Intentionality Effects on Event Boundaries

Ariel Mathis (apmathis@sas.upenn.edu)
Department of Linguistics, University of Pennsylvania
3401-C Walnut St. Suite 300, C-Wing, Philadelphia, PA 19104

Anna Papafragou (anna4@sas.upenn.edu)
Department of Linguistics, University of Pennsylvania
3401-C Walnut St. Suite 300, C-Wing, Philadelphia, PA 19104

Abstract

Theories of event cognition have hypothesized that the boundaries of events are characterized by change, including a change in the agent's goal, but the role of higher-order goal information on the placement of event boundaries has not been addressed experimentally. We tested whether goals can affect how viewers determine event boundaries. Participants read a context sentence stating an agent's goal (e.g., "Jesse wants to eat the orange with her breakfast" vs. "Jesse wants to use the orange as a garnish"). Participants then saw an image of an event outcome (e.g., a partly peeled orange) and were asked to identify whether the event had occurred ("Did she peel the orange?"). Participants were more likely to respond *Yes* to a partly complete outcome if the outcome satisfied the agent's goal. Our results offer the first direct evidence in support of the conclusion that higher-order intentionality information affects the way events are conceptualized.

Keywords: events, aspect, telicity, goals, intentionality, perfective

Introduction

Daily life, from morning routines and daily commutes to getting ready for bed at night, can be viewed as a series of events. From this point of view, "completing your morning routine" can be thought of as an event by itself, but this overarching event can also be broken down into smaller subevents such as making a pot of coffee, cooking breakfast, or making the bed (Zacks, Tversky, & Iyer, 2001). One definition of an event that aligns with this description states that events are characterized as temporal entities unfolding at a specific time and location and having a beginning and ending point (Zacks & Tversky, 2001).

According to a prominent model of how people parse dynamic events (Event Segmentation Theory, or EST; Zacks, Speer, Swallow, Braver, & Reynolds, 2007), event comprehension is an ongoing process that is facilitated by the use of multiple simultaneous event models that are maintained in working memory. Event models are used to predict near future occurrences and are adjusted when there is an increase in prediction errors. These increases in prediction error, as indicated by transient changes in neural activity, correspond to the placement of boundaries during both active event segmentation and passive viewing (Zacks et al., 2001). An increase in prediction errors, and the corresponding detection of event boundaries, has been found to correspond with points of change in the stream of input (Speer, Zacks, & Reynolds, 2004; 2007). These changes can be perceptual such as changes in movement (e.g., a car turns; Magliana, Kopp, McNerney, Radvansky, & Zacks, 2012). Additionally, top-down processes such as inferences about an agent's intentions have also been argued to influence the placement of event boundaries (Zacks, 2004). For instance, adults have been shown to adjust their level of segmentation (i.e., the density of event boundaries) depending on their familiarity with the actor's intent – a stream of actions will be divided into smaller units when viewers are uncertain about the goal of making these actions (Newtson, 1973; Vallacher & Wegner, 1987; Wilder, 1978). Relatedly, very young infants can parse everyday actions by placing boundaries at the points where a goal is achieved (Baldwin, Baird, Saylor, & Clark, 2001; cf. also Brandone & Wellman, 2009; Luo & Johnson, 2008; Woodward, 1998).

In the line of research just reviewed, the roles of perceptual and conceptual/intentional cues to event boundaries are typically intertwined, with changes in the latter often initiating or corresponding with changes in the former (Tversky, Zacks, & Hard, 2008). This close relationship creates a limitation when it comes to understanding how intentionality contributes to event boundaries (and thus event conceptualization) more generally. For instance, in many of the classic event segmentation studies, it is not possible to completely isolate changes in the goals of an event agent from co-occurring spatiotemporal cues as participants view and segment film clips. In the event of making a pot of coffee, for instance, a change in the agent's goal (e.g., completing the goal of filling the coffee maker with water and then deciding to turn it on) is also accompanied by a distinct change in movement (the change in motion away from the reservoir towards the 'on' button). One study by Levine, Hirsch-Pasek, Pace, and Golinkoff (2017) attempted to eliminate such spatiotemporal cues in a film segmentation task by playing motion clips in reverse. Participants shown the reversed film continued to segment similarly to those shown the original film. However, as the authors note, while reversing the film reduced the available spatiotemporal cues, the cohesion of the agent's movements was not eliminated.

