Generating event boundaries in memory without prediction error

Sunjae Shim[1], Franck B. Mugisho[2], Lila Davachi[2,3], & Christopher Baldassano[2]
1. Department of Psychology, Stanford University
2. Department of Psychology, Columbia University
3. Nathan Kline Institute for Psychiatric Research

Word Count: 5205
Data and code available at: https://github.com/sjshim/event_boundary_wm

Abstract: Continuous experiences are segmented into discrete long-term memories through the generation of event boundaries. A leading theory of event segmentation proposes that event boundaries are triggered by prediction errors caused by unexpected stimuli. However, recent studies have raised doubts about whether prediction error is necessary for event segmentation. In this study, we tested an alternative account: that event boundaries in memory reflect the temporal structure of working memory during perception and can occur even in the absence of prediction error. In experiment 1, participants were asked to detect repeats within sequences of random images. The switch to a new sequence could be predictable, with a continuous display of the number of images remaining in each sequence, or unpredictable, with no prior indication to participants that a sequence was about to end. We found boundary-related effects on temporal order memory in both cases, with higher accuracy for the within-sequence comparisons even when the event boundary between sequences was predictable. In experiments 2a and 2b, event boundaries were always predictable, and participants performed either the (event-related) repeat-detection task from experiment 1 or a (non-event-related) 2-back task. We observed event-boundary effects on order memory only when the working memory task was event-related. Both of these experiments support an alternative theory of event segmentation, in which boundaries are critically related to working memory dynamics rather than prediction error.

Public significance statement: Although our lives unfold steadily over time, our memories are organized as a library of separate, distinct events. Current theories disagree about how the boundaries between these events are determined, with a prominent proposal arguing that we only start a new event memory when we experience a surprise. We show in our experiments that, in fact, event boundaries can occur even for changes that are entirely predictable, and are instead driven by moments at which we stop keeping track of information about the recent past.

Keywords: Event boundaries, Event segmentation, Prediction error, Working memory

Event segmentation, the cognitive process of dividing ongoing activities into smaller events, has been studied for decades in cognitive psychology (Lichtenstein & Brewer, 1980; Trabasso & Van Den Broek, 1985; Kurby & Zacks, 2008). Event boundaries have immediate effects during encoding, visible in behavioral measures (Ongchoco et al., 2023; Swallow et al., 2011; Zwaan et al., 1995) and neuroimaging (Baldassano et al., 2017), and impact the organization of experiences in long-term memory (Clewett et al., 2019; Ezzyat & Davachi, 2011; Heusser et al., 2018; V. Wang et al., 2023). However, despite the fact that observers tend to broadly agree about which moments during a stimulus constitute event boundaries (Sasmita & Swallow, 2022), there is still substantial controversy over the mechanisms by which event boundaries are generated (Clewett et al., 2019; Güler et al., 2024; Nolden et al., 2024; Shin & DuBrow, 2021; Y. C. Wang et al., 2024).

A prominent theory by Zacks and colleagues (2007), referred to as Event Segmentation Theory (EST), proposes that boundaries in human perception and comprehension are primarily guided by prediction error. According to EST, we consistently maintain a model of the current event in working memory and use this model to make predictions about the near future. When these predictions are incorrect, this creates an event boundary and triggers an update to our event models. There is substantial empirical support for the idea that prediction errors can generate event boundaries (Rouhani et al., 2020; Zacks et al., 2009, 2011), and recent work using continuous paradigms has found evidence for ongoing and automatic predictive processes, based on predictive eye movements (Huang et al., 2023) and neural signals (Lee et al., 2021).

Although these prior papers provide evidence that prediction error may be *sufficient* for event segmentation, it is not known whether it is *necessary*. The widely used Ezzyat-Dubrow-Davachi paradigm (Buonomano et al., 2023) contains highly predictable event transitions (with a fixed number of items per event) and the transitions between events have been shown to evoke the behavioral markers of event boundaries (Davachi & DuBrow, 2015; DuBrow & Davachi, 2014, 2016; Heusser et al., 2016, 2018). Furthermore, reaction-time effects at boundaries do not diminish with repeated exposure to the same event-structured list (Bein & Davachi, 2022). Using a statistical learning paradigm, it is possible to produce event clustering among items with shared community structure while exactly controlling for transition probability at boundary vs. non-boundary transitions (Schapiro et al., 2013). Even one of the classic paradigms for generating boundaries, moving through a physical doorway (Pettijohn & Radvansky, 2018; Radvansky et al., 2011; Radvansky & Copeland, 2006; Seel et al., 2019), seems at odds with the prediction-error account; walking in the front door of one's own home should generally not be a surprising experience. In fact, even *imagining* walking through a doorway causes event-boundary effects in working memory (Lawrence & Peterson, 2016), despite the fact that there is no new perceptual information which could cause a perceptual prediction error.

Here we test an alternative account: that perception of continuous stimuli requires active management of our working memory, and event boundaries are moments at which information is strategically removed from working memory. When using our limited working-memory resources to track the most relevant information for our current tasks and goals, boundaries are moments at which recent stimuli suddenly become less relevant. This change in relevance could be triggered by prediction error (prompting a revision of our current event model), but could also come from a completely expected change in task, goals, location, or event type. This theory is consistent with the recent framing of event segmentation as a working memory process (Güler et al., 2024; Jafarpour et al., 2022; Radvansky, 2017), as well as proposals that event boundaries reflect changes in tasks (Clewett et al., 2019; Y. C. Wang & Egner, 2022) or inferred event type (Shin & DuBrow, 2021).

In this study we contrasted the working-memory account against the prediction-error account by separately manipulating whether boundaries were predictable and whether the task required the maintenance of working memory during events. After participants were shown a sequence of images, we tested for the presence of event boundaries in episodic memory using a temporal-order memory task, since order judgments that cross an event boundary are known to result in lower accuracy (DuBrow & Davachi, 2013; Heusser et al., 2018). In experiment 1, participants were asked to detect repeated images within each "segment," with transitions between segments occurring when a clock was shown with its hand pointing up. These transitions could be made predictable by having the clock hand continuously visible as it rotated during each segment, or be made unpredictable by having the clock hand only appear at the beginning of each segment. We found boundary-related effects on temporal order memory in both the predictable- and unpredictable-boundary conditions, showing that a working-memory task is *sufficient* to predict event boundaries, even in the absence of prediction error. In experiments 2a and 2b, event boundaries were always predictable, and participants performed either the (event-related) repeat-detection task from experiment 1 or a (non-event-related) 2-back task. We observed event-boundary effects in memory only when the working-memory task was event related, suggesting that maintaining items in working memory during an event is also *necessary* for the generation of event boundaries. Both of these results are inconsistent with the prevailing theory of event segmentation and support an alternative view in which boundaries are critically related to working memory dynamics rather than prediction error.

