Published in final edited form as:

Curr Dir Psychol Sci. 2007 April ; 16(2): 80–84. doi:10.1111/j.1467-8721.2007.00480.x.

EVENT SEGMENTATION

Jeffrey M. Zacksⁱ and Khen M. Swallow

Washington University in Saint Louis

Abstract

One way to understand something is to break it up into parts. New research indicates that segmenting ongoing activity into meaningful events is a core component of ongoing perception, with consequences for memory and learning. Behavioral and neuroimaging data suggest that event segmentation is automatic and that people spontaneously segment activity into hierarchically organized parts and sub-parts. This segmentation depends on the bottom-up processing of sensory features such as movement, and on the top-down processing of conceptual features such as actors' goals. How people segment activity affects what they remember later; as a result, those who identify appropriate event boundaries during perception tend to remember more and learn more proficiently.

Keywords

event perception; segmentation; motion; intentions

EVENT SEGMENTATION

Look at the scene depicted in Figure 1. Though most viewers will eventually figure out what is shown, many will have an easier time understanding the alternative version shown in Figure 2. What's the difference? The first picture fractures the scene in a way that obscures its natural part structure, whereas the second respects that structure. For quite a while psychologists have known that in order to recognize or understand an object people often segment it into its spatial parts (e.g., Biederman, 1987). A new body of research has shown that just as segmenting in *space* is important for understanding objects, segmenting in *time* is important for understanding events.

Event segmentation is the process by which people parse a continuous stream of activity into meaningful events. Recent developments in perceptual psychology and cognitive neuroscience have provided new insights into the role of event segmentation in human cognition. We will review three. First, event segmentation appears to be an automatic, ongoing component of human perception. Second, segmentation during perception scaffolds later memory and learning. Finally, specialized neural mechanisms identify event boundaries by tracking significant changes in physical and social features.

SEGMENTATION IS AUTOMATIC

Much of the research on event segmentation has used variants of a procedure developed by Newton (1976). In this procedure participants watch a movie of some activity and press a button whenever—in their judgment—one meaningful event ends and another event begins. This task produces event boundary judgments that are reliable across viewers and within

ⁱ Address correspondence to Jeffrey M. Zacks, Department of Psychology, Washington University, Saint Louis MO 63130. jzacks@artsci.wustl.edu.

viewers across time (Newtson, 1976; Speer, Swallow, & Zacks, 2003). The boundaries they identify tend to be readily nameable “chunks,” corresponding to subgoals an actor is attempting to satisfy in order to fulfill the larger goal of the activity. For example, boundaries in “putting up a tent” might be placed when the tent is staked out and when the rain fly is attached (see Figure 3). Event boundaries are hierarchically structured, such that fine-grained events are clustered into larger coarse-grained events. For example, observers segmenting a movie of a woman making a bed tend to identify events such as removing individual pillowcases and also to identify the removal of all the pillowcases as a larger event (Zacks, Tversky, & Iyer, 2001). The consistency and structure of these results suggest that the Newtson procedure taps into ongoing naturally occurring perceptual processing. But there is a problem: The task requires that observers attend to event boundaries and make decisions about where they occur. Such task demands may change the nature of the perceptual processing involved.

Stronger evidence that event segmentation is naturally ongoing comes from indirect measures, particularly functional neuroimaging studies. In one experiment (Zacks et al., 2001), participants viewed a series of movies of everyday activities (e.g. washing dishes, fertilizing a houseplant) while changes in brain activity were recorded with functional MRI. After first passively viewing the movies, participants segmented the same movies, twice, to identify event boundaries at two temporal grains. *Fine* boundaries marked the smallest events the participants found natural and meaningful, and *coarse* boundaries marked the largest events they found natural and meaningful. These boundaries were then used as markers to analyze the brain activity data from the initial passive viewing session. During passive viewing, regions of posterior and frontal cortex showed transient increases in activity that began several seconds before each event boundary and peaked several seconds after the boundary. Responses were larger for coarse than fine boundaries. Because the critical brain data were acquired before participants learned of the segmentation procedure, these changes cannot be attributed to overt or covert performance of a laboratory-specific task. These results strongly imply that brain processes correlated with event segmentation are a normal part of ongoing perception.

