Event Boundedness Affects Attention Allocation during Online Event Processing

Jamie Yuen (jkyuen@sas.upenn.edu)
Department of Linguistics, University of Pennsylvania
Philadelphia, Pennsylvania, USA

Sarah Hye-yeon Lee (sarahhl@sas.upenn.edu)
Department of Linguistics, University of Pennsylvania
Philadelphia, Pennsylvania, USA

Anna Papafragou (anna4@sas.upenn.edu)
Department of Linguistics, University of Pennsylvania
Philadelphia, Pennsylvania, USA

Abstract

The human mind can segment continuous streams of activity in the world into meaningful, discrete units known as events. However, not all events are created equal. We draw a distinction between *bounded* events (e.g., folding a handkerchief) that have a predictable structure that develops in distinct stages (i.e., a beginning, middle, and end) and a well-defined endpoint, and *unbounded* events (e.g., waving a handkerchief) that lack such a well-defined structure and endpoint. We predict that event boundedness (bounded vs. unbounded) will affect attention allocation patterns over the course of the event. Here, we tested this prediction using a dwell time paradigm by measuring the time participants spent on each still frame of an activity. We found that event endpoints attracted increased attention compared to midpoints; importantly, this increase was significantly greater when people viewed bounded events, compared to unbounded events. In addition, event endpoints attracted increased attention compared to event beginnings, but this pattern also interacted with event boundedness. We conclude that abstract internal event structure (specifically, event boundedness) affects attention allocation during online event apprehension.

Keywords: event cognition; boundedness; attention; dwell time; prediction

Introduction

Humans are surrounded with rich, dynamic, and continuous input, yet are able to chunk this input into discrete *events*. This ability has been studied by having observers press a key to dynamically mark boundaries, or breakpoints, between events in a video of naturalistic human or human-like action (Newtson, 1973; Newtson & Engquist, 1976). Evidence suggests that there is robust agreement on event boundaries across individuals (e.g., Newtson, 1973, Zacks et al., 2001), and event segmentation appears to occur spontaneously as people process sequences of ongoing activity (Saylor & Baldwin, 2004; Baldwin & Pederson, 2016; Huff et al., 2012; Zacks et al., 2001).

Dwell Times as a Window onto Event Structure

Previous studies have used the dwell time paradigm to probe mechanisms of event segmentation and representation (Hard et al., 2011, 2019; Kosie & Baldwin, 2019a, 2019b, 2021). In a dwell time experiment, activity sequences are converted into image slideshows, and observers advance through them at their own pace. Their dwell times –time spent looking ("dwelling") at each slide– provide a measure of attention, with greater time spent looking on a slide indicating greater attention. In a pioneering study, Hard et al. (2011) presented to participants four familiar, potentially hierarchically organized activities (e.g., cleaning a dorm room, eating breakfast) in the form of slideshows and measured dwell times at each slide as participants proceeded through them at their own pace. Then participants viewed the same activity in video-form and segmented it into smaller events by marking where one event ended and another began (i.e., the event boundary). Dwell times were analyzed on slides that each participant later identified as being within-event, or at an event boundary. Increased dwell times were found at event boundary slides compared to within-event slides, indicating that event boundaries elicit greater attention. Such enhanced attention to event boundaries is known as the 'boundary advantage.' The boundary advantage has been found across different event exemplars (Sage & Baldwin, 2014), and even in preschoolers (Kosie & Baldwin, 2021), and appears to be linked to enhanced subsequent memory for event units (Gold, Zacks, & Flores, 2017; McGatlin, Newberry, & Bailey, 2018; see Yousif et al., 2024 for related findings with larger-scale stimuli from outside the lab). More broadly, a variety of existing studies suggests that event boundaries – especially event endpoints – are salient within the representation of an event (Huff et al., 2012; Ongchoco & Scholl, 2019; Pettijohn & Radvansky, 2016; Schwan & Garsoffky, 2004; Swallow et al., 2009; Zacks et al., 2007; Zacks & Tversky, 2001).

