional images were collected using a multiband echo-planar imaging (EPI) sequence with the following parameters: repetition time (TR) = 1.5 s, echo time (TE) = 39 ms, flip angle = 50°, multiband acceleration factor = 4, and slice shift = 3. Each volume consisted of 60 oblique axial slices acquired in an interleaved order, with an isotropic voxel size of 2.0 mm. High-resolution T1-weighted structural images were obtained with an MPRAGE sequence (1.0 mm isotropic resolution) for anatomical reference and registration. Additionally, field maps were collected for registration (40 oblique axial slices, 3.0 mm isotropic resolution). The fMRI scanning session consisted of three experimental runs, during which participants viewed clips from the Intact, SFix, and SRnd conditions.

fMRI Preprocessing

Data preprocessing and registration were performed using tools from FEAT, FLIRT, and command-line functions in FSL (http://fsl.fmrib.ox.ac.uk/fsl/). The first three EPI volumes of each run were discarded to allow for T1 signal equilibration. Preprocessing steps included brain extraction, motion correction, high-pass temporal filtering (cutoff period = 140 s), and spatial smoothing (3-mm FWHM Gaussian kernel). Field maps were processed according to the FSL FUGUE user guide; magnitude images were averaged and skull-stripped, while phase images were converted to rad/s and smoothed (with a 2-mm Gaussian kernel). The resulting phase and magnitude images were used to unwrap the functional images in the preprocessing step of FEAT analysis to reduce distortions and aid registration to anatomical space. Functional data were then registered to each participant's anatomical scan, followed by normalization to standard Montreal Neurological Institute (MNI) space. After preprocessing, functional time series for each run were divided into volumes according to individual movie clip viewings, with the initial four volumes of each clip viewing removed to account for hemodynamic lag. Finally, the timeseries for each clip was z-scored within each participant to ensure comparability across runs and participants.

fMRI Visualization

Functional data were projected from volume space to cortical surface space using Nilearn for visualization purposes only; all analyses were conducted in volume space using custom Python scripts and libraries.

Statistical Analysis

Scoring of event memory

Free recall data were scored by counting the number of details reported by each participant for each movie clip. Details included characters, dialogue, actions, and perceptual details such as the

appearance of characters and the spaces through which they moved. Reported details were checked against the movie clips for veracity. Points were only given for accurate details.

Filtering for Regions with Meaningful Timescales Across Clip Viewings

Event Segmentation Model

We used the event segmentation Hidden Markov Model (HMM) described by Baldassano et al., 2017 to identify event boundaries for each searchlight in the brain (**Fig. 2**). The HMM operates under the following assumptions: (1) a brain region's response to a dynamic stimulus consists of a sequence of discrete event states, and (2) each event is represented in a brain region by a unique spatial activity pattern. We fit the HMM to each searchlight in the brain, independently for each presentation of each clip, varying the number of events from 2 to 10 to capture event structure at slower timescales (few events) and faster timescales (more events). This resulted in nine models for each of the six viewings for each of the three clips. These event numbers covered the range of plausible event timescales for meaningful events in the video (from ~8 seconds to 45 seconds). Each model provided the location of event boundaries for the specified number of events (2-10 across models) for each searchlight and each viewing of each clip. These event boundaries were used in subsequent analyses, described below. This approach deviates from prior work in that we do not summarize each brain region's activity with a single timescale, nor do we assume that the voxel pattern representing each event remains stable across movie repetitions.

Within- vs. Between-Event Similarity

To evaluate the quality of event segmentation, we computed the similarity of spatial activity patterns within versus across HMM-defined event boundaries, using only held-out (testing) data. Event boundaries were defined by fitting the HMM (described above) to the group-averaged BOLD activity timecourse from a randomly selected half of the participants and identifying the timepoints at which there was a change in the most-probable event label. These boundaries were then applied to the group-averaged BOLD activity timecourse from the held-out participants. For each clip, we generated a timepoint-by-timepoint similarity matrix showing the correlation between brain activity patterns across time. To minimize the influence of BOLD autocorrelation, we calculated correlations at a fixed lag of 5 TRs. Within-event correlations were defined as pairs of timepoints separated by this lag that fell within the same event, while between-event correlations spanned an event boundary. The Within- vs. Between-Event Similarity value was calculated as the difference between the average within-event and between-event correlations, providing a measure of how well the event structure learned from training data generalized to stable neural activity patterns in unseen data.

Figure 2. **Computing Within- vs. Between-Event Similarity for Each Clip Presentation.** We used a Hidden Markov Model (HMM) to identify event boundaries in each clip presentation and then assessed Within-Event vs. Between-Event pattern similarity in an independent group of participants. The analysis pipeline is visualized here with 3 events, but the model was fit for 2-10 events to assess event structure at multiple timescales. **(A)** The analysis is based on identifying stable patterns of activity within each searchlight in the brain, which correspond to a neural "event". For each clip presentation, a Timepoint-by-Timepoint Similarity Matrix visualizes the correlation between the spatial patterns of activity for two timepoints in the clip. Dashed white outlines indicate the event boundaries identified by the HMM when the model segments the neural data into 3 events. The strength of these model-identified event boundaries was measured by computing the difference in pattern similarity for pairs of timepoints within the same HMM-defined event (blue rectangles) versus between events (red rectangles). To minimize the effects of BOLD autocorrelation, we examine Within-Event and Between-Event correlations at a delay of 5 TRs. **(B)** To produce an unbiased estimate of Within- vs. Between-Event Similarity, data from participants were randomly split into training and testing groups. Averaged time series data from the training group were used to identify event boundaries with the HMM (here, 3 events; the process is conducted for event counts 2-10). Within- vs. Between-Event Similarity around these training-group HMM boundaries was calculated for neural patterns in the testing group and used for subsequent analyses. This approach is repeated 5 times, with different participant splits.