Experiment 1

Method

Participants

80 participants (female = 27, male = 50, unreported = 3, M_{age} = 27.81, age range: 21-54 years) were recruited through the Prolific online platform. All participants were fluent in English, gave informed consent through an online questionnaire, and received monetary compensation for their time. The experimental protocols for this experiment and all subsequent experiments were approved by the Institutional Review Board of Columbia University.

Materials

432 images were selected from the Bank of Standardized Stimuli (Brodeur et al., 2014) and THINGS database (Hebart et al., 2019) to create 2 image stimuli sets, each consisting of 216 items. Each set of 216 was randomly divided into six 36-item lists, and each list was divided into 6 segments: two 5-item segments, two 6-item segments, and two 7-item segments. 20% of items were then replaced by a previous item from the same segment (creating within-segment repeats), and 10% were replaced with a previous item from a different segment (creating across-segment repeats). The full randomization procedure was carried out 19 times, and participants were randomly assigned to one of these 19 versions of the stimulus order. Stimulus presentation and data collection was carried out using PsychoPy/Pavlovia (Peirce et al., 2019).

Procedure

During the first phase of each run, participants were presented with a series of 36 random images along with an image of a clock (Figure 1, left). In the Predictable Boundary condition, the arm of the clock started pointing upwards (at the 12 o'clock position), and the clock was visibly divided into 5, 6, or 7 wedges. After each image, the arm advanced one wedge around the clock. When the arm came back to the vertical position, a new *segment* began, which could again consist of 5, 6, or 7 images. Participants were asked to respond whether the presented image was new or repeated from within the current segment. Only the images within the current segment were relevant to this task; if an image had been seen before but in a previous segment, the participant was instructed to respond to these as "new." The start of a new segment could always be exactly anticipated by the participants, since they could always see how many wedges (images) remained in the current segment. The images were presented for a fixed time of 2.5 seconds (s) during which the participants had to respond, followed by a fixed 2 s inter-trial interval.

In the Unpredictable Boundary condition the participants performed the same task, but the arm of the clock was visible pointing up only at the beginning of a new segment, and the wedges (indicating segment length) were never shown. Because the segments varied in length, this meant that participants could only partially predict when each segment would end, and were cued to start the next segment by the sudden appearance of the vertical arm.

In the second phase of each run, we tested temporal order memory for the presented images (Figure 1, right). Participants were shown two images from the presentation phase and were asked to select which of the images they had seen first. Each pair was chosen from items that were never repeated during the presentation phase and were 3 items apart from each other (i.e. had 2 intervening items). For each block, 6 pairs of images were tested: 3 pairs used images that came from the same segment, and 3 pairs used images that came from two different segments (i.e. a segment boundary occurred between them). The trials were self-paced and advanced once a response was given. There was a fixed .5 s inter-trial interval between test trials.

Participants completed a practice run of both conditions before participating in the main task. The practice task provided detailed instructions about the task and a run-through of both phases of the task, consisting of 6 segments of item presentation and 6 pairs of temporal order accuracy judgment. Stimuli used in the practice task were not included in the main task. In the main task, each participant completed 6 of these runs in one condition (Predictable or Unpredictable Boundaries) followed by 6 runs in the other condition, with condition order counterbalanced across subjects. Items were never repeated across different runs.

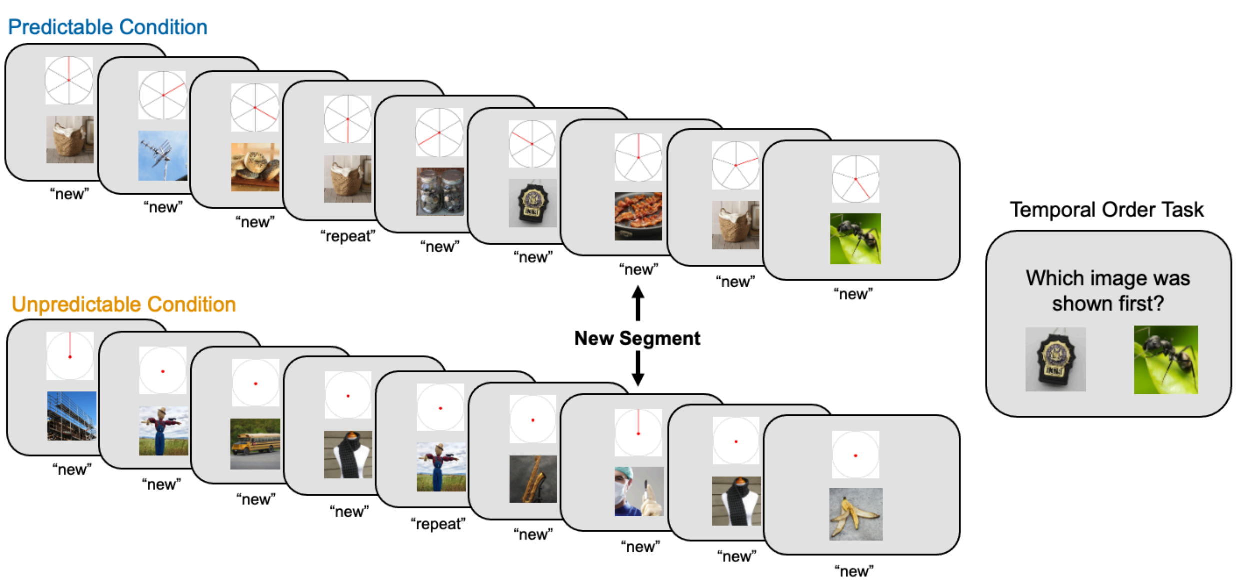

Figure 1: Experiment 1 task design. Participants were presented with a sequence of 36 images (each paired with a clock image), and asked to determine when an image was repeated within a segment. A new segment began whenever the clock arm pointed straight up, which could either be trivially anticipated (in the Predictable Condition) or occur without warning (in the Unpredictable Condition). Participants then performed 6 trials of a temporal order memory task, to identify which of two images was presented first. This full procedure was repeated 12 times, 6 in each condition (blocked by condition, with condition order counterbalanced across participants).

