TOPICS
TOPICS IN COGNITIVE SCIENCE

Topics in Cognitive Science (2020) 1–22
© 2020 Cognitive Science Society, Inc. All rights reserved.
ISSN:1756-8765 online
DOI: 10.1111/tops.12505

This article is part of the topic "Event-Predictive Cognition: From Sensorimotor via Conceptual to Language-Based Structures and Processes," Martin V. Butz, David Bilkey, and Alistair Knott (Topic Editors). For a full listing of topic papers, see http://onlinelibrary.wiley.com/journal/10.1111/(ISSN)1756-8765/earlyview

Structuring Memory Through Inference-Based Event Segmentation

Yeon Soon Shin,[a] Sarah DuBrow[b]

[a]*Princeton Neuroscience Institute, Princeton University*
[b]*Department of Psychology, University of Oregon*

Received 29 March 2020; received in revised form 29 March 2019; accepted 14 April 2020

Abstract

Although the stream of information we encounter is continuous, our experiences tend to be discretized into meaningful clusters, altering how we represent our past. Event segmentation theory proposes that clustering ongoing experience in this way is adaptive in that it promotes efficient online processing as well as later reconstruction of relevant information. A growing literature supports this theory by demonstrating its important behavioral consequences. Yet the exact mechanisms of segmentation remain elusive. Here, we provide a brief overview of how event segmentation influences ongoing processing, subsequent memory retrieval, and decision making as well as some proposed underlying mechanisms. We then explore how beliefs, or inferences, about what generates our experience may be the foundation of event cognition. In this inference-based framework, experiences are grouped together according to what is inferred to have generated them. Segmentation then occurs when the inference changes, creating an event boundary. This offers an alternative to dominant theories of event segmentation, allowing boundaries to occur independent of perceptual change and even when transitions are predictable. We describe how this framework can reconcile seemingly contradictory empirical findings (e.g., memory can be biased toward both extreme episodes and the average of episodes). Finally, we discuss open questions regarding how time is incorporated into the inference process.

Correspondence should be sent to Sarah DuBrow, Department of Psychology, University of Oregon, Eugene, OR 97403. E-mail: sdubrow@uoregon.edu or Yeon Soon Shin, Princeton Neuroscince Institute, Princeton University, Princeton, NJ 08544. E-mail: yshin@princeton.edu

Keywords: Event segmentation; Temporal context; Situation models; Latent cause inference; Episodic sampling; Episodic memory

1. Introduction

Humans have a natural tendency to segment the continuous stream of incoming information we experience into discrete events, broadly defined as units of activity with an identifiable beginning and end (Zacks, Speer, Swallow, Braver, & Reynolds, 2007). People can detect the transitions between events, called event boundaries, while reading (Speer & Zacks, 2005) and watching films (Newtson, 1973), as well as during more naturalistic first-person experiences (Magliano, Radvansky, Forsythe, & Copeland, 2014). Event segmentation may occur spontaneously, as evidenced by longer processing times at boundaries in the absence of a segmentation task (e.g., Hard, Recchia, & Tversky, 2011; Speer & Zacks, 2005). In addition, neuroimaging data have shown that the brain responds to boundaries during passive viewing (Baldassano et al., 2017; Ben-Yakov & Henson, 2018; Speer, Zacks, & Reynolds, 2007; Zacks, Braver, et al., 2001), representing information within an event more similarly than across events (Chen et al., 2017). Segmentation behavior tends to be consistent across people both behaviorally (Jeunehomme & D'Argembeau, 2018; Newtson, 1973; Zacks, Tversky, & Iyer, 2001) and neurally (Baldassano et al., 2017; Ben-Yakov & Henson, 2018; Speer, Swallow, & Zacks, 2003), suggesting that event structure is construed in a systematic way. Moreover, segmentation has been shown to have important behavioral consequences such as enhancing memory for items encountered at boundaries (Heusser, Ezzyat, Shiff, & Davachi, 2018; Swallow, Zacks, & Abrams, 2009) and warping time perception such that intervals with boundaries are estimated as longer in memory (Ezzyat & Davachi, 2014; Lositsky et al., 2016). Boundaries may also provide an opportunity to consolidate information from the previous event (Bilkey & Jensen, 2020). Here, we first review the behavioral effects of event segmentation, then we review potential mechanisms of event segmentation, and finally we explore the role of inference as a framework for event cognition.

2. Why do we segment events?

Segmenting experience into distinct events has been shown to have extensive psychological consequences (see Table 1). In this section, we review three domains affected by event segmentation: (a) in-the-moment processing of *ongoing* experiences, (b) memory organization of *past* experiences, and (c) making decisions that best serve the current situation. The effects of event segmentation observed in these domains demonstrate a benefit for adaptive behavior.

2.1. Facilitation of ongoing event processing

When encountering incoming information, event segmentation can facilitate processing by increasing access to the event that is currently being experienced. One measure that

Table 1
Behavioral effects of event segmentation by boundary manipulations

Boundary Manipulation	Study	Behavioral Effect(s)
		Reading time
Narrative changes	Zwaan, Magliano, and Graesser (1995)	Temporal and causal change > no change
	Zwaan, Radvansky, Hilliard, and Curiel (1998)	More situational change > less
	Rinck and Weber (2003)	Temporal and protagonist change > no change
	Zacks et al. (2009)	More situational change > less
	Radvansky and Copeland (2010)	Temporal change > no change
	McNerney, Goodwin, and Radvansky (2011)	Causal and character change > no change > spatial and temporal change
	Pettijohn and Radvansky (2016)	Unexpected change > expected or no change
Narrative time changes	Zwaan (1996)	Temporal change > no change
	Speer and Zacks (2005)	Temporal change > no change
		Memory access
Narrative time changes	Zwaan (1996)	Within > across recognition (online and delayed)
	Speer and Zacks (2005)	Within > across recognition (online)
	Ezzyat and Davachi (2011)	Within > across cued recall
Activity change (video)	Swallow et al. (2009)	Within > across nonboundary recognition (online); Across > within boundary recognition (online)
	Swallow et al. (2011)	Within > across nonboundary recognition (online); Across > within boundary recognition (online)
Task and category change	DuBrow and Davachi (2013)	Within > across order memory and serial recall
	DuBrow and Davachi (2014)	Within > across order memory
	DuBrow and Davachi (2016)	Within > across serial recall
Virtual room change	Horner et al. (2016)	Within > across sequence recognition
Background color change	Heusser et al. (2018)	Within > across order memory
Turns in virtual navigation	Brunec et al. (2020)	Within > across order memory; Across > within duration discrimination
		Spontaneous clustering
Narrative changes	Zwaan, Magliano, et al. (1995)	Within > across verb clustering
Task change	Polyn et al. (2009b)	Within > across recall transitions
Background color change	Heusser et al. (2018)	Within > across recall transitions
		Prediction
Narrative changes	Zacks et al. (2009)	Within > across perceived predictability
	Pettijohn and Radvansky (2016)	Within > across expectedness ratings
Activity change (video)	Zacks et al. (2011)	Within > across prediction accuracy

has been used to assess facilitated processing is how long it takes to read narrative passages with and without event boundaries. Reading time increases have been observed in studies that explicitly signal time shifts (e.g., “an hour/day later” vs. “a moment later,” Speer & Zacks, 2005; Zwaan, 1996; “the next morning,” Pettijohn & Radvansky, 2016) as well as protagonist changes (Rinck & Weber, 2003). These results suggest that narrative comprehension is facilitated within events versus across event boundaries, perhaps because within-event content is predicted by learned event schemas that involve a sequences of states (Franklin, Norman, Ranganath, Zacks, & Gershman, 2019; Radvansky & Zacks, 2011) or situation models (Zwaan & Radvansky, 1998). Zacks, Speer, and Reynolds (2009) explicitly probed event segmentation and predictability and found that narrative passages rated low in predictability showed longer reading times and more identified boundaries. Moreover, consistent with a role for predictability in mediating event segmentation, Pettijohn and Radvansky (2016) showed that reading times do not slow down at shifts when they are foreshadowed (i.e., predicted). These data suggest that predictable content that belongs to the same event facilitates rapid comprehension.