A related limitation in prior work lies with the difficulty of presenting an agent's goals explicitly and naturally within a

film segment. Ideally, in order to address the hypothesis that a salient goal might affect placement of event endpoints, the goal information presented with the event must be both explicit and amenable to manipulation in isolation from the physical components of the event. In the present experiment, we addressed these limitations in a novel paradigm exploring the role of intentionality in event cognition.

Current Study

In the current study we tested whether contextually supplied knowledge about the goals of an agent within an event affects how viewers place event endpoints. We presented participants with a context sentence that introduced an overarching goal of an agent (e.g., either "Jesse wants to eat the orange with her breakfast", or "Jesse wants to use the orange as a garnish"). Then people were shown images of objects depicting an event at a stage of partial completion (e.g., a partly-peeled orange) and had to answer a question about the event (e.g., "Did she peel the orange?"). Critically, the event in the test question was an intermediate step (or subevent) in fulfilling the agent's overarching goal. We were interested in whether participants would be more likely to give non-culmination responses (e.g., to deny that the agent peeled the orange) when the agent's stated goals involved a higher degree of subevent development (as in *eat the orange*, that requires that all of the orange be peeled) as opposed to a lower degree of development (as in *use the orange as a garnish*, where even a small piece of the skin is enough).

The present study bears on theories of event cognition, especially those that have hypothesized that higher-order considerations such as intentionality affect event boundaries (Baldwin et al., 2001; Newtson, 1973; Vallacher & Wegner, 1987; Wilder, 1978; Zacks, 2004). As previously discussed, many of the methods utilized in prior work on event cognition were insufficient to isolate the role of intentionality from that of visual features of the input stream. The solution employed in the current experiment was to use a combination of narrative text and static images of event outcomes depicting various stages of completion. The use of a partially narrative format to present goal information allows for the agent's goal to be made explicit while also allowing the manipulation of the goal in isolation from other cues (cf. also Madden & Zwaan, 2003). The use of narrative is justified by studies of event segmentation that have found similar patterns of event boundary placement (Magliano et al., 2012) and similar neural activation around event boundaries (Speer et al., 2007; Zacks et al., 2001) regardless of whether an event was presented visually or in a narrative text. Similarly, the use of a static image allows for the manipulation of visual cues to event progression and culmination. The choice of static pictures is further justified by evidence that event information can be reliably extracted from a single event snapshot (e.g., Hafri, Papafragou, & Trueswell, 2013), even when the snapshot involves only the event outcome (specifically, the object affected by the action) in the absence of the agent (Hindy, Altmann, Kalenik, & Thompson-Schill, 2012;

Solomon, Hindy, Altmann, Thompson-Schill, 2015; Ünal & Papafragou, 2019).

Finally, the present study connects to discussions of event culmination in the linguistic literature. When asked whether someone colored a picture (perfective aspect), people are often likely to say *yes* even when the coloring is not complete (Arunachalam & Kothari, 2011; van Hout, 2018). This phenomenon has been observed in different languages (Jeschull, 2007; Li & Bowerman, 1998; Schulz & Penner, 2002; Weist, Wysocka, & Lyytinen, 1991) and characterizes both adults' and children's responses (van Hout, 2018; Jeschull, 2007; Schulz & Penner, 2002), but its origins are poorly understood. Previous linguistic studies have suggested that, even for perfective descriptions of events such as *color a picture* that have a definitive endpoint (as opposed to, say, *color pictures*), contextual factors seem to play a role in whether people think that the description applies to a half-finished event (Arunachalam & Kothari, 2011; van Hout, 2018). Such contextual effects are more likely to occur when the event is not totally incomplete or totally complete. Our study bears on this literature because it directly probes contextual effects on the interpretation of perfective event descriptions.

Experiment

Participants

Forty-three native English speakers were recruited from the Psychological and Brain Sciences department subject pool at the University of Delaware. Participation in the study fulfilled a course requirement.