Results

Working-memory task accuracy and response time

Participants performed well in the presentation phase of the experiment with an average accuracy of 95.9% (SD = 3.76%) in labeling images as new or repeated within the current segment. For subsequent analyses, we excluded 9 participants with accuracies below 90% on this task, resulting in a final N=71. Performance was comparable for all positions within a segment, with near-ceiling performance for the first item which was by definition always “new” in that segment (Supplementary Figure 1).

Although participants were not explicitly instructed to make speeded responses, we observed reaction time differences related to segment boundaries in both conditions. The average response time for boundary items (i.e. first item shown in a new segment) was significantly longer than for non-boundary items (i.e. all other items) for both unpredictable and predictable conditions (Predictable: $t_{70}=3.19$, $p<0.01$, 95% CI of RT difference [10ms, 60ms]; Unpredictable: $t_{70}=6.88$, $p<0.001$, 95% CI of RT difference [50ms, 100ms]). Note that this task is *trivial* for boundary items (which are “new” for that segment), suggesting that the slowing in reaction times for these items is not due to the decision-making process itself, but is instead related to the processing of an event transition.

Effect of perceptual boundaries on temporal order memory

In order to determine whether items within a segment were bound together into an event memory, we compared the accuracy of temporal order judgments for pairs of items that had been presented 3 images apart. Higher accuracy for within-event versus between-event pairs was used to measure the strength of event boundaries in memory (Heusser et al., 2018). As shown in Figure 2 (left), we observed a significant accuracy advantage for within-event judgments in both conditions (Predictable: t_{70}=3.02, p<0.01, 95% CI of accuracy difference [0.03,0.12]; Unpredictable: t_{70}=2.62, p<0.05, 95% CI of accuracy difference [0.01,0.11]). Plotting the difference in this accuracy effect between conditions for each participant (Figure 2, right) we find that the 95% confidence interval for the population mean difference is tightly constrained around 0 ([-0.03, 0.06]); since we have 18 temporal order trials in each condition, this corresponds to a population effect that is weaker than 0.61 trials in favor of the Unpredictable condition or 1.14 trials in favor of the Predictable condition.

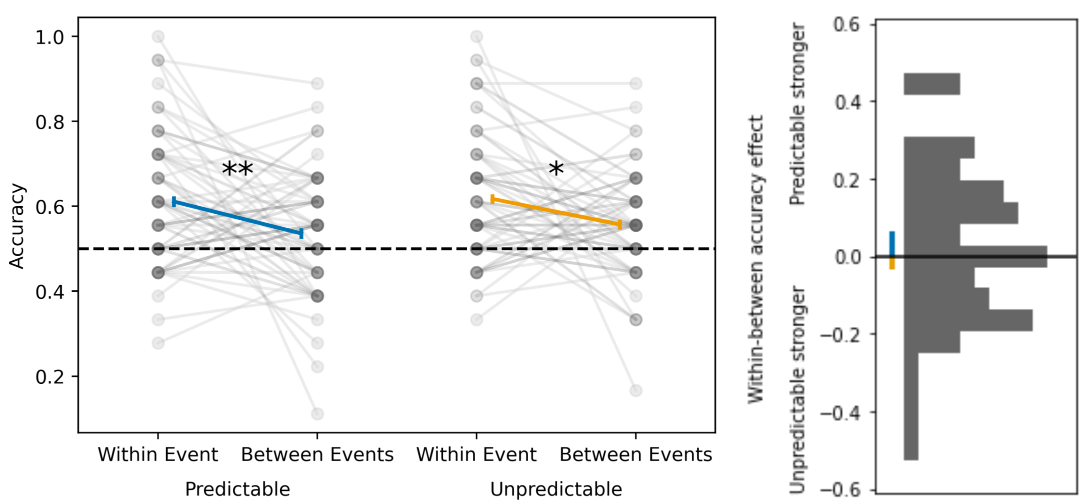

Figure 2: Temporal order accuracy for Experiment 1. (left) Participants were more accurate for temporal order judgments when both items came from the same event rather than adjacent events. This event boundary effect was present even when boundaries were entirely predictable. (right) Across participants, the strength of the boundary effect in the two conditions was on average the same in both conditions, with a 95% confidence interval for the population mean tightly constrained around 0. * p<0.05, ** p<0.01

Discussion

We found that an event-structured working memory task was sufficient to generate event boundary effects in order memory, even when event boundaries were entirely predictable. This result demonstrates that prediction error is not critically necessary to generate event boundaries in memory. In the Predictable Boundary condition, participants could perfectly anticipate the moments when they would need to update the contents of their working memory (flushing out previous items), suggesting that these working memory dynamics were driving event segmentation rather than prediction error.

We therefore designed a second experiment to further test the hypothesis that dynamics in working memory drive event segmentation. This experiment also sought to address some limitations of Experiment 1. First, participants in the Predictable Boundary condition do receive a small amount of new information at the beginning of a new segment, since they can see whether the new segment will be five, six, or seven images long (according to the number of wedges presented on the clock). While this setup was necessary to add ambiguity about the length of events to the task, this extra information could be considered to cause an increased "prediction error" at boundaries since the number of wedges could not be predicted in advance. We addressed this issue in Experiment 2 by making all segments six images long, ensuring that boundaries were entirely predictable and boundary items did not contain any extra information compared to non-boundary items. Second, we made the Experiment 2 stimulus sequences identical across conditions, controlling for the influence of perceptual information and ensuring that any observed differences in order memory were solely due to the internal dynamics of memory updating.

Experiment 2

Methods

Participants

For Experiment 2a, 137 participants fluent in English and living in the United States (female = 55, male = 57, unreported = 25, M_{age} = 34.12, age range: 20-75 years) were recruited from Prolific and from a university in the United States: 55 participants for the segment memory condition and 82 participants for the 2-back condition. For replication study Experiment 2b, 132 participants fluent in English and living in the United States (female = 69, male = 56, unreported = 7, M_{age} = 35.22, age range: 19-69 years) were recruited from Prolific: 68 participants for the segment memory condition and 64 participants for the 2-back condition. All participants gave informed consent through an online questionnaire and received monetary compensation for their time.

