

Segmentation in the perception and memory of events

Christopher A. Kurby and Jeffrey M. Zacks

Department of Psychology, Washington University, 1 Brookings Drive, St Louis, MO 63130, USA

People make sense of continuous streams of observed behavior in part by segmenting them into events. Event segmentation seems to be an ongoing component of everyday perception. Events are segmented simultaneously at multiple timescales, and are grouped hierarchically. Activity in brain regions including the posterior temporal and parietal cortex and lateral frontal cortex increases transiently at event boundaries. The parsing of ongoing activity into events is related to the updating of working memory, to the contents of long-term memory, and to the learning of new procedures. Event segmentation might arise as a side effect of an adaptive mechanism that integrates information over the recent past to improve predictions about the near future.

Making sense by segmenting

Imagine walking with a friend to a coffee shop. If asked to describe this activity in more detail you might list a few of the events that make it up. The events listed could be broken up by changes in the physical features of the activity, such as location: ‘We started out by going down to the laboratory. We grabbed our coats and put them on. Then we walked out of the building to the corner by the subway station. . .’ Or, they could be broken up by changes in conceptual features, such as your goals: ‘We started our walk talking about how much construction is going on. When the topic turned to the new building with the coffee shop we decided to head over there to give it a try. . .’ Such descriptions are typical of how people talk about events, and they illustrate something important about perception: people make sense of a complex dynamic world in part by segmenting it into a modest number of meaningful units. Recent research on event perception reveals that, as an ongoing part of normal perception, people segment activity into events and subevents. This segmentation is related to core functions of cognitive control and memory encoding, and is subserved by isolable neural mechanisms.

Events and their boundaries

By ‘event’ we mean a segment of time at a given location that is conceived by an observer to have a beginning and an end [1]. In particular we focus on the events that make up everyday life on the timescale of a few seconds to tens of minutes – things like opening an envelope, pouring coffee into a cup, changing the diaper of a baby or calling a friend on the phone. Event Segmentation Theory (EST) [2] (see Glossary) proposes that perceptual systems spontaneously

segment activity into events as a side effect of trying to anticipate upcoming information (see Box 1). When perceptual or conceptual features of the activity change, prediction becomes more difficult and errors in prediction increase transiently. At such points, people update memory representations of ‘what is happening now’. The processing cascade of detecting a transient increase in error and updating memory is perceived as the subjective experience that a new event has begun.

Segmentation tasks

How can a researcher discover when a person perceives that a new event has begun? One simple but surprisingly powerful answer is simply to ask them, usually by having them press a button [3]. Viewers tend to identify event boundaries at points of change in the stimulus, ranging from physical changes, such as changes in the movements of the actors, to conceptual changes, such as changes in goals or causes. To investigate movements of actors, Newton *et al.* [4] asked participants to segment movies of an actor conducting everyday activities. The physical pose of the actor was coded at one-second intervals. Event boundaries tended to be marked at larger changes in pose. Zacks [5] extended this qualitative finding to a quantitative analysis of movement variables in a set of studies using simple animated stimuli. Event boundaries were predicted by changes in movement parameters including the acceleration of the objects and their location relative to one another (see also Ref. [6].) Larger physical changes have been studied using commercial cinema as a stimulus, in which the locations and temporal setting of characters can change from shot to shot. In one study, such changes were found to predict where viewers segmented Hollywood movies [7]. Conceptual changes that are correlated with event boundaries include changes in goals of actors, in causal relations and in interactions amongst characters [8,9].

Automatic event segmentation

Does asking people for conscious judgments about event boundaries really tell us anything about ongoing

Glossary

Event model: an actively maintained representation of the current event, which is updated at perceptual event boundaries.

Event segmentation: the perceptual and cognitive processes by which a continuous activity is segmented into meaningful events.

Temporal grain: events can be perceived on a range of temporal grains, or timescales, from a second or less to tens of minutes.

Corresponding author: Zacks, J.M. (jzacks@artsci.wustl.edu).

Box 1. One step ahead of the game

Reacting quickly to changes is good – but anticipating them is better. Humans and many other animals make predictions about the near future when perceiving and acting. For example, when watching someone wrap a present you can make precise predictions about the trajectories of the arms based on the recent paths of the arms combined with previously learned information about how bodies typically move. Or, you can make broader predictions about sequences of intentional actions based on the recent sequence of actions combined with knowledge about present-wrapping and inferences about the current goals of the actor. Sequential predictions depend in part on a neural pathway involving the anterior cingulate cortex, which might represent the disparity between predictions and outcomes [54], and subcortical catecholamine nuclei, which might broadcast error signals throughout the brain [55,56]. Event Segmentation Theory (EST, [2]) proposes that the perception of event

boundaries arises as a side effect of using prediction for perception. According to EST, perceivers form working memory representations that capture ‘what is happening now’, called event models. When predictions are accurate, event models are maintained in a stable state to guide prediction, integrating information over the recent past. When prediction errors transiently increase relative to their current baseline, event models are updated based on currently present information (Figure 1). This updating affects what information is currently actively maintained, and the transient increase in processing during updating produces more robust encoding into long-term memory. The error-based updating mechanism has been modeled in a set of connectionist simulations [57]. If the environment contains sequences of inputs that recur in the training set, the updating mechanism can use such predictable sequential structure to improve prediction.

