Viewers' Sensitivity to Abstract Event Structure

Yue Ji (Jiyue@Udel.Edu)
Department of Linguistics & Cognitive Science, University of Delaware, 125 E. Main Street
Newark, DE 19716 USA

Anna Papafragou (Apapafragou@Udel.Edu)
Department of Psychological & Brain Sciences, University of Delaware, 105 The Green
Newark, DE 19716 USA

Abstract

Bounded and unbounded events differ in whether they include an inherent endpoint (Bach, 1986). Even though this distinction can be important for the way events are identified and processed, the literature on event cognition has not focused on such abstract aspects of event structure. In the present study, we asked whether viewers are sensitive to the distinction between bounded and unbounded events in a category learning task. Our results show that people were more successful in forming the category of bounded events than that of unbounded events. We discuss implications of this finding for event cognition.

Keywords: event structure; endpoint; boundedness

Introduction

Our experience of the world is intrinsically dynamic. To make sense of the complex flow of changes in our environment, we break continuous streams of experience into separate entities and classify such entities into different types.

Much work has focused on how people segment continuous experience into discrete units, i.e. events. The term "event" refers broadly to a temporal segment that has "a beginning and an ending" (Zacks & Tversky, 2001). People identify the boundaries of an event through tracking changes in perceptual features such as direction, location, or speed of action (e.g. an arrow hitting a target); more importantly, people encode events based on conceptual features, especially the goal-directedness or causal structure of the corresponding experience (e.g. a person on diet hitting a target; Zacks & Swallow, 2007). Event boundaries have a privileged status in memory and provide anchors for later learning and describing (Swallow, Zacks, & Abrams, 2009). In particular, the endpoint is conceptualized as a critical event component. For instance, when comparing two events, the resultant state (e.g. whether a ball knocked over the whole tower or just a few blocks) has more psychological weight than other perceptual features (e.g. the moving direction of the ball) (He & Arunachalam, 2016). In the well-studied domain of motion events, the goal of motion is more accurately encoded in both language and memory as opposed to other components such as the source (Lakusta & Landau, 2005, 2012; Papafragou, 2010; Regier & Zheng, 2007; Wagner, 2009). In addition, people tend to fill the gap between successive events within a causal chain by generating rapid inferences about the endpoint of the first event. In a study by Strickland and Keil (2011), after watching videos of someone launching an object (e.g. kicking a soccer ball) followed by the object's directed motion (e.g. the ball flying into the goal), participants mistakenly reported that they saw the moment of contact, i.e. the endpoint of the launching event, even when it was actually omitted from the display.

Despite the richness of the literature on event segmentation and the salience of endpoints in event perception, the nature of event endpoints has been less discussed. In most event-segmentation studies, the stimuli are actions by an intentional actor (e.g. a person putting up a tent) and the endpoint is taken to be obvious and well-defined (e.g. the moment the tent is put up). In studies of motion events, the endpoint appears similarly self-evident (and is typically the moment that a moving entity reaches the goal). However, across a broad range of events, the notion of endpoint is not always straightforward. Consider the following situations described by the two sentences in (1):

(1) a. The child played the Moonlight Sonata.
    b. The child played the piano.

There is subtle difference between (1a) and (1b). The event in (1a) comes to an end when the last note of the sonata was played. In contrast, it is hard to specify how or when the situation in (1b) ends — the child could stop playing at any point. The endpoint is inherent in the former event but is arbitrary in the latter. Such contrasts have been discussed extensively in the linguistic literature on aspect (i.e., the linguistic encoding of the internal temporal profile of events). In this literature, the distinction between the two sentences in (1) is captured by assuming that (1a) encodes an experience as a "bounded" event but (1b) encodes it as an "unbounded" event (Bach, 1986; Harley, 2003; Jackendoff, 1991). Bounded events have an internal structure with a "built-in terminal point" (Comrie, 1976), "climax" (Vendler, 1957) or "culmination" (Parsons, 1990), while unbounded events are homogenous, lacking internal development (Krifka, 1998). This linguistic distinction presumably has a non-linguistic counterpart in the way events are perceived and understood but to date, this connection has not been explored in detail.