Materials

For Experiment 2a and 2b, the same pool of 432 images from Experiment 1 was used to create 4 stimulus sequences, each consisting of 324 items. For each of these 4 sets, items were chosen from the image pool and organized into 6-image segments. We biased the placement of item repeats to ensure that 2-back repeats were at least partially dissociable from same-segment repeats; this process resulted in items that could be completely novel and unrepeated (51.4%), repeated from exactly 2 images ago (17.6%), repeated from within the current segment (24.9%), or repeated from previous segments (13.7%). Note that an image can belong to more than one of these categories; for example, if an item at position 5 in a segment was a repeat of the item at position 3, then it would be both a within-segment repeat and a 2-back repeat. Each participant was randomly assigned to use one of the four image sequences, regardless of task condition. Since the same sequences were used for the two tasks, this ensured that any differences between conditions were unrelated to the perceptual dynamics of the image sequences.

Procedure

As in Experiment 1, each experimental run consisted of a memory task on a sequence of 36 images and a subsequent temporal order judgment task (Figure 3). On each trial of the memory task, an image was presented with a colored background border that was unrelated to the content of the image. The border cycled through six colors, which were selected to be equidistant around the perceptually-uniform CIELAB hue wheel, therefore providing a constant amount of hue change between each image and the subsequent image. The blue border denoted the start of each six-item sequence.

Participants were assigned to one of two experimental conditions. One group of participants performed the same task as in Experiment 1 (the *segment memory* condition), determining whether the presented image was new or repeated within the current six-item segment (i.e. since the last appearance of a blue-colored border). The images were presented for 2.5 s and were followed by a fixed 2 s inter-trial interval. Another group of participants performed a control working-memory task which was unrelated to the segment boundaries (the *2-back* condition). In this condition, participants were given the same sequence of stimuli as the background condition task, but were instructed to respond “repeat” only if they saw an image repeated from exactly 2 images ago. The participants were told that they would see a colored background along with the image, but that this color was not relevant to the task they would be performing.

Participants in both conditions then performed a temporal order task identical to the task in Experiment 1. Each participant completed a total of 9 runs, each with an item presentation phase and a temporal order task phase, and novel items were used in each run. For Experiment 2a, participants completed a practice task for their assigned condition that was similarly structured as Experiment 1. For Experiment 2b, to ensure that participants were understanding the task clearly, we implemented a performance threshold before participants moved on to the main task. The same practice block was repeated until participants reached an accuracy of 70% in the first phase of the task. All participants included in the sample met this threshold.

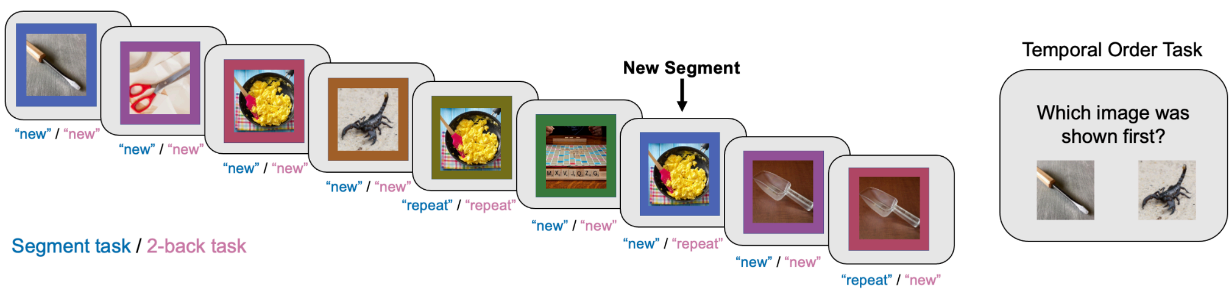

Figure 3: Experiment 2a and 2b task design. As in Experiment 1, participants were presented with a sequence of 36 images, and then performed six trials of the temporal order task. This procedure was repeated nine times. Segments were indicated by a predictable sequence of colors starting with blue. Unlike Experiment 1, only half of the participants performed the segment memory task (detecting repeats within each segment), while the other participants performed a 2-back task (detecting repeats two items apart). For the example sequence shown, the fifth and ninth images are within-segment repeats, while the fifth and seventh images are 2-back repeats.

Participant exclusions

Average presentation phase accuracy for Experiment 2a was 90.81%. However, we observed that a number of participants had high rates of false alarms for items that did not match their task instructions. Specifically, we examined critical items that were 2-back repeats but not within-segment repeats (these images spanned an event boundary) and within-segment repeats that were not 2-back repeats (these images were more or less than 2 images apart). We excluded 24 participants from the segment memory condition who false alarmed to more than 50% of the images that were 2-back repeats but not within-segment repeats (which they should have labeled as “new”) and excluded 29 participants from the 2-back condition who false alarmed to more than 50% of the images that were within-segment repeats but not 2-back repeats (which they should have labeled as new). As in Experiment 1, we also excluded 3 and 7 participants from the segment memory and 2-back conditions, respectively, for having overall accuracy below 90%. This resulted in 28 and 46 participants left in each condition for subsequent analyses.

Since these additional exclusion criteria were decided post-hoc, we ran an additional replication Experiment 2b in which we applied these exclusions to a new dataset to ensure that our results would generalize. We additionally made minor improvements to the task instructions, and added additional practice trials as described above. Overall accuracy during the presentation phase was 92.60%. Applying the exclusion criteria from Experiment 2a based on false-alarms to condition-critical items, we excluded 34 and 12 participants in the segment memory and 2-back conditions, respectively, and also excluded 1 and 10 more participants, respectively, for overall presentation phase accuracy below 90%. This resulted in 33 and 42 participants left in each condition for subsequent analyses.

Results

Working memory task performance

Following exclusions, presentation-phase accuracy for Experiment 2a and 2b was 96.81% and 96.48%, respectively. Overall, the accuracy for items in each position of a segment was high for both segment memory and 2-back conditions. As in Experiment 1, we observed longer average reaction times for boundary versus non-boundary items in both Experiment 2a (t_{27}=2.17, $p<0.05$, 95% CI of RT difference [0ms, 70ms]) and Experiment 2b (t_{32}=4.50, $p<0.001$, 95% CI of RT difference [30ms, 80ms]), and when combining both datasets (t_{60}=4.39, $p<0.001$, 95% CI of RT difference [20ms, 70ms]), confirming that slowed reaction times at boundaries are present even for predictable boundaries. Unexpectedly, we also observed a reaction time difference for boundary items in the 2-back condition, in which participants were told to ignore the cycling background colors (2a: t_{45}=5.70, $p<0.001$, 95% CI of RT difference [30ms, 70ms]; 2b: t_{41}=4.65, $p<0.001$, 95% CI of RT difference [20ms, 60ms]; Combined: t_{87}=7.36, $p<0.001$, 95% CI of RT difference [30ms, 60ms]). It is possible that participants (implicitly or explicitly) recognized when the color sequence repeated, and this served to structure the timing of their responses.