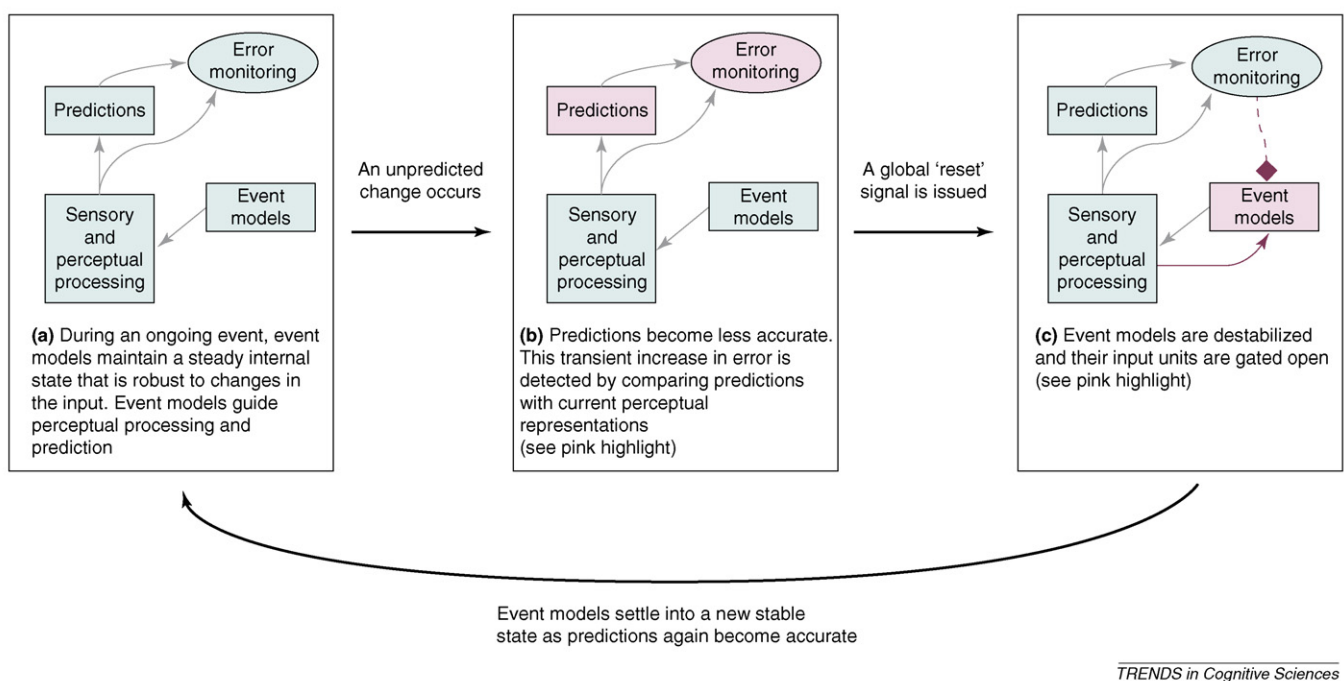


Figure 1. A schematic depiction of how event segmentation emerges from perceptual prediction and the updating of event models. (a) Most of the time, sensory and perceptual processing leads to accurate predictions, guided by event models that maintain a stable representation of the current event. Event models are robust to moment-to-moment fluctuations in the perceptual input. (b) When an unexpected change occurs, prediction error increases and this is detected by error monitoring processes. (c) The error signal is broadcast throughout the brain. The states of event models are reset based on the current sensory and perceptual information available; this transient processing is an event boundary. Prediction error then decreases and the event models settle into a new stable state.

perception outside the laboratory? Event segmentation tasks have good intersubjective agreement [10] and reliability [11], which suggests they tap into ongoing processing. Nonetheless, a basic limitation of directly applying segmentation tasks is that they might interfere with the ongoing perceptual processes they attempt to measure. Stronger evidence that event segmentation is automatic comes from implicit behavioral measures and from neurophysiological measures that require no overt task.

Reading-time evidence for automaticity of event segmentation

Studies of the pace of reading during narrative text comprehension indicate that readers slow down at event boundaries. These studies arise mostly from discourse comprehension theories, which propose that readers construct a series of mental models of a situation described by a narrative. In the Structure Building Framework [12],

new mental models are initiated when the text refers to a new set of people, places or things, when the action in the text changes its temporal or spatial location, or when a new causal sequence is initiated. In the Event Indexing Model [13], new mental models are initiated when there is a change in space, time, protagonist, objects, goals or causes. When the changes identified by these models occur, readers have been found to read more slowly. This is consistent with the processing load hypothesis put forward by the Event Indexing Model, which states that changes in situational features increase the difficulty of integrating newly encountered information into the current mental model. In one such study [14], participants read two literary narratives one sentence at a time on a computer screen, and reading time was recorded. The texts were coded to identify changes in time, space and causal contingency. Reading time increased consistently at changes in time and cause; reading time sometimes increased at spatial changes, but

this depended on the previous knowledge of the reader and task goals. More recent work has found that reading time increases at shifts in characters and their goals [15,16]. Such results support the notion that reading times increase at event boundaries because, as noted previously, event boundaries are associated with changes in time, space, causes, characters and goals. Although most reading-time studies have not measured event segmentation, two recent studies have directly compared event segmentation and narrative reading time. In one [17], cues to event boundaries were experimentally manipulated by changing the temporal contiguity of events (Figure 1). Phrases marking a temporal discontinuity ('an hour later') increased the likelihood that a clause would be identified as an event boundary during segmentation, and slowed reading during self-paced comprehension. The second study looked at the relation between event segmentation and reading time correlationally [8]. Two groups of participants read narratives about a young boy. One group segmented them into events; the other read them one clause at a time while reading time was recorded. Clauses in which more readers identified event boundaries produced longer reading times. These results support the view that when comprehenders encounter boundaries between events they perform extra processing operations.

Neurophysiological evidence for automaticity of event segmentation

Converging evidence for the automaticity of event segmentation comes from noninvasive neurophysiological measures, including functional magnetic resonance imaging (fMRI) and electroencephalography (EEG). These

studies have been motivated by the hypothesis that if the brain is undertaking processing that relates to event boundaries, transient changes in brain activity should be observed at those points in time corresponding to event boundaries – whether or not you are attending to event boundaries – or not you are attending to event boundaries. In one study [18], participants passively viewed short movies of everyday activities while their brain activity was recorded with fMRI. During the initial viewing and fMRI data recording, participants were asked simply to watch the movies and try to remember as much as possible. In the second phase of the experiment these participants segmented the movies into events. Event boundaries were associated with increases in brain activity during passive viewing in bilateral posterior occipital, temporal and parietal cortex and right lateral frontal cortex. The posterior activation included the MT complex [11], an area associated with the processing of motion [19]. A subsequent study using simple animations of geometric objects found that activity in the MT complex was correlated with the speed of motion of objects, and with the presence of event boundaries [20].