The preceding data make a compelling case that event segmentation is automatic. However, segmentation still may be affected by observers’ attention and goals. Indeed, there is evidence that observers can adapt their performance of the button-pressing segmentation task based on situational needs. For example, observers adjust the temporal grain of their segmentation based on explicit instructions, the sort of information they are trying to learn from a stimulus, and how much they know about the activity they are watching (see Zacks & Tversky, 2001 for a review). An important question for future research is whether these variations in overt task-related behavior reflect changes in ongoing perception or changes in the decision processes that are specific to the button-pressing task.

SEGMENTATION GUIDES MEMORY AND LEARNING

One important consequence of perceptual segmentation is that the resulting segments can form the basis of memory and learning. When it comes to remembering the details of what has recently happened, new data indicate that those individuals who are better able to segment an activity into events are better able to remember it later (Zacks, Speer, Vettel, & Jacoby, 2006). In this study, older adults segmented movies of everyday events (e.g., setting up a tent, planting a flower bed). Each person’s segmentation was compared to the segmentation of the other observers to determine whether they identified event boundaries that were similar to those of the others or placed their event boundaries in idiosyncratic locations. Although there is no gold standard for establishing whether one has segmented a movie correctly, observers generally agree on the locations of event boundaries. So, if a

particular observer's segmentation deviates from this norm, chances are something is amiss. All participants later completed a test requiring them to discriminate still pictures taken from the movies from similar pictures that were not from the movies. Those participants who segmented "well" showed better memory for the visual contents. Importantly, this strong relationship was observed above and beyond the effects of individual differences in cognitive ability and the presence of senile dementia.

Segmenting an activity well is not simply a matter of identifying the right event boundaries; it also requires tracking how sets of fine-grained events group together into larger meaningful units. Recent studies suggest such grouping is important for learning new activities (Hard, Lozano, & Tversky, in press; Lozano, Hard, & Tversky, in press). In these studies participants watched movies demonstrating simple assembly tasks and segmented them at both a fine and coarse grain. They then were asked to perform the assembly task they had just watched. The experimenters analyzed the degree to which participants grouped events into hierarchical units. A number of experimental manipulations affected the degree of hierarchical segmentation, and in all cases segmenting hierarchically was associated with better performance of the learned task. Together, these data suggest that event boundaries form anchors for long-term memory, and that interventions that encourage people to identify appropriate event boundaries can improve memory for what has happened and the learning of new skills.

NEURAL AND INFORMATION PROCESSING MECHANISMS

The previous sections argue that event segmentation is an automatic component of normal perception that shapes how people remember and learn. How does the brain do this segmentation? Evidence indicates that the brain and mind track features of one's environment, and when a salient feature changes unpredictably an event boundary is perceived (Zacks, Speer, Swallow, Braver, & Reynolds, in press). The critical features may include sensory features, such as color, sound, and movement, and conceptual features, such as cause-and-effect interactions and actors' goals. Sensory features likely are processed in a primarily bottom-up fashion, where the nature of the processing is driven primarily by perceptual input. Processing conceptual features, however, likely relies on top-down processing that integrates an observer's representation of the current event with previously stored knowledge. For example, segmenting events based on an actor's goals requires maintaining a representation of those goals over time, and often will depend on prior knowledge about the actor's dispositions and abilities.

Sensory Features: The Movement Of Objects and People

One hint at the sensory features that are important for event segmentation came from the neuroimaging study described previously (Zacks et al., 2001). In that study, the most active sites of transient brain responses at event boundaries appeared to be in visual processing areas known to process movement information. These areas collectively are called the human MT complex (MT+), because they are thought to be homologs of motion-sensitive areas in the medial temporal cortex of the macaque monkey (hence MT). A follow-up to the Zacks et al. (2001) study identified MT+ in individual observers and confirmed that the areas activated by event boundaries did include this area (Speer, Swallow, & Zacks, 2003), suggesting that people use movement cues to identify those boundaries. These results motivated recent experiments exploring the quantitative relationship between movement features and event segmentation.

In one set of experiments (Zacks, 2004), participants viewed simple animations depicting two objects moving around the screen and segmented them to identify fine or coarse event boundaries. The animations were generated either by asking people to play a video game in

which one player controlled each object and tried to achieve some goal (e.g., chase the other object and catch it), or by a random algorithm that produced animations with similar velocities and accelerations. The animations were analyzed to quantify movement features such as the speed and acceleration of each object. In all experimental conditions people tended to segment at points when movement features changed—for example, when objects were accelerating quickly, or when they reached a point of being maximally close to each other and turned away. (See also Hard, Lozano, & Tversky, in press). Thus, movement changes are strongly related to event segmentation.