Effect of memory task on temporal order memory

To test our hypothesis that event boundary effects rely on working memory dynamics, we measured within- versus between-event temporal order memory for participants in both task conditions. We found that participants performing the segment memory task showed a significant boundary effect on temporal order memory, replicating our results from Experiment 1 (Experiment 2a: t_{27}=4.21, $p<0.001$, 95% CI of accuracy difference [0.05, 0.17]; Experiment 2b:

t_{32}=4.64, p<0.001, 95% CI [0.08, 0.20]; Combined: t_{60}=5.98, p<0.001, 95% CI [0.08, 0.17]). When participants performed the 2-back task on an identical sequence of stimuli, however, there was no significant difference between within- and between-segment temporal order memory (Experiment 2a: t_{45}=1.39, p>0.05, 95% CI [-0.01, 0.06]; Experiment 2b: t_{41}=-0.22, p>0.05, 95% CI [-0.05, 0.04]; Combined: t_{87}=0.70, p>0.05, 95% CI [-0.02, 0.04]). The boundary effect in the segment memory condition was significantly greater than the effect in the 2-back condition (Experiment 2a: t_{72}=2.63, p<0.05, 95% CI of difference in accuracy difference [0.02, 0.14]; Experiment 2b: t_{73}=3.87, p<0.001, 95% CI [0.07, 0.22]; Combined: t_{153}=4.74, p<0.001, 95% CI [0.07, 0.16]).

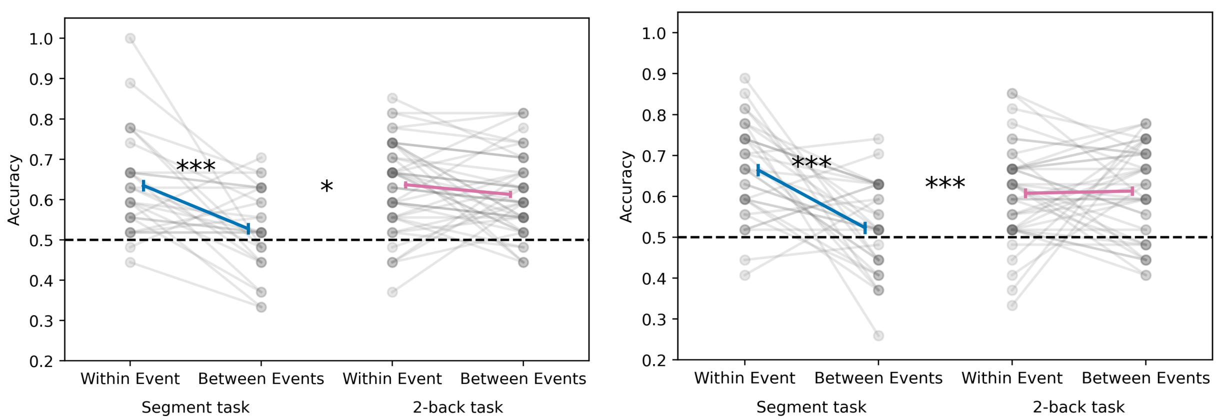

Figure 4: Temporal order accuracy for Experiments 2a and 2b. In both Experiment 2a (left) and 2b (right), participants performing the segment memory task exhibited a within- versus between-event order memory effect, replicating our results from the predictable-boundary condition in Experiment 1. Participants performing a 2-back task on identical stimuli, however, showed no boundary-related effect in temporal memory. * p<0.05, *** p<0.001

Discussion

The results from Experiments 2a and 2b replicated our findings from Experiment 1, showing that event boundaries in memory do not require a spike in prediction error and can arise solely through the resetting of working memory load. Having a working-memory task that was coupled to the event boundaries was critical for producing event-structured memories; when the same sequence of items was instead experienced while performing a 2-back task, there were no boundary-related effects in the temporal order task.

Event boundaries are to some extent a property of a stimulus itself, and computational models can partially predict event boundary structure for stimuli such as narratives (Kumar et al., 2023). Our findings, however, argue for a view of event segmentation as an active process occurring in the mind, which can result in task-dependent event boundaries. Unlike work in which event boundaries were shifted by manipulating attention to specific dimensions of a stimulus (Bailey et al., 2017; De Soares et al., 2023), here participants were performing qualitatively similar tasks (detecting repeated images) that differed in their dynamics: images were held in working memory either in event units (that expanded and then reset) or in a continuously-moving window (of the last two items). Event segmentation in episodic memory therefore emerged as a consequence of the working-memory strategy employed by participants (Güler et al., 2024; Radvansky, 2017), rather than participants simply “detecting” bottom-up event boundaries. Interestingly we did find that reaction times were slightly slowed when the background color cycled restarted, regardless of memory condition, potentially suggesting that this measure is more sensitive to the perceptual dynamics of the stimulus rather than the event structure of the memory traces being formed.

General Discussion

In all three of our Experiments, we found boundary-related effects in temporal order memory for task conditions in which switches to new events were entirely predictable. These results are incompatible with a straightforward prediction-error account of how event boundaries are generated. Instead, we found that event boundaries were generated at moments when working memory was reset. This provides compelling evidence for an alternative account of event segmentation: that event boundaries in memory correspond to moments at which information about the recent past is dropped from working memory. Understanding and contextualizing our current sensory input generally requires maintaining information about the recent past (Hasson et al., 2015), but at moments when the latent structure of the world shifts, these past observations may suddenly become less relevant.

This view explains why event boundaries can occur even at entirely-predictable transitions. For example, when crossing a doorway into another (familiar) room, the details that we had just been observing in the previous room are no longer relevant; to make the best use of our limited working memory, we should create a new event model for our current environment which could include, for example, the specific locations of items in this room that we'd like to interact with. Similarly, when the countdown on New Year's Eve reaches 0, this highly predictable moment still results in a strong event boundary, since our task and attentional sets suddenly change (e.g. from counting down in sync with the broadcast to grabbing a drink for a toast). There is a sense in which our account extends the Event Segmentation Theory of Zacks et al. (2007); rather than proposing that event models are discarded from memory *only* due to spikes in prediction error, we argue that event models can be preemptively "scheduled" for retirement (and removed from working memory). Baldwin & Kosie (2021) have made similar arguments based on dwell-time studies, showing that viewers actively anticipate upcoming event boundaries and that these boundaries become *stronger* with repeated exposure.