The finding that event boundaries yield greater attention aligns with Event Segmentation Theory (EST) (Zacks et al., 2007). EST proposes that people are constantly predicting what will happen next and, when happenings become increasingly unpredictable and prediction errors emerge, an event boundary occurs, indicating where one event ends and another begins (Zacks et al., 2007). According to EST, unpredictable, critical information is being relayed at event boundaries and warrants increased attention. Consistent with

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

this idea, dwell times tend to increase just before, rather than in response to, the completion of an action (Hard et al., 2011). As viewers anticipate that an action is about to be completed, they increasingly direct more attention to the action. With increased attention, viewers can better process information at the boundary region (Hard et al., 2019).

Boundedness: Beyond a Single Notion of 'Event Structure' and 'Event Boundary'

The literature on event segmentation and processing reviewed earlier has largely treated all events as similar entities characterized by the same attentional signatures (including the 'boundary advantage'). Here, we adopt the perspective that there are, in fact, foundational ontological differences in event types captured by the notion of *boundedness* (see Ji & Papafragou, 2020a, b). *Bounded* events have a non-homogenous structure with distinct stages that lead to a well-defined endpoint (or culmination). A girl folding a handkerchief is an example of a bounded event: its endpoint is defined as the point at which the handkerchief is entirely folded. By contrast, *unbounded* events have a homogenous structure that lacks a well-defined endpoint and therefore may terminate at any arbitrary moment (Ji & Papafragou, 2020a, b). A girl waving a handkerchief is an example of an unbounded event: this event could be considered to reach an end whenever the waving action stops. The bounded-unbounded distinction in the event domain is akin to the object-substance distinction in the object domain: bounded events resemble objects (e.g., a book) because they possess well-defined boundaries, and behave like canonical individuals; unbounded events, on the other hand, resemble substances (e.g., paper) because they lack well-defined boundaries and are not canonical individuals (they cannot be individuated or counted; see Papafragou & Ji, 2024; Lee, Ji & Papafragou, 2024, for empirical evidence of the event-object parallel; cf. also Kuhn et al., 2021; Wellwood, Hespos, & Rips, 2018).

Within event cognition, boundedness should be understood as a perspective on events imposed by the human mind and not an objective feature of the world (Vurgun, Ji & Papafragou, 2023; Wagner & Carey, 2003). Boundedness is reflected in human language (via the expression of *telicity*; Filip, 1993; Folli & Harley, 2006; Krifka, 1992, 1998; van Hout, 2016). Most relevantly for present purposes, viewers have been found to be sensitive to boundedness: for instance, both adults and young children can classify novel visual events as being either bounded or unbounded (Ji & Papafragou, 2020a; 2020b). Furthermore, viewers process these two types of events differently in event cognition – for instance, by treating structural disruptions of the content of events as more problematic if the event is bounded than unbounded (Ji & Papafragou, 2020b). Sensitivity to boundedness arises spontaneously, even when it is irrelevant to task demands (Ji & Papafragou, 2022).

If boundedness is a foundational property of event architecture that affects attention, it follows that bounded and unbounded events should elicit systematically different dwell time patterns. That is, viewers' dwell times throughout an event should reveal effects of event structure beyond a generalized 'boundary advantage'. Such findings would suggest that attention shifts do not occur only during segmentation of multiple sequential activities, but also *within* an event. Event boundedness allows an analysis of fine-grained event internal structure (beginning, middle, and ending) beyond boundary phenomena. More importantly, it allows us to treat varieties of 'event endpoints' differently depending on the representation they belong to.