Stimuli

A total of 54 images were included in the experiment. Of these, 36 were filler items: 18 depicted Incomplete (Visual Outcomes and 18 Complete Visual Outcomes for events (see Figure 1). The remaining 18 were target items and depicted Partly Complete Visual Outcomes (Figure 2). Images were assigned to Outcome types based on a prior norming study that asked a separate group of 25 participants to assess the percentage of event completion after each image was presented with the Neutral Goal context. For Incomplete Visual Outcomes, this estimated percentage was very low (M=7.91%), for Complete Visual Outcomes, it was very high (M=92.78%) and for Partly Complete Visual Outcomes, it was intermediate-low (M=27.02%). Unlike Incomplete Outcomes, Partly Complete Outcomes were cases where the event had been initiated

Each Visual Outcome image was followed by a test question (in perfective aspect) requiring a *Yes* or *No* answer (e.g., "Did she peel the orange?", for the target item in Figure 2). Each Visual Outcome image was preceded by a Context that stated the agent's goal (i.e., a sentence that included verbs such as *want, plan*, etc.). For each filler, there was a single Context (Figure 1). For each target (Partly Complete Visual Outcome), there were 3 possible Contexts: (a) Low Goal Contexts introduced an overarching goal that

could be satisfied even by a relatively modest degree of progress along the subevent timeline (e.g., "Jesse wants to use the orange as a garnish", where a small amount of peeling an orange can yield enough for a garnish); (b) High Goal Contexts introduced a goal for which a greater development of the subevent was needed (e.g., "Jesse wants to eat the orange for her breakfast", where the orange needs to be completely or almost completely peeled to be eaten); (c) Neutral Goal Contexts simply included the information later found in the test question ("Jesse wants to peel the orange").

Procedure

Visual Outcome (Incomplete, Complete, and Partly Complete) and - for Partly Complete Visual Outcomes only - Context (Low Goal, Neutral Goal, and High Goal) were within-subjects variables. Three lists were created by counterbalancing the Contexts for Partly Complete Outcomes so all participants saw a total of 54 trials: 18 involved Incomplete Outcomes, 18 Complete Outcomes, and 18 Partly Complete Outcomes (with 6 Partly Complete Outcomes in each of the 3 types of Context: Low Goal, High Goal, and Neutral Goal). Contexts and Partly Complete Outcome pairs varied across the lists so each participant saw only 1 instance of a given Outcome. The experiment was programmed and administered in OpenSesame. Trial order was randomized separately for each participant within the OpenSesame software.

Participants were asked to "read the following scenarios, look at the accompanying image, and answer each question" prior to beginning the experiment. Each trial began with a fixation point and participants moved on by pressing the spacebar on the keyboard. The Context sentence was then shown in the upper quarter of the screen. Participants were instructed to press the spacebar after reading the sentence. Next the Visual Outcome appeared below the Context sentence. The test question and response options ("Yes"/ "No") automatically appeared below the Visual Outcome after an additional 500ms. The Context sentence, Visual Outcome, and test question remained on screen until a response was given by pressing "d" for "Yes" and "k" for "No".
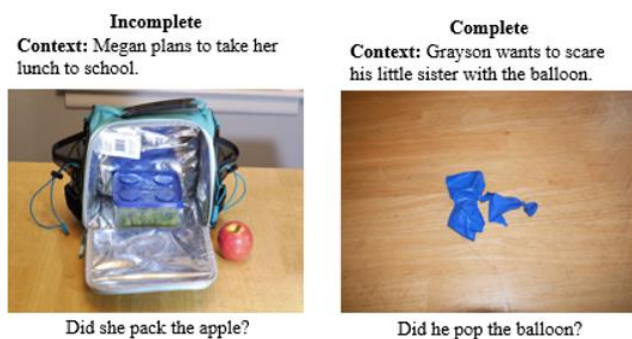

Figure 1. Example of an Incomplete and a Complete Visual Outcome and corresponding Contexts (filler items).

Figure 2. Example of a Partly Complete Visual Outcome and corresponding Contexts (target item).

Results

Response data were first submitted to a logit model with Visual Outcome contrasts (Partly Complete vs. Incomplete; Partly Complete vs. Complete) as predictors (see Figure 3). Responses were coded on a binary scale (yes = 1; no = 0). We collapsed across Contexts for the Partly Complete items for this analysis. Unsurprisingly, the observable degree of completion of the event affected whether viewers assessed that the event boundary had been reached: Partly Complete Visual Outcomes elicited *Yes* responses significantly more often than Incomplete Outcomes (z = -5.80, p < .001), and Complete Outcomes elicited *Yes* responses significantly more often than Partly Complete Outcomes (z = 6.32, p < .001; see Table 1). Visual information, therefore, clearly affected the placement of event boundaries.