Similar to the reading time data, when presented with still frames of action sequences, people tend to dwell longer on transition frames, suggesting increased processing demands at boundaries (Hard et al., 2011; reviewed in Baldwin & Kosie, 2020). Interestingly, however, predictability generally enhances this dwell time effect. That is, the dwell time difference between actions at boundaries and within-event actions increases as a function of experience (Hard, Meyer, & Baldwin, 2019; Kosie & Baldwin, 2019). This exaggeration may be driven both by reduced processing time within events as the predictability of actions increases, as well as by longer dwell times at boundaries as the anticipation of change triggers viewers to gather more information (see Baldwin & Kosie, 2020, for an information optimization account of ongoing event processing).

When processing ongoing information, it may be helpful to have selective access to currently relevant information. Event segmentation can help prioritize information relevant to the currently active event, while making the information from previous events less accessible. This has been demonstrated in paradigms that interleave narrative reading or movie watching with recognition memory tests. When there is an event boundary between the encoding of the probe and the recognition test, people recognize the probe more slowly (Speer & Zacks, 2005; Zwaan, 1996) and less accurately (Speer & Zacks, 2005; Swallow et al., 2011). In a similar set of studies in which people experienced event boundaries by walking through doorways, items learned in a previous room became less accessible (Radvansky & Copeland, 2006; Radvansky, Krawietz, & Tamplin, 2011). This suggests that when an event ends, the information that was learned within that event drops out of active working memory because it is no longer relevant. Updating an active event in working memory (i.e., event model) in this way is one of the principles of the Event Horizon framework (Radvansky, 2012; Radvansky & Zacks, 2017) and has beneficial effects for online processing (e.g., reduced reading time and increased working memory access) of incoming information within the same event. Dropping previous information out of working memory may have the added benefit of preventing that irrelevant information from intruding on the current processing. This may also protect information in the previous event from retroactive interference, increasing the accuracy of later reconstruction (Gershman, Radulescu, Norman, & Niv, 2014).

2.2. Memory organization

Segmenting events can also support the encoding and retrieval of episodic memories. One line of evidence that event segmentation helps episodic memory is that, when items belong to the same event, they show mutually facilitated recognition memory. For instance, when people read multiple sentences, some of which share a location, they are faster to recognize the sentences that share a common location compared to sentences that do not (Radvansky & Zacks, 1991; Radvansky, Zwaan, Federico, & Franklin, 1998). This benefit may be driven by recent activation of a sentence that could prime retrieval of other episodes in the same event. One study that supports incidental retrieval of within-event items used a subset of a previously studied group of words as a target for an unrelated memory task, and showed that a lure from the same group was more likely to be falsely recognized than a lure from a different group (Hoskin, Bornstein, Norman, & Cohen, 2018). While this demonstrates that it can come at a cost at times, having segmented structure in memory can keep related information together, facilitating later retrieval.

Memory studies have also directly probed whether people are more likely to retrieve items that belong to the same event together. Analogous to the suppression of previous event information observed during ongoing processing, retrieving events from episodic memory that have a boundary between them can be more challenging than retrieving information from the same event. Zwaan (1996) used a cued recognition paradigm where a sentence that was followed by a time shift signal phrase served as a cue to facilitate recognition speed of the next target sentence. Event segmentation was manipulated by shift magnitudes ("a day/hour/moment later"). Consistent with the prediction that event segmentation enhances retrieval of items in the same event and diminishes retrieval of items in a different event, people were slower to recognize the target sentence after a larger shift. Using a similar time shift signal ("an hour later" vs. "a moment later"), Ezzyat and Davachi (2011) asked people to recall what came after a cued sentence. In the large shift condition, recall performance was lower when a pre-shift sentence was used as a cue as compared with when a post-shift sentence served as a cue. However, a difference was not observed when the time shift was small. Complimenting the online predictability effects discussed in the previous section, these results suggest a mechanism by which retrieving an episode from long-term memory can cue the next episodes that occurred within the same event, guiding predictions of what will happen next.

Relatedly, studies that probe the temporal order of items that either belong to the same event or different events provide additional evidence for facilitated within-event retrieval. In DuBrow and Davachi (2013, 2014), subjects judged the relative recency of two items within or across boundaries that were created by switching stimulus categories and their associated tasks (e.g., male/female judgment for the face category; bigger/smaller than shoebox judgment for the object category). Recency judgments between two studied items were less accurate when the intervening sequence contained boundaries. To examine whether this performance drop was due to retrieval failure for the intervening items, recognition memory for those intervening items was tested immediately after (i.e., primed by) the order judgments. Consistent with the aforementioned studies that looked at cued recall and cued recognition, when people made a correct order judgment, the speed at which they recognized the intervening items was faster when there was no event boundary, suggesting greater within-event access to item sequences. The within-event versus across-events difference in temporal order memory has also been shown in more recent studies that used perceptual boundaries (background color changes, Heusser et al., 2018) and spatial boundaries (room changes, Horner, Bisby, Wang, Bogus, & Burgess, 2016; turns in navigation, Brunec et al., 2020). These experiments show that temporal information is better preserved within an event than across events, potentially via better reconstruction of a study sequence.

Another way to test how individuals reconstruct sequences from memory is to examine transition probabilities in verbal recall. In one study, people were asked to recall items in the same order that they were studied. Accurate serial transitions between recalled items were found to be more common within than across category boundaries, providing additional support for the better reconstruction of a study sequence within an event (DuBrow & Davachi, 2016). Similarly, in unconstrained free recall studies, event-level clustering (Polyn, Norman, & Kahana, 2009b) and a tendency toward more forward serial transitions within events compared to across boundaries (Heusser et al., 2018) have been observed. Since recall order was unconstrained, these results suggest that event-level organization may be a fundamental property of recall. That is, event structure may provide a scaffold for spontaneously recalling past experiences in their sequential order.

2.3. *Adaptive decision making*

Interestingly, this sequential recall closely resembles sequential reactivation in the hippocampus (Foster & Wilson, 2007), in which event structure has been observed while an experience unfolds. In particular, Gupta, van der Meer, Touretzky, and Redish (2012) showed that hippocampal activation reflects the segment of the environment that is currently being navigated (i.e., the event model), disproportionately representing paths ahead within the segment in the beginning and paths behind within the segment as they approach its end. This activation of currently relevant information has implications for decision making, where generalizing relevant past experiences can guide decisions for unknown possible futures. Indeed, an extensive literature on rodent navigation and decision making has shown that hippocampal activation of forward trajectories occurs *preceding* decisions about where to go next (Johnson & Redish, 2007; Pfeiffer & Foster, 2013; for reviews, see Ólafsdóttir, Bush, & Barry, 2018; Pezzulo, van der Meer, Lansink, & Pennartz, 2014). Similar neural reinstatement effects, both during rest (Momennejad, Otto, Daw, & Norman, 2018; Schuck & Niv, 2019) and prospectively at decision times (Doll, Duncan, Simon, Shohamy, & Daw, 2015), have also been shown to influence subsequent decision making in humans.