The training and testing sets were then swapped, and the Within- vs. Between-Event Similarity values from both directions were averaged. To avoid overfitting and ensure robustness, we averaged Within- vs. Between-Event Similarity values from five different participant splits for each model. Because the movie clips used in the Scrambled-Fixed and Scrambled-Random conditions were counterbalanced across participants, the training/testing split for these clips was done within each counterbalancing group separately, and the Within- vs. Between-Event Similarity measures were averaged across the counterbalancing groups. For each clip, this procedure

produced 54 Within- vs. Between-Event Similarity values (nine event counts for each of six clip presentations).

To validate the reliability of the Within- vs. Between-Event Similarity metric across participants, we measured the stability of this metric across subsets of the Intact dataset. Specifically, we randomly divided participants into two independent datasets. Within each of these datasets, we carried out the procedure described above: We identified HMM event boundaries in half of the participants (training group) and then used those boundaries to compute Within- vs. Between-Event Similarity in the other half of participants (testing group). As before, training and testing sets were then swapped, and results were averaged across five different participant splits. We repeated this procedure across event counts, viewings, and searchlights (limited to those showing meaningful event structure across all viewings; see *Statistical Significance Testing*, below). Finally, we examined between-group reliability in the Within– vs. Between-Event Similarity values. High agreement (mean Pearson $r = 0.48$, range = 0.36–0.59 for initial viewing; mean Pearson $r = 0.59$, range = 0.42–0.69 for subsequent viewings) between the two independently derived Within- vs. Between-Event Similarity values demonstrated that the measure is a stable measure across subsets of participants (**Supplementary Figs. 1 and 2**).

Whole-Brain Searchlight Analysis

We conducted a whole-brain analysis using spherical searchlights generated by Lee et al. (2021). The searchlights were spaced evenly throughout the MNI volume (radius = 5 voxels; stride = 5 voxels). We only analyzed searchlights containing at least 20 voxels within a standard MNI brain mask and with valid data from at least 15 participants for all viewings. The Within- vs. Between-Event Similarity value of each voxel was computed as the average Within- vs. Between-Event Similarity values of all searchlights including that voxel, separately for each of the 6 viewings for each of the 3 clips.

Statistical Significance Testing

We tested whether Within- vs. Between-Event Similarity values were statistically greater than zero using a permutation-based null hypothesis approach. Null event boundaries were generated by randomly shuffling the order of the event segment lengths from the trained HMM model, producing boundaries that had the same spacings as in the real analysis but were placed on arbitrary timepoints. Within- vs. Between-Event Similarity values for these null boundaries were then computed on the testing data. Importantly, permutations were constrained such that the shuffled sequence never matched the real sequence of event neural patterns to ensure a valid null comparison. For each event count, we conducted up to 50 permutations, calculating the Within- vs. Between Event Similarity value for each permutation. Note that for event counts below 5, there were fewer than 50 possible permutations, and so we instead used all possible shuffled permutations. We then fit a normal distribution to all of these values (across event counts) to create a single null distribution for each searchlight, clip, and viewing. The p-value was then

calculated by computing the area under the normal distribution that surpassed the real Within- vs. Between-Event Similarity for each event count.

For subsequent analyses, we only kept searchlights for which, for each clip, every viewing had at least one event count that produced a Within- vs. Between-Event Similarity value above zero and significantly greater than the null distribution (p > 0.05, one-tailed). This ensured that only searchlights with reliable event structures for all viewings proceeded, enabling subsequent identification of changes in event structure.

Identifying Changes in Event Structure

We quantified changes in event structure across repetitions of each clip at both slow (2-event) and fast (10-event) timescales. Change was defined as the difference between the average Within- vs. Between-Event Similarity across repeated viewings (viewings 2–6) vs. the first viewing, separately for 2 events and 10 events (**Fig. 3**). We focused our statistical comparisons on the first viewing (no prior familiarity with the clips) vs. subsequent viewings (varying levels of familiarity with the clips) because our primary hypotheses were about the effects of familiarity generally rather than a specific rate of change as a function of the amount of familiarity. We nevertheless visualize results for each clip viewing separately for interested readers. Likewise, we limited our statistical analyses to the slowest (2-event) and fastest (10-event) timescales that we examined, to reduce the number of statistical comparisons and convey results in a concise manner. Results are nevertheless visualized across all event counts for transparency.

To determine statistical significance for changes in event structure across viewings, we used a permutation-based null hypothesis testing method. Null datasets were created by randomly shuffling the order of six brain activity timecourses for each participant, corresponding to the six presentations of a given clip. We ran this analysis pipeline 51 times: once on the actual (unpermuted) dataset and 50 times on the null (permuted) datasets. This allowed us to examine the effect of first viewing vs. repeated viewings on the shuffled order. A two-tailed p-value was obtained by fitting a normal distribution to the null *change* in Within- vs. Between-Event Similarity values and computing the area under the normal distribution that exceeded the real *change* in Within- vs. Between-Event Similarity value for each searchlight, for each clip. Significant changes in Within- vs. Between-Event Similarity (q<0.05) in searchlights for a given clip were determined after applying the Benjamini-Hochberg FDR correction, as implemented in AFNI (Cox, 1996).

