



Cognitive Science (2011) 1–29

Copyright © 2011 Cognitive Science Society, Inc. All rights reserved.

ISSN: 0364-0213 print / 1551-6709 online

DOI: 10.1111/j.1551-6709.2011.01202.x

The Impact of Continuity Editing in Narrative Film on Event Segmentation

Joseph P. Magliano,^a Jeffrey M. Zacks^b

^a*Department of Psychology, Northern Illinois University*

^b*Department of Psychology, Washington University in Saint Louis*

Received 13 May 2010; received in revised form 9 November 2010; accepted 11 March 2011

Abstract

Filmmakers use continuity editing to engender a sense of situational continuity or discontinuity at editing boundaries. The goal of this study was to assess the impact of continuity editing on how people perceive the structure of events in a narrative film and to identify brain networks that are associated with the processing of different types of continuity editing boundaries. Participants viewed a commercially produced film and segmented it into meaningful events, while brain activity was recorded with functional magnetic resonance imaging (MRI). We identified three degrees of continuity that can occur at editing locations: edits that are continuous in space, time, and action; edits that are discontinuous in space or time but continuous in action; and edits that are discontinuous in action as well as space or time. Discontinuities in action had the biggest impact on behavioral event segmentation, and discontinuities in space and time had minor effects. Edits were associated with large transient increases in early visual areas. Spatial-temporal changes and action changes produced strikingly different patterns of transient change, and they provided evidence that specialized mechanisms in higher order perceptual processing regions are engaged to maintain continuity of action in the face of spatiotemporal discontinuities. These results suggest that commercial film editing is shaped to support the comprehension of meaningful events that bridge breaks in low-level visual continuity, and even breaks in continuity of spatial and temporal location.

Keywords: Film comprehension; Event comprehension; Event segmentation; Situation models; Mental models; fMRI

1. Introduction

On the face of it, it is surprising how easy it is to comprehend a popular narrative film given how complicated its structure can be. Feature films typically are constructed from

Correspondence should be sent to Joe Magliano, Department of Psychology, Northern Illinois University, DeKalb, IL 60115. E-mail: jmagliano@niu.edu

hundreds to thousands of individual camera shots (continuous runs of the camera) that are edited together to create a perceptually continuous sequence of events that comprise a narrative. These camera shots are often filmed at different times and locations and may have little feature overlap, and as such, flow of information across shots often bears little resemblance to the perceptual flow of information as we interact in the real world (Cutting, 2005; Hochberg & Brooks, 1978). For example, a common way to establish the location of a scene within a building is the use of an establishing shot sequence that is comprised of a shot of the outside of a building and a shot of an interior room within a building. Comprehension of such sequences improves with experience (Schwan & Ildirar, 2010), but most of the time viewers readily infer that the interior room is inside the building despite never experiencing such an abrupt perceptual sequence in real-world interactions.

How do filmmakers manage this trick and do it so well most of the time? Filmmakers use a common, well-documented system of editing together camera shots called *continuity editing*, a system of editing that is designed to maintain a continuous and clear narrative action (Bordwell, 1985; Bordwell & Thompson, 2006; Cutting, 2005). It relies on a variety of technical procedures designed to ensure shots are matched across boundaries, along perceptual features, spatiotemporal (ST) relations, and actions. Continuity editing enables viewers to perceive some sequences of shots as depicting a continuously unfolding event, but to perceive other transitions between shots as transitions from one episode to another (Magliano, Miller, & Zwaan, 2001). Although there has been considerable discussion of continuity editing in film theory (Bordwell, 1985, 2002; Bordwell & Thompson, 2006; Dancyger, 2007), there is a dearth of understanding of the higher level cognitive (e.g., processes involved in constructing mental models) and neurological systems that make it work so well in the minds of the viewers.