A second analysis was run only on Partly Complete (i.e., target) Visual Outcomes using Context contrasts (Neutral Goal vs. Low Goal; Neutral Goal vs. High Goal) as the predictors (see Figure 4). Partly Complete Visual Outcomes paired with a Low Goal context elicited more *Yes* responses compared to those paired with a Neutral Goal context (z = 3.93, p < .001). No significant difference was found between Neutral Goal and High Goal contexts (z = -0.21; p > .05; Table 2).
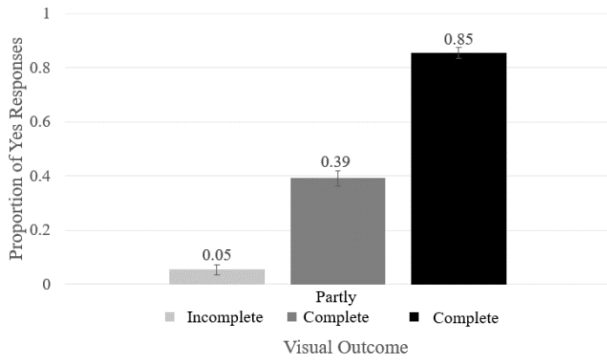
413

Figure 3. Proportion of Yes responses by Visual Outcome.

Table 1: Fixed effect estimates for the multi-level model of Yes responses by Visual Outcome.

Effect	Estimate	SE	z value
(Intercept)	-0.58	0.22	-2.65**
Visual Outcome (Partly Complete vs. Incomplete)	-3.32	0.57	-5.80***
Visual Outcome (Partly Complete vs. Complete)	3.41	0.54	6.32***

Formula in R: Response ~ 1 + Visual Outcome + (1|Subject) + (1|Item) + (0 + Visual Outcome | Subject)
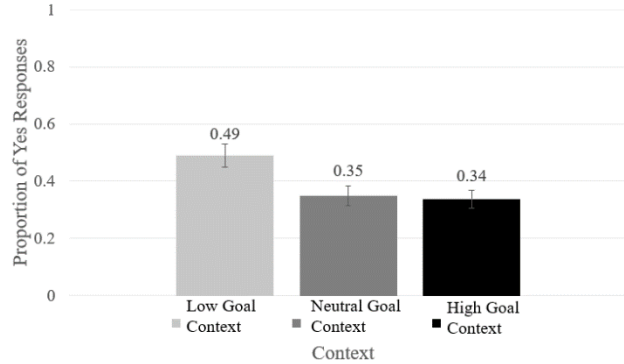
* $p < .05$, ** $p < .01$, *** $p < .001$

Figure 4. Proportion of Yes responses to Partly Complete Visual Outcomes by Context.

Table 2: Fixed effect estimates for the multi-level model of Yes responses to Partly Complete Visual Outcomes by Context.

Effect	Estimate	SE	z value
(Intercept)	-0.63	0.34	-1.85
Context (Neutral Goal vs. Low Goal)	0.86	0.22	3.93***
Context (Neutral Goal vs. High Goal)	-0.05	0.22	-0.21

Formula in R: Response ~ 1 + Goal Context + (1|Subject) + (1|Item)

* $p < .05$, ** $p < .01$, *** $p < .001$

Discussion

Streams of events make up every component of our daily lives, from making a cup of coffee in the morning to getting ready for bed at night. It has long been recognized that one cue that helps us to recognize when one event ends and another begins is the knowledge that an agent's goal has changed (Zacks & Tversky, 2001). Nevertheless, the role of higher-order goal information on event representations has not been addressed experimentally in detail. Similarly, psycholinguistic research has found that both children and adults sometimes accept sentences with perfective aspect ("She colored the picture") even for incomplete events, and has pointed to the possible role of extra-linguistic, contextual factors in determining event culmination, even though it has not examined such factors directly (van Hout, 2018; Arunachalam & Kothari, 2011). Here, using a novel paradigm, we asked whether prior knowledge of an agent's goals can affect viewers' placement of event boundaries (as assessed by viewers' answers to perfective questions, e.g., "Did she do X?").

We found that goal information affected endpoint placement for Partly Complete Visual Outcomes. Specifically, participants were more likely to place an event boundary at a Partly Complete Outcome if the Outcome satisfied the agent's goal. For target items, event boundaries did not change between Neutral Goal and High Goal contexts but did between Neutral Goal and Low Goal contexts. Given that Low Goal contexts differed from the Neutral baseline, while High Goal contexts did not, our findings suggest that Low Goal contexts shifted the placement of event boundaries that would otherwise be placed later in the event timeline. Put simply, events that are less-than-halfway complete (such as the one in Figure 2) are not considered culminated; however, in contexts where a contextual standard is satisfied by such an incomplete degree of event unfolding, viewers are more likely to consider that a culmination point has been reached.