Whereas most of the neurophysiological research on event perception involves visual events, a new study by Sridharan *et al.* [21] investigated the perception of event structure in music. This study examined the extent to which musically untrained listeners use transitions between movements to segment classical pieces into coarse-grained events. A particular advantage in studying musical movements is that they provide objective, normative events. As Figure 2 illustrates, two dissociable networks in the right hemisphere were selectively responsive at transitions between movements: an early-responding

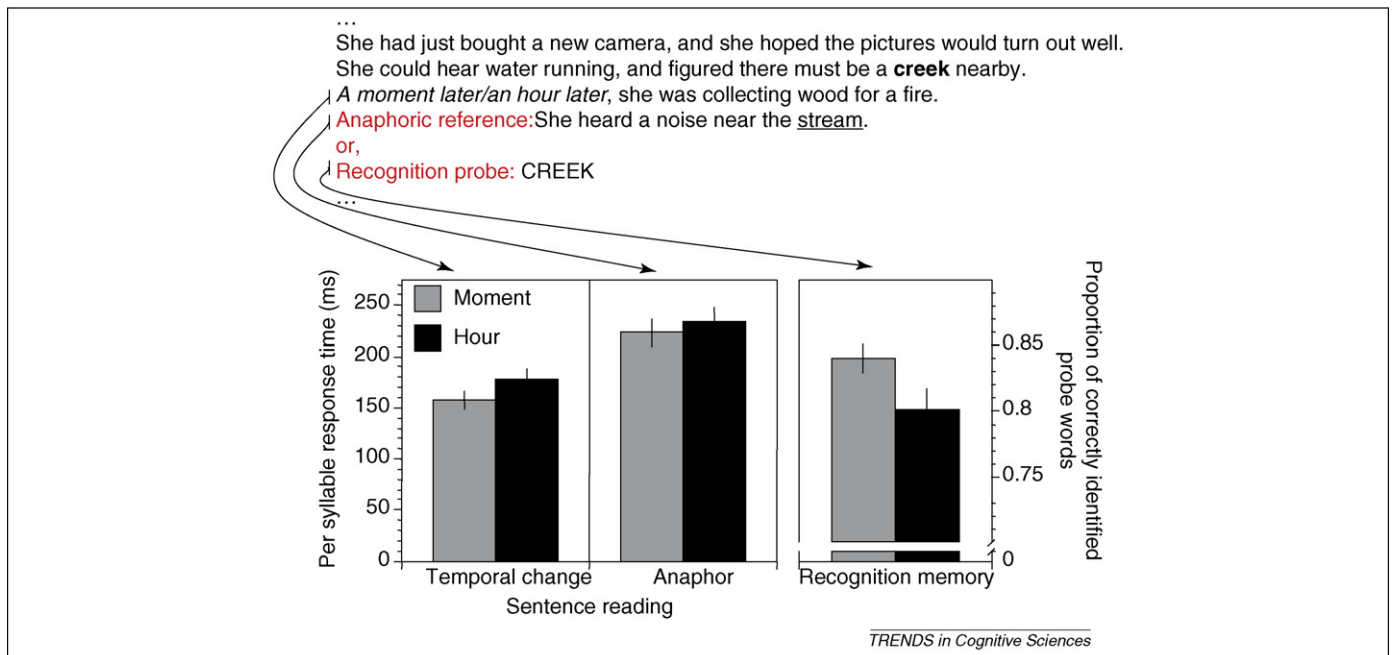


Figure 1. Temporal changes are perceived as event boundaries, and this affects working memory [17]. Participants read narratives containing sentences (marked here in italics) that either indicated a significant interval of time had passed ('an hour later') or not ('a moment later'). These temporal reference sentences often were identified as event boundaries – particularly in the 'hour later' condition. Subsequent experiments assessed memory for objects mentioned before the temporal shift (creek, in this example). Memory was tested in two ways: by measuring the time it took to read an anaphoric reference to the previously mentioned object, or by asking directly whether the participant recognized the word as having been read. Bottom left: Sentences in the 'hour later' condition were read more slowly than the nearly identical 'moment later' versions. Bottom middle: Anaphoric references to previously mentioned objects were read more slowly in the 'hour later' condition. Bottom right: Recognition memory for previously mentioned objects was less accurate in the 'hour later' condition. Such results indicate that working memory is updated at event boundaries. Error bars are standard errors.