Although traditionally decision making has been viewed as relying on a representation of incrementally learned value that is independent of episodic memory (Knowlton, Mangels, & Squire, 1996; cf. Poldrack et al., 2001), replay of past experiences at decision points suggests the potential contribution of episodic memories. Indeed, a growing literature has shown that adaptive decisions are influenced by episodic memory retrieval (Shadlen & Shohamy, 2016; Shohamy & Daw, 2015). For example, people are more likely to choose previously encountered items that were associated with high values than low values only when they remember such associations (Murty, FeldmanHall, Hunter, Phelps, & Davachi, 2016), and they may rely more on individual episodes over summary values for decision making following memory retrieval (Duncan, Semmler, & Shohamy, 2019). Linking the roles of episodic memory and event segmentation, Bornstein and Norman (2017) showed that, when reminded of an image experienced in a previous event, people's decisions are biased by the summary of their experiences in that event, not just the specific experience associated with the reminded image. This result suggests that event segmentation can support the interaction between episodic memory and decision making, by guiding retrieval of past decision outcomes from previous events that most closely match the current situation.

Together, these studies suggest that organizing information into event structures can have important benefits for online processing and retrieval of relevant information that, in turn, can help guide adaptive decision making. Given the widespread effects of segmentation, it is crucial to appropriately segment our everyday experiences into events in a way that promotes later utilization. One of the major challenges in doing so is the inherent ambiguity of when an event starts and ends. Many of the studies reviewed showing the benefits for memory and decision making of event segmentation cannot address this challenge, as they imposed stark changes in spatial contexts, perceptual features, task sets, and so forth to manipulate segmentation (see Table 1). Experiments that use narratives to induce event boundaries are more ambiguous because there can be multiple changes along different dimensions (cf. event indexing theory; Zwaan, Langston, & Graesser, 1995), but they still often contain signals for boundaries such as time shifts or scene changes. When event boundaries are not overtly signaled, how do we segment events? Below we review theories of event segmentation under ambiguity. We focus on the traditional prediction error account and our proposed framework that identifies boundaries based on changes in inference rather than extrinsic change.

3. Potential mechanisms of event segmentation

3.1. Prediction error

The dominant account of event segmentation is that "prediction error," the difference between one's experience-based expectation and the currently observed outcome, signals the end of events and induces event boundaries (Zacks et al., 2007). To test how explicit predictions are related to event segmentation, Zacks, Kurby, Eisenberg, and Haroutunian (2011) showed people movie clips of everyday activities, and occasionally paused the clip to ask them to predict what would happen in 5 s. When there was an intervening event boundary, as identified by independent observers, the prediction accuracy dropped, and there was greater activation in the substantia nigra, a region traditionally associated with dopaminergic responses to reward prediction errors. In line with this idea, a neural network model that implements perceptual prediction error as a gating signal to update event representations can identify simulated event boundaries (Reynolds, Zacks, & Braver, 2007). However, there are major challenges regarding the precise relationship between prediction errors and event segmentation.

First, prediction errors may not always signal meaningful changes in event perception, particularly when the environment is uncertain (O'Reilly, 2013). That is, while it can be ideal to draw boundaries in order to discount the past when the underlying statistics of the world change abruptly, disregarding the broader environmental context would be suboptimal when the environment is noisy, as frequent high prediction errors would lead to over-segmentation. Instead, when prediction errors are frequent, boundaries should only be drawn following *unexpected* deviations, as expected deviations are not reflective of meaningful changes in the underlying event structure. This can be accomplished by scaling down the degree to which prediction error updates expectations (i.e., the learning rate) for these expected deviations. Empirically, boundary-related memory effects that occur when changes are infrequent are absent when changes occur frequently (Siefke, Smith, & Sederberg, 2019). This idea is further reflected in studies where people's learning rates dynamically adjust to the uncertainty of the environment, reducing the effects of prediction errors in noisy environments while enhancing them in stable environments, in which a relatively high prediction error indicates meaningful change (Behrens, Woolrich, Walton, & Rushworth, 2007; Nassar et al., 2012; Nassar, Wilson, Heasly, & Gold, 2010; Pearce & Hall, 1980). This suggests that the magnitude of the update (i.e., prediction error $\times$ learning rate) may be a better indicator of event boundaries than prediction error alone, given the non-stationary nature of everyday experiences.

Second, prediction errors may not be necessary to create an event boundary. For example, in narrative reading, when an event shift is foreshadowed, readers still respond more slowly to the pre-shift memory probes even though they no longer experience surprise nor slow their reading (Pettijohn & Radvansky, 2016). This suggests that even expected change, when sufficiently meaningful, can drive event segmentation in memory. One demonstration of this is statistical learning (Saffran, Aslin, & Newport, 1996). In one study, Schapiro, Rogers, Cordova, Turk-browne, and Botvinick (2013) found that people could identify boundaries in a series of stimuli based on their learned temporal transition statistics without ever experiencing prediction errors. In their experiments, most stimuli were exclusively followed by other stimuli within the same group, while some stimuli served as entry/exit points where a transition across groups was available. Notably, the individual transition probability from a group's exit point to a neighboring group's entry point was equal to the probability of each within-group transition, and thereby did not induce prediction errors (see also Richmond & Zacks, 2017, for discussion). After learning, participants successfully indicated the transitions between groups (i.e., the entry points) in the sequence of stimuli. This highlights the importance of segmenting experience even at predictable transitions between events. Indeed, predictable transitions may enhance segmentation effects by shifting attentional resources to boundaries (Baldwin & Kosie, 2020).

As reviewed here, prediction error defined as the difference between one's expectation and the observed outcome may not be sufficient or necessary for segmenting events. This calls for an alternative account of the mechanisms underlying event segmentation.

3.2. *Inference of event types*

No two events can be exactly the same even when they are from the same category, as at least one dimension, time, is always unique. For example, when items occur in two different instances of the same category (i.e., ABA structure), there is an intervening change in event type, and the two visits to A are idiosyncratic in terms of time. When comparing this scenario to a single continuous instance of an event (AAA structure), segmentation effects are observed despite comparing two items of the same type (DuBrow & Davachi, 2013, 2014). Radvansky et al. (2011) used a more naturalistic boundary manipulation with the same structure by having people walk around a series of rooms and then testing their memory. When returning to the same room after having been to a different room (ABA), their memory was worse than when they never left the room at all (AAA). Thus, staying in the same event instance (no event boundaries) has a memory advantage over merely having shared context (room A). This suggests that boundaries affect memory above and beyond the effects of changing context.

Time-sensitive and idiosyncratic *event instances* do interact with the structure of knowledge formed by multiple encounters with similar situations. In Radvansky et al. (2011), people also performed better when they returned to the same room (ABA) compared to when they went to another room (ABC). Thus, the *event type* that represents a class of experiences (called event schemata in Event Segmentation Theory; Zacks et al., 2007) is a key factor in understanding segmentation (for reviews, see Radvansky & Zacks, 2011; Zacks et al., 2007). For example, in the statistical learning study described in the previous section, stimuli that tended to be experienced in close temporal proximity became discrete clusters (event types) in memory (Schapiro et al., 2013). When revisited, these clusters could be recognized even though different paths were experienced across learning instances due to the probabilistic structure. These studies raise the possibility that event types can be a useful basis for promoting generalization of learning within type, allowing us to extrapolate our previous experiences in an event to a new instance.