Figure 3. Computing Changes in Event Structure at Slow and Fast Timescales Across Clip Presentations. (A) Example timepoint-by-timepoint pattern similarity matrices with event structure at different timescales. The white boxes overlaid on each matrix denote event boundaries identified by the HMM analysis, capturing clusters of timepoints that are highly similar to each other. The left matrix shows a relatively coarse event segmentation, as indicated by a small number of large white boxes (events). The right matrix shows a relatively fine event segmentation, as indicated by a larger number of small white boxes (events). A brain region's event segmentation may change with clip repetition, becoming either finer or coarser. Event representations become finer if event structure is lost at slow timescales, gained at fast timescales, or both—leading to more, shorter events. Conversely, event representations become coarser if event structure is gained at slow timescales, lost at fast timescales, or both—resulting in fewer, longer events. **(B)** We fit HMMs to each presentation of each clip, measuring the strength of event structure for each viewing (within- minus between-event similarity) while varying the number of events in the models. In this example, the initial clip presentation (yellow) shows event structure at slow timescales (fewer events), whereas subsequent clip presentations (green) show event structure at faster timescales (more events). With repeated viewing, events therefore become fine-tuned in this example, both because event structure is lost at slow timescales and gained at fast timescales. **(C)** Following the same conventions of (B), in this example, the initial clip presentation (yellow) shows event structure at fast timescales (more events), whereas subsequent clip presentations (green) show event structure at slow timescales (fewer events). With repeated viewing, events therefore become coarser in this example, both because event structure is gained at slow timescales and lost at fast timescales. Subsequent analyses focus on changes in event structure (within- vs. between-event similarity) at the lowest and highest event counts (2 and 10), although data will be visualized across the full range of event lengths.

Combined Analysis of Movie Clips

We tested whether brain regions' changes in event structure across repeated viewings were consistent across different movie clips. Specifically, we sought to identify which brain regions showed a consistent change in event structure at slow (2-event) and fast (10-event) timescales

with repeated exposure to a movie clip, regardless of the specific content of that movie clip. Because each clip underwent an initial filtering step (described above) to include only searchlights with meaningful event structure across all repetitions, we first identified the set of overlapping searchlights that passed this filter for all three clips. This ensured that only overlapping searchlights with reliable event structure were included in the cross-clip analysis.

Next, within this overlapping set of searchlights, we calculated the change in the Within- vs. Between-Event Similarity value from the first viewing to the average of viewings 2–6, separately for slow (2-event) and fast (10-event) timescales. This procedure was first done for each clip separately. We then grouped searchlights based on whether they showed a consistent increase or decrease in the Within- vs. Between-Event Similarity value across all three clips—meaning the direction of change was the same (gain or loss of event structure) in every clip. For each clip and each type of change (gain or loss of event structure) at each timescale (2-event and 10-event), we performed two-tailed statistical tests and applied FDR correction at $\sqrt[3]{0.05}$. such that the conjunction of effects across all 3 clips would be FDR corrected at q < 0.05. Finally, we identified searchlights that overlapped across the FDR-corrected maps for each clip, resulting in separate maps for increases vs. decreases in the Within- vs. Between-Event Similarity value with repeated viewing, at both the slow (2-event) and fast (10-event) timescales.

Linking Changes in Event Structure to Memory

We examined whether changes in Within- vs. Between-Event Similarity predicted memory for the three clips. To do so, we focused on the searchlights from the conjunction analysis (i.e., regions that showed significant and consistent changes in event structure across all three clips). For each searchlight, we obtained each individual participant's Within vs. Between-Event Similarity value (as described below) at the event timescale that exhibited change in event structure with repeated viewing for that searchlight at the group level. We also obtained each participant's average narrative recall (across clips). Finally, we obtained the across-participant correlation between these values using a robust bootstrap linear regression model (10,000 iterations). Significant brain-behavior correlations were determined after applying multiple comparisons corrections across searchlights.

To obtain individual Within vs. Between-Event Similarity values, we modified the original pipeline. For each clip, we calculated per-participant Within- vs. Between-Event similarity values by randomly splitting the participants within each counterbalancing group into two halves. The HMM model was trained on one half and tested individually on the participants in the other half, providing Within- vs. Between-Event similarity values for all event counts for each presentation for each participant.

Code and resource availability

The Python code used to reproduce all the results in this study is available at https://github.com/narjes-alzahli/EventTimescaleChanges. The results in MNI space can be accessed at https://neurovault.org/collections/21744/.

Data availability

The analyses in this study used a publicly available dataset, accessible at OpenNeuro: https://openneuro.org/datasets/ds001545/versions/1.1.1.

Results

Changes in Event Structure in The Brain with Repeated Movie Viewing

To identify changes in event structure in the brain, we examined TR-by-TR brain activity patterns during each of the six viewings of each movie clip (**Figs. 1-3**). To assess the strength of event structure at different timescales, we computed the similarity of spatial activity patterns within versus across HMM-defined event boundaries (Within vs. Between-Event Similarity) at slow and fast timescales (2 and 10 events) across the six viewings of each clip. We first conducted method validation, confirming that our approach of measuring Within vs. Between-Event Similarity provided reliable estimates of event structure across datasets/participants (see Methods/Within- vs. Between-Event Similarity and **Supplementary Fig. 1**). Having shown that our measure is reliable, we next turned to examining our main question about *changes* in event structure. Change in event structure at a specific timescale (slow or fast) was defined as the difference between the average Within- vs. Between-Event Similarity values across viewings 2-6 and the Within- vs. Between-Event Similarity value during the first viewing.

Our analysis of the Intact clip revealed that the majority of brain regions (with reliable event structure for at least one timescale) showed no significant changes in event structure at either slow or fast timescales over movie repetitions (**Fig. 4**). Among regions that did show event structure changes across repeated viewings, the direction and magnitude of these shifts varied, such that we observed evidence for both finer event structure (loss of slow timescale structure and/or gain of fast timescale structure with repeated viewings) and coarser event structure (gain of slow timescale structure and/or loss of fast timescale structure with repeated viewings) across different brain regions. Regions showing finer event representations with repeated viewing included fusiform gyrus, which showed a loss of event structure at a slow timescale (**Fig. 4, searchlight 1**), and temporoparietal junction, which gained event structure at a fast timescale (**Fig. 4, searchlight 2**). Regions showing coarser event representations with repeated viewing included superior temporal sulcus, which gained event structure at a slow timescale and lost event structure at a fast timescale (**Fig. 4, searchlight 3**), and inferior frontal cortex, which gained event structure at a slow timescale (**Fig. 4, searchlight 4**).