Inspired by the rich linguistic research on how event endpoints are encoded in language (see Filip, 2004; Krifka, 1998, etc.), one can further identify two major types of consideration that determine whether an event is bounded or not. First, intuitions about boundedness may be due to the nature of the action. In particular, some actions lead to a change of state in the affected object, such that the endpoint is the resultant state (bounded events); other actions do not affect the object in a perceptible way or the change lacks a well-defined resultant state (unbounded events). The contrast is shown in the following example:

(2) a. The child dressed the teddy bear.
    b. The child patted the teddy bear.

(2a) describes a bounded event—the teddy bear was dressed when the child finished. (2b) describes an unbounded event—no predictable result followed from the child's patting. Although both events involve the same object, the difference in actions leads to the contrast in boundedness.

Second, intuitions about boundedness may be due to the nature of the affected object. Sometimes, there is a homomorphism between the affected object and the time course of the event (Dowty, 1991; Krifka, 1989), such that the changes in the object track or "measure out" the way the event develops (Tenny, 1987). When the object itself is quantified, the event is bounded. The contrast can be illustrated by the example below:

(3) a. The child ate a pretzel.
    b. The child ate cheerios.

(3a) depicts a bounded event— the event unfolds as the pretzel changes and it ends at the moment when the pretzel is gone. (3b) depicts an unbounded event that lacks an inherent endpoint—the child could stop at any time.

To sum up, bounded and unbounded events differ in whether they have an inherent endpoint. Two major components, i.e. the nature of the action and the affected object, might determine whether an inherent endpoint is available. So far the literature on event perception has not explored the role of boundedness in determining event boundaries, and little is known about whether viewers are sensitive to such abstract aspects of event cognition. One suggestive piece of evidence comes from work focusing on how events are counted. Bounded events are naturally counted in terms of how many inherent endpoints have been achieved. Lacking an inherent endpoint, unbounded events are counted according to spatio-temporal criteria. Returning to the example in (1), imagine that the child paused for a break and then resumed her playing in both situations. The event of playing the Moonlight Sonata still occurred once, but the child played the piano twice. When counting events like (1a), adults look for the inherent endpoint regardless of the pauses (Barner, Wagner & Snedeker, 2008) but 3-to-5-year-olds tend to over-generalize spatio-temporal criteria and count the number of pauses (Wagner, 2006; Wagner & Carey, 2003).

In the present paper, we explore viewers' sensitivity to the distinction between bounded and unbounded events (defined in terms of the availability of an inherent endpoint, as in (2) and (3) above). Specifically, we ask whether viewers can group events into the bounded vs. unbounded category in a category learning task. Drawing on the linguistic literature in which the category of unbounded events is definitionally dependent on the category of bounded events (such that boundedness and unboundedness form a positive-negative pair), we ask whether there is an asymmetric relation between the two types of event in non-linguistic cognition. If so, the category of bounded events might be learned by observers more easily compared to that of unbounded events.

Experiment 1

Experiment 1 was a category learning task. Participants were exposed to minimal pairs of bounded and unbounded events (defined by the availability of an inherent endpoint) and had to extract what was shared by different events of the same category and extend this information to new events.

Method

Participants Forty adults participated in the experiment. All were undergraduates at the University of Delaware and received course credit for participation. Data from an additional group of 2 adults were collected but excluded because these adults were color-blind and failed to identify an important test feature (a red frame) consistently.

Stimuli Twenty pairs of videos were created, such that each pair showed a bounded and an unbounded event (see Table 1). Within each pair, the videos had the same duration (range: 4.5s-13s; M = 7.98s) and involved the same actor but differed minimally from each other in one of two ways that involved boundedness. For half of the pairs, the bounded and unbounded events within a pair involved the same object but differed in terms of the nature of the action performed on the object: the bounded event displayed an action that caused a clear and temporally demarcated change of state in the object (e.g. fold up a handkerchief) while its unbounded counterpart did not involve such a change (e.g. wave a handkerchief). For the other half of the pairs, the bounded and unbounded events within a pair involved the same action but differed in terms of the nature of the affected object: the bounded event involved a single object (e.g. draw a circle) but its unbounded counterpart involved either an unspecified plurality of objects or a mass quantity (e.g. draw circles).