1.1. *The role of continuity editing in conveying event structure in film*

Most narrative film is made with the goal to construct a clear and coherent event structure so that the viewer can readily understand the sequence of events depicted in it (Bordwell, 1985, 2002). Film continuity is managed in the service of telling a story and thus creating a sense of continuity and discontinuity of events within the story world. To this end, filmmakers rely on formal devices for editing together the distinct camera shots that make up a film. Some of these devices are quite dramatic, such as the fade-out, fade-in, or dissolve, and are intended to signal a shift in scenes. However, approximately 95% of editing boundaries are *cuts* (Cutting, 2005), which constitutes the splicing together of two camera frames. The juxtaposition of the content of two edited frames can be jarring and contain little feature overlap, but most often, a cut is intended to convey a continuous flow of events and be “invisible” to the viewer (Murch, 2001). Some authors have argued that cuts are invisible because they correspond to visual interruptions that occur naturally due to the movements of the eyes, in particular *blinks* and *saccades* (Cutting, 2005; Murch, 2001). However, there are no significant differences in the patterns of blinks and saccades between locations of a film that fall within camera shots or at editing boundaries (Smith & Henderson, 2008).

An alternative reason that cuts are not perceptually salient stems from the nature of continuity editing. The purpose of continuity editing is to tell a story by creating a spatially and temporally coherent sequence of events and actions (Bordwell, 1985; Bordwell & Thompson, 2006) with the end result of enabling the viewer to perceive a sense of causal cohesion across cuts. One could contend that many cuts are invisible because they are created such that viewers perceive a strong sense of situational cohesion across the cut boundary.

Continuity editing is important for the management of perception of spatial and temporal ellipsis. Films rarely depict all of the subevents that make up a larger event. For example, one shot may show an actor approaching the stairs of a building and the next shot may show the actor entering that building. In such cases, viewers are intended to perceive these events as being roughly continuous in space and time. The editing technique most likely to be used in this situation would be the cut. In contrast, filmmakers may use fade-ins, fade-outs, or dissolves across shots to indicate that there is a significant amount of missing narrative time between the two shots.

Bordwell (1985; Bordwell & Thompson, 2006) has described how the sense of spatial and temporal continuity is largely managed by following the *180° system*. Specifically, spatial continuity across cuts is maintained as long as any scene's action (e.g., entities navigating, social interactions, physical confrontations) operates within a discernable and predictable 180° range, which creates a notional "stage" containing the action. The axis defining this 180° range is often created by an establishing shot, which provides information regarding the spatial location of a scene and the entities that will interact during it. Reliance on the 180° system across shots also helps to maintain a sense of temporal continuity because shots that follow this rule will likely be temporally contiguous. However, there are other conventions that can be used to create a sense of temporal continuity, such as cross-cutting between different locations, which will create a sense of temporal concurrence in spite of breaking the 180° rule (Magliano et al., 2001). It is often the case that filmmakers want to depict a narrative event in a matter of seconds that would otherwise take a much longer period of time to complete in the real world. In these instances, filmmakers may use cuts in conjunction with continuity in the soundtrack (music, sounds occurring within the story world, dialog) across the cuts to create the perception of temporal continuity when the visual input and general knowledge of the world suggest otherwise.

Thus, there is a class of edits that maintain continuity by preserving spatial and temporal relations. However, creating a sense of causal continuity in action is arguably a primary function of continuity editing (Bordwell, 1985; Bordwell & Thompson, 2006), and action can be perceived as continuous even when space and time are not. Continuity of action is largely maintained by having an action start in one shot and continue in the shot following the cut. Adherence to the 180° system ensures that the motion across shots will match in term of its flow across the frames that comprise the shots (i.e., crossing the 180° boundary will flip the image on a frame, as is illustrated in Bordwell & Thompson, 2006). Cuts that happen during periods of high motion may be masked by that motion (Smith & Henderson, 2008). When the action tracks human actors, mechanisms for processing biological motion may create perceptual continuity across changes in viewpoint (Alaerts, Swinnen, & Wenderoth, 2009; Troje, Westhoff, & Lavrov, 2005).

Filmmakers can create a variety of combinations of continuities and discontinuities in space, time, and action across cuts. For example, interweaving actions that occur in different locations across cuts should lead to the perception that the events in the different locations are occurring concurrently in the story world (Bordwell, 1985; Bordwell & Thompson, 2006) and will likely converge in the story world (Magliano, Dijkstra, & Zwaan, 1996). On the other hand, breaks in space and time across a cut boundary should lead to the perception that there is a break in story time and space (Magliano et al., 2001; Zacks, Speer, & Reynolds, 2009).