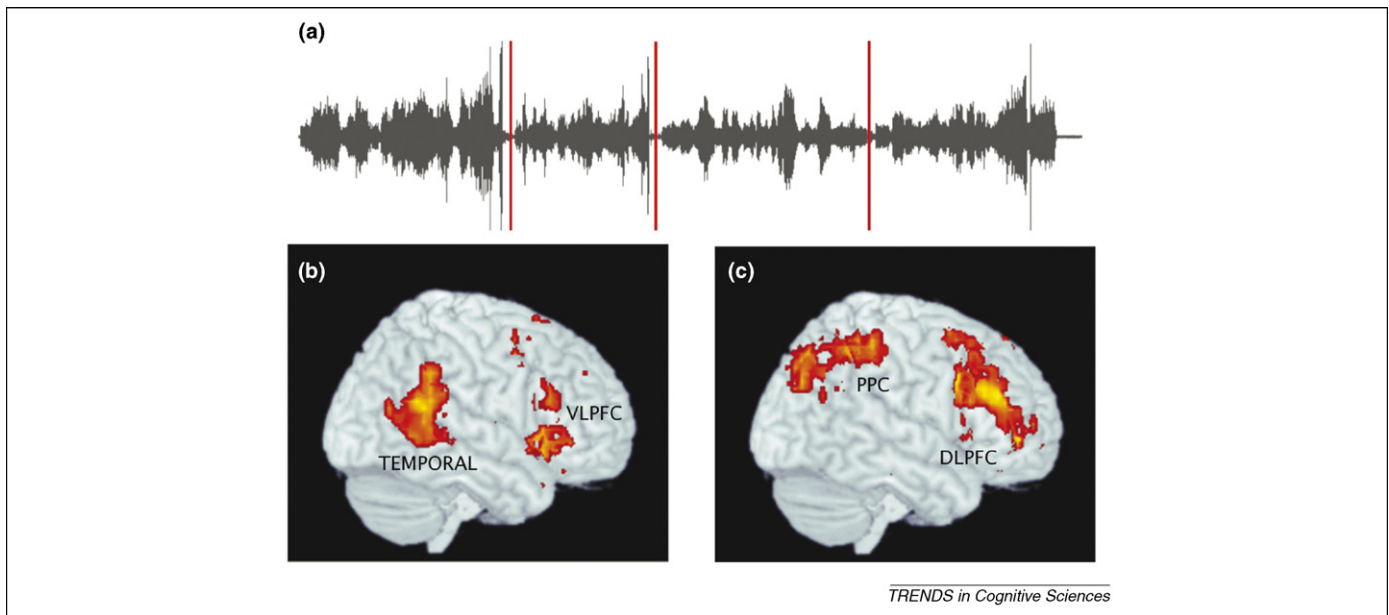


Figure 2. The boundaries between musical movements elicited increased activity at a network of brain regions. (a) In Western concert music, symphonies are made up of movements. The sound wave is plotted, and the breaks between movements are illustrated with red lines. In [21], musically untrained participants listened to two 8–10-min segments of symphonies by William Boyce, consisting of movements lasting an average of 1 min and 10 s, while their brain activity was recorded using fMRI. (b) A ventral network including the ventrolateral prefrontal cortex (VLPFC) and posterior temporal cortex (TEMPORAL) increased in activity first at movement boundaries. (c) A dorsal network including the dorsolateral prefrontal cortex (DLPFC) and posterior parietal (PPC) increased in activity slightly later. Both networks were lateralized to the right hemisphere. The authors interpreted these data by suggesting that the response of the ventral network reflected the processing of violations of musical expectancy, whereas the response of the dorsal network reflected consequent top-down modulation of the processing of new musical information. Reproduced, with permission, from [21].

ventral network that included the ventrolateral prefrontal cortex and posterior temporal cortex, and a late-responding dorsal network that included the dorsolateral prefrontal cortex and posterior parietal cortex.

fMRI data also have provided evidence for automatic event segmentation during reading. Speer *et al.* [22] got participants to read narrative texts, one word at a time, while brain activity was recorded with fMRI. Participants subsequently segmented the texts into events. Brain activity increased at the later-identified event boundaries in a set of regions that corresponded substantially with the regions that increased at event boundaries in movies; these included bilateral regions in medial posterior temporal, occipital and parietal cortex, in the lateral temporal-parietal and anterior temporal cortex, and in the right posterior dorsal frontal cortex.

Converging with the fMRI data, recent experiments using EEG indicate that perceptual processing is modulated at event boundaries on an ongoing basis. In one set of experiments [23], participants viewed movies of goal-directed activities while undergoing EEG recording, after which they segmented the movies into events. Evoked responses were detected at frontal and parietal electrode sites, and these responses were modulated by whether or not participants were familiar with the movies themselves or with the activities they depicted.

Finally, measurements of pupil diameter provide further evidence that information processing increases transiently after the perception of an event boundary. Pupil diameter provides an online measure of cognitive processing load, as demonstrated in studies using a variety of motor and cognitive paradigms [24]. In one study [25], participants viewed movies while their pupil diameters

and eye movements were recorded. They then segmented the movies into events. Pupil diameter transiently increased following those points that were later identified as event boundaries. (Saccades also became more frequent around event boundaries, which might reflect the reorienting of viewers at the beginning of a new event.)

Together, the reading-time, neurophysiological and oculomotor data strongly suggest that there are transient changes in brain activity correlated with the subjective experience that one event has ended and another has begun. We can be confident these effects are independent of task demands such as the conscious intention to segment activity or covert attention to the location of event boundaries, because for all the studies discussed here the crucial reading-time or physiological data were collected before participants were introduced to the event segmentation task. An important question for ongoing research is: to what extent do such effects reflect neural processing that has a causal role in event segmentation, and to what extent do they reflect processing that is correlated with the presence of event boundaries but not a cause or consequence of event segmentation as such?

Hierarchical event perception

Events can be identified at a range of temporal grains, from brief (fine-grained) to extended (coarse-grained). In goal-directed human activity it is natural to think of such events as being hierarchically organized, with groups of fine-grained events clustering into larger units; this is partly because actions are hierarchically organized by goals and subgoals [1]. For example, the activity of making a sandwich includes the coarse-grained events of (i) removing ingredients from the refrigerator, and (ii) assembling

them. The event of assembling the ingredients, in turn, might include subevents such as adding meat, adding cheese and spreading mayonnaise. When understanding activities, people seem spontaneously to track the hierarchical grouping of events. Evidence for hierarchical grouping has been found by asking a viewer or reader to segment an activity twice, on different occasions, once to identify coarse-grained event boundaries and once to identify fine-grained boundaries. If coarse-grained events subsume a group of finer-grained events, it would be expected that each coarse event boundary would fall slightly later in time than the fine event boundary to which it is closest. That is exactly what is observed [26]. Hierarchical dividing of coarse events into finer-grained events also would predict that each coarse event boundary should fall near one of the fine boundaries; this is also the case [27].