To ensure that all video stimuli would illustrate the contrast in boundedness presented in Table 1, a new group of 18 adults from the same population was asked to watch a subset of the clips and describe what happened in a full English sentence. For this norming task, the events in Table 1 were split into 2 lists, such that each list included only one member of each pair and an equal number of bounded and unbounded events. Each of the 18 participants was randomly assigned to one of the two lists. Their descriptions were coded for the verb used to describe the action and the noun phrase used to describe the affected object(s). As expected, differences in boundedness within a pair that were due to the nature of the action were reflected in verb choices: bounded stimuli elicited verbs of change of state (e.g. "dress a teddy bear") 98.3% of the time and unbounded stimuli elicited verbs denoting activities (e.g. “pat a teddy bear”) 93.1% of the time. Similarly, differences in boundedness within a pair that were due to the nature of the object were reflected in noun phrase choices: bounded events elicited count nouns with definite or indefinite articles (e.g. “eat the/a pretzel”) 98% of the time and unbounded events elicited bare plurals, mass nouns, or related devices (e.g. “eat cheerios”) 92.4% of the time.

For purposes of Experiment 1, the video stimuli were arranged into three basic lists corresponding to the three phases of the experiment (see Table 1). For the initial learning phase, we selected 8 pairs of events (4 in which boundedness was due to the Action and 4 in which boundedness was due to the Affected Object) and arranged them into a pseudorandomized presentation list in which a single video was played in the center of the screen and the two videos within a pair appeared in immediate succession (the order of bounded-unbounded events within pairs was counterbalanced within the list).

For the later testing phase, we arranged 8 of the remaining pairs of videos (see Table 1) into 2 lists. Each list contained one video from each pair. We counterbalanced whether the event was bounded or unbounded and whether source of boundedness was the action or the object across lists.

For the final (short) surprise testing phase, we used the last 4 pairs of videos, arranged into 2 lists. The same counterbalancing was used as in the (main) testing phase.

Procedure Participants were randomly assigned to one of two conditions. In the Bounded condition, the videos of bounded events shown in the learning phase were given a red frame while their unbounded counterparts were given a black frame. In the Unbounded condition, the reverse assignment occurred.

In the learning phase for both conditions, participants were asked to watch a few videos and to pay attention to those appearing within a red frame. Their task was to figure out what kind of videos were given the red frame and to decide whether a new video could have the red frame or not.

In the testing phase, participants saw a new set of videos and for each one they were asked: “Could the video have a red frame or not?” (test question) In the surprise testing phase, participants were unexpectedly asked: “Could the video have a black frame or not?” (surprise question) This question was included to probe whether participants formed any hypothesis about the secondary event category present within the experiment, even though it was not the target of the study.

After the end of the session, participants were asked to write down what kind of videos could have a red frame. This was used as an additional source of information about the category that participants had just formed.

Table 1: Videos used in Experiment 1.

Phase	Boundedness Source	No.	Bounded Events	Unbounded Events
Learning	Nature of Action	1	fold up a handkerchief	wave a handkerchief
		2	put up one’s hair	scratch one’s hair
		3	pile up a deck of cards	shuffle a deck of cards
		4	group pawns based on color	mix pawns of two colors
	Nature of Affected Object	5	draw a balloon	draw circles
		6	tie a knot	tie knots
		7	eat a pretzel	eat cereal
		8	flip a postcard	flip pages
Testing	Nature of Action	9	dress a teddy bear	pat a teddy bear
		10	roll up a towel	twist a towel
		11	fill a glass with milk	shake a bottle of milk
		12	scoop up yogurt	stir yogurt
	Nature of Affected Object	13	peel a banana	crack peanuts
		14	blow a balloon	blow bubbles
		15	tear a paper towel	tear slices off paper towels
		16	paint a star	paint stuff
Surprise Testing	Nature of Action	17	close a fan	use a fan for oneself
		18	crack an egg	beat an egg
	Nature of Affected Object	19	cut a ribbon in half	cut ribbon from roll (into many pieces)
		20	stick a sticker	stick stickers

Results

An ANOVA was performed on the proportion of correct responses to all questions with Source of Boundedness (i.e. Nature of Action vs. Nature of the Affected Object) as a within-subjects factor. No significant difference was found (F (1, 39) = .042, p = .838). Therefore, answers to questions targeting the two sources of boundedness were collapsed for further analysis.