In sum, continuity editing uses multiple techniques to establish continuity across successive shots. Some of these may require little special perceptual processing on the part of the viewer to comprehend: A cut that respects the 180° system, maintains continuous sound, and matches on action results in a stimulus that may not be too different from what the viewer would experience during a saccade. However, in addition to continuity edits, there are edits in which the filmmakers do not attempt to maintain continuity of space, time, or action. These edits occur most frequently when a significant scene boundary takes place, and filmmakers generally conceive of these as corresponding to structural units (Bordwell, 1985). Such edits may require specialized mechanisms to comprehend. One possibility is that such cuts may “work” because human perceptual and comprehension systems are set up to segment ongoing activity into discrete events (Kurby & Zacks, 2008; Newton & Engquist, 1976; Zacks & Swallow, 2007) and to build mental models of the depicted activity that are individuated based on this segmentation (Zwaan & Radvansky, 1998).

1.2. Event segmentation and narrative comprehension

How do perceivers segment ongoing activity? Event segmentation theory (Zacks, Speer, Swallow, Braver, & Reynolds, 2007) provides one explanation. The theory assumes that people construct and maintain a mental representation of the current unfolding event in working memory. This mental representation provides a basis for predicting how activity will unfold—that a bicycle will continue rolling down a hill, or that someone being asked a question will soon begin to speak an answer. When predictions about feature continuities are violated (e.g., a hand stops moving forward), viewers perceive the onset of a new event. According to this theory, event boundaries should tend to occur when features in an activity change; because in general, changes are less predictable than stasis (though there are some exceptions). Interestingly, the emphasis on predictive processes in event segmentation theory is consistent with that of film theorists who have speculated why continuity editing works. Specifically, Bordwell (1985; Bordwell & Thompson, 2006) has argued that film viewers expect the continuity of events in a film and that continuity editing works by substantiating (or violating) those expectations.

What sorts of changes matter for event segmentation? The event-indexing model, which was proposed to account for mental model construction in the context of narrative experiences, provides an answer to this question (Magliano, Zwaan, & Graesser, 1999; Zwaan, Langston, & Graesser, 1995; Zwaan & Radvansky, 1998). The event-indexing model assumes that comprehenders monitor the extent that the focal event is continuous

with the event representation in working memory along dimensions of situational continuity, such as causality, time, and space. Changes in one or more of these dimensions indicate that the situation has changed or is no longer relevant, and thus the mental model in episodic memory must be updated.

There is considerable evidence that comprehenders monitor dimensions of situational continuity in the context of text (see Zwaan & Radvansky, 1998) and film comprehension (Magliano, Taylor, & Kim, 2005; Magliano et al., 2001; Zacks et al., 2009). Recent data suggest that the processing of situational discontinuities is causally involved in the perception of event boundaries. Zacks, Speer, Swallow, and Maley (2010) had participants watch a narrative film (*The Red Balloon*; Lamorisse, 1956), while brain activity was recorded with functional magnetic resonance imaging (fMRI), and then segment the film. Replicating previous results, transient changes in brain activity were observed throughout the cortex: in the lateral and medial posterior cortex, mid-cingulate cortex, and lateral frontal cortex. Many of these regions also showed substantial changes in response to situational discontinuities. Most important, the brain response to situational discontinuity partially mediated the brain response to event boundaries. This supports the hypothesis that the processing of unexpected changes in the situation leads to the perception of an event boundary. Similar results have been reported for segmentation of narrative texts (Speer, Reynolds, Swallow, & Zacks, 2009; Whitney et al., 2009).

Zacks et al. (2009) also made a preliminary assessment of the impact of cuts on event segmentation while viewing the film. Event boundaries tended to occur at cuts; however, when the presence of situational discontinuities was controlled, the direction of this effect reversed. Hence, the reason that cuts were positively associated with event boundaries was that they tended to be points of situational discontinuity. These results are consistent with those of Magliano et al. (2001), who assessed the impact of continuity editing on event segmentation judgments and suggested that event perception in narrative film is, in part, controlled by continuity editing.