Dwell Times and Signatures of Event Boundedness

The present study adopts the dwell time paradigm to directly test for signatures of boundedness in how attention is allocated during the unfolding of an event. Unlike past dwell time studies that presented relatively long, hierarchically organized activity sequences (e.g., Hard et al., 2011: 156-247 slides; 2019: 2,800 slides; Kosie & Baldwin, 2019a: 66-137 slides; 2019b: 57-112 slides; 2021: 20 slides), we focus on single events instantiated via very short image sequences. We hypothesize that fine-grained event structure lives within the edges of even such short exposures: in particular, bounded events – which have predictable and well-defined internal structure – should elicit more dramatic shifts in attention over time compared to unbounded events – which have a homogenous structure. This hypothesis leads to two novel predictions. First, attention at the ends of events should be greater than at the middles of events – but this change in attention should be greater for bounded than unbounded events. Second, attention at the ends of events should increase compared to beginnings – but again this difference should be greater for bounded events. We test these predictions in the experiment that follows.

Experiment

Method

Participants Forty-three adult native speakers of English (19 female, 24 male, M_{Age} = 43.2 years, SD = 13.4) were recruited via Prolific. The experiment took about 4 minutes and participants were compensated $0.55 for completing the study.

Materials Twelve pairs of videos (*M*=7.76 s; range: 6.1 s-9.21 s) were adapted from Ji and Papafragou (2020a). Each pair included a closely matched bounded and an unbounded event with the same duration. Each video showcased the same actor performing an action on a tabletop against the same simple background (see Table 1 for a complete list). These videos belonged to a larger group of videos that had been normed by Ji and Papafragou (2020a) to ensure the contrast in boundedness between the paired events: (a) bounded events were considered to have "a beginning, middle and ending" 87.2% of the time compared to 20.3% of the time for unbounded events ($t(1, 39) = 20.05$, $p < .001$); (b) bounded events prompted descriptions with change-of-state verbs (e.g., *fold a handkerchief*) or quantified count noun phrases (e.g., *tie a knot*) 98.2% of the time, while unbounded events prompted descriptions with verbs of activity (e.g., *wave a handkerchief*) or unquantified noun phrases (bare plurals or mass nouns: e.g., *tie knots*) 92.8% of the time. These videos were also normed to be equal in terms of intentionality, speed of the action, and visual similarity within pairs of bounded and unbounded events (ibid.; cf. also Vurgun et al., 2024).

Table 1. List of test event stimuli.

	Bounded events	Unbounded events
1	cut a piece of paper in half	cut pieces from a paper roll
2	roll up a towel	twist a towel
3	fill a glass with milk	shake a bottle of milk
4	crack an egg	beat an egg
5	fold a handkerchief	wave a handkerchief
6	tie a knot	tie knots
7	draw a balloon	draw doodles
8	put up one's hair	scratch one's hair
9	tear a paper towel	tear slices off paper towels
10	stack a deck of cards	shuffle a deck of cards
11	blow a balloon	blow bubbles
12	group pawns based on color	mix pawns of two colors

Dwell time slideshows were created by selecting 9 still images from each video within this set (see Figure 1 for an example). The very first and last frame were extracted; seven additional still images were extracted at equal intervals between the first and last frames (at the rate of 10 frames per second). Because the video pairs varied in length, these seven images were sampled at slightly different rates across pairs (*M*=1.22 fps; range: 0.98 fps-1.91 fps). Overall, though, these sampling rates were similar to those in previous dwell time studies (Hard et al., 2011: 1fps, Kosie and Baldwin, 2019a: 1fps; Kosie & Baldwin, 2019b: 1fps and 2fps).[1] The total set of 9 images per slideshow chosen by our proportion method sufficiently and comprehensibly showcased the unfolding of each event.

The twelve pairs of slideshows were divided into two presentation lists. Each list contained only one member from a bounded-unbounded event pair and had the same total number of bounded and unbounded events (6 each), arranged in a pseudo-randomized order. For both lists, we also created the same two practice slideshow trials with additional events, one unbounded (stir yogurt) and the other bounded (stick a sticker; both also drawn from Ji & Papafragou, 2020a).