One might find the lack of difference between Neutral and High Goal contexts puzzling. However, Neutral Goal contexts were not expected to represent a halfway point between the Low and High Goal contexts; indeed, much of prior literature has assumed that, even in the absence of a specific context, an event such as peel the orange would culminate at its natural endpoint (Filip, 2017). It seems that the High Goal contexts did not shift that expectation any further. We are currently exploring further interactions between context types and visual outcomes.

Our findings offer the first direct piece of evidence in support of the conclusion that higher-order intentionality information affects how viewers place event endpoints. Furthermore, given that our materials involved assigning boundaries to events denoted by telic (perfective) sentences, our results support theories of aspectual interpretation that

allow for a strong contribution of pragmatic-contextual factors to culmination interpretations of perfectivity (cf. Depraetere, 2007; Borer, 2005; Filip, 2017; cf. Arunachalam & Kothari, 2011).

Implications and Extensions

In important regards, our results are in line with the claim that change is an indicator of event boundaries (Speer, Zacks, & Reynolds, 2004; 2007; Zacks & Tversky, 2001) - more specifically, the idea that salient goals can affect the placement of event boundaries (e.g., Levine et al., 2017; Speer et al., 2004; 2007; Newtson, 1973; Vallacher & Wegner, 1987; Wilder, 1978). However, our results go beyond prior work that has focused on event segmentation and did not always decouple the contribution of goal changes from that of spatiotemporal cues to event boundaries. Most broadly, our results support the conclusion that event boundaries are determined by a variety of considerations, some of which may be very abstract (cf. also Zacks & Tversky, 2001; Ji & Papafragou, in press; Cohn & Paczynski, 2019).

The current study explored the effect of goal information on in-the-moment placement of event boundaries but can be extended to shed light on the role of event boundaries in other domains of cognition. Prior research has demonstrated that event boundaries can act as a marker for event memory. For example, event boundaries appear to facilitate the updating of event information in working, long term, and procedural memory (Kurby & Zacks, 2008); furthermore, objects located at event boundaries are remembered better than those located outside of an event boundary (Swallow, Zacks, & Abrams, 2009). Our findings raise the tantalizing possibility that conceptual cues to event boundaries such as an agent's goals or intentions (alongside perceptual cues to event boundaries such as motion) might influence event memory.

Concluding Remarks

In conclusion, we found that knowledge of an agent's goal affected whether an event would be considered as culminated or not: viewers were more likely to accept a Partly Complete Outcome as indicating culmination if the agent's goal was satisfied. This result strongly suggests that higher-order goal information affects the way events are conceptualized.

References

Arunachalam, S., & Kothari, A. (2011). An experimental study of Hindi and English perfective interpretation. *Journal of South Asian Linguistics*, *4*(1), 27-42.

Baldwin, D. A., Baird, J. A., Saylor, M. M., & Clark, M. A. (2001). Infants parse dynamic action. *Child Development*, *72*(3), 708–717. https://doi.org/10.1111/1467-8624.00310

Borer, H. (2005). *Structuring sense volume 2: The normal course of events.* Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199263929.001.0001

Brandone, A. C., & Wellman, H. M. (2009). You can't always get what you want. *Psychological Science*, *20*(1), 85–91. https://doi.org/10.1111/j.1467-9280.2008.02246.x

Cohn, N., & Paczynski, M. (2019). The neurophysiology of event processing in language and visual events. In R. Truswell (Ed.), *The Oxford Handbook of Event Structure*. Oxford: OUP.

Depraetere, I. (2007). (A)telicity and intentionality. *Linguistics*, *45*(2), 243-269. https://doi.org/10.1515/LING.2007.008

Filip, H. (2017). The semantics of perfectivity. Italian journal of linguistics, 29(1), 167-200. https://doi.org/10.26346/1120-2726-107

Hafri, A., Papafragou, A., & Trueswell, J. C. (2013). Getting the gist of events: Recognition of two-participant actions from brief displays. *Journal of Experimental Psychology: General*, *142*(3), 880-905. https://dx.doi.org/10.1037%2Fa0030045

Hindy, N. C., Altmann, G. T. M., Kalenik, E., & Thompson-Schill, S. L. (2012). The effect of object state-changes on event processing: Do objects compete with themselves? *Journal of Neuroscience, 32*, 5795–5803. https://doi.org/10.1523/JNEUROSCI.6294-11.2012

van Hout, A. (2018). On the acquisition of event culmination. In K. Syrett & S. Arunachalam (Eds.), *Semantics in Language Acquisition*. Amsterdam: John Benjamins. https://doi.org/10.1075/tilar.24.05hou

Jeschull, L. (2007). The pragmatics of telicity and what children make of it. *Proceedings of the 2nd Conference on Generative Approaches to Language Acquisition North America (GALANA)* (pp. 180-187). Cascadilla Press.