Events on different temporal grains can be sensitive to different features of activities. Running descriptions by viewers of coarse- and fine-grained events in movies has provided evidence for such differences [27]. Descriptions of coarse-grained events focused on objects, using more precise nouns and less precise verbs. Descriptions of fine-grained events focused on actions on those objects, using more precise verbs but specifying the objects less precisely. If different temporal grains depend on different features, you would expect that segmentation of fine- or coarse-grained events could be selectively impaired. Selective impairments of coarse-grained segmentation have been found in patients with frontal lobe lesions [28] and in patients with schizophrenia [29] (Box 2).

Parsing ongoing activity into hierarchically organized events and subevents might be important for relating ongoing perception to knowledge about activities. Studies of memory for events and story comprehension suggest that people use hierarchically organized event representations (scripts or schemata) to understand a particular activity in relation to similar previously experienced activities (see, e.g., Refs [30,31]). Infants as young as 12 months seem to be sensitive to the hierarchical organization of behavior [32,33] (Box 3) and non-human primates seem to be sensitive to hierarchical organization in the behavior of conspecifics such that they can use such organization to learn manual skills [34]. Recent research has found that learning of a new procedure by adults can be facilitated by explicitly representing the hierarchical structure of the activity, or impaired by misrepresenting that structure [35]. Such learning might be related to the mechanisms by which people form 'chunks' in semantic knowledge [36]. In sum, people seem to track the hierarchical structure of activity during perception because this enables them to use prior knowledge for understanding, and to adapt that prior knowledge to learning new skills.

Memory for events

Event boundaries relate systematically to both the online maintenance of information (working memory) and the permanent storage of information for later retrieval (long-term memory). According to EST [2], this is because at event boundaries people update representations of the current event, which frees information from working

Box 2. Event segmentation and neuropathology

Recent results suggest that event segmentation is impaired as a result of frontal lobe lesions, schizophrenia and neuropathology associated with aging. In a pair of studies by Zalla and colleagues [27,28], participants segmented brief movies of everyday activities and then completed recall and recognition memory tests. Compared with controls, patients with frontal lobe lesions [27] and patients with schizophrenia [28] were less likely to identify normative coarse-grained units. The fine-grained segmentation of the patient groups did not differ from that of the controls.

Zacks *et al.* [52] investigated the effect of aging and Alzheimer-type dementia on event segmentation and event memory. Healthy younger adults, healthy older adults and older adults with mild dementia segmented movies of everyday activities. They then completed a recognition memory test for pictures from each movie and a test of order memory. Segmentation of older adults differed more from group norms than did that of younger adults. This was particularly pronounced for those with dementia, and held both for fine- and coarse-grained segmentation. Older adults (particularly those with dementia) also performed less well on the memory tests. Moreover, those older adults who segmented less well also remembered less.

One possibility is that damage to prefrontal cortex is the common cause of impaired event segmentation in all three of these groups. This could reflect selective damage to the neural substrate of working memory representations of events (see Box 1). Alternatively, the event segmentation impairments of these groups could reflect damage to other mechanisms that participate in event segmentation, such as perceptual processing or error monitoring. Further studies including other conditions such as Parkinson disease and obsessive-compulsive disorder will help to improve our understanding of the causes of impairments in event segmentation.

Box 3. Developing into an adult event segmentation system

How does event segmentation emerge in early development? There is evidence that some components of this ability are present at 10-11 months of life. The relevant studies have used two infant-friendly paradigms: habituation paradigms, in which infants tend to look longer at stimuli they perceive as more different from what came just before; and cross-modal matching paradigms, in which infants tend to look longer at visual stimuli that they perceive as synchronized with auditory stimuli.

In one habituation experiment [58], infants as young as 6 months were familiarized with displays in which a puppet jumped two or three times; they then looked longer when tested on sequences with a different number of jumps. In another habituation experiment [59], 10-11-month-old infants were familiarized with 4-s movies of everyday actions, and then shown movies with 1.5-s still frames inserted at different moments. They looked longer if the still frame interrupted an intentional action than if it came at the completion of the action.

In a study using the cross-modal matching paradigm [60], 10-11-month-old infants simultaneously viewed two visual displays; in each, an actor performed sequences of simple actions involving objects in, on and next to containers. While both action displays were simultaneously displayed, a series of tones sounded that were synchronized to the boundaries between actions in one of the visual displays. Infants looked longer at the display whose action boundaries matched the tones.

By 12 months, infants seem to be sensitive to how sequences of actions are grouped to achieve higher-level goals. In one habituation study [32], infants were familiarized to action sequences involving pulling objects toward the experimenter. They then looked longer at actions that were physically similar but not goal-appropriate, compared with sequences that were physically dissimilar but were goal-appropriate (see also Refs [61,62]). By 24 months, toddlers clearly code action in terms of hierarchical goals, and this affects memory recall [63].

memory, and orient to incoming perceptual information, which encodes it particularly strongly for long-term memory. Evidence for working memory updating at event boundaries comes from the comprehension of text narratives, picture stories and cinema.

Working memory for events in text

Studies of memory access during text reading generally have not directly measured event boundaries, but have found evidence that those cues associated with event boundaries reduce memory access. Memory for recently mentioned objects is poorer after reading sentences that indicate a shift in time or space [37–41]. As noted previously, one study [17] experimentally manipulated the presence of event boundaries in narratives by varying the passage of time indicated in an auxiliary phrase. As is shown in Figure 1, implicit and explicit measures of memory indicated that information became less accessible following a time shift.

Working memory for events in pictures

Gernsbacher [42] got participants to identify episode boundaries in a picture story. New participants then viewed the picture story, and from time to time were probed to discriminate pictures that had recently been presented in the story from pictures that were left–right reversed. Recognition was better for images from within the current episode than for images from the previous episode, indicating that some of the surface information in the pictures was less available once a new event had begun.