Results from Experiment 1 are shown in Figure 1. The proportion of correct responses to test questions was significantly higher in the Bounded (M = 92.50%) than in the Unbounded condition (M = 76.25%) (t (38) = 3.563, p = .001). No significant difference in the proportion of correct responses to the surprise questions in the two conditions was found (t (38) = -.831, p = .411).

An ANOVA was conducted on the proportion of correct responses with Question Type (Test vs. Surprise) as a within-subjects factor and Condition (Bounded vs. Unbounded) as a between-subjects factor. There was a significant effect of Question Type (F (1, 38) = 19.795, p < 0.0001), no significant effect of Condition (F (1, 38) = 2.247, p = .142), and an interaction between the two factors (F (1, 38) = 7.833, p = .008). The participants were more accurate in test questions than in the surprise questions in the Bounded condition (t (19) = 6.114, p < .00001) but not in the Unbounded condition (t (19) = 1.022, p = .320).

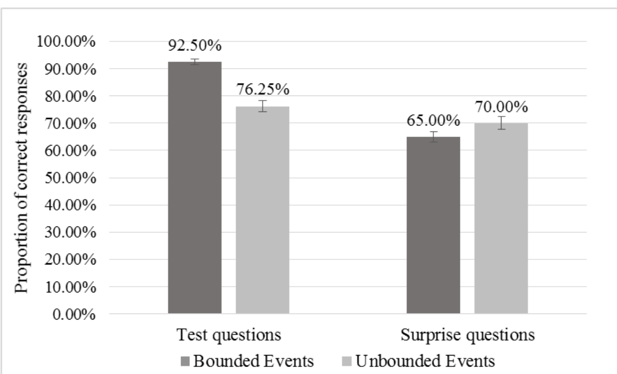

Figure 1: Proportion of correct responses in Experiment 1. Error bars represent standard error.

Answers to the last open question asking about the target category focused on 3 aspects of the stimuli—organization, neatness and intention. Organization was the most frequent hypothesis (29 out of the 40 answers). Specifically, modifiers such as "organized", or "structured" were used to describe bounded events while "unorganized", or "lacking structure" were given for unbounded events. Neatness was the second most frequent hypothesis (15 out of the 40 answers). Words used for bounded events included "neat", "tidy" and "clean" while those for unbounded events included "messy" and "untidy". Lastly, intention was mentioned in 9 out of the 40 answers. Bounded events were depicted as aiming "to achieve a goal", or being "on purpose" while unbounded events were "lacking an end or purpose", "random".

Discussion

Performance in test questions directly showed that, given the same contrastive examples in the learning phase, the participants were better at forming the category of bounded events compared to that of unbounded events. Furthermore, in the Bounded condition, learning was focused, with participants being less successful in the surprise compared to the test questions; however, no such asymmetry was found in the Unbounded condition. Further intuitions about boundedness were found in answers to the last open question about the nature of the target (red-frame) stimuli. The most frequent hypotheses referred to the organization of the stimuli. This suggests that participants attended to the internal structure of events when forming hypotheses about the meaning of the to-be-acquired category.

Experiment 2

In Experiment 1, participants might have benefited from the presentation of paired videos in the learning phase. By showing 2 successive videos with minimal differences, the contrast between bounded and unbounded events was highlighted. Experiment 2 asked whether the category of bounded or unbounded events could be efficiently extracted in a less supportive learning context.

Method

Participants A new group of forty undergraduates at the University of Delaware were recruited. Data from an additional adult were collected but excluded because he failed to understand the task and did not finish all the questions.

Stimuli and Procedure The stimuli and procedure were identical to those in Experiment 1 with one exception. In the learning phase, the sequence of the 16 videos was pseudo-randomized such that any 2 videos within a pair were separated by at least 5 other videos. This made it impossible to detect the contrast between bounded and unbounded events by simply comparing 2 consecutive videos.

Results

As in Experiment 1, no difference in the proportion of correct responses was found between the two sources of boundedness (F (1, 39) = 1.595, p = .214). The answers were thus collapsed in the following analysis.

Results from Experiment 2 are shown in Figure 2. Test questions elicited a significantly higher proportion of correct responses in the Bounded (M = 84.38%) than in the Unbounded condition (M = 68.13%) (t (38) = 3.365, p = .002). There was no significant difference between the two conditions in the proportion of correct responses to the surprise questions (t (38) = -1.129, p = .266).