Intact Clip

(A) Change at Slow Timescale

Coarser — gaining slow timescale
Finer — losing slow timescale

-0.4   0   0.4

Change in Within vs. Between-Event
Similarity from Repeated Clip Viewing

(B) Change at Fast Timescale

Coarser — losing fast timescale
Finer — gaining fast timescale

-0.4   0   0.4

Change in Within vs. Between-Event
Similarity from Repeated Clip Viewing

(C)

Events Become Finer
losing slow timescale and/or gaining fast timescale

Events Become Coarser
gaining slow timescale and/or losing fast timescale

first viewing
average of later viewings
later viewings

Figure 4. Changes in Event Structure at Slow and Fast Timescales with Repeated Viewing of the Intact Clip. Brain maps depict how repeated viewing of the Intact clip changes within- vs. between-event similarity at **(A)** a slow timescale (two events) and **(B)** a fast timescale (ten events). The color gradient from blue to red indicates the direction of change from the initial viewing to subsequent viewings: blue represents a shift toward coarser/slower events, and red indicates a shift toward finer/faster events. Only searchlights with event structure for at least one timescale for each presentation are visualized. Significant searchlights are outlined with a black border; significance was determined using a permutation test with false discovery rate (FDR) correction (q < 0.05). Overall, events became finer in lateral visual cortex, and became coarser in ventral visual and lateral temporal cortex. **(C)** Four sample searchlights (1–4) were selected post hoc for illustration. Following the conventions in Figure 3, within- vs. between-event similarity at event counts 2-10 is shown for the first viewing (yellow line), subsequent viewings (thin green lines), and the average of subsequent viewings (thick green line). Red and blue arrows highlight the direction of change in within- vs. between-event similarity from the first viewing (thick yellow line) to the average of later viewings (thick green line).

We next repeated this analysis for the Scrambled-Fixed clip. Similarly to the Intact clip, the majority of brain regions with reliable event structure showed no significant changes in event structure across repetitions of the Scrambled-Fixed clip (**Fig. 5**). The few regions that showed significant shifts showed changes only at the slow timescale. Superior temporal gyrus displayed finer event structure with repeated viewing, by losing event structure at a slow timescale (**Fig. 5, searchlight 1**). Superior temporal sulcus displayed coarser event representations with repeated viewing by gaining event structure at a slow timescale (**Fig. 5, searchlight 2**).

Figure 5. Changes in Event Structure at Slow and Fast Timescales with Repeated Viewing of the Scrambled-Fixed Clip. Following the conventions of Figure 4, brain maps depict how repeated viewing changes within- vs. between-event similarity at **(A)** a slow timescale (two events) and **(B)** a fast timescale (ten events). Significant searchlights are outlined with a black border (q < 0.05). Only subtle changes in timescales were observed for the Scrambled-Fixed condition, primarily in lateral temporal cortex. **(C)** Two sample searchlights (1–2) were selected post hoc for illustration. Red and blue arrows highlight the direction of change in within- vs. between-event similarity from the first viewing (thick yellow line) to the average of later viewings (thick green line).

Finally, we turned to the Scrambled-Random condition. As with previous clips, most brain regions with reliable event structure did not show significant changes in event structure across repetitions (**Fig. 6**). However, several regions exhibited significant shifts. Superior occipital cortex,

cuneus, and collateral sulcus demonstrated finer event representations with repeated viewing, due to weaker slow-event structure in later viewings (**Fig. 6, searchlights 1 and 2**). Most regions, however, exhibited coarser event representations with repeated viewing, often due to a gain of event structure at a slow timescale. These regions included middle temporal gyrus and inferior frontal lobe, which exhibited both stronger slow-event structure and weaker fast-event structure in later viewings (**Fig. 6, searchlights 3 and 4**). Overall, for the Scrambled-Random clip, 39 out of 45 searchlights with significant timescale changes (87%) became coarser with repeated viewing, compared with 12/23 (52%) for the Intact clip and 1/3 (33%) for the Scrambled-Fixed clip.

Scrambled-Random Clip

Figure 6. Changes in Event Structure at Slow and Fast Timescales with Repeated Viewing of the Scrambled-Random Clip. Following the conventions of Figures 4 and 5, brain maps depict how repeated viewing changes within- vs. between-event similarity at **(A)** a slow timescale (two events) and **(B)** a fast timescale (ten events). Significant searchlights are outlined with a black border (q < 0.05). Repeated viewing of the Scrambled-Random clip led to coarser event timescales throughout the lateral temporal cortex. **(C)** Four sample searchlights (1–4) were selected post hoc for illustration. Red and blue arrows highlight the direction of change in within- vs. between-event similarity from the first viewing (thick yellow line) to the average of later viewings (thick green line).

Examining across-clip overlap in timescale changes

To identify brain regions that exhibited consistent changes in event structure across different movie clips, we performed a conjunction analysis across the Intact, Scrambled-Fixed, and Scrambled-Random conditions (**Fig. 7**). The analysis revealed changes in event structure that generalized across movie content. Some regions, such as lateral occipital cortex, showed a consistent loss of event structure at the slow timescale across all clips, reflecting finer event structure with repeated viewing (**Fig. 7, searchlight 1**). Other regions, across superior and middle temporal gyri, exhibited a consistent gain of event structure at the slow timescale, reflecting coarser event structure over repetitions (**Fig. 7, searchlight 2**). These findings suggest that certain brain regions adapt the timescale of event segmentation in a systematic manner across varied stimulus contexts.