Together with the preceding section, these results suggest that the distinction between continuity edits and edits that create event boundaries is critical for film perception. When an edit creates an event boundary, as at a major scene change, the mechanisms of event segmentation work to establish a new model of the situation (Gernsbacher, 1990; Zacks et al., 2007; Zwaan & Radvansky, 1998). Given the hierarchical structure of narratives (and events in general), viewers will likely infer situational relationships between the event models that comprise a mental model for a narrative (e.g., causal, motivational; Trabasso, van den Broek, & Suh, 1989; Trabasso & Suh, 1993; Zwaan & Radvansky, 1998), but inferring those relationships will require the events to unfold in the new scene. As such, the cognitive and perceptual system does not initially bridge the gap between what came before the edit boundary and what came after, and no special kind of processing is required.

1.3. Mechanisms for bridging discontinuity in film

What processing mechanisms might be responsible for bridging at continuity edits and the suppression of that bridging at major scene changes? At least two kinds of relevant

neural mechanisms can be distinguished: *stimulus-driven increases* of processing and *attention-driven regulation*. Measures of brain activity obtained using fMRI are a valuable source of data for studying the operation of both mechanisms.

By stimulus-driven increases, we mean additional neural activity induced by the presentation of novel information. In the visual system, presentation of novel stimulus features is associated with dishabituation of the relevant neurons and therefore transient increases in fMRI signal. This can be observed, for example, with fMRI adaptation paradigms (also called repetition suppression), in which the fMRI response evoked to the initial presentation of a stimulus feature is greater than the response when that feature is repeated (Grill-Spector & Malach, 2001). Similar mechanisms may operate for the processing of ST features and features of action categories. Thus, when one dimension is changing, activity in regions associated with processing that dimension may increase. Discontinuities that occur at event boundaries may stimulate additional processing in a bottom-up fashion that allows the perceiver to assimilate the change.

By attention-driven regulation, we mean a change in activity due to endogenous processes, also called top-down processing. Attention-driven regulation can produce increases (up-regulation) or decreases (down-regulation) in activity. In visuospatial processing, attending to a particular location in space can increase the responsivity of brain regions that process stimuli at that location (Reynolds, Pasternak, & Desimone, 2000), and it can also increase the baseline level of activity in such regions (Luck, Chelazzi, Hillyard, & Desimone, 1997). Similar effects have been found for attending to particular stimulus dimensions such as color, orientation, and motion direction (Maunsell & Treue, 2006). Film cues are likely to produce attention-driven increases in activity. For example, in the famous scene in *The Wizard of Oz* (Fleming, 1939) in which the color palette transitions from grayscale to color, color processing would likely be up-regulated. Film editing also may produce attention-driven decreases, when an edit renders a stimulus dimension momentarily less relevant or meaningful. In particular, when action discontinuities indicate a major scene change, it is not helpful to attempt to bridge the visual discontinuity of the edit. The adaptive response is to form a new mental model of the activity and begin to populate it. Thus, whereas continuity edits might give rise to stimulus-driven increases, the addition of action discontinuity might cue attention-driven down-regulation that would suppress such responses at major scene changes.

1.4. Overview of current studies

The current studies aimed to characterize continuity editing in terms of current cognitive and neurophysiological theories of event perception by testing two key proposals. First, we wanted to test the hypothesis that the meaningful units of activity in a narrative film correspond to the scene boundaries proposed by the continuity editing system—those points at which viewers do not need to bridge across cuts to maintain continuity. This is a proposal on which contemporary cognitive psychology, film theory, and film practice would appear to converge, but to our knowledge it has never been tested empirically. Second, we wanted to test the novel neurophysiological hypothesis that we developed in the previous section:

that neural mechanisms of perception are selectively engaged when processing continuity edits to bridge low-level discontinuities.

To test these hypotheses, we re-analyzed previously collected behavioral (Zacks et al., 2009; Study 2; Zacks et al., 2010) and fMRI data (Zacks et al., 2010). In these studies, participants watched *The Red Balloon* (Lamorisse, 1956), a 33-min feature film. In one study, participants segmented the film into meaningful events (Zacks et al., 2009). In the other study, participants passively viewed the movie while brain activity was recorded with fMRI, and then subsequently segmented the movie (Zacks et al., 2010). In both studies, the film was presented in four segments of 7–10 min each, and participants marked off event boundaries at both a coarse and fine grain by segmenting the movie twice with grain counterbalanced across participants. (Continuity editing accounts do not make specific predictions about differences between different grains of segmentation, but these did occasionally occur and we will describe those instances.) In the previous studies, *The Red Balloon* was coded for changes in the situational dimensions specified by the event-indexing model: spatial location, time, the characters present, and the interactions between characters, causes, and goals. Event segmentation theory proposes that observers should tend to perceive event boundaries at situation changes, because such changes lead to increases in prediction error. This was the case. In the fMRI study, large transient brain responses were observed at event boundaries in the parietal and frontal cortex. Event segmentation theory proposes that these responses should be mediated by the presence of situational changes; this was also the case. Both results replicate previous studies using text materials (Zacks et al., 2009; Study 1; Speer, Reynolds, & Zacks, 2007).