There are multiple kinds of mechanisms that could lead to this observed relationship between working memory dynamics and event segmentation. One possibility is that holding items in working memory binds them together in relation to each other, with the critical work of creating events primarily occurring during the event itself as items are added to a linked sequence or an associative web (DuBrow & Davachi, 2014; Heusser et al., 2016). Another possibility is that the moments at which working memory is reset trigger specific consolidation processes that encode information into long-term memory, resulting in items being bound into events only at the end of an event (Baldassano et al., 2017; Ben-Yakov & Dudai, 2011; Reagh et al., 2020; Silva et al., 2019; Wu et al., 2022). Güler and colleagues (2024) describe these two possible roles of working memory as “accumulation” and “reactivation” processes, and it is still unclear which of these processes best describes how events are constructed (or if both play a role, either simultaneously or in different contexts).

In this study we directly manipulated working memory dynamics through an explicit memory task, but in more naturalistic settings we would expect the contents of working memory to be strategically managed to optimize comprehension. Moments when the future and past are largely independent given our present sensory input are the optimal times to drop information from working memory (Baldassano, 2023); at these moments, retaining recent information is no longer valuable for upcoming perception or decision-making, and it can be either forgotten (if it is unlikely to ever be relevant) or stored solely in long-term memory (if it may need to be re-activated in the future). Determining when our past and future are relatively disconnected is of course inherently challenging, since the future has not yet been observed, but in many cases we can rely on predictive scripts for familiar event sequences (Abelson, 1981; Baldassano et al., 2018) to anticipate what information will be irrelevant in the near future. In addition, there are a number of heuristics that could indicate when our previous situation model should be reset, including not only prediction error (surprisal) (Richmond & Zacks, 2017), but also salient changes in context (Clewett & Davachi, 2017) or sudden changes in our predictions about the next stimulus (Kumar et al., 2023).

The neuronal mechanisms by which sequences are represented in working memory is still an open question, but electrophysiological work has suggested that theta oscillations serve to organize stimuli into sequences. Specifically, cell assemblies firing at gamma frequencies represent specific items, and these item representations activate in an ordered sequence throughout the phases of a theta oscillation (Bahramisharif et al., 2018; Heusser et al., 2016). Our results suggest that this kind of theta-gamma representation of recent items does not occur automatically in response to a sequence of images, but is only maintained when these recent images are relevant for an ongoing task. Interestingly, our segment memory task (detecting repeats within a segment) did not require maintenance of sequential information about the order of the previous items in a segment, but we still observed improved order memory for these items; this suggests that a sequential code is a default representational format for items within an event.

Overall, our results are inconsistent with theories in which event boundaries can *only* arise due to prediction error, showing that items simultaneously held in working memory can be bound into events even when the switches between events are entirely predictable. Rather than only reacting to prediction errors, the cognitive process of event segmentation can be more proactive and strategic, discretizing experience into meaningful units that correspond to known temporal dynamics of the environment.

Constraints on Generality

Our participant pool consisted only of English-speaking residents of the United States who had access to the Prolific online platform, and focused primarily on young adults, so we are unable to assess whether there are differences across cultures or ages in the strength of event boundaries evoked by these conditions. Since we used controlled image sequences as stimuli, we do not have direct evidence that these same processes drive event segmentation in other sensory modalities, or precisely how these effects arise in more naturalistic settings.

References

Abelson, R. P. (1981). Psychological status of the script concept. *American Psychologist*, *36*(7), 715–729. https://doi.org/10.1037/0003-066X.36.7.715

Bahramisharif, A., Jensen, O., Jacobs, J., & Lisman, J. (2018). Serial representation of items during working memory maintenance at letter-selective cortical sites. *PLOS Biology*, *16*(8), e2003805. https://doi.org/10.1371/journal.pbio.2003805

Bailey, H. R., Kurby, C. A., Sargent, J. Q., & Zacks, J. M. (2017). Attentional focus affects how events are segmented and updated in narrative reading. *Memory & Cognition*, *45*(6), 940–955. https://doi.org/10.3758/s13421-017-0707-2

Baldassano, C. (2023). Studying waves of prediction in the brain using narratives. *Neuropsychologia*, *189*, 108664. https://doi.org/10.1016/j.neuropsychologia.2023.108664

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

Baldassano, C., Hasson, U., & Norman, K. A. (2018). Representation of Real-World Event Schemas during Narrative Perception. *The Journal of Neuroscience*, *38*(45), 9689–9699. https://doi.org/10.1523/JNEUROSCI.0251-18.2018

Baldwin, D. A., & Kosie, J. E. (2021). How Does the Mind Render Streaming Experience as Events? *Topics in Cognitive Science*, *13*(1), 79–105. https://doi.org/10.1111/tops.12502

Bein, O., & Davachi, L. (2022). *Event integration and temporal differentiation: How hierarchical knowledge emerges in hippocampal subfields through learning* [Preprint]. Neuroscience. https://doi.org/10.1101/2022.07.18.500527

Ben-Yakov, A., & Dudai, Y. (2011). Constructing Realistic Engrams: Poststimulus Activity of Hippocampus and Dorsal Striatum Predicts Subsequent Episodic Memory. *Journal of Neuroscience*, *31*(24), 9032–9042. https://doi.org/10.1523/JNEUROSCI.0702-11.2011

Brodeur, M. B., Guérard, K., & Bouras, M. (2014). Bank of Standardized Stimuli (BOSS) Phase II: 930 New Normative Photos. *PLoS ONE*, *9*(9), e106953. https://doi.org/10.1371/journal.pone.0106953

Buonomano, D. V., Buzsáki, G., Davachi, L., & Nobre, A. C. (2023). Time for Memories. *The Journal of Neuroscience*, *43*(45), 7565–7574. https://doi.org/10.1523/JNEUROSCI.1430-23.2023

Clewett, D., & Davachi, L. (2017). The ebb and flow of experience determines the temporal structure of memory. *Current Opinion in Behavioral Sciences*, *17*, 186–193. https://doi.org/10.1016/j.cobeha.2017.08.013

Clewett, D., DuBrow, S., & Davachi, L. (2019). Transcending time in the brain: How event memories are constructed from experience. *Hippocampus*, *29*(3), 162–183. https://doi.org/10.1002/hipo.23074