Figure 7. Conjunction Analysis Showing Consistent Changes in Event Structure Across All Clip Conditions (Intact, Scrambled-Fixed, Scrambled-Random). Brain regions from a conjunction analysis designed to identify areas showing significant changes in within- vs. between-event similarity across repeated viewings of all three clip conditions. Top row shows a sample searchlight for which events become finer with repetition across all clips, due to a loss of event structure at the slow timescale of two events (downward red arrows). Bottom row shows a sample searchlight for which events become

coarser with repetition across all clips, due to a gain of event structure at the slow timescale of two events (upward blue arrows). SFix = Scrambled-Fixed condition; SRnd = Scrambled-Random condition.

Relating changes in event structure to event memory

We next aimed to determine the behavioral relevance of event structure changes in the brain by examining the correlation between these changes and subsequent memory. We focused on brain regions that showed a consistent change in event structure across all three clips (as identified in the conjunction analysis) and tested if these changes predicted average recall across the clips in an individual differences analysis.

Figure 8. Memory Recall Examples. (A) Sample frames are shown from the Intact clip. Above and below these frames (connected by lines) are, respectively, examples of better and worse recall for the corresponding events. Responses from the worse recall example (bottom row, purple borders) are characterized by a relatively small number of vague, less detailed, or less accurate recollections of the events depicted in the clip. Responses from the better recall example (top row, yellow borders) are

characterized by a relatively large number of detailed, accurate, and contextually rich recollections of the clip's content. The worse recall example is shown in its entirety; only a subset of the better recall example is shown due to space constraints. **(B)** Recall scores for the Intact, Scrambled-Fixed (SFix), and Scrambled-Random (SRnd) conditions. Recall was superior for the Intact vs. Scrambled conditions. The better and worse recall examples for the Intact condition, shown in **(A)**, are highlighted in **(B)** with a triangle and square, respectively. * p < 0.05; ** p < 0.01.

We first examined overall memory performance by quantifying the number of details recalled for each clip (Fig. 8; see Materials and Methods/Scoring of event memory). Memory for the Intact clip (M = 25.75, SD = 11.24) was better than memory for the Scrambled-Fixed clip (M = 21.45, SD = 12.62) (t_{29} = 2.17, p = .04, 95% CI: 0.24– 8.36) and the Scrambled-Random clip (M = 19.93, SD = 8.82) (t_{29} = 3.36, p = .002, 95% CI: 2.28–9.36). Recall scores for the Scrambled-Fixed and Scrambled-Random clips were not significantly different (t_{29} = 0.66, p = .51, 95% CI:−3.17–6.20).

To relate changes in event structure to memory recall, we first estimated the change in Within- vs. Between-Event similarity values for each participant for a given searchlight at the timescale (slow or fast) that exhibited a significant group-level effect in that region. This was done for each clip separately and then averaged across clips to yield a single value for each participant. Finally, we obtained the correlation between this value and participants' average recall across the three clips, separately for each searchlight.

Figure 9. Across-Participant Correlation between Changes in Within- vs. Between-Event Similarity and Recall. For two brain regions identified via conjunction analysis across all three clip conditions (Intact, SFix, SRnd), each individual participant's average recall score across all clips is plotted against their average change in event structure for 2 events across all clips. Left panels show a region with group-level loss of slow timescale structure with repetition of all three clips (red arrow at 2 events). Right panels show a region with a group-level gain in slow timescale structure with repetition of all three clips (blue arrow at 2 events). For both regions, participants with greater loss of slow-timescale structure had significantly better recall (p < .05).

Lateral occipital cortex (**Fig. 9, left panel**, r=-0.324, p=0.0238) and middle temporal gyrus (**Fig. 9, right panel**, r=-0.310, p=0.0341) showed significant correlations with recall performance. In both regions, participants with greater loss of slow-timescale event structure had significantly better recall. Interestingly, this effect was observed despite contrasting group-level changes in these

brain regions, with lateral occipital cortex exhibiting group-level loss of slow timescale structure with clip repetition and middle temporal gyrus exhibiting group-level gain of slow timescale structure with clip repetition. This suggests that increasingly fine event structure over clip repetitions supports recall of narrative details, regardless of the overall tendency of a brain region to form finer vs. coarser event representations with repeated viewing.

Discussion

Summary of findings

We found both stability and flexibility in the brain's event timescales. While many regions showed consistent event timescales across repeated viewings of a narrative movie, others exhibited changes in the structure of their event representations. These included regions whose event representations became finer with repeated viewing, due to gaining event structure at a fast timescale and/or losing structure at a slow timescale; and regions whose event representations became coarser, due to gaining event structure at a slow timescale and/or losing structure at a fast timescale. The Intact clip—with preserved narrative coherence—elicited topographically organized timescale changes: in visual and auditory regions, such as lateral occipital cortex and middle temporal gyrus, repeated viewing led to finer event segmentation (loss of slow-timescale structure). In anterior regions involved in higher-order processing, such as inferior frontal cortex, event representations became coarser (gaining slow-timescale structure). This spatial organization of finer vs. coarser shifts was less apparent in the Scrambled-Fixed condition and largely absent in the Scrambled-Random condition, where coarsening was dominant. Thus, narrative coherence supports more spatially organized timescale changes with increasing familiarity, whereas disorder in narrative structure prompts a more consistent shift toward coarser event representations. Coarser representations for narratively incoherent events (in the Scrambled conditions) may aid in integrating information over longer timescales—perhaps as the brain attempts to impose structure where little exists. Finally, increasingly fine-tuned event representations in lateral occipital cortex and middle temporal gyrus were associated with more detailed memory. Thus, the brain's event timescales can flexibly change with experience, and these changes are correlated with event memory.