Working memory for events in cinema and virtual reality

Converging with the results from text and picture stories, experiments looking at recognition memory for objects in cinema and virtual reality indicate that working memory is updated at event boundaries. In one recent study [43], participants navigated in a virtual reality environment, and memory for recently appearing objects was tested. Memory was reduced after walking through a door into an adjacent room – a probable event boundary. In another recent set of studies [44], participants viewed excerpts from movies that were occasionally interrupted by recognition memory tests for objects that had been on the screen five seconds previously. The information that participants could retrieve about these objects differed systematically depending on whether an event boundary had occurred during those five seconds. Further, neuroimaging data suggested that the basis for responding also differed: retrieval of information about objects from within the current event selectively activated brain areas including the bilateral occipital and lateral temporal cortices and the right inferior lateral frontal cortex, whereas retrieval of information about objects from the previous event selectively activated brain areas including the medial temporal cortex and medial parietal cortex. The medial temporal regions correspond to the hippocampal formation, which is known to have a key role in long-term memory storage and retrieval [45]. This is consistent with the hypothesis that participants depended preferentially on working memory

for within-event retrieval and shifted to long-term memory for across-event retrieval – although the information to be retrieved was only a few seconds old.

Long-term memory for events

Of course, the core function of long-term memory is to support access to information over much longer delays. Evidence from tests in which participants retrieve information over delays of minutes to hours indicates that event boundaries serve as anchors in long-term memory. Recognition memory for pictures drawn from event boundaries has been found to be better than memory for pictures drawn from points between the boundaries [46]. Manipulations that affect the perception by viewers of the locations of event boundaries in a film also affect their memory for the film when tested later. Boltz [47] got participants to watch feature films with or without embedded commercials. The commercials were inserted at event boundaries, or between event boundaries. Memory for the activity was better for movies without commercials and for movies in which commercials were inserted at event boundaries. Similarly, Schwan and Garsoffky [48] got participants to view movies of everyday events with or without deletions. The deletions were either of segments of time surrounding an event boundary, or segments of time within an event. The researchers found that recall was better for events when there were no deletions and when the deletions were of segments within an event, preserving the event boundaries. Segmentation grain has also shown an impact on memory: recall for details is better after fine-grained segmentation than after coarse-grained segmentation [49–51].

If event boundaries serve as anchors for long-term memory encoding, then individuals who segment an activity effectively should have better later memory for it. As described in Box 2, recent evidence supports this hypothesis [52].

How does segmentation help?

Why do people segment ongoing activity into events? Segmentation results from the continual anticipation of future events. This anticipation enables you adaptively to encode structure from the continuous perceptual stream, to understand what an actor will do next, and to select your own future actions [53]. Segmentation simplifies, enabling you to treat an extended interval of time as a single chunk. If you segment well, this chunking saves on processing resources and improves comprehension. However, to segment well, you must identify the correct units of activity – the correct events. One possible mechanism for identifying events is to monitor your ongoing comprehension and break activity into units when comprehension begins to falter (Box 1). Once events have been individuated you can start to learn to recognize sequences of events and plan reactions based on such sequences. Grouping fine-grained events hierarchically into larger units enables you to learn not just rote one-after-the-other relations but more complex ones. One such relation is partial ordering, which is ubiquitous in problem-solving and planning. Think of baking a cake, where there are several different orders

Box 4. Questions for future research

Recent work on event segmentation raises a number of interesting challenges for future studies. These relate to both the neural mechanisms of perception and to broader aspects of cognition and memory.

Changes have been observed in the central nervous system and in the oculomotor system at those points identified by observers as event boundaries. Which changes reflect neural processing that has a causal role in event segmentation, and which reflect processing that is correlated with the presence of event boundaries but is not a cause or consequence of event segmentation as such?

Segmenting activity hierarchically is associated with better learning of procedural skills. Is hierarchical understanding of an activity necessary to learn it, or does hierarchical encoding simply tend to occur in the same learning circumstance that produces good performance?

How do prior knowledge and expectation affect which points are identified as event boundaries? There is some evidence that participants segment familiar activities more at a coarser grain than unfamiliar activities [6], although this effect is not consistent in the literature [27]. Although the mechanisms underlying the role of event knowledge in segmentation remain unclear, one possibility is that possessing a schema or script for a type of activity might reduce prediction error during perception of that activity.

How is event segmentation related to the understanding and planning of self-initiated events? Ideomotor theories of action control [64–66] propose that actions are coded in terms of the goals they are intended to satisfy rather than motor movements themselves. Similar to the predictive nature of event perception, actions are preceded by the anticipation of a desired perceptual effect [53,67]. A recent theory of perception and action planning expanded on this idea by proposing that the perception of events and production of events (actions) share a representational medium [68], possibly subserved by a mirror system that maps the actions of others onto your own action representations [69]. In such accounts, the events that you are currently producing constrain the events that you can concurrently perceive.

in which you could mix the ingredients, but all those steps must be complete before the cake goes in the oven.

Events segmented during perception also can form the units for memory encoding, enabling you to store compact representations of extended activities. Identifying the correct events facilitates memory, much as it is easier to remember a sequence of vocalizations if it comes from a language you know, enabling you to segment it into words, clauses and sentences.

Event segmentation is automatic and important for perception, comprehension, problem-solving and memory (See Box 4 for open questions). None of this would be true if the structure of the world were not congenial to segmentation. If sequential dependencies were not predictable, if activity were not hierarchically organized, there would be no advantage to imposing chunking and grouping on the stream of behavior. In this regard, as in many others, human perceptual systems seem to be specialized information-processing devices that are tuned to the structure of their environment.

Acknowledgements

Preparation of this article was supported in part by grants R01-MH070674 and T32 AG000030–31 from the National Institutes of Health. We thank Devarajan Sridharan for providing figure materials, and thank the following colleagues for thoughtful comments on the manuscript: Dare Baldwin, Joe Magliano, G.A. Radvansky, Stephan Schwan, Ric Sharp, Khen Swallow and Barbara Tversky.