Procedure The experiment was hosted on an online experiment-building platform, PennController IBEX (Zehr & Schwarz, 2018). Participants received the following instructions: "In this experiment, you will see 12 videos. You will advance through the videos at your own pace by pressing the spacebar. You can spend as much time looking at each screen as you wish. Please pay careful attention to each video because you will be asked questions throughout the study about what you saw." Participants were randomly assigned to one of the two presentation lists for the dwell time task. Dwell

Figure 1: Sample dwell time trials: slideshow for a bounded event (fill a glass with milk; Panel a) and an unbounded event (shake a bottle of milk; Panel b).

[1] Hard et al.'s (2011) study selected the middle frame of 30 fps slideshows at each 1-s intervals. However, the videos adapted for our study were significantly shorter in comparison and homed in on a singular event per trial (as opposed to several events in one sequence). We note that dwell-time patterns, such as the boundary advantage, seem robust to slideshow resolution (Kosie & Baldwin, 2019b).

times were calculated per still image as the latency between the timepoint of image presentation and participants' keypress to advance the display.

As a means to keep participants engaged, four memory trials were included throughout the experiment. Those appeared between 2-4 dwell time trials. For these, participants were shown an image from an event they might have seen from the pool of slideshows within that series. (We selected events that were presented in the first or second dwell time trial in each 2-4-trial set, in order to avoid recency effects.) In each memory trial, an image from the halfway point of a video (i.e., the fifth image of a dwell time sequence) was shown. The participants answered "Yes" if they thought they had seen it and "No" if they did not. The memory trials depicted two unbounded and two bounded events, which were counterbalanced such that participants in any group had seen only one of each event type during the dwell time trials. Participants in both groups were asked the same memory questions. Prior to the main experiment, we presented participants with the same two practice dwell time trials, and one practice memory question to familiarize them with the task.

Results

An initial analysis showed respectable accuracy on the memory trials (M = 77%). With this in mind, we turned to our main measure of interest.

Prior to our main dwell time data analysis, dwell times faster than 80 ms and slower than 2500 ms were excluded from analysis. In accordance with procedures established by Hard et al. (2011, 2019) and Kosie and Baldwin (2019a, 2019b), dwell times greater than 3 standard deviations (SD) from the overall group mean were also excluded. Additionally, one participant garnered an overall mean dwell time greater than 3 SD away from the means of other participants and was excluded from the analysis. Overall, 6.5% of the original data was affected.

We divided the nine total slides evenly into thirds to represent three stages (beginning, middle, and end) within each event. Following past research (Hard et al., 2011; Kosie & Baldwin, 2019a, 2021), we then excluded data from the very first slide of each event because it always elicited very high dwell times compared to the rest ($M_{bounded}$ = 0.838 s, $M_{unbounded}$ = 0.854 s). This is known to occur because participants are adjusting to the task. Figure 2 presents average dwell times at each event stage by event type for the remaining slides.

Event middles vs. ends To test our first prediction, we analyzed dwell times from the middle and end of bounded versus unbounded events. We used Linear Mixed Effects models with Event Type (contrast-coded, Unbounded=0.5, Bounded=-0.5), Event Stage (contrast-coded, End=0.5, Middle=-0.5) and the Event Stage by Event Type interaction as fixed effects. As random effects, we entered intercepts for subjects and items, in addition to by-subject and by-item random slopes for the effects of Event Type, as well as Event Stage and their interaction, when justified by model comparison: Random effects began fully crossed and fully specified with by-subject and by-item slopes for relevant main effects. They were reduced (starting with by-item effects) through model comparison such that only random effects that contributed significantly to the model (p<.05) were included (Baayen et al., 2008). Analyses were conducted using the *lme4* package (version 1.1.35.2) (Bates et al., 2015) and *lmerTest* (version 3.1.3) (Kuznetsova et al., 2017) in the R software environment (R Development Core Team, 2023).