Relation to prior work

Our findings build on evidence that the brain's timescales are flexible and adjust to the statistics of incoming stimuli (Lerner et al., 2014; Baumgarten et al., 2021; Çatal et al., 2024). Lerner et al. (2014) reported that dynamics in early auditory cortex and higher-order areas rescaled in time to match compressed or dilated speech. The timescale of neural responses lengthened and shortened with slowed and sped-up sensory information, respectively. Baumgarten et al. (2021) also found that timescales adapt to sensory input rates. When predicting tone sequences played

at half, normal, or double speed, the brain integrated a fixed amount of information—a constant number of tones—rather than a fixed time window. Our observations align with such adaptive capacity, but extend it to repeated naturalistic viewing—rather than observing how timescales adapt to *external* sensory manipulations, we observe timescale changes due to increasing familiarity with a narrative, reflecting *internal* adaptation.

The changes we observed in the brain's structure of event representations may be related to changes in stimulus predictability. According to a popular theory of event segmentation, event boundaries are triggered when incoming information deviates from predictions, forcing an update of the current "event model" (Kurby & Zacks, 2009). On initial viewing, many moments may be surprising, leading to frequent prediction errors and fine segmentation. With repetition, events become expected; a reduction in surprise could allow the brain to merge events into longer segments (Hard et al., 2006). Simultaneously, increasing knowledge of the plot might enable the detection of subtle event transitions (e.g., noticing minor details or foreshadowing). Under alternative theories of event segmentation, which argue for a primary role of contextual instability (rather than prediction error) in shaping event boundaries (Güler et al., 2025; Shin & DuBrow, 2020; Shim et al., 2023), this enhanced sensitivity to previously-ignored details could result in additional fine boundaries that were initially overlooked. Thus, experience can recalibrate the brain's event structure at multiple scales.

Our findings provide insights into how temporal and narrative structure influence neural timescale changes. For the Intact clip, predictability and cohesion of the storyline allowed the brain to fine-tune the narrative into meaningful sub-events with experience. In contrast, the Scrambled-Random condition, which lacked coherent structure, showed widespread coarsening in event segmentation with experience. This coarsening may be related to participants' ability to gradually extract the global narrative across repeated viewings, as suggested by their subsequent recall of the overall plot. However, due to the scrambled order of segments, shorter-term predictability was limited, and fine-tuning of event boundaries was less pronounced. The Scrambled-Fixed condition, with stable event order but an incohesive narrative, exhibited both coarsening and fine-tuning—suggestive of the brain balancing prediction of upcoming events with gradual extraction of narrative structure. However, relatively limited changes in event structure in this condition warrant caution in overinterpreting these results. Together, our results showcase the brain's adaptive capacity to balance the extraction of global structure and encoding of detailed information. They underscore the dynamic nature of neural event segmentation and highlight the importance of considering stimulus characteristics when interpreting neural timescale changes.

We speculate that changes in the brain's event representations with experience may influence, or reflect, the detection of larger-scale structure and/or meaningful sub-events. Yet, this flexibility poses a methodological challenge. Repetition of dynamic stimuli (e.g., Huth et al., 2016; Golland et al., 2007; Hasson et al., 2009) may qualitatively alter temporal event structure in neural

responses. Thus, researchers should use caution when assuming that repeated presentations of dynamic stimuli leave the underlying neural event structure unchanged.

Limitations and alternative explanations

We interpret our results as qualitative changes in the structure of neural event representations with experience, but might they be due to other factors? One concern is regression to the mean: regions with initially strong event structure at a particular timescale might appear to shift toward intermediate values on later viewings. However, some regions with initially strong structure at the slow timescale exhibited an additional increase in this structure with repeated viewing, inconsistent with regression to the mean (e.g., searchlight 4 in **Fig. 4**). Further, many regions showed stable temporal structure with repeated viewing, which would be unexpected if our results were a statistical artifact. Additionally, we only considered regions with statistically reliable event structure for each clip viewing, reducing the risk of observing changes due purely to noise. Finally, losing event structure at the slow timescale predicted better memory (**Fig. 9**), arguing against the neural changes being artifactual.

Another concern is participant disengagement—viewers might become bored during later clip presentations, and increasing disengagement may drive changes in neural event structure. We think this is unlikely to explain our results. First, disengagement should lead to a loss of both fine and coarse event structure: idiosyncratic mind wandering or distraction should produce a general and widespread loss of consistency in event representations across people, but our analyses require similar event representations across individuals. Further, we observed *selective* changes in event structure with movie repetition: changes were not widespread across all timescales, they were not widespread across the brain, they varied by stimulus type, and the direction of changes reversed across brain regions and across fast and slow timescales within a region. Further, if the loss of event structure with repeated viewing was due to distraction or mind wandering, one would not expect it to be associated with *better* memory (**Fig. 9**).

Finally, we used the number of recalled details as an index of memory, which could align with finer event segmentation (parsing events more finely might produce more event fragments to describe). Though widely used in narrative recall research ([Martinez, 2024](#); [Chen et al., 2017](#); [Silva et al., 2019](#)), future studies should include complementary metrics (e.g., narrative recall coherence; recognition memory accuracy) to ensure that the relationship between neural event timescales and memory generalizes beyond detail quantity.

Conclusion and future directions

Brain regions' event representations are flexible and can be reshaped by experience. Regions dynamically adjust how they segment continuous information, either compressing or expanding event durations as events become familiar. Neural timescales are therefore not static properties of each region—they can adapt to the observer's knowledge.