Davachi, L., & DuBrow, S. (2015). How the hippocampus preserves order: The role of prediction and context. *Trends in Cognitive Sciences*, *19*(2), 92–99.

https://doi.org/10.1016/j.tics.2014.12.004

De Soares, A., Kim, T., Mugisho, F., Zhu, E., Lin, A., Zheng, C., & Baldassano, C. (2023). *Top-down attention shifts behavioral and neural event boundaries in narratives with overlapping event scripts*. https://doi.org/10.1101/2023.08.08.552465

DuBrow, S., & Davachi, L. (2013). The influence of context boundaries on memory for the sequential order of events. *Journal of Experimental Psychology: General*, *142*(4), 1277–1286. https://doi.org/10.1037/a0034024

DuBrow, S., & Davachi, L. (2014). Temporal Memory Is Shaped by Encoding Stability and Intervening Item Reactivation. *The Journal of Neuroscience*, *34*(42), 13998–14005. https://doi.org/10.1523/JNEUROSCI.2535-14.2014

DuBrow, S., & Davachi, L. (2016). Temporal binding within and across events. *Neurobiology of Learning and Memory*, *134*, 107–114. https://doi.org/10.1016/j.nlm.2016.07.011

Ezzyat, Y., & Davachi, L. (2011). What Constitutes an Episode in Episodic Memory? *Psychological Science*, *22*(2), 243–252. https://doi.org/10.1177/0956797610393742

Güler, B., Adıgüzel, Z., Uysal, B., & Günseli, E. (2024). Discrete memories of a continuous world: A working memory perspective on event segmentation. *Current Research in Behavioral Sciences*, *6*, 100145. https://doi.org/10.1016/j.crbeha.2023.100145

Hasson, U., Chen, J., & Honey, C. J. (2015). Hierarchical process memory: Memory as an integral component of information processing. *Trends in Cognitive Sciences*, *19*(6), 304–313. https://doi.org/10.1016/j.tics.2015.04.006

Hebart, M. N., Dickter, A. H., Kidder, A., Kwok, W. Y., Corriveau, A., Van Wicklin, C., & Baker, C. I. (2019). THINGS: A database of 1,854 object concepts and more than 26,000 naturalistic object images. *PLOS ONE*, *14*(10), e0223792. https://doi.org/10.1371/journal.pone.0223792

Heusser, A. C., Ezzyat, Y., Shiff, I., & Davachi, L. (2018). Perceptual boundaries cause mnemonic trade-offs between local boundary processing and across-trial associative binding. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(7), 1075–1090. https://doi.org/10.1037/xlm0000503

Heusser, A. C., Poeppel, D., Ezzyat, Y., & Davachi, L. (2016). Episodic sequence memory is supported by a theta–gamma phase code. *Nature Neuroscience*, *19*(10), 1374–1380. https://doi.org/10.1038/nn.4374

Huang, J., Velarde, I., Ma, W. J., & Baldassano, C. (2023). Schema-based predictive eye movements support sequential memory encoding. *eLife*, *12*, e82599. https://doi.org/10.7554/eLife.82599

Jafarpour, A., Buffalo, E. A., Knight, R. T., & Collins, A. G. E. (2022). Event segmentation reveals working memory forgetting rate. *iScience*, *25*(3), 103902. https://doi.org/10.1016/j.isci.2022.103902

Kumar, M., Goldstein, A., Michelmann, S., Zacks, J. M., Hasson, U., & Norman, K. A. (2023). Bayesian Surprise Predicts Human Event Segmentation in Story Listening. *Cognitive Science*, *47*(10), e13343. https://doi.org/10.1111/cogs.13343

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

Lawrence, Z., & Peterson, D. (2016). Mentally walking through doorways causes forgetting: The location updating effect and imagination. *Memory*, *24*(1), 12–20. https://doi.org/10.1080/09658211.2014.980429

Lee, C. S., Aly, M., & Baldassano, C. (2021). Anticipation of temporally structured events in the brain. *eLife*, *10*, e64972. https://doi.org/10.7554/eLife.64972

Lichtenstein, E. H., & Brewer, W. F. (1980). Memory for goal-directed events. *Cognitive Psychology*, *12*(3), 412–445. https://doi.org/10.1016/0010-0285(80)90015-8

Nolden, S., Turan, G., Güler, B., & Günseli, E. (2024). Prediction error and event segmentation in episodic memory. *Neuroscience & Biobehavioral Reviews*, *157*, 105533. https://doi.org/10.1016/j.neubiorev.2024.105533

Ongchoco, J. D. K., Walter-Terrill, R., & Scholl, B. J. (2023). Visual event boundaries restrict anchoring effects in decision-making. *Proceedings of the National Academy of Sciences*, *120*(44), e2303883120. https://doi.org/10.1073/pnas.2303883120

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. https://doi.org/10.3758/s13428-018-01193-y

Pettijohn, K. A., & Radvansky, G. A. (2018). Walking through doorways causes forgetting: Active and passive interaction. *Journal of Cognitive Psychology*, *30*(8), 771–777. https://doi.org/10.1080/20445911.2018.1540492

Radvansky, G. A. (2017). Event segmentation as a working memory process. *Journal of Applied Research in Memory and Cognition*, *6*(2), 121–123. https://doi.org/10.1016/j.jarmac.2017.01.002

Radvansky, G. A., & Copeland, D. E. (2006). Walking through doorways causes forgetting: Situation models and experienced space. *Memory & Cognition*, *34*(5), 1150–1156. https://doi.org/10.3758/BF03193261

Radvansky, G. A., Krawietz, S. A., & Tamplin, A. K. (2011). Walking through Doorways Causes Forgetting: Further Explorations. *Quarterly Journal of Experimental Psychology*, *64*(8), 1632–1645. https://doi.org/10.1080/17470218.2011.571267

Reagh, Z. M., Delarazan, A. I., Garber, A., & Ranganath, C. (2020). Aging alters neural activity at event boundaries in the hippocampus and Posterior Medial network. *Nature Communications*, *11*(1), 3980. https://doi.org/10.1038/s41467-020-17713-4

Richmond, L. L., & Zacks, J. M. (2017). Constructing Experience: Event Models from Perception to Action. *Trends in Cognitive Sciences*, *21*(12), 962–980. https://doi.org/10.1016/j.tics.2017.08.005

Rouhani, N., Norman, K. A., Niv, Y., & Bornstein, A. M. (2020). Reward prediction errors create event boundaries in memory. *Cognition*, *203*, 104269.