Whereas the situation change coding implemented by Zacks et al. (2009, 2010) examined the narrative content of the film, for this study we conducted a new content analysis to classify editing boundaries along different levels of continuity specified by theories of continuity editing (Bordwell, 1985). In the present study, we first identified locations where editing boundaries occurred and then coded whether each edit boundary maintained continuity of the scene. Following film practice, we coded for continuity in terms of space, time, and action (Bordwell, 1985). We then classified the edit boundaries hierarchically into three classes: (a) continuity edits that were continuous in space, time, and action, (b) spatial-temporal discontinuity edits that were discontinuous in space or time but continuous in action, and (c) action discontinuity edits that were discontinuous in action as well as space or time. Fig. 1 shows examples of each type of boundary. As can be seen in the continuity edit boundary example, the frames around the edit show different viewpoints of the same event, namely the boy talking to the streetcar conductor. The spatial-temporal discontinuity edit example is part of an editing sequence depicting the boy navigating through the small town. Although the shots before and after the edit boundary are in different locations, the boy continues his journey home. The shot prior to the edit boundary in the action discontinuity example depicts the boy closing the window of his apartment and the shot following the boundary depicts a man opening a storefront. The differences between the three classes allow one to isolate three potential classes of processing mechanism: processing common to all edit boundaries, processing specific to ST discontinuity, and processing specific to action



Fig. 1. Examples of continuity, spatial-temporal discontinuity, and action discontinuity editing boundaries.

discontinuity. We modeled these three hypothesized classes of processing using regression models that coded for each type of discontinuity.

The analysis of a film provided by film theoretical continuity features has similarities and differences to the analysis given by the event-indexing model, used in Zacks et al. (2009). Edits classified as spatial-temporal discontinuities in the current coding often corresponded to changes in space and time in the event-indexing coding. However, the relationship between edit type and the remaining event-indexing features was weak.¹ In sum, whereas the event-indexing coding captures what happens in the story of the film, the continuity analysis captures one dimension of how the story was told. There is some degree of correspondence between the two, but it is by no means perfectly overlapping.

According to continuity editing theory, successful continuity edits should be perceived differently than action discontinuities: Even if there are major discontinuities in low-level sensory information and ST frame, viewers should tend to bridge these discontinuities and perceive a continuous narrative event. This proposal is specific to the techniques of narrative film and goes beyond general cognitive theories of event perception. In Study 1, we tested this proposal using behavioral event segmentation data. As described above, event segmentation theory and the event-indexing model predict that the likelihood of perceiving event

boundaries should increase as a function of perceived discontinuity. Continuity editing theory makes the further prediction that physical discontinuities in the stimulus (edits) and ST discontinuities should be less important than action discontinuities (see also Murch, 2001), which make up the scene breaks that filmmakers intend viewers to perceive as meaningful boundaries.

If viewers actively bridge sensory and spatial-temporal continuities at continuity edits, this should be evident in patterns of brain activity. Study 2 investigated the mechanisms of bridging at continuity edits using fMRI. We focused on the hypothesis that specific neural mechanisms would be associated with processing visual discontinuities that were not accompanied by action discontinuity.

2. Study 1: Assessing the impact of continuity editing on event segmentation

Study 1a reanalyzed the behavioral event segmentation data obtained in Study 2 of Zacks et al. (2009); Study 1b reanalyzed the behavioral segmentation data from Zacks et al. (2010). We coded the editing boundaries in *The Red Balloon* and assessed the relationship between editing discontinuities and the perception of event boundaries.

2.1. Method

The methodology used by Zacks et al. (2009, 2010) is summarized here.