Future research can leverage our new methodological approach, which quantifies event structure at multiple timescales simultaneously within a brain region, rather than assigning it a single preferred timescale. Tracking how both slow and fast event structures evolve *within a brain region* with repeated experience allows researchers to characterize flexible, multi-level segmentation dynamics across the brain. This framework could be applied to study learning, schema formation, or narrative processing. Additionally, future studies can explore how changes in event timescales relate to the perception of elapsed time (Matthews, 2011; Eaglement, 2010; Rose & Summers, 1995; Sherman & Yousif, 2025)—for example, whether gaining slow-timescale structure in higher-order regions is associated with the perception of time slowing with repetition. Finally, do timescale changes persist long-term, or does the brain revert to its initial segmentation patterns after a delay? Does persistence of timescale changes predict durable memory, whereas reversion to initial segmentation patterns predicts forgetting? Addressing these questions will elucidate how the brain flexibly constructs and updates event representations, and how this flexibility supports learning, re-interpreting, and remembering real-world experiences.

Together, our findings lay important groundwork for studying the brain's representation of events, revealing that even the temporal "building blocks" of experience, events themselves, can reorganize with experience. This opens the door to understanding how the brain's dynamic balancing of stability and flexibility supports our ability to understand and remember complex narratives.

Acknowledgements

This work was funded by a National Institutes of Health Research Project Grant (R01EY034436 to M.A. and C.B). We thank the Alyssano Group for valuable feedback on this project.

Author Contributions

Narjes Al-Zahli: Conceptualization, Data Curation, Formal analysis, Methodology, Project Administration, Software, Validation, Visualization, Writing—original draft, Writing—review and editing; **Mariam Aly**: Conceptualization, Methodology, Funding Acquisition, Resources, Project Administration, Supervision, Writing—original draft, Writing—review and editing; **Chris Baldassano:** Conceptualization, Methodology, Funding Acquisition, Resources, Project Administration, Supervision, Writing—original draft, Writing—review and editing.

References

Aly, M., Chen, J., Turk-Browne, N. B., & Hasson, U. (2018). Learning Naturalistic Temporal Structure
in the Posterior Medial Network. Journal of Cognitive Neuroscience, 30(9), 1345-1365.
https://doi.org/10.1162/jocn_a_01308

Aly, M., Chen, J., Turk-Browne, N.B., Hasson, U. (2019). Learning Naturalistic Temporal Structure in
the Posterior Medial Network. OpenNeuro. [Dataset].
https://doi.org/10.18112/openneuro.ds001545.v1.1.1

Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., & Norman, K. A. (2017).
Discovering event structure in continuous narrative perception and memory. *Neuron*,
95(3), 709-721.e5. https://doi.org/10.1016/j.neuron.2017.06.041

Baumgarten, T. J., Maniscalco, B., Lee, J. L., Flounders, M. W., Abry, P., & He, B. J. (2021). Neural
integration underlying naturalistic prediction flexibly adapts to varying sensory input rate.
Nature Communications, 12(1). https://doi.org/10.1038/s41467-021-22632-z

Çatal, Y., Keskin, K., Wolman, A., Klar, P., Smith, D., & Northoff, G. (2024). Flexibility of intrinsic
neural timescales during distinct behavioral states. *Communications Biology*, *7*(1).
https://doi.org/10.1038/s42003-024-07349-1

Chaudhuri, R., Knoblauch, K., Gariel, M., Kennedy, H., & Wang, X. (2015). A Large-Scale circuit
mechanism for hierarchical dynamical processing in the primate cortex. *Neuron*, *88*(2),
419–431. https://doi.org/10.1016/j.neuron.2015.09.008

Chen, J., Leong, Y. C., Honey, C. J., Yong, C. H., Norman, K. A., & Hasson, U. (2016). Shared
memories reveal shared structure in neural activity across individuals. Nature
Neuroscience, 20(1), 115–125. https://doi.org/10.1038/nn.4450

De Soares, A., Kim, T., Mugisho, F., Zhu, E., Lin, A., Zheng, C., & Baldassano, C. (2024). Top-down
attention shifts behavioral and neural event boundaries in narratives with overlapping
event scripts. *Current Biology*. https://doi.org/10.1016/j.cub.2024.09.013

Eagleman, D. M. (2008). Human time perception and its illusions. Current Opinion in
Neurobiology, 18(2), 131–136. https://doi.org/10.1016/j.conb.2008.06.002

Geerligs, L., Gözükara, D., Oetringer, D., Campbell, K. L., Van Gerven, M., & Güçlü, U. (2022). A
partially nested cortical hierarchy of neural states underlies event segmentation in the
human brain. *eLife*, *11*. https://doi.org/10.7554/elife.77430

Golland, Y., Bentin, S., Gelbard, H., Benjamini, Y., Heller, R., Nir, Y., Hasson, U., & Malach, R. (2006). Extrinsic and Intrinsic Systems in the Posterior Cortex of the Human Brain Revealed during Natural Sensory Stimulation. Cerebral Cortex, 17(4), 766–777. https://doi.org/10.1093/cercor/bhk030

Güler, B., Serin, F., & Günseli, E. (2025). Prediction error is out of context: The dominance of contextual stability in structuring episodic memories. Psychonomic Bulletin & Review. https://doi.org/10.3758/s13423-025-02723-4

Günseli, E., & Aly, M. (2020). Preparation for upcoming attentional states in the hippocampus and medial prefrontal cortex. *eLife*, *9*. https://doi.org/10.7554/elife.53191

Hard, B. M., Tversky, B., & Lang, D. S. (2006). Making sense of abstract events: Building event schemas. Memory & Cognition, 34(6), 1221–1235. https://doi.org/10.3758/bf03193267