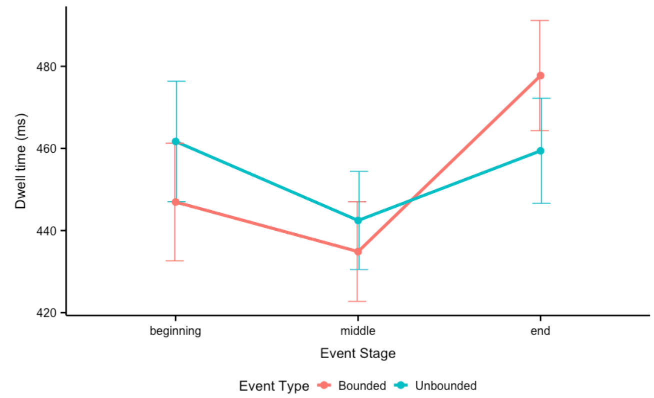

Figure 2: Average dwell times by Event Stage and Event Type

The statistical analysis is reported in Table 2. The analysis revealed a main effect of Event Stage on dwell times: participants paid significantly more attention to event ends than to event middles (p<.001). Importantly, as we predicted, we found an interaction between Event Stage and Event Type (p<.05): there was a greater change in dwell time from the middle to the end of an event in bounded events compared to unbounded events.

Table 2: Results of the lmer model for dwell times at the middles and ends of bounded vs. unbounded events.

Effect	Estimate	SE	df	t value	p value
Intercept	455.39	47.56	42.21	9.58	<.001***
Event Type	-8.65	6.12	2899.18	-1.41	.16
Event Stage	35.55	7.18	40.37	4.95	<.001***
Event Type:Event Stage	-30.91	12.22	2896.71	-2.53	.012*

Note. Formula in R: Dwell Time ~ Event Type * Event Stage + (1 + Event Stage | participant) + (1 | item)
* *p* < .05, ** *p* < .01, *** *p* < 0.001

Event beginnings vs. ends To test our second prediction, we conducted a separate, similar statistical analysis on dwell times at the beginnings and ends of unbounded and bounded events, reported in Table 3. In line with our prediction, the analysis revealed a significant interaction between Event Type and Event Stage (p<.01): participants paid more attention to event ends compared to event beginnings for bounded events (p<.01) but showed no significant difference between event ends and beginnings for unbounded events.

Table 3: Results of the lmer model for dwell times at the beginnings and ends of bounded vs. unbounded events.

Effect	Estimate	SE	df	t value	p value
Intercept	464.52	46.92	42.99	9.90	<.001***
Event Type	-3.42	16.96	12.44	-.20	.84
Event Stage	18.21	11.38	23.39	1.60	.12
Event Type:Event Stage	-42.09	13.79	2327.92	-3.05	.0023**

Note. Formula in R: Dwell Time ~ Event Type + (1 | participant) + (1 | item)
* $p < .05$, ** $p < .01$, *** $p < 0.001$

Discussion

We examined how viewers' attention (measured by dwell times) changed as bounded and unbounded events unfolded over time. Two major findings emerged. First, viewers' attention was significantly higher at event endings compared to event middles but this increase in attention was modulated by boundedness: the midpoint-endpoint difference was more pronounced when viewing bounded than unbounded events. Second, and relatedly, attention increased between beginnings and endings for bounded events, but no such difference emerged for unbounded events. These results replicate earlier findings reporting an event boundary advantage in the dwell time literature (Hard et al., 2011, 2019; Kosie & Baldwin, 2019a, 2019b, 2021) but go beyond past work by showing that this advantage depends on the structure of distinct, abstract event types defined by the notion of boundedness.