If the processing of movement information in MT+ contributes to the segmentation of activity into events, then one should expect MT+ to track both movement information and event boundary locations when people view animations such as these. A recent neuroimaging study indicates that this is the case (Zacks, Swallow, Vettel, & McAvoy, 2006). In this experiment participants viewed simple animations while brain activity was recorded. Separate scans were used to identify MT+ in individual observers. Activity in MT+ increased as the velocity of the objects increased and also at event boundaries. Thus, the processing of movement information appears to be well situated to play a causal role in the detection of event boundaries.

Conceptual Features: Actors' Goals

Although movement features account for a substantial part of where people segment activity, they are not the whole story. Movement is more strongly tied to segmentation when viewers identify fine-grained units or segment random animations than when they identify coarse-grained units or segment animations depicting intentional activity (Zacks, 2004). This suggests that people depend on other sources of information, such as inferences about actors' intentions and goals, to understand the larger structure of activity. One piece of direct evidence for the importance of goals in event segmentation comes from infant perception (Baldwin, Baird, Saylor, & Clark, 2001). In this study, infants were familiarized with one of two movies depicting a woman cleaning a kitchen. Each movie depicted a salient goal-directed action (e.g., replacing a fallen dishtowel or storing an ice cream container in the freezer). After the familiarization phase, infants were presented with excerpts with one-second pauses inserted into the movie. The pauses were placed either at the moment when the woman achieved the action's goal, or several seconds before. The infants looked longer at the excerpts when the pauses were placed before the goal completions, suggesting that they found those more disruptive.

Another piece of evidence for the importance of goals comes from a neuroimaging study of events in texts (Speer, Reynolds, & Zacks, in press). In this experiment participants read narrative texts describing the activities of a small boy while brain activity was recorded, and then segmented the texts into events. The narratives were coded to localize changes in a number of features, including changes in the characters' spatial locations and in the characters' goals. Event boundaries in the narratives were associated with brief increases in brain activity that were similar in timing and location to those for live-action movies. Many of these areas also responded to changes in the narrative features—and the brain responses to event boundaries could be entirely accounted for by the responses to the narrative features. This suggests that both physical movement features (changes in location) and changes in actors' goals play important roles in the segmentation of activity into events.

These studies begin to provide the database for a mechanistic account of how observers segment ongoing activity into events. However, the available data afford only the barest outlines of such an account. We regard the detailed characterization of the relation between bottom-up and top-down processing in event segmentation as one important goal for future research. Further, we believe that a number of little-studied features, from purely sensory to

purely conceptual, must be important for event segmentation. Toward the sensory end are features such as sound, lighting, and contact between actors and objects. Toward the conceptual end are features such as goals and social conventions. In the middle are features such as statistical dependencies amongst events. The systematic exploration of these bases for segmentation is a second important research goal.

IN THE COURSE OF EVENTS

The previous sections have reviewed recent evidence supporting three conclusions about event perception. First, event segmentation is an automatic component of ongoing perceptual processing. Second, how people segment activity online has significant effects on how they remember it later: Events form the units of memory encoding, so identifying the right events leads to good memory and learning but identifying the wrong events leads to poor memory and learning. Finally, there are specialized neural systems that process features including movement and goals in order to use changes in those features to identify event boundaries. These findings have implications for education and for clinical practice. For education, they suggest that interventions that help people appropriately segment events will help them remember and learn from those events (see Segmentation Guides Learning and Memory, above). For clinical practice, these findings suggest that some cognitive deficits may reflect impaired event segmentation. A small number of studies indicate that event segmentation is impaired in patients with lesions to prefrontal cortex, schizophrenia, and mild Alzheimer-type dementia (see Zacks, Speer, Swallow, Braver, & Reynolds, in press for a review). As discussed previously, the last of these studies established a link between event segmentation and later memory, raising the possibility that it may be possible to remediate some memory deficits by improving segmentation.