References

- Zacks, J.M. and Tversky, B. (2001) Event structure in perception and conception. *Psychol. Bull.* 127, 3–21
- Zacks, J.M. et al. (2007) Event perception: A mind/brain perspective. *Psychol. Bull.* 133, 273–293
- Newton, D. (1973) Attribution and the unit of perception of ongoing behavior. *J. Pers. Soc. Psychol.* 28, 28–38
- Newton, D. et al. (1977) Objective basis of behavior units. *J. Pers. Soc. Psychol.* 12, 847–862
- Zacks, J.M. (2004) Using movement and intentions to understand simple events. *Cognit. Sci.* 28, 979–1008
- Hard, B.M. et al. (2006) Making sense of abstract events: Building event schemas. *Mem. Cognit.* 34, 1221–1235
- Magliano, J.P. et al. (2001) Indexing space and time in film understanding. *Appl. Cognit. Psychol.* 15, 533–545
- Speer, N.K. et al. (2004) Perceiving narrated events. In *Proceedings of the 26th Annual Meeting of the Cognitive Science Society, Chicago, IL* (Forbus, K. et al., eds), p. 1637, Cognitive Science Society
- Zacks, J.M. et al. (2006) The human brain's response to change in cinema. *Abstr. Psychon. Soc.* 11, 9
- Newton, D. (1976) Foundations of attribution: the perception of ongoing behavior. In *New Directions in Attribution Research* (Harvey, J.H. et al., eds), pp. 223–248, Erlbaum
- Speer, N.K. et al. (2003) Activation of human motion processing areas during event perception. *Cogn. Affect. Behav. Neurosci.* 3, 335–345
- Gernsbacher, M.A. (1990) *Language Comprehension as Structure Building*, Erlbaum
- Zwaan, R.A. et al. (1995) The construction of situation models in narrative comprehension: An event-indexing model. *Psychol. Sci.* 6, 292–297
- Zwaan, R.A. et al. (1995) Dimensions of situation model construction in narrative comprehension. *J. Exp. Psychol. Learn. Mem. Cogn.* 21, 386–397
- Rinck, M. and Weber, U. (2003) Who when where: an experimental test of the event-indexing model. *Mem. Cognit.* 31, 1284–1292
- Zwaan, R.A. et al. (1998) Constructing multidimensional situation models during reading. *Sci. Stud. Read.* 2, 199–220
- Speer, N.K. and Zacks, J.M. (2005) Temporal changes as event boundaries: processing and memory consequences of narrative time shifts. *J. Mem. Lang.* 53, 125–140
- Zacks, J.M. et al. (2001) Human brain activity time-locked to perceptual event boundaries. *Nat. Neurosci.* 4, 651–655
- Sekuler, R. et al. (2002) Perception of visual motion. In *Steven's handbook of experimental psychology: Sensation and perception* (Vol. 1, 3rd edn) Pashler, H. and Yantis, S., eds In pp. 121–176, John Wiley & Sons
- Zacks, J.M. et al. (2006) Visual movement and the neural correlates of event perception. *Brain Res.* 1076, 150–162
- Sridharan, D. et al. (2007) Neural dynamics of event segmentation in music: Converging evidence for dissociable ventral and dorsal networks. *Neuron* 55, 521–532
- Speer, N.K. et al. (2007) Human brain activity time-locked to narrative event boundaries. *Psychol. Sci.* 18, 449–455
- Sharp, R.M. et al. (2007) Electrophysiological correlates of event segmentation: how does the human mind process ongoing activity? In *Proceedings of the Annual Meeting of the Cognitive Neuroscience Society, New York*, p. 235, Cognitive Neuroscience Society
- Beatty, J. and Lucero-Wagoner, B. (2000) The pupillary system. In *Handbook of Psychophysiology* (Cacioppo, J.T. et al., eds), pp. 142–162, Cambridge University Press
- Swallow, K.M. and Zacks, J.M. (2004) Hierarchical grouping of events revealed by eye movements. *Abstr. Psychon. Soc.* 9, 81
- Hard, B.M. et al. (2006) Hierarchical encoding of behavior: translating perception into action. *J. Exp. Psychol. Gen.* 135, 588–608
- Zacks, J.M. et al. (2001) Perceiving, remembering, and communicating structure in events. *J. Exp. Psychol. Gen.* 130, 29–58
- Zalla, T. et al. (2003) Perception of action boundaries in patients with frontal lobe damage. *Neuropsychologia* 41, 1619–1627
- Zalla, T. et al. (2004) Perception of dynamic action in patients with schizophrenia. *Psychiatry Res.* 128, 39–51