An ANOVA conducted with Question Type as a within-subjects factor and Condition as a between-subjects factor showed a significant effect of Question Type (F (1, 38) = 7.095, p = .011), no significant effect of Condition (F (1, 38) = .646, p = .427), and an interaction between Question Type and Condition (F (1, 38) = 7.839, p = .008). The participants performed better in test questions than in surprise questions in the Bounded condition (t (19) = 4.174, p = .001), but not in the Unbounded condition (t (19) = -.093, p = .927).

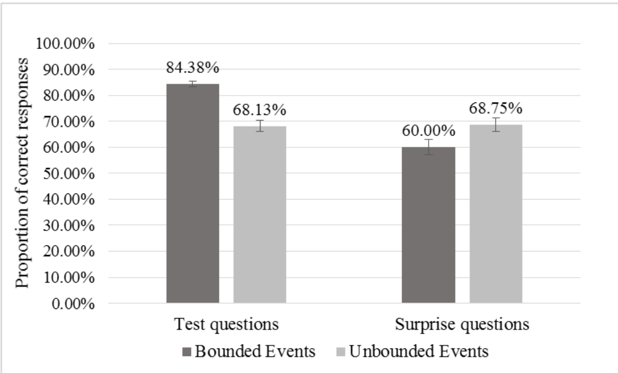

Figure 2: Proportion of correct responses in Experiment 2. Error Bars represent standard error.

Answers to the open question about the nature of the target category still mainly referred to organization, neatness and intention. These were mentioned in 16, 8 and 7 out of the 40 answers respectively. In addition, repetition was used to describe unbounded events in 5 answers. Completion appeared in 3 answers about bounded events.

As is clear from Figures 1-2, performance on the test questions was better in Experiment 1 than in Experiment 2. This was confirmed in an ANOVA that used the proportion of correct responses on the test questions as the dependent measure, and included Condition (Bounded vs. Unbounded) and Experiment (1 vs. 2) as between-subjects factors. The analysis showed main effects of Condition, (F (1, 76) = 23.940, p < .0001), and Experiment (F (1, 76) = 5.986, p = .017), and no interaction between the two factors (F (1, 76) = .000, p = 1.000). (Results were similar when accuracy on both test and surprise questions was used as the dependent measure.)

Discussion

Results from Experiment 2 showed a learning advantage for the category of bounded compared to unbounded events. This pattern was similar to Experiment 1, even though performance in Experiment 2 was worse compared to the earlier study, presumably because of the lack of direct contrast between bounded and unbounded events during the learning phase.

General Discussion

Our findings provide direct evidence for viewers' sensitivity to the abstract feature of boundedness in event cognition. In that sense, the present data go beyond prior work on how bounded and unbounded events are individuated and counted (Barner, Wagner & Snedeker, 2008; Wagner, 2006; Wagner & Carey, 2003). Furthermore, our results demonstrate that there is an asymmetry between bounded and unbounded events, such that it is easier to form the category of bounded compared to unbounded events. Our results raise the possibility that unboundedness is asymmetrically dependent on boundedness during event perception and apprehension, and that bounded - but not unbounded - events form a natural class.

The present data leave several directions open for further research. An important direction concerns the exact nature of the conjectures underlying participants' groupings of events into boundedness categories. The notion of boundedness is broad and can be subject to more abstract considerations than the present discussion has suggested. For instance, the inherent endpoint that defines bounded events can provided by a salient intention (Depraetere, 2007). To take an isolated example, even though the action of warming a soup does not have a clearly defined endpoint, it is often construed as culminating at the point at which the soup has reached someone's favorite temperature. In our study, it seems unlikely that intentionality was the feature responsible for participants' success in the Bounded condition. We asked a new group of 10 people to rate the degree of intentionality for all the videos used in the experiments on a scale from 1 (totally unintentional) to 7 (intentional). There was no significant difference between scores for bounded events (M = 5.829) and unbounded events (M = 5.704) (t (9) = 1.059, p = .330).

Finally, a number of researchers has drawn close parallels between object and event systems from a semantic perspective, such that the property of boundedness in the domain of events has been linked to the issue of quantification in the domain of objects (Bach, 1986; Jackendoff, 1991). In our study, the quantification of the affected object served as a cue for distinguishing bounded events from unbounded ones. It is possible that viewers are better at forming the category of bounded events because it is easier to track a single object compared with an unindividuated substance or objects of a variable number. An interesting further question is whether there is a common notion of boundedness underlying cognitive representations of both events and objects (see Wellwood, Hespos & Rips, in press) and how distinctions in one domain might generalize to the other.