Hasson, U., Yang, E., Vallines, I., Heeger, D. J., & Rubin, N. (2008). A hierarchy of temporal receptive windows in human cortex. *Journal of Neuroscience*, *28*(10), 2539–2550. https://doi.org/10.1523/jneurosci.5487-07.2008

Hasson, U., Avidan, G., Gelbard, H., Vallines, I., Harel, M., Minshew, N., & Behrmann, M. (2009). Shared and idiosyncratic cortical activation patterns in autism revealed under continuous real-life viewing conditions. Autism Research, 2(4), 220–231. https://doi.org/10.1002/aur.89

Honey, C. J., Thesen, T., Donner, T. H., Silbert, L. J., Carlson, C. E., Devinsky, O., Doyle, W. K., Rubin, N., Heeger, D. J., & Hasson, U. (2012). Slow Cortical Dynamics and the Accumulation of Information over Long Timescales. *Neuron*, *76*(2), 423–434. https://doi.org/10.1016/j.neuron.2012.08.011

Huth, A. G., De Heer, W. A., Griffiths, T. L., Theunissen, F. E., & Gallant, J. L. (2016). Natural speech reveals the semantic maps that tile human cerebral cortex. Nature, 532(7600), 453–458. https://doi.org/10.1038/nature17637

Kanter, B. R., Lykken, C. M., Polti, I., Moser, M., & Moser, E. I. (2025). Event structure sculpts neural population dynamics in the lateral entorhinal cortex. Science, 388(6754). https://doi.org/10.1126/science.adr0927

Kurby, C. A., & Zacks, J. M. (2008). Segmentation in the perception and memory of events. Trends in Cognitive Sciences, 12(2), 72–79. https://doi.org/10.1016/j.tics.2007.11.004

Lee, C. S., Aly, M., & Baldassano, C. (2021). Anticipation of temporally structured events in the brain. *eLife*, *10*. https://doi.org/10.7554/elife.64972

Lerner, Y., Honey, C. J., Silbert, L. J., & Hasson, U. (2011). Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *The Journal of Neuroscience*, *31*(8), 2906–2915. https://doi.org/10.1523/jneurosci.3684-10.2011

Martinez, D. (2024). Scoring story recall for individual differences research: Central details, peripheral details, and automated scoring. Behavior Research Methods, 56(8), 8362–8378. https://doi.org/10.3758/s13428-024-02480-7

Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., & Wang, X. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*, *17*(12), 1661–1663. https://doi.org/10.1038/nn.3862

Matthews, W. J. (2011). Stimulus repetition and the perception of time: The effects of prior exposure on temporal discrimination, judgment, and production. PLoS ONE, 6(5), e19815. https://doi.org/10.1371/journal.pone.0019815

Rose, D., & Summers, J. (1995). Duration illusions in a train of visual stimuli. Perception, 24(10), 1177–1187. https://doi.org/10.1068/p241177

Sherman, B. E., & Yousif, S. R. (2025). An illusion of time caused by repeated experience. Psychological Science. https://doi.org/10.1177/09567976251330290

Shim, S., Mugisho, F. B., Davachi, L., & Baldassano, C. (2014). Generating event boundaries in memory without prediction error (Preprint). PsyArXiv. https://doi.org/10.31234/osf.io/4bdp6

Shin, Y. S., & DuBrow, S. (2020). Structuring memory through Inference‑Based event segmentation. *Topics in Cognitive Science*, *13*(1), 106–127. https://doi.org/10.1111/tops.12505

Silva, M., Baldassano, C., & Fuentemilla, L. (2019). Rapid memory reactivation at movie event boundaries promotes episodic encoding. Journal of Neuroscience, 39(43), 8538–8548. https://doi.org/10.1523/jneurosci.0360-19.2019

Stephens, G. J., Honey, C. J., & Hasson, U. (2013). A place for time: the spatiotemporal structure of neural dynamics during natural audition. Journal of Neurophysiology, 110(9), 2019–2026. https://doi.org/10.1152/jn.00268.2013

Tarder-Stoll, H., Baldassano, C., & Aly, M. (2024). The brain hierarchically represents the past and future during multistep anticipation. *Nature Communications*, *15*(1). https://doi.org/10.1038/s41467-024-53293-3

Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: A mind-brain perspective. *Psychological Bulletin*, *133*(2), 273–293. https://doi.org/10.1037/0033-2909.133.2.273

Supplementary Figures

Reliability of Within- vs. Between-Event Similarity for Initial Viewing

Supplementary Figure 1. Validation of Within- vs. Between-Event Similarity (Intact Clip). To assess the reliability of the Within- vs. Between-Event Similarity measure during participants' initial viewing of the movie clip, we examined how correlated this measure was across searchlights for two independent participant samples. Participants were first randomly divided into two independent groups of 15 participants each. Then, within each group separately, participants were randomly divided into training and testing subsets. We identified HMM event boundaries for each searchlight in the training subset and then used those boundaries to compute Within- vs. Between-Event Similarity in the testing subset. Training and testing subsets were then swapped, and results were averaged across five different participant splits within each independent group. Finally, we examined the correlation in Within- vs. Between-Event Similarity values across the two independent groups. Reliability for Within- vs. Between-Event Similarity values was consistently high across all event counts, demonstrating that this measure is relatively stable across subsets of participants.

Reliability of Within- vs. Between-Event Similarity for Subsequent Viewings

Supplementary Figure 2. Validation for Subsequent Viewings (Intact Clip). The same split-half reliability procedure was repeated for participants' subsequent viewings of the narrative. Results again showed strong reliability across all event counts, indicating that the Within- vs. Between-Event Similarity measure remains robust across repeated viewings.