Event types can also stabilize event segmentation hierarchically such that, in parsing events, low-level perceptual changes that are not relevant to the event type carry less weight than changes along dimensions that are pertinent to the event type. In a neural network application of this idea, the REPRISE model stabilizes low-level perceptual and motor information to be consistent with the current event type (called a "context vector") when executing event-level control (Butz, Bilkey, Humaidan, Knott, & Otte, 2019). Empirically, hierarchical structures in event segmentation are consistently observed in behavioral (Zacks, Tversky, et al., 2001) and brain imaging studies (Baldassano et al., 2017; Hasson, Yang, Vallines, Heeger, & Rubin, 2008; Lerner, Honey, Silbert, & Hasson, 2011), supporting the idea that event types are utilized in segmenting experience.

It is important to note that these event types do not exist in isolation and are continuously created and updated based on event instances. How do we dynamically update and create new event types? To inform this, we turn to category learning, where a similar set of challenges exists. Categories, like event types, help facilitate understanding of their individual members through generalization of category properties. However, as in event perception, we do not know how many categories exist in the world, and we need to update the existing categories' summary statistics based on their members, while being open to creating a completely new category. The rational model of categorization addresses these challenges by proposing that a stimulus is more likely to be a member of a prolific category with many members, and yet a new category can be inferred at any point when the existing categories are not a good match (Anderson, 1991; also called the Chinese Restaurant Process [CRP], Aldous, 1985). In discovering clusters such as categories from observations, this rich-gets-richer process reins in the tendency to create too many clusters and thus keeps the model simple while still allowing for new cluster formation. The adaptable nature of the model aids predictions within categories by extrapolating features and functions within the categories, whether or not an explicit label was given to a new category (Anderson, 1991). Thanks to its flexibility, this model can be generalized to two popular models of categorization where instances of a category can either (a) be lumped together such that one summary value can account for the category (prototype model; Reed, 1972) or (b) be perfectly preserved to be later compared with other instances (exemplar model; Medin & Schaffer, 1978; Nosofsky, 1986), by varying a single parameter that governs creation of a new category (Sanborn, Griffiths, & Navarro, 2010). The rational model provides a unified framework that successfully predicts categorization behavior in humans across domains (Anderson, 1991; Sanborn et al., 2010).

The model is particularly useful in situations where there is a latent variable (i.e., a cluster of experiences such as a category or event type) that needs to be discovered to properly generalize across experiences. For example, the model can explain how people effectively generalize between tasks by forming clusters of task sets (Collins & Frank, 2013). In reinforcement learning, the model can explain how the latent variable can guide reward predictions, and how a new latent variable can be inferred when existing ones do not predict the outcome (Courville, Daw, & Touretzky, 2005). This model performs better than classic reinforcement learning models in explaining how compound reward cues are flexibly represented (Courville et al., 2005; Soto, Gershman, & Niv, 2014) and why conditioned responses come back after extinction (Gershman, Blei, & Niv, 2010). In event perception, the latent variable would correspond to event types. That is, we can identify an event type from a time-specific event instance, updating the type based on the instance, or create a new event type when an instance does not fit with any of the existing event types' properties. The latent variable model would also predict that, due to its rich-gets-richer property, we are more likely to infer an event type that we encounter often (i.e., high prior probability) when it sufficiently explains the current episode (i.e., high likelihood), rather than inferring a new event type de novo. In this framework, event segmentation is likely to occur when the distribution over inferred event types diverges from the distribution at the previous time point. Event segmentation based on distributional changes can explain how experiences are clustered into units even when there is no prediction error between experiences, as in statistical learning-based clustering described above.

A demonstration of this inference process is depicted in Fig. 1. As we watch a movie, we learn that the character Julie recently lost her husband Patrice, a famous composer, in an accident. Later, Julie discovers that Patrice's assistant Olivier is trying to complete one of Patrice's unfinished scores *and* that Patrice had a mistress. When we see Julie approaching Olivier (Fig. 1.a-1), we do not know whether she will confront him about the score or the mistress, and the probabilities of the two event types are uniformly distributed (Fig. 1.c *leftmost panel*). When Julie and Olivier have a conversation about the unfinished score (Fig. 1.a-2) and then the movie jumps to Olivier playing the song (Fig. 1.a-3), the same "unfinished score" event type is active despite the perceptual input and spatial context having changed (denoted by the change in background color in Fig. 1.c). Note that the probability of event segmentation is low despite the location change (denoted by a gray dotted line in Fig. 1.c,d). Conversely, event segmentation can happen without big perceptual changes (denoted by black dotted lines in Fig. 1.c,d). In *scene 4*, the prior for the "unfinished score" is higher than "mistress conversation" because of the greater number of previous instances of the "unfinished score" event type (denoted by the number of boxes in Fig. 1.b left) and the model's rich-gets-richer property. However, despite occurring at the same location, as Julie diverts her attention away from the piano to ask Olivier whether he knew about the mistress, the likelihood of the "unfinished score" event becomes lower while the likelihood of the "mistress conversation" becomes higher (Fig. 1.b center). Thus, the updated posterior probability shifts toward "mistress conversation" (Fig. 1.b right). In the proposed framework, the divergence between the updated probability distribution and that of the previous time point increases the probability of event segmentation (Fig. 1.d).

Inference-based segmentation can be useful when there is high uncertainty about the upcoming event type, as it allows an event to end without introducing a new event, thereby protecting the event that just ended from further interference. In the example, when Olivier asks what Julie wants to do about the mistress (Fig. 1.a-5), the screen fades out (Fig. 1.a-6). With no visual input, we can still make predictions about what will happen in the next scene. It turns out that the next scene is the continuation of *scene 5* where Julie answers Olivier that she wants to meet the mistress (Fig. 1.a-7). Note, the probability distribution at *scene 7* is similar to the one in *scene 5*. However, event segmentation is still likely to occur between *5* and *7* because the prediction at *scene 6* became highly uncertain and thus the probability distributions diverged.

One key aspect of the framework is that the inference process makes use of previously created event types. For example, after meeting the mistress, we again see Julie visiting Olivier (Fig. 1.a-9), and we are again unsure about what the topic of their conversation will be. Thus, the probabilities of the two event types are uniformly distributed. When we see that they start working on Patrice's score together (Fig. 1.a-10), instead of