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

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

Levine, D., Hirsh-Pasek, K., Pace, A., & Michnick Golinkoff, R. (2017). A goal bias in action: The boundaries adults perceive in events align with sites of actor intent. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(6), 916–927. https://doi.org/10.1037/xlm0000364

Li, P., & Bowerman, M. (1998). The acquisition of lexical and grammatical aspect in Chinese. *First Language, 18*(54), 311-350. https://doi.org/10.1177/014272379801805404

Luo, Y., & Johnson, S. C. (2009). Recognizing the role of perception in action at 6 months. *Developmental Science*, *12*(1), 142–149. https://doi.org/10.1111/j.1467-687.2008.00741.x

Madden, C.J., Zwaan, R.A. (2003). How does verb aspect constrain event representations?. *Memory & Cognition, 31,* 663–672. https://doi.org/10.3758/BF03196106

Magliano, J., Kopp, K., McNerney, M. W., Radvansky, G. A., & Zacks, J. M. (2012). Aging and perceived event structure as a function of modality. *Aging, Neuropsychology, and Cognition, 19*(1-2), 264–282. https://doi.org/10.1080/13825585.2011.633159

Newtson, D. (1973). Attribution and the unit of perception of ongoing behavior. *Journal of Personality and Social Psychology*, *28*(1), 28–38. https://doi.org/10.1037/h0035584

Schulz, P., & Penner, Z. (2002). How you can eat the apple and have it too: Evidence from the acquisition of telicity in German. In J. Costa & M.J. Freitas (Eds.), *Proceedings of the GALA 2001 Conference on Language Acquisition* (pp. 239–246).

Solomon, S. H., Hindy, N. C., Altmann, G. T., & Thompson-Schill, S. L. (2015). Competition between mutually exclusive object states in event comprehension. *Journal of Cognitive Neuroscience, 27*, 2324–2338. https://doi.org/10.1162/jocn_a_00866

Speer, N. K., Zacks, J. M., & Reynolds, J. R. (2007). Human brain activity time-locked to narrative event boundaries: Research article. *Psychological Science*, *18*(5), 449–455. https://doi.org/10.1111/j.1467-9280.2007.01920.x

Speer, N. K., Zacks, J. M., & Reynolds, J. R (2004). Perceiving narrated events. *Proceedings of the 26th Annual Meeting of the Cognitive Science Society*, *26*(26), 1637.

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

Tversky, B., Zacks, J. M., & Hard, B. M. (2008). The structure of experience. In T. F. Shipley & J. M. Zacks (Eds.) *Understanding events*. New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195188370.003.0019

Ünal, E., & Papafragou, A. (2019). How children identify events from visual experience. *Language Learning and Development, 15*, 138-156. https://doi.org/10.1080/15475441.2018.1544075

Vallacher, R. R., & Wegner, D. M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review*, *94*(1), 3–15. https://doi.org/10.1037/0033-295X.94.1.3

Weist, R. M., Wysocka, H., & Lyytinen, P. (1991). A cross-linguistic perspective on the development of temporal systems. *Journal of Child Language*, *18*(01), 67. https://doi.org/10.1017/S0305000900013301

Wilder, D. A. (1978). Predictability of Behaviors, Goals, and Unit of Perception. *Personality and Social Psychology Bulletin*, *4*(4), 604–607. https://doi.org/10.1177/014616727800400422

Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, *69*(1), 1–34. http://dx.doi.org/10.1016/S0010-0277(98)00058-4

Zacks, J. M. (2004). Using movement and intentions to understand simple events. *Cognitive Science*, *28*(6), 979-1008. https://doi.org/10.1207/s15516709cog2806_5

Zacks, J. M., Braver, T. S., Sheridan, M. A., Donaldson, D. I., Snyder, A. Z., Ollinger, J. M., … Raichle, M. E. (2001). Human brain activity time-locked to perceptual event boundaries. *Nature Neuroscience*, *4*(6), 651–655. https://doi.org/10.1038/88486

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. 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(1), 29–58. https://doi.org/10.1037/0096-3445.130.1.29