General Discussion

Dwell Times Reveal Signatures of Event Boundedness

We present evidence that the internal structure of events affects attention allocation during online event processing. Using a dwell time paradigm, we found that viewers allocated greater attention to boundaries (endpoints) compared to non-boundaries (middles), replicating the boundary advantage that has been attested in the literature (Hard et al., 2011, 2019; Kosie & Baldwin, 2019a, 2019b, 2021). However, our data show that not all event boundaries are equal. The change in attention as viewers proceeded from the beginning, middle, to the end of an event interacted with boundedness: While viewers proceeded through the event, they were sensitive to the event's structure and paid more attention as they neared the endpoint, with this effect being more pronounced in bounded events. The effect of boundedness also significantly interacted with attention to the beginnings vs. endings of events, which could indicate that an event's internal structure is recognized almost immediately and processed accordingly.

These findings add to existing evidence that boundedness is a cognitive feature of cognitive event architecture (Papafragou & Ji, 2024; Lee et al., 2024; see also Kuhn et al., 2021; Wellwood et al., 2018). Importantly, they also show that boundedness rapidly and systematically affects the online allocation of attention throughout an event.

Revisiting the 'Boundary Advantage'

Why are endings of bounded events more privileged than endings of unbounded events? One possibility is that it is relatively easier to predict when a bounded event is going to end compared to unbounded events. Take, for example, a bounded event like filling a glass with milk. As long as the speed of the pouring remains consistent throughout the event, the fullness of the cup perfectly corresponds to the progression of the filling-a-glass-with-milk event. When the cup is 90% full, the event is typically 90% completed, and viewers can anticipate that the event is nearing its ending. Because viewers are better able to anticipate when they are approaching the endpoint, they can deploy increased attention for the upcoming event boundary.

Conversely, endings of unbounded events are less predictable. Take, for example, an unbounded event like shaking a bottle of milk. Perhaps we can, to a limited extent, predict when a shaking-a-bottle-of-milk event might end (from what we know about typical shaking-a-bottle events in the world.) However, these predictions are less reliable. The shaking may only last for a very short time, or it may go on for longer than one would expect. Furthermore, the endpoint of this type of event might depend on the agent's (or another person's) unobservable goals or intentions, a factor that would make the predictability of such endpoints less well-defined (on the role of intentionality on event boundaries, see Mathis & Papafragou, 2022). This means that in viewing unbounded events, it would be more difficult to anticipate when the event is going to end, and viewers would have decreased ability to efficiently increase their attention near the event boundary.

Another, perhaps related possible explanation for why endings of bounded events are more privileged than endings of unbounded events has to do with the content of the endings: the endings of bounded events contain important information for event understanding, since they correspond to high degrees of (predictable) change in the object involved in the event; by contrast, endings of unbounded events do not contain privileged information, at least in terms of content, because they do not involve major changes. This line of reasoning is consistent with findings from an eye tracking study showing that, at the video offset, people pay more attention to the action and the object in resultative events that involve a high degree of (presumably predictable) change in an affected object compared to events with a less salient change of state (Sakarias & Flecken, 2019).

These specific patterns are broadly compatible with the idea that predictability within an event determines attention allocation around boundaries (see EST; Zacks et al., 2007). Beyond this idea, however, the boundedness effect is not anticipated by current theories focusing broadly on the role of event boundaries in cognition (e.g., Zacks et al., 2007) but falls out from theories that recognize the role of the abstract internal structure of different ontological types of events in how events are processed and understood (e.g., Ji & Papafragou, 2023). Unlike classic notions of 'event boundary' in event segmentation or dwell time studies that do not specify whether the event has reached its natural endpoint or has simply stopped, the notion of boundedness introduced here offers a more fine-grained notion of boundary: even though both bounded and unbounded events can and do come to an end, the notion of 'endpoint' represents something different in each case (reaching an inherent endpoint, i.e. culmination, as opposed to mere cessation respectively). Similarly, the notion of beginning is different in each case, with beginnings in bounded events – but not unbounded events - being lawfully connected to a projected, structurally organized whole and eventual endpoint.