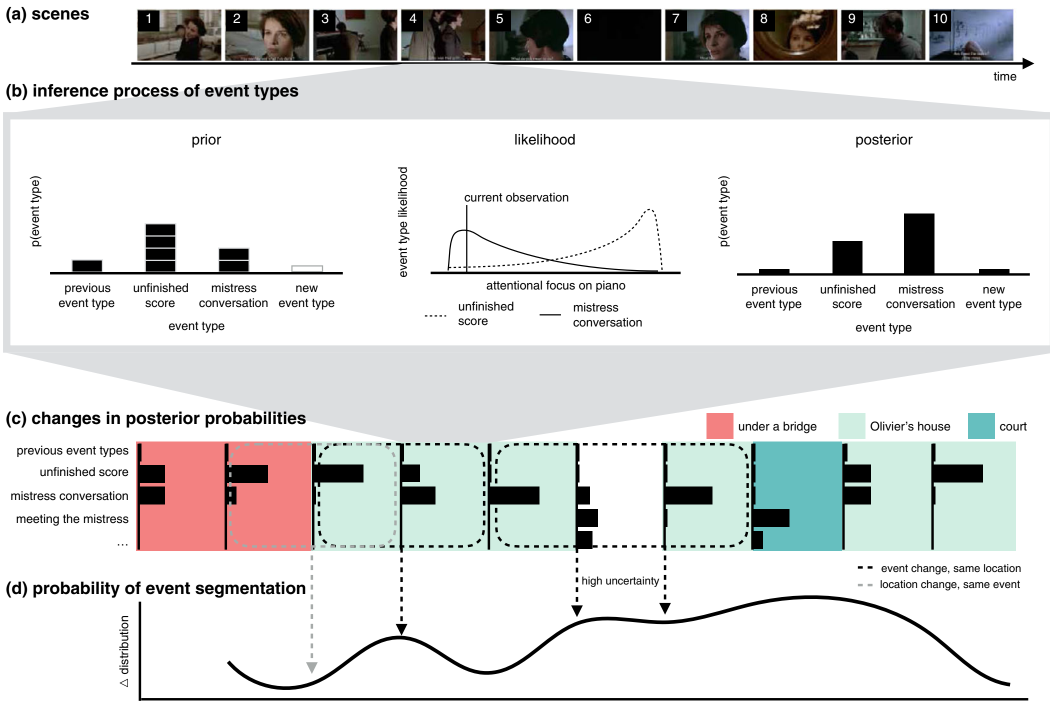

Fig. 1. An example of inference-based segmentation: (a) scenes from Krzysztof Kieslowski's movie *Three colors: Blue*. In this movie, Julie, a recent widow of a famous composer, tries to finish her husband's score and tracks down his mistress. (b) Inference process of event types. *Left*: Prior probability of each event type reflects the popularity of each event type (indicated by number of boxes constituting the bars). Note that there is a small chance of creating a new event type that has not yet been observed. *Center*: The likelihood of event type "mistress conversation" is higher than the likelihood of "unfinished score" given the current observation (low attention on the piano). *Right*: Combining the prior distribution and likelihoods results in the current posterior probability distribution over event types. Here, posterior probability is higher for the "mistress conversation" than the "unfinished score" event type, although the prior was higher for "unfinished score." (c) Changes in posterior probabilities. The posterior distribution over event types changes as observations change. Note that changes in the location (denoted by the background color) are not always predictive of posterior distribution changes. (d) Probability of event segmentation. Events are more likely to be segmented when changes in the probability distribution from the previous time point are large. Again, event segmentation (solid dotted lines) and location change (gray dotted lines) do not always correspond.

creating a whole new event type, we can reactivate the "unfinished score" event type and update the event with new information. Again, notice that the probability distributions for *scenes 3* and *10* are similar, but segmentation has occurred between those two event instances. Overall, this example illustrates the inference process and three key features of the model: (a) event segmentation occurs when the inferred events, rather than observed features, change, (b) common event types are more likely to be inferred (rich-gets-richer property), and (c) previous event types can be revisited and updated, enabling generalization.

3.3. Support for the latent variable account

Beyond online inference during event perception, the latent variable account also makes specific predictions about memory retrieval related to its cluster organization of individual episodes. As an analogy to memory retrieval, imagine looking back at some photo albums of a recent trip to Hawaii during which you went surfing many times and hiking just once (Fig. 2). You may look at the two photo albums “Surfing in Hawaii” and “Hiking in Hawaii” and summarize your Hawaii trip as “surfing and hiking.” That is, summarizing experiences by sampling at the cluster (photo album) level leads to an over-representation of rare episodes relative to how often they were actually experienced. In this example, cluster-level sampling will bias the memory toward equal weighting of surfing and hiking. By contrast, episodic sampling of the individual images will accurately represent their relative frequency and reflect that surfing was the main activity. This strong prediction of cluster-level sampling was recently tested by modeling the inference of a latent variable based on a set of observations. By manipulating the distribution of those observations, Shin and Niv (2020) showed that more clusters are inferred when fewer and more variable values are observed, which in turn biases summary estimates

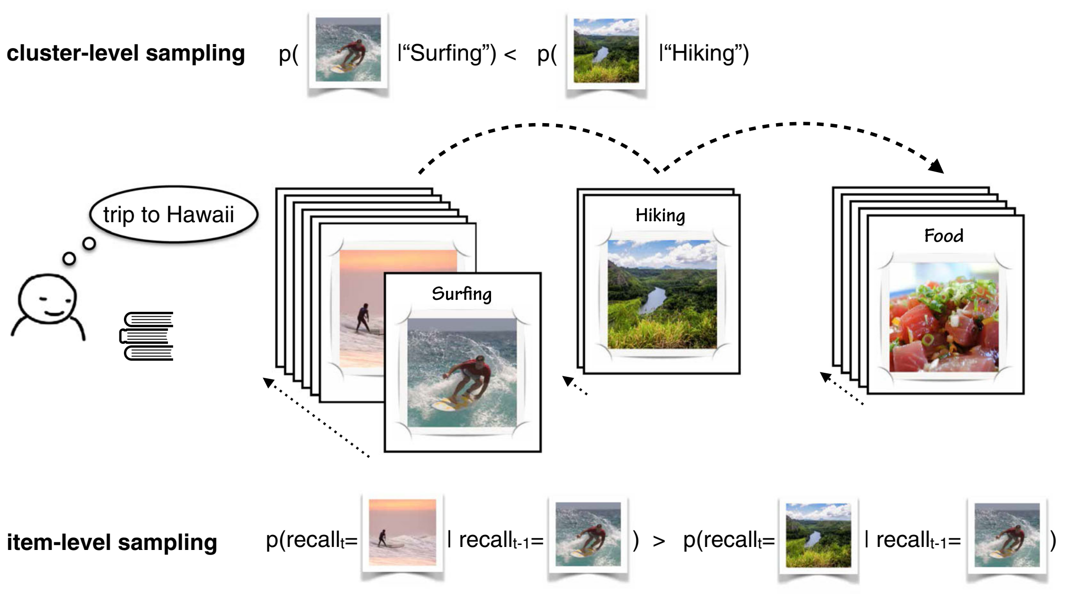

Fig. 2. Schematic for cluster-level and item-level sampling. As an analogy to memory retrieval, imagine looking back at some photo albums of a recent trip to Hawaii. If you summarize the trip, you may look at a few photos from each album (cluster-level sampling; top dashed arrow). In this case, any individual photo from a thinner album (e.g., “Hiking”) would be more likely to be picked than any individual photo from a thicker album (e.g., “Surfing”), gaining prominence in the overall summary. On the other hand, you may want to retrieve detailed experiences, looking at each photo (item-level sampling; bottom dotted arrows). In this case, a photo will be more likely to be followed by another photo from the same album (e.g., “Surfing”) than a different album (e.g., “Hiking”).

toward those rarer and more extreme events. This suggests that people can and do use cluster-level sampling to summarize experiences.

Relying on latent structure for summarizing experiences can also explain systematic memory distortions toward gist—the summary statistics of the latent variable. For instance, in a study where the organizing structure was imposed by item categories (e.g., lamp), color memory for individual members of a category was distorted toward the center of the color distribution of its category (Brady, Schacter, & Alvarez, 2018). At its extreme, the latent variable can even create false memories such that words that are not a part of the studied list yet exist at the conceptual/semantic center of the list are falsely recalled as one of the studied items (Deese, 1959; Roediger & McDermott, 1995). These types of gist biases can be useful in summarizing experiences, albeit at the cost of accuracy for each individual episode.