Segmentation is a powerful perceptual operation. By reducing a continuous flux of activity to a modest number of discrete events a perceiver can achieve terrific economy of representation for perception and later memory. Segmentation not only is economical—it also allows one to think about different things in relation to each other, generatively, which is notoriously difficult with continuous, unsegmented representations. For this reason, people generally perceive space as consisting not of continuous gradations of color and texture but of spatially coherent objects. The same holds in time: Just as much as our everyday perceptual world is made up of objects, it is made up of events.

Acknowledgments

Preparation of this article was supported by National Institute of Mental Health Grant MH70674.

REFERENCES

- Baldwin DA, Baird JA, Saylor MM, Clark MA. Infants parse dynamic action. *Child Development*. 2001; 72(3):708–717. [PubMed: 11405577]
- Biederman I. Recognition-by-components: A theory of human image understanding. *Psychological Review*. 1987; 94(2):115–117. [PubMed: 3575582]
- Hard BM, Lozano SC, Tversky B. Hierarchical encoding of behavior: Translating perception into action. *Journal of Experimental Psychology: General*. (in press).
- Lozano SC, Hard BM, Tversky B. Perspective-taking promotes action understanding and learning. *Journal of Experimental Psychology: Human Perception & Performance*. (in press).
- Newton, D. Foundations of attribution: the perception of ongoing behavior. In: Harvey, JH.; Ickes, WJ.; Kidd, RF., editors. *New directions in attribution research*. Hillsdale, New Jersey: Lawrence Erlbaum Associates; 1976. p. 223–248.
- Speer NK, Reynolds JR, Zacks JM. Human brain activity time-locked to narrative event boundaries. *Psychological Science*. (in press).

- Speer NK, Swallow KM, Zacks JM. Activation of human motion processing areas during event perception. *Cognitive, Affective & Behavioral Neuroscience*. 2003; 3(4):335–345.
- Zacks JM. Using movement and intentions to understand simple events. *Cognitive Science*. 2004; 28(6):979–1008.
- Zacks JM, Braver TS, Sheridan MA, Donaldson DI, Snyder AZ, Ollinger JM, et al. Human brain activity time-locked to perceptual event boundaries. *Nature Neuroscience*. 2001; 4(6):651–655.
- Zacks JM, Speer NK, Swallow KM, Braver TS, Reynolds JR. Event perception: A mind/brain perspective. *Psychological Bulletin*. (in press).
- Zacks JM, Speer NK, Vettel JM, Jacoby LL. Event understanding and memory in healthy aging and dementia of the Alzheimer type. *Psychology & Aging*. 2006; 21(3):466–482. [PubMed: 16953710]
- Zacks JM, Swallow KM, Vettel JM, McAvoy MP. Visual movement and the neural correlates of event perception. *Brain Research*. 2006; 1076(1):150–162. [PubMed: 16473338]
- Zacks JM, Tversky B. Event structure in perception and conception. *Psychological Bulletin*. 2001; 127(1):3–21. [PubMed: 11271755]
- Zacks JM, Tversky B, Iyer G. Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology: General*. 2001; 130(1):29–58. [PubMed: 11293458]

RECOMMENDED READINGS

- Baldwin, DA.; Baird, JA. Action analysis: A gateway to intentional inference. In: Rochat, P., editor. *Early social cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates; 1999. p. 215-240.
- Newton D. 1976 See references.
- Zacks JM, Braver TS, Sheridan MA, Donaldson DI, Snyder AZ, Ollinger JM, et al. 2001 See references.
- Zacks JM, Speer NK, Swallow KM, Braver TS. (in press). See references.
- Zacks JM, Tversky B. 2001 See references.

Figure 1.
How easily can you describe this scene? This may be difficult since it is broken up in a way that obscures its natural part structure.

Figure 2.
This version may be easier to recognize. Though the scene is still broken up, the natural parts are preserved.

Figure 3.

Example of event boundaries. These frames show the six coarse-grained event boundaries selected most frequently by a group of younger and older adults (Zacks, Speer, Vettel, & Jacoby, 2006, exp. 2). These boundaries marked the ends of events that could be described as: 1) Put down the tent. 2) Spread it out. 3) Insert the front tent pole. 4) Stake out the ends of the tent. 5) Stake out the sides. 6) Attach the rain fly. [Boundaries were identified by estimating the continuous probability of segmentation over time using a gaussian smoothing kernel, bandwidth = 5 s, and selecting the six highest local maxima in the resulting probability distribution.]