Expanding the Dwell Time Paradigm in Event Cognition

Our experiment validates the use of dwell times as an implicit, rapid and systematic index of online event processing. We now show for the first time that the dwell time paradigm can reveal not simply event structure in general but finely articulated sub-types of event ontology (and corresponding notions of what counts as an 'event boundary').

In the current study, systematic dwell time patterns emerged even with single isolated events that were considerably shorter and simpler than the activity sequences tested in earlier studies. In this sense, our findings indicate that viewers extract internal temporal event structure even for short, simple, already segmented everyday events and this structure, in turn, rapidly affects the allocation of event-internal attention. Future uses of the dwell time paradigm (e.g., with naturalistic stimuli obtained beyond the lab) have the potential to further illuminate how people process internal event structure and thus advance theories of event cognition.

Acknowledgments

This work is supported by the NSF Grant BCS-2041171 awarded to AP and the Penn College of Arts and Sciences Alumni Society's Millstein Family Undergraduate Research Grant to JY.

References

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. https://doi.org/10.1016/j.jml.2007.12.005

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using Ime4. *Journal of Statistical Software, 67*(1), 1-48. https://doi.org/10.18637/jss.v067.101

Gold, D. A., Zacks, J. M., & Flores, S. (2017). Effects of cues to event segmentation on subsequent memory. *Cognitive Research: Principles and Implications*, *2*(1), 1. https://doi.org/10.1186/s41235-016-0043-2

Hard, B. M., Meyer, M., & Baldwin, D. (2019). Attention reorganizes as structure is detected in dynamic action. *Memory & Cognition*, *47*(1), 17–32. https://doi.org/10.3758/s13421-018-0847-z

Hard, B. M., Recchia, G., & Tversky, B. (2011). The shape of action. *Journal of Experimental Psychology: General*, *140*(4), 586–604. https://doi.org/10.1037/a0024310

Hard, B. M., Tversky, B., & Lang, D. S. (2006). Making sense of abstract events: Building event schemas. *Memory & Cognition, 34,* 1221–1235. https://doi.org/10.3758/BF03193267

Huff, M., Papenmeier, F., & Zacks, J. M. (2012). Visual target detection is impaired at event boundaries. *Visual Cognition*, *20*(7), 848–864. https://doi.org/10.1080/13506285.2012.705359

Ji, Y., & Papafragou, A. (2020a). Is there an end in sight? Viewers' sensitivity to abstract event structure. *Cognition*, *197*, 104197. https://doi.org/10.1016/j.cognition.2020.104197

Ji, Y., & Papafragou, A. (2020b). Midpoints, endpoints and the cognitive structure of events. *Language, Cognition and Neuroscience*, *35*(10), 1465–1479. https://doi.org/10.1080/23273798.2020.1797839

Ji, Y., & Papafragou, A. (2022). Boundedness in event cognition: Viewers spontaneously represent the temporal texture of events. *Journal of Memory and Language, 127*, 104353. https://doi.org/10.1016/j.jml.2022.104353

Ji, Y., & Papafragou, A. (2024). Boundedness supports children's event representations. *Developmental Psychology, 60(8),* 1457–1473. https://doi.org/10.1037/dev0001782

Kosie, J. E., & Baldwin, D. (2019a). Attention rapidly reorganizes to naturally occurring structure in a novel activity sequence. *Cognition*, 182, 31-44, https://doi.org/10.1016/j.cognition.2018.09.004

Kosie, J. E., & Baldwin, D. (2019b). Attentional profiles linked to event segmentation are robust to missing information. *Cognitive Research: Principles and Implications*, *4*(1), 8. https://doi.org/10.1186/s41235-019-0157-4

Kosie, J. E., & Baldwin, D. A. (2021). Dwell times showcase how goal structure informs preschoolers' analysis of unfolding motion patterns. *Child Development*, *92*, 2235–2243. https://doi.org/10.1111/cdev.13661