While the latent variable model can summarize experiences at the cluster level, this does not preclude sampling individual episodes during retrieval. Indeed, episodic sampling from clusters makes additional predictions for memory biases. Returning to the photo album analogy, imagine choosing a few photos from each photo album. Any given photo from an album with fewer photos would be more likely to be chosen. In a similar sense, when searching memory for past experiences by going through latent variables (events) and sampling from each of them (cluster-level sampling; Fig. 2 top), the episodes that have fewer companions in a cluster are more likely to be sampled, and thus are more likely to be retrieved. This idea is consistent with work by Alves et al. (2015), where people were found to be better at recognizing words that have fewer close conceptual neighbors in the same studied list. Sampling in this way can also account for memory biases like the von Restorff effect in which distinctive items are better remembered (Hunt, 1995; von Restorff, 1933) and the cue overload effect in which memory is better for items whose retrieval cue has fewer associated items (Watkins & Watkins, 1976).

What determines how many episodes each cluster contains when the clusters need to be dynamically inferred? One answer is extreme events that deviate from the average of previously existing clusters in terms of the relevant feature dimensions (e.g., a very loud noise that is not commonly experienced). During inference, the latent variable model is more likely to create an entirely new cluster for an extreme event. Since any given episode in a small cluster is more likely to be sampled, cluster-level sampling can account for memory and decision biases toward extreme episodes. For example, when people directly experience risky outcomes, they become risk seeking for gains and risk averse for losses (Hertwig, Barron, Weber, & Erev, 2004; Ludvig & Spetch, 2011). These decision biases are mediated by a memory bias in which extreme outcomes are remembered better and judged to have occurred more frequently (Madan, Ludvig, & Spetch, 2014, 2017). Similarly, episodes are better remembered when they elicit high reward prediction errors, regardless of their valence (Rouhani, Norman, & Niv, 2018). Recent work by Lieder, Griffiths, and Hsu (2014, 2018) also demonstrates that extreme events are more likely to be sampled from memory, although they propose a different sampling mechanism.

When trying to search for specific episodes, it would be useful to skim over a range of events (e.g., skimming the cover of photo albums) and go deeper only once we find the relevant event (e.g., open a photo album and look through the album). Supporting this idea, neural replay patterns measured with magnetoencephalography have been shown to skip between events, reinstating only the entry points of events in sequence (Michelmann, Staresina, Bowman, & Hanslmayr, 2019). When retrieval continues (e.g., choosing photos one after another), however, episodic samples would be more likely to transition within events than across events (i.e., item-level sampling; Fig. 2 bottom). That is, when we are motivated to retrieve details of episodic memories, rather than providing a quick summary of the entire experiences, it is easier to hold onto one album and to go through individual photos within an album than going back and forth between albums. This idea is supported by empirical data that show transitions within events are more likely than across events in free recall (Heusser et al., 2018) and serial recall (DuBrow & Davachi, 2016). The latent variable model would explain such recall behavior by modulating recall probability according to the similarity between posterior distributions over latent variables. Linking the latent variable model to episodic memory, Socher et al. (2009) showed that a variant of the model can predict within-event transitions in human recall behavior. In their variant, the probability of recalling a specific word was determined based on the mixture of the latent topic structure (semantic context) and the temporal adjacency (episodic context) active at a given time. This model could predict recall transitions better than models with purely semantic or episodic context, suggesting that both conceptual and temporal structures are critical features of clusters in memory.

4. Future directions

Theories of episodic memory organization provide additional insight into how temporal information may play a role in structuring events. For instance, the Temporal Context Model (Howard & Kahana, 2002) and the Context Maintenance and Retrieval model (Polyn, Norman, & Kahana, 2009a) have emphasized how storing a separate temporal representation may provide a scaffold for organizing memories. While these models have been highly successful in predicting memory recall, the way in which time interacts with ongoing event encoding and segmentation needs further investigation. Rather than having to store an independent representation of temporal information (Socher et al., 2009), a nonparametric Bayesian model in which time is incorporated into the process of inferring latent variables could be more parsimonious. That is, the model could be sensitive to the recency of events without having to separately track time, as the probability of a previous event would decay over time since it was last active. In addition, by assuming that recently encountered event types have a higher chance of producing the current observation, the model can provide temporal stability in the inference process. One candidate model is a latent variable model that utilizes distance in the prior probability of events, called the distance-dependent Chinese Restaurant Process (ddCRP; Blei & Frazier, 2011). Instead of relying on cluster popularities as the standard rational model of categorization and CRP do, the ddCRP prior can assign probabilities according to the temporal distance between previous and current observations. Similarly, a simpler variant of this model in which a currently active cluster gets an extra boost, called a sticky CRP (Fox, Sudderth, Jordan, & Willsky, 2011), has been used to account for stable event perception (Franklin et al., 2019; Gershman et al., 2014). Another intriguing direction for event segmentation research is incorporating time into the ddCRP hierarchically (cf. Ghosh, Ungureanu, Sudderth, & Blei, 2011) for event types to perform a stabilizing function that can improve generalization across instances.

The exact mechanisms by which event types are inferred and how the inferred event type interacts with event segmentation need further exploration. One possibility is that events are segmented when there is a change in the distribution over inferred event types (Fig. 1). This would explain the event segmentation effects observed when boundaries are imposed by time shift signals (e.g., "a while later") before the next event begins. That is, such phrases or other signals that the previous event has ended and is no longer relevant (e.g., the fade-to-black sixth scene in Fig. 1.a) would increase uncertainty over event types, thereby changing the probability distribution and inducing an event boundary. A conceptual parallel can be found in a neural network model called the Connectionist Temporal Classification model (Graves, Fernández, Gomez, & Schmidhuber, 2006). In this model, boundaries can be created at moments of high uncertainty (e.g., silence) where the likelihoods of any existing labels (event types) are low. This approach of treating high uncertainty as a non-labeled state increases flexibility in terms of how long an event lasts by allowing an event to end before the next one begins. These models differ from Event Segmentation Theory, which assumes that a new event begins as soon as the previous event ends (Kurby & Zacks, 2008). Experiments that test specific hypotheses based on the proposed framework (e.g., event boundaries will occur at the offset of an event instance as well as at the onset) would provide further insight as to how events are segmented, being guided by, and guiding, predictions.

There are remaining questions regarding how transitions between event types are learned and represented. In its current form, the latent variable framework does not directly address how transition probabilities would be incorporated in event segmentation. Neural network approaches have begun to investigate how transition models between low-level event types (behavioral primitives) can be learned such that the history of previous event transitions would inform the subsequent prediction at an event boundary (Gumbsch, Butz, & Martius, 2019). Another important question pertains to the hierarchical structure of event types. That is, when one event type is repeatedly followed by another, would those two event types continue to be recognized as distinct or would they ultimately be merged into a single, more complex event type? In either case, the mechanisms by which transition probabilities (within and across events) are represented and utilized in event cognition will need further examination.

5. Conclusion

A large body of work now suggests that event segmentation is a fundamental process that emerges naturally and is remarkably consistent across individuals. Research on how segmentation influences memory demonstrates its adaptive utility in increasing access to relevant information and reducing interference. In order to better understand the cognitive operations that support event segmentation, we must examine patterns of behavior when the answer is not clear (i.e., during ambiguous transitions) and model the internal processes that infer change based on ambiguous input. In particular, we propose that latent variable inference provides a useful framework for characterizing how we identify when an event is no longer relevant and select among alternatives. This framework accounts for existing data on the consequences of event segmentation for online processing, memory, and decision making, and generates new predictions that can guide future research and model development.

Acknowledgments

We thank Martin Butz, J. Benjamin Hutchinson, Nina Rouhani, and the anonymous reviewers for comments on previous versions of the manuscript. We are also grateful to Yael Niv, Kennenth A. Norman, and members of the DuBrow lab for helpful discussions.

References

Aldous, D. J. (1985). *Exchangeability and related topics*. École d'Été de Probabilités de Saint-Flour XIII—1983, Volume 1117. Berlin: Springer.