Acknowledgments

This material is based upon work supported by the National Science Foundation under grant no. 1632849.

References

Bach, E. (1986). The algebra of events. *Linguistics and Philosophy, 9,* 5-16.

Barner, D., Wagner, L. & Snedeker, J. (2008). Events and the ontology of individuals: Verbs as a source of individuating mass and count nouns. *Cognition, 106,* 805-832.

Comrie, B. (1976). *Aspect*: *An Introduction to the Study of Verb Aspect and Related Problems.* Cambridge: CUP.

Dowty, D. R. (1991). Thematic proto-roles and argument selection. *Language*, *67*, 547-619.

Depraetere, I. (2007). (A)telicity and intentionality. *Linguistics, 45*, 243-269.

Filip, H. (2004). The telicity parameter revisited. *SALT, 14,* 92-109.

Harley, H. (2003). How do verbs get their names? Denominal verbs, manner incorporation and the ontology of verb roots in English. In N. Erteschik-Shir, & T. Rapoport (Eds.), *The Syntax of Aspect*. Oxford: OUP.

He, X., & Arunachalam, S. (2016). How event endstates are conceptualized [pdf document]. Retrieved from http://tedlab.mit.edu/~mekline/ELC2016/25_ELC2016-He&Arunachalam.pdf

Jackendoff, R. (1991). Parts and boundaries. *Cognition, 41*, 9-45.

Krifka, M. (1989). Nominal reference, temporal constitution and quantification in event semantics. In R. Bartsch, J. van Benthem, & P. van Emde Boas (Eds.), *Semantics and Contextual Expression, Groningen-Amsterdam Studies in Semantics 11*. Dordrecht: Foris Publications.

Krifka, M. (1998). The origins of telicity. In S. Rothstein (ed.), *Events and Grammar*. Dordrecht: Kluwer.

Lakusta, L., & Landau, B. (2005). Starting at the end: the importance of goals in spatial language. *Cognition, 96*, 1-33.

Lakusta, L., & Landau, B. (2012). Language and memory for motion events: Origins of the asymmetry between source and goal. *Cognitive Science*, *36*, 517-544.

Papafragou, A. (2010). Source-goal asymmetries in motion representation: Implications for language production and comprehension. *Cognitive Science, 34,* 1064-1092.

Parsons, T. (1990). *Events in the Semantics of English*: *A Study in Subatomic Semantics*. Cambridge, MA: MIT Press.

Regier, T., & Zheng, M. (2007). Attention to endpoints: A cross-linguistic constraint on spatial meaning. *Cognitive Science, 31,* 705-719.

Strickland, B., & Keil, F. (2011). Event completion: Event based inferences distorts memory in a matter of seconds. *Cognition, 121,* 409-415.

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

Tenny, C. (1987). *Grammaticalizing Aspect and Affectedness* (Doctoral dissertation). MIT, Cambridge, MA.

Vendler, Z. (1957). Verbs and times. *The Philosophical Review, 66*, 143-160. Reprinted as Chapter 4 of Z. Vendler (Ed.), *Linguistics in Philosophy* (1967).

Wagner, L. (2006) Aspectual bootstrapping in language acquisition: Telicity and transitivity. *Language Learning and Development, 2*, 51-76.

Wagner, L. (2009). Manners and goals in pre-linguistic thought: the origins of aspectual construal. *Proceedings of the 33th Annual BU Conference on Language Development* (pp. 599-610). Somerville: Cascadilla Press.

Wagner, L., & Carey, S. (2003). Individuation of objects and events: a developmental study. *Cognition, 90*, 163-191.

Wellwood, A., Hespos, S. J., & Rips, L. (in press). The object : substance :: event : process analogy. In T. Lombrozo, S. Nichols, & J. Knobe (Eds.) *Oxford Studies on Experimental Philosophy*. New York, NY: OUP.

Zacks, J., & Swallow, M. (2007). Event segmentation. *Current Directions in Psychological Science*, *16,* 80-84.

Zacks, J., & Tversky, B. (2001). Event structure in perception and conception. *Psychological Bulletin*, *127*, 3-21.