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software, 82*(13), 1–26. https://doi.org/10.18637/jss.v082.i13

Lee, S. H., Ji, Y., & Papafragou, A. (2024). Signatures of individuation across objects and events. *Journal of Experimental Psychology: General*, *153*(8), 1997–2012. https://doi.org/10.1037/xge0001581

Madden, C. & Zwaan, R. (2003). How does verb aspect constrain event representations? *Memory & Cognition, 31*, 663-672. https://doi.org/10.3758/BF03196106
McGatlin, K., Newberry, K. & Bailey, H. (2018). Temporal Chunking Makes Life's Events More Memorable. *Open Psychology*, *1*(1), 94-105. https://doi.org/10.1515/psych-2018-0007
Newtson, D., & Engquist, G. (1976). The perceptual organization of ongoing behavior. *Journal of Experimental Social Psychology, 12*(5), 436–450. https://doi.org/10.1016/0022-1031(76)90076-7
Ongchoco, J. D. K., & Scholl, B. J. (2019). Did that just happen? Event segmentation influences enumeration and working memory for simple overlapping visual events. *Cognition, 187,* 188–197. https://doi.org/10.1016/j.cognition.2019.01.002
Papafragou, A., & Ji, Y. (2023). Events and objects are similar cognitive entities. *Cognitive Psychology, 143,* 101573. https://doi.org/10.1016/j.cogpsych.2023.101573
Pettijohn, K. A., & Radvansky, G. A. (2016). Walking through doorways causes forgetting: Event structure or updating disruption? *Quarterly Journal of Experimental Psychology, 69*(11), 2119–2129. https://doi.org/10.1080/17470218.2015.1101478
R Core Team (2023) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna. https://www.R-project.org/
Sage, K. D., & Baldwin, D. (2014). Looking to the hands: Where we dwell in complex manual sequences. *Visual Cognition, 22*(8), 1092–1104. https://doi.org/10.1080/13506285.2014.962123
Sakarias, M., & Flecken, M. (2019). Keeping the Result in Sight and Mind: General Cognitive Principles and Language-Specific Influences in the Perception and Memory of Resultative Events. *Cognitive Science*, *43*(1), 10.1111/cogs.12708. https://doi.org/10.1111/cogs.12708
Saylor, M., & Baldwin, D. (2004). Action analysis and change blindness: Possible links. In D. T. Levin (Ed.), *Thinking and seeing: Visual metacognition in adults and children* (pp. 37–56). Greenwood/Praeger.
Schwan, S., & Garsoffky, B. (2004). The cognitive representation of filmic event summaries. *Applied Cognitive Psychology, 18*, 37–55. https://doi.org/10.1002/acp.940
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. https://doi.org/10.1037/a0015631
Vurgun, U., Ji, Y., & Papafragou, A. (2024). Aspectual processing shifts visual event apprehension. *Cognitive Science*, *48*(6), e13476. https://doi.org/10.1111/cogs.13476
Wagner, L., & Carey, S. (2003). Individuation of objects and events: a developmental study. *Cognition, 90*(2), 163–191. https://doi.org/10.1016/S0010-0277(03)00143-4
Yousif, S. R., Lee, S. H. Y., Sherman, B. E., & Papafragou, A. (2024). Event representation at the scale of ordinary experience. *Cognition, 249*, *105833*. https://doi.org/10.1016/j.cognition.2024.105833
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
Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological Bulletin*, *127*(1), 3–21. https://doi.org/10.1037/0033-2909.127.1.3
Zacks, J. M., Tversky, B., & Iyer, G. (2001). Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology: General, 130*, 29–58. https://doi.org/10.1037/0096-3445.130.1.29
Zehr, J., & Schwarz, F. (2018). PennController for Internet Based Experiments (IBEX). https://doi.org/10.17605/OSF.IO/MD832