https://doi.org/10.1016/j.cognition.2020.104269

Sasmita, K., & Swallow, K. M. (2022). Measuring event segmentation: An investigation into the stability of event boundary agreement across groups. *Behavior Research Methods*, *55*(1), 428–447. https://doi.org/10.3758/s13428-022-01832-5

Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience*, *16*(4), 486–492. https://doi.org/10.1038/nn.3331

Seel, S. V., Easton, A., McGregor, A., Buckley, M. G., & Eacott, M. J. (2019). Walking through doorways differentially affects recall and familiarity. *British Journal of Psychology*, *110*(1), 173–184. https://doi.org/10.1111/bjop.12343

Shin, Y. S., & DuBrow, S. (2021). 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. *The Journal of Neuroscience*, *39*(43), 8538–8548. https://doi.org/10.1523/JNEUROSCI.0360-19.2019

Swallow, K. M., Barch, D. M., Head, D., Maley, C. J., Holder, D., & Zacks, J. M. (2011). Changes in Events Alter How People Remember Recent Information. *Journal of Cognitive Neuroscience*, *23*(5), 1052–1064. https://doi.org/10.1162/jocn.2010.21524

Trabasso, T., & Van Den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language*, *24*(5), 612–630. https://doi.org/10.1016/0749-596X(85)90049-X

Wang, V., Ongchoco, J. D. K., & Scholl, B. J. (2023). Here it comes: Active forgetting triggered even just by anticipation of an impending event boundary. *Psychonomic Bulletin & Review*. https://doi.org/10.3758/s13423-023-02278-2

Wang, Y. C., Adcock, R. A., & Egner, T. (2024). Toward an integrative account of internal and external determinants of event segmentation. *Psychonomic Bulletin & Review*, *31*(2), 484–506. https://doi.org/10.3758/s13423-023-02375-2

Wang, Y. C., & Egner, T. (2022). Switching task sets creates event boundaries in memory. *Cognition*, *221*, 104992. https://doi.org/10.1016/j.cognition.2021.104992

Wu, X., Viñals, X., Ben-Yakov, A., Staresina, B. P., & Fuentemilla, L. (2022). Post-encoding Reactivation Is Related to Learning of Episodes in Humans. *Journal of Cognitive Neuroscience*, *35*(1), 74–89. https://doi.org/10.1162/jocn_a_01934

Zacks, J. M., Kurby, C. A., Eisenberg, M. L., & Haroutunian, N. (2011). Prediction Error Associated with the Perceptual Segmentation of Naturalistic Events. *Journal of Cognitive Neuroscience*, *23*(12), 4057–4066. https://doi.org/10.1162/jocn_a_00078

Zacks, J. M., Speer, N. K., & Reynolds, J. R. (2009). Segmentation in reading and film comprehension. *Journal of Experimental Psychology: General*, *138*(2), 307–327. https://doi.org/10.1037/a0015305

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

Zwaan, R. A., Magliano, J. P., & Graesser, A. C. (1995). Dimensions of situation model construction in narrative comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(2), 386–397. https://doi.org/10.1037/0278-7393.21.2.386

Supplementary

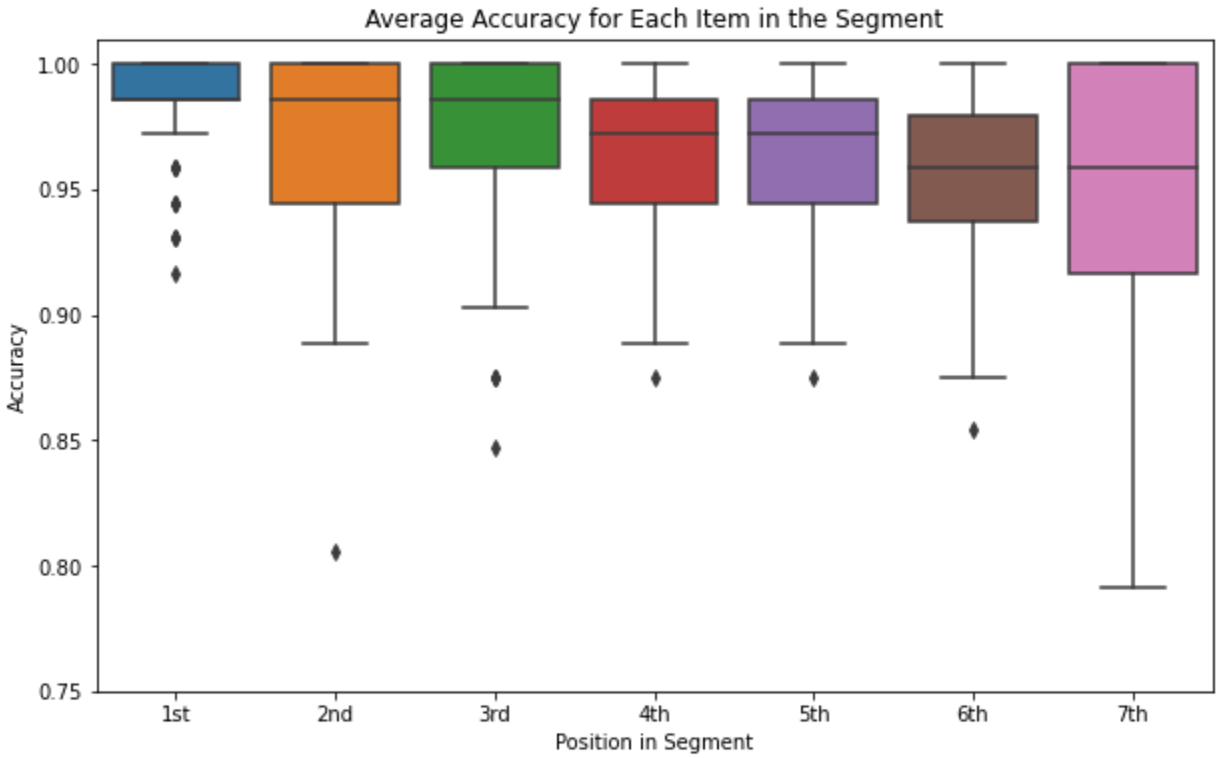

Supplementary Figure 1: Participant accuracies on new/repeat task in Experiment 1, separated by position within each segment. Performance was similar across positions 2-7, with very high performance at position 1 (for which the correct response was always "new"). Note that positions 6 and 7 were only present in 2/3 and 1/3 of segments, respectively, leading to increased variability in the accuracy estimates for these positions.