2.1.1. Participants

For Study 1a, 41 participants (10 male, 31 female, ages 18–22) were recruited from the Washington University community. For Study 1b, 24 participants (12 male and 12 females, ages 19–28) were recruited from the same community.

2.1.2. Stimuli and tasks

For Studies 1a and 1b, participants viewed the feature film *The Red Balloon* (Lamorisse, 1956), a narrative film about a boy interacting with an apparently sentient red balloon. While viewing the film, participants performed an *event segmentation* task (Newtson, 1973). In one session, participants viewed the film and were asked to press a button to identify the largest units of activity that were natural and meaningful to them (*coarse* segmentation). In another session, they again performed this task, but were asked to identify the smallest units that were natural and meaningful to them (*fine* segmentation). Order of segmentation grain was counterbalanced across participants. *The Red Balloon* was always presented in four clips, with short breaks between each clip. Before performing the coarse and fine segmentation tasks, participants were given the opportunity to practice the task using a brief (155 s) movie of a man building a model with a construction toy. During this practice, the number of event boundaries identified was recorded. If participants gave too few responses (three for coarse grain instructions and six for fine grain), they were given feedback to consider identifying more boundaries.

For Study 1b, participants viewed the film once prior to the segmentation task, during fMRI scanning. Participants were instructed to comprehend the film, but they were not given instructions to segment it. The segmentation session followed the scanning session by 1–6 days.

2.1.3. Continuity editing coding

We identified all the editing boundaries (cuts, fade-ins, fade-outs, dissolves) that occurred in the film and performed a continuity analysis to determine the extent that there was continuity in space, time, or action. Procedures developed by Magliano et al. (2001) were used as a basis for determining whether there was continuity along these dimensions (see Fig. 1), which were inspired by Bordwell's analysis of continuity editing techniques (Bordwell, 1985; Bordwell & Thompson, 2006). If there was continuity for a given variable at a cut point, it was coded as a 0 and if there was discontinuity it was coded as a 1. With respect to space, there were two criteria; adherence to the 180° system and feature overlap in the frames around a boundary that indicated that the events across a cut were in the same location. Determining adherence to the 180° rule is a matter of judgment and is informed by a consideration of the axis of action created in the establishing shot of a scene (the first shot that establishes the spatial layout of the locations and the entities within it). Typically, the camera will not cross this boundary because doing so will cause the orientation of the image to flip on the frame.² If the shot after an editing boundary was within the 180° field established in the scene, or if the content of the frames around the cut clearly indicated that the events were in the same location, then it was coded as continuous. *The Red Balloon* was edited in a fairly conventional manner and locational feature overlap after a boundary was largely diagnostic of spatial continuity/discontinuity. However, consideration of both criteria may be useful to code editing boundaries in less conventional films. The analysis of temporal continuity focused on an assessment of the presence of significant temporal ellipsis. Editing techniques such as fade-outs and dissolves are conventions often used to indicate these. There are instances in which a relatively short temporal ellipsis (say a matter of minutes) can be conveyed with a cut. In these instances, an assessment of the continuity of sounds, music, and action was used to determine whether cut boundaries depicted continuity or discontinuity in time. Finally, continuity in action was determined by an assessment of whether the action after an editing boundary was the continuation of an action established in a prior shot. If so, the editing boundary was considered to be continuous.

The coding was used to identify three different types of continuity situations that can occur at editing boundaries: continuity edits, spatial-temporal discontinuity edits, and action discontinuity edits. Continuity edits occurred when there was continuity in space and time. Spatial-temporal discontinuity edits occurred when there was continuity in action, but discontinuity in space or time. We were not able to distinguish spatial and temporal discontinuities because they were highly correlated in this film. Virtually all of the temporal discontinuities co-occurred with a spatial discontinuity, although spatial discontinuities sometimes occurred without temporal discontinuity. (This occurred in montage sequences that interlaced events happening at different locations, which is a convention for conveying that the events are temporally continuous.) Finally, action discontinuity edits contained a

discontinuity in action and a discontinuity in space or time. There were 211 edits in the film. Of these, 111 were continuity edits, 81 were spatial-temporal discontinuities, and 19 were action discontinuities.