Alves, H., Unkelbach, C., Burghardt, J., Koch, A. S., Krüger, T., & Becker, V. D. (2015). A density explanation of valence asymmetries in recognition memory. *Memory & Cognition*, *43*(6), 896–909.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*(3), 409–429.

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.

Baldwin, D. A., & Kosie, J. (2020). How does the mind render streaming experience as events. *Topics in Cognitive Science*. https://doi.org/10.1111/tops.12502

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221.

Ben-Yakov, A., & Henson, R. N. (2018). The hippocampal film editor: Sensitivity and specificity to event boundaries in continuous experience. *The Journal of Neuroscience*, *38*(47), 10057–10068.

Bilkey, D., & Jensen, C. (2020). Neural markers of event boundaries. *Topics in Cognitive Science*. https://doi.org/10.1111/tops.12470

Blei, D. M., & Frazier, P. I. (2011). Distance dependent Chinese restaurant processes. *Journal of Machine Learning Research*, *12*, 2461–2488.

Bornstein, A. M., & Norman, K. A. (2017). Reinstated episodic context guides sampling-based decisions for reward. *Nature Neuroscience*, *20*(7), 997–1003.

Brady, T., Schacter, D. L., & Alvarez, G. (2018). The adaptive nature of false memories is revealed by gist-based distortion of true memories. *PsyArXiv*.

Brunec, I. K., Ozubko, J. D., Ander, T., Guo, R., Moscovitch, M., & Barense, M. D. (2020). Turns during navigation act as boundaries that enhance spatial memory and expand time estimation. *Neuropsychologia*, *141*, 107437.

Butz, M. V., Bilkey, D., Humaidan, D., Knott, A., & Otte, S. (2019). Learning, planning, and control in a monolithic neural event inference architecture. *Neural Networks*, *117*, 135–144.

Chen, J., Leong, Y. C., Honey, C. J., Yong, C. H., Norman, K. A., & Hasson, U. (2017). Shared memories reveal shared structure in neural activity across individuals. *Nature Neuroscience*, *20*(1), 115–125.

Collins, A. G. E., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychological Review*, *120*(1), 190–229.

Courville, A. C., Daw, N. D., & Touretzky, D. S. (2005). Similarity and discrimination in classical conditioning: A latent variable account. *Advances in Neural Information Processing Systems*, *17*, 313–320.

Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, *58*(1), 17–22.

Doll, B. B., Duncan, K. D., Simon, D. A., Shohamy, D., & Daw, N. D. (2015). Model-based choices involve prospective neural activity. *Nature Neuroscience*, *18*(5), 767–772.

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.

DuBrow, S., & Davachi, L. (2014). Temporal memory is shaped by encoding stability and intervening item reactivation. *The Journal of Neuroscience*, *34*(42), 13998–14005.

DuBrow, S., & Davachi, L. (2016). Temporal binding within and across events. *Neurobiology of Learning and Memory*, *134*, 107–114.

Duncan, K., Semmler, A., & Shohamy, D. (2019). Modulating the use of multiple memory systems in value-based decisions with contextual novelty. *Journal of Cognitive Neuroscience*, *31*(10), 1455–1467.

Ezzyat, Y., & Davachi, L. (2011). What constitutes an episode in episodic memory? *Psychological Science*, *22*(2), 243–252.

Ezzyat, Y., & Davachi, L. (2014). Similarity breeds proximity: Pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron*, *81*(5), 1179–1189.

Foster, D. J., & Wilson, M. A. (2007). Hippocampal theta sequences. *Hippocampus*, *17*(11), 1093–1099.

Fox, E. B., Sudderth, E. B., Jordan, M. I., & Willsky, A. S. (2011). A sticky HDP-HMM with application to speaker diarization. *The Annals of Applied Statistics*, *5*(2A), 1020–1056.

Franklin, N., Norman, K. A., Ranganath, C., Zacks, J. M., & Gershman, S. J. (2019). Structured event memory: A neuro-symbolic model of event cognition. *bioRxiv*, 541607.

Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, *117* (1), 197–209.

Gershman, S. J., Radulescu, A., Norman, K. A., & Niv, Y. (2014). Statistical computations underlying the dynamics of memory updating. *PLoS Computational Biology*, *10*(11), e1003939.

Ghosh, S., Ungureanu, A. B., Sudderth, E. B., & Blei, D. M. (2011). Spatial distance dependent Chinese restaurant processes for image segmentation. *Advances in Neural Information Processing Systems*, *24*, 1476–1484.

Graves, A., Fernández, S., Gomez, F., & Schmidhuber, J. (2006). Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd International Conference on Machine Learning* (pp. 369–376). Pittsburgh, PAL: Association for Computing Machinery.

Gumbsch, C., Butz, M. V., & Martius, G. (2019). Autonomous identification and goal-directed invocation of event-predictive behavioral primitives. *IEEE Transactions on Cognitive and Developmental Systems*, 1–1.

Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S., & Redish, A. D. (2012). Segmentation of spatial experience by hippocampal theta sequences. *Nature Neuroscience*, *15*(7), 1032–1039.

Hard, B. M., Meyer, M., & Baldwin, D. (2019). Attention reorganizes as structure is detected in dynamic action. *Memory & Cognition*, *47*(1), 17–32.

Hard, B. M., Recchia, G., & Tversky, B. (2011). The shape of action. *Journal of Experimental Psychology: General*, *140*(4), 586–604.

Hasson, U., Yang, E., Vallines, I., Heeger, D. J., & Rubin, N. (2008). A hierarchy of temporal receptive windows in human cortex. *The Journal of Neuroscience*, *28*(10), 2539–2550.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539.

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.

Horner, A. J., Bisby, J. A., Wang, A., Bogus, K., & Burgess, N. (2016). The role of spatial boundaries in shaping long-term event representations. *Cognition*, *154*, 151–164.

Hoskin, A. N., Bornstein, A. M., Norman, K. A., & Cohen, J. D. (2018). Refresh my memory: Episodic memory reinstatements intrude on working memory maintenance. *Cognitive, Affective, & Behavioral Neuroscience*, *19*(2), 338–354.

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*(3), 269–299.

Hunt, R. R. (1995). The subtlety of distinctiveness: What von Restorff really did. *Psychonomic Bulletin & Review*, *2*(1), 105–112.

Jeunehomme, O., & D'Argembeau, A. (2018). Event segmentation and the temporal compression of experience in episodic memory. *Psychological Research Psychologische Forschung*, *157*(7), 1–10.

Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *The Journal of Neuroscience*, *27*(45), 12176–12189.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, *273*(5280), 1399–1402.

Kosie, J. E., & Baldwin, D. (2019). Attentional profiles linked to event segmentation are robust to missing information. *Cognitive Research: Principles and Implications*, *4*(1), 8.

Kurby, C. A., & Zacks, J. M. (2008). Segmentation in the perception and memory of events. *Trends in Cognitive Sciences*, *12*(2), 72–79.

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.

Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review*, *125*(1), 1–32.

Lieder, F., Hsu, M., & Griffiths, T. L. (2014). The high availability of extreme events serves resource-rational decision-making. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *36*. Retrieved from https://escholarship.org/uc/item/5433c2jb

Lositsky, O., Chen, J., Toker, D., Honey, C. J., Shvartsman, M., Poppenk, J. L., Hasson, U., & Norman, K. A. (2016). Neural pattern change during encoding of a narrative predicts retrospective duration estimates. *eLife*, *5*, 255.

Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS ONE*, *6*(6), e20262.

Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin & Review*, *21*(3), 629–636.

Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2017). The role of memory in distinguishing risky decisions from experience and description. *Quarterly Journal of Experimental Psychology*, *70*(10), 2048–2059.

Magliano, J. P., Radvansky, G. A., Forsythe, J. C., & Copeland, D. E. (2014). Event segmentation during first-person continuous events. *Journal of Cognitive Psychology*, *26*(6), 649–661.

McNerney, M. W., Goodwin, K. A., & Radvansky, G. A. (2011). A novel study: A situation model analysis of reading times. *Discourse Processes*, *48*(7), 453–474.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85* (3), 207–238.

Michelmann, S., Staresina, B. P., Bowman, H., & Hanslmayr, S. (2019). Speed of time-compressed forward replay flexibly changes in human episodic memory. *Nature Human Behaviour*, *3*(2), 143–154.

Momennejad, I., Otto, A. R., Daw, N. D., & Norman, K. A. (2018). Offline replay supports planning in human reinforcement learning. *eLife*, *7*, e32548.

Murty, V. P., FeldmanHall, O., Hunter, L. E., Phelps, E. A., & Davachi, L. (2016). Episodic memories predict adaptive value-based decision-making. *Journal of Experimental Psychology: General*, *145*(5), 548–558.

Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasly, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*, *15*(7), 1040–1046.

Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *The Journal of Neuroscience*, *30*(37), 12366–12378.

Newtson, D. (1973). Attribution and the unit of perception of ongoing behavior. *Journal of Personality and Social Psychology*, *28*(1), 28–38.

Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, *115*(1), 39–57.

Ólafsdóttir, H. F., Bush, D., & Barry, C. (2018). The tole of hippocampal replay in memory and planning. *Current Biology*, *28*(1), R37–R50.

O'Reilly, J. X. (2013). Making predictions in a changing world-inference, uncertainty, and learning. *Frontiers in Neuroscience*, *7*, 105.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning – Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*(6), 532–552.

Pettijohn, K. A., & Radvansky, G. A. (2016). Narrative event boundaries, reading times, and expectation. *Memory & Cognition*, *44*(7), 1064–1075.

Pezzulo, G., van der Meer, M. A. A., Lansink, C. S., & Pennartz, C. M. A. (2014). Internally generated sequences in learning and executing goal-directed behavior. *Trends in Cognitive Sciences*, *18*(12), 647–657.

Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, *497*(7447), 74–79.

Poldrack, R. A., Clark, J., Pare-Blagoev, E. J., Shohamy, D., Moyano, J. C., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. *Nature*, *414*(6863), 546–550.

Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009a). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, *116*(1), 129–156.

Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009b). Task context and organization in free recall. *Neuropsychologia*, *47*(11), 2158–2163.

Radvansky, G. A. (2012). Across the event horizon. *Current Directions in Psychological Science*, *21*(4), 269–272.

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

Radvansky, G. A., & Copeland, D. E. (2010). Reading times and the detection of event shift processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*(1), 210–216.

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.

Radvansky, G. A., & Zacks, J. M. (2011). Event perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(6), 608–620.

Radvansky, G. A., & Zacks, J. M. (2017). Event boundaries in memory and cognition. *Current Opinion in Behavioral Sciences*, *17*, 133–140.

Radvansky, G. A., & Zacks, R. T. (1991). Mental models and the fan effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*(5), 940–953.

Radvansky, G., Zwaan, R. A., Federico, T., & Franklin, N. (1998). Retrieval from temporally organized situation models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(5), 1224–1237.

Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, *3*(3), 382–407.

Reynolds, J. R., Zacks, J. M., & Braver, T. S. (2007). A computational model of event segmentation from perceptual prediction. *Cognitive Science*, *31*(4), 613–643.

Richmond, L. L., & Zacks, J. M. (2017). Constructing experience: Event models from perception to action. *Trends in Cognitive Sciences*, *21*(12), 962–980.

Rinck, M., & Weber, U. (2003). Who when where: An experimental test of the event-indexing model. *Memory & Cognition*, *31*(8), 1284–1292.

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(4), 803–814.

Rouhani, N., Norman, K. A., & Niv, Y. (2018). Dissociable effects of surprising rewards on learning and memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(9), 1430–1443.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926–1928.

Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: Alternative algorithms for category learning. *Psychological Review*, *117*(4), 1144–1167.

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.

Schuck, N. W., & Niv, Y. (2019). Sequential replay of nonspatial task states in the human hippocampus. *Science*, *364*(6447), eaaw5181.

Shadlen, M. N., & Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, *90*(5), 927–939.

Shin, Y. S., & Niv, Y. (2020). Biased evaluations emerge from inferring hidden causes. *PsyArXiv*.

Shohamy, D., & Daw, N. D. (2015). Integrating memories to guide decisions. *Current Opinion in Behavioral Sciences*, *5*, 85–90.

Siefke, B. M., Smith, T. A., & Sederberg, P. B. (2019). A context-change account of temporal distinctiveness. *Memory & Cognition*, *79*(2), 1–15.

Socher, R., Gershman, S., Sederberg, P., Norman, K., Perotte, A. J., & Blei, D. M. (2009). A Bayesian analysis of dynamics in free recall. *Advances in Neural Information Processing Systems*, *22*, 1714–1722.

Soto, F. A., Gershman, S. J., & Niv, Y. (2014). Explaining compound generalization in associative and causal learning through rational principles of dimensional generalization. *Psychological Review*, *121*(3), 526–558.

Speer, N. K., Swallow, K. M., & Zacks, J. M. (2003). Activation of human motion processing areas during event perception. *Cognitive, Affective, & Behavioral Neuroscience*, *3*(4), 335–345.

Speer, N. K., & Zacks, J. A. (2005). Temporal changes as event boundaries: Processing and memory consequences of narrative time shifts. *Journal of Memory and Language*, *53*(1), 125–140.

Speer, N. K., Zacks, J. M., & Reynolds, J. R. (2007). Human brain activity time-locked to narrative event boundaries. *Psychological Science*, *18*(5), 449–455.

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.

Swallow, K. M., Zacks, J. M., & Abrams, R. A. (2009). Event boundaries in perception affect memory encoding and updating. *Journal of Experimental Psychology: General*, *138*(2), 236–257.

von Restorff, H. (1933). Über die Wirkung von Bereichsbildungen im Spurenfeld. *Psychologische Forschung*, *18*(1), 299–342.

Watkins, M. J., & Watkins, O. C. (1976). Cue-overload theory and the method of interpolated attributes. *Bulletin of the Psychonomic Society*, *7*(3), 289–291.

Zacks, J. M., Braver, T. S., Sheridan, M. A., Donaldson, D. I., Snyder, A. Z., Ollinger, J. M., Buckner, R. L., & Raichle, M. E. (2001). Human brain activity time-locked to perceptual event boundaries. *Nature Neuroscience*, *4*(6), 651–655.

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.

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.

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.

Zacks, J. M., Tversky, B., & Iyer, G. (2001). Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology: General*, *130*(1), 29–58.

Zwaan, R. A. (1996). Processing narrative time shifts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(5), 1196–1207.

Zwaan, R. A., Langston, M. C., & Graesser, A. C. (1995). The construction of situation models in narrative comprehension: An event-indexing model. *Psychological Science*, *6*(5), 292–297.

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.

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, *123*(2), 162–185.

Zwaan, R. A., Radvansky, G. A., Hilliard, A. E., & Curiel, J. M. (1998). Constructing multidimensional situation models during reading. *Scientific Studies of Reading*, *2*(3), 199–220.