- 30 Lichtenstein, E.D. and Brewer, W.F. (1980) Memory for goal-directed events. *Cognit. Psychol.* 12, 412–445
- 31 Rumelhart, D.E. (1977) Understanding and summarizing brief stories. In *Basic Processes in Reading: Perception and Comprehension* (Laberge, D. and Samuels, S.J., eds), pp. 265–303, Erlbaum
- 32 Sommerville, J.A. and Woodward, A.L. (2005) Pulling out the intentional structure of action: the relation between action processing and action production in infancy. *Cognition* 95, 1–30
- 33 Bauer, P.J. and Mandler, J.M. (1989) One thing follows another: effects of temporal structure on 1- to 2-year-olds' recall of events. *Dev. Psychol.* 25, 197–206
- 34 Byrne, R.W. (2002) Seeing actions as hierarchically organized structures: great ape manual skills. In *The Imitative Mind: Development, Evolution, and Brain Bases* (Meltzoff, A. and Prinz, W., eds), pp. 122–142, Cambridge University Press
- 35 Zacks, J.M. and Tversky, B. (2003) Structuring information interfaces for procedural learning. *J. Exp. Psychol. Appl.* 9, 88–100
- 36 Gobet, F. et al. (2001) Chunking mechanisms in human learning. *Trends Cogn. Sci.* 5, 236–243
- 37 Glenberg, A.M. et al. (1987) Mental models contribute to foregrounding during text comprehension. *J. Mem. Lang.* 26, 69–83
- 38 Bower, G.H. and Rinck, M. (2001) Selecting one among many referents in spatial situation models. *J. Exp. Psychol. Learn. Mem. Cogn.* 27, 81–98
- 39 Morrow, D.G. et al. (1989) Updating situation models during narrative comprehension. *J. Mem. Lang.* 28, 292–312
- 40 Morrow, D.G. et al. (1987) Accessibility and situation models in narrative comprehension. *J. Mem. Lang.* 26, 165–187
- 41 Rinck, M. and Bower, G. (2000) Temporal and spatial distance in situation models. *Mem. Cognit.* 28, 1310–1320
- 42 Gernsbacher, M.A. (1985) Surface information loss in comprehension. *Cognit. Psychol.* 17, 324–363
- 43 Radvansky, G.A. and Copeland, D.E. (2006) Walking through doorways causes forgetting: situation models and experienced space. *Mem. Cognit.* 34, 1150–1156
- 44 Swallow, K.M. et al. (2007) Perceptual events may be the “episodes” in episodic memory. *Abstr. Psychon. Soc.* 12, 25
- 45 Squire, L.R. and Zola-Morgan, S. (1991) The medial temporal lobe memory system. *Science* 253, 1380–1386
- 46 Newton, D. and Engquist, G. (1976) The perceptual organization of ongoing behavior. *J. Exp. Soc. Psychol.* 12, 436–450
- 47 Boltz, M. (1992) Temporal accent structure and the remembering of filmed narratives. *J. Exp. Psychol. Hum. Percept. Perform.* 18, 90–105
- 48 Schwan, S. and Garsoffky, B. (2004) The cognitive representation of filmic event summaries. *Appl. Cognit. Psychol.* 18, 37–55
- 49 Hanson, C. and Hirst, W. (1989) On the representation of events: A study of orientation, recall, and recognition. *J. Exp. Psychol. Gen.* 118, 136–147
- 50 Lassiter, G.D. (1988) Behavior perception, affect, and memory. *Soc. Cognit.* 6, 150–176
- 51 Lassiter, G.D. et al. (1988) Memorial consequences of variation in behavior perception. *J. Exp. Soc. Psychol.* 24, 222–239
- 52 Zacks, J.M. et al. (2006) Event understanding and memory in healthy aging and dementia of the Alzheimer type. *Psychol. Aging* 21, 466–482
- 53 Schütz-Bosbach, S. and Prinz, W. (2007) Prospective coding in event representation. *Cogn. Process.* 8, 93–102
- 54 Holroyd, C.B. and Coles, M.G. (2002) The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol. Rev.* 109, 679–709
- 55 Schultz, W. and Dickinson, A. (2000) Neuronal coding of prediction errors. *Annu. Rev. Neurosci.* 23, 473–500
- 56 Bouret, S. and Sara, S.J. (2005) Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends Neurosci.* 28, 574–582
- 57 Reynolds, J.R. et al. (2007) A computational model of event segmentation from perceptual prediction. *Cognit. Sci.* 31, 613–643
- 58 Wynn, K. (1996) Infants' individuation and enumeration of actions. *Psychol. Sci.* 7, 164–169
- 59 Baldwin, D.A. et al. (2001) Infants parse dynamic action. *Child Dev.* 72, 708–717
- 60 Saylor, M.M. et al. (2007) Infants' on-line segmentation of dynamic human action. *J. Cognit. Dev.* 8, 113–128
- 61 Woodward, A.L. (1998) Infants selectively encode the goal object of an actor's reach. *Cognition* 69, 1–34
- 62 Woodward, A.L. and Sommerville, J.A. (2000) Twelve-month-old infants interpret action in context. *Psychol. Sci.* 11, 73–77
- 63 Travis, L.L. (1997) Goal-based organization of event memory in toddlers. In *Developmental Spans on Event Comprehension and Representation: Bridging Fictional and Actual Events* (van den Broek, P.W. et al., eds), pp. 111–139, LEA
- 64 Greenwald, A.G. (1970) A choice reaction time test of ideomotor theory. *J. Exp. Psychol.* 6, 20–25
- 65 Greenwald, A.G. (1970) Sensory feedback mechanisms in performance control: with special reference to ideo-motor mechanism. *Psychol. Rev.* 77, 73–99
- 66 James, W. (1890) *The Principles of Psychology*, Holt
- 67 Schütz-Bosbach, S. and Prinz, W. (2007) Perceptual resonance: action-induced modulation of perception. *Trends Cogn. Sci.* 11, 349–355
- 68 Hommel, B. et al. (2001) The theory of event coding (TEC): a framework for perception and action planning. *Behav. Brain Sci.* 24, 849–937
- 69 Rizzolatti, G. and Craighero, L. (2004) The mirror-neuron system. *Annu. Rev. Neurosci.* 27, 169–192

How to re-use Elsevier journal figures in multimedia presentations

It's easy to incorporate figures published in *Trends*, *Current Opinion* or *Drug Discovery Today* journals into your multimedia presentations or other image-display programs.

1. Locate the article with the required figure on ScienceDirect and click on the 'Full text + links' hyperlink
2. Click on the thumbnail of the required figure to enlarge the image
3. Copy the image and paste it into an image-display program

Permission of the publisher is required to re-use any materials from *Trends*, *Current Opinion* or *Drug Discovery Today* journals or from any other works published by Elsevier. Elsevier authors can obtain permission by completing the online form available through the Copyright Information section of Elsevier's Author Gateway at <http://authors.elsevier.com>. Alternatively, readers can access the request form through Elsevier's main website at:

www.elsevier.com/locate/permissions