To assess the relations between editing and behavioral event segmentation, we divided the film into 5 s intervals. We coded each interval as to whether it contained a continuity edit, a spatial-temporal discontinuity, or an action discontinuity. These variables were hierarchically nested: All spatial-temporal discontinuities were also edits, and all action discontinuities were also spatial-temporal discontinuities. If an interval contained multiple edits, it was coded according to the most discontinuous edit, with action discontinuity being designated as such. There were 170 time bins that contained an edit and 285 that did not. Of the time bins with edits, 85 contained continuity edits, 67 contained spatial-temporal discontinuities, and 18 contained action discontinuities.

2.2. Results and discussion

For each participant, we recorded the intervals where they identified coarse and fine event boundaries. In Study 1a, the mean proportion of intervals coded as coarse boundaries was 0.16 ($SD = 0.10$), and the mean proportion of intervals coded as fine boundaries was 0.44 ($SD = 0.20$). In Study 1b, the mean proportion of intervals coded as coarse boundaries was 0.13 ($SD = 0.10$), and the mean proportion of intervals coded as fine boundaries was 0.34 ($SD = 0.20$).

We conducted logistic regressions for each participant in which the hierarchically nested predictor variables were as follows: the presence of an edit, the presence of a spatial-temporal discontinuity, and the presence of an action discontinuity. To test whether each of these features was significantly related to segmentation, the regression weights were compared to zero using t -tests with participants as the random variable. This approach takes advantage of the hierarchical nesting of the presence of edits and discontinuities in space/time and action. The test of the coefficient for edits gives a pure estimate of the effect of an edit on segmentation, the test of the coefficient for spatial-temporal discontinuity gives a pure estimate of the additional effect of such discontinuities, and the coefficient for action discontinuity gives a pure estimate of the additional effect of such discontinuities.

For ease of interpretation, we visualize the data using segmentation frequencies rather than logistic regression weights. Fig. 2 shows the mean proportion of 5 s intervals identified as event boundaries in Study 1a as a function of whether the interval had no edits, a continuity edit, a spatial-temporal discontinuity, or an action discontinuity. The no-edit interval provided a comparison for the impact of the different types of editing boundaries. As can be seen in the figure, the presence of continuity edits, spatial-temporal discontinuities, and action discontinuities all were associated with an increase in the frequency of fine segmentation. All three of these increases were statistically significant [edit: mean regression weight = 0.26 ($SD = 0.42$), $t(1, 40) = 3.93$, $p < .001$; spatial-temporal discontinuity: mean regression weight = 0.13 ($SD = 0.36$), $t(1, 40) = 2.47$, $p = .02$; action discontinuity: mean regression weight = 0.35 ($SD = 0.80$), $t(1, 40) = 2.82$, $p = .007$]. For coarse segmentation, edits were not associated with a significant increase in segmentation frequency [mean




Fig. 2. Mean proportion of event segmentation judgments for time bins with no cuts, continuity edit, spatial/temporal discontinuity edit, and action discontinuity edit for Study 1 (error bars reflect standard error).

regression weight = 0.06 ($SD = 0.39$), $t(1, 40) = 0.98$, $p = .33$]. Spatial-temporal discontinuities were associated with a small but significant *decrease* in segmentation frequency [mean regression weight = -0.28 ($SD = 0.60$), $t(1, 40) = -2.99$, $p = .004$]. Action discontinuities were associated with a larger and significant increase in segmentation [mean regression weight = 0.96 ($SD = 1.03$), $t(1, 40) = 5.99$, $p < .001$].

As can be seen in Fig. 3, Study 1b closely replicated the results of Study 1a. For fine segmentation, the logistic regression analyses found statistically significant increases in segmentation with the presence of an edit [mean regression weight = 0.28 ($SD = 0.28$), $t(1, 23) = 4.84$, $p < .001$] and an action discontinuity [mean regression weight = 1.04 ($SD = 0.56$), $t(1, 23) = 9.10$, $p < .001$]. Unlike Study 1a, the effect of spatial-temporal discontinuities was not significant [mean regression weight = -0.07 ($SD = 0.55$, $p = .59$]. For coarse segmentation, there was again no significant effect of edits on segmentation [mean regression weight = 0.17 ($SD = 0.49$), $t(1, 23) = 1.65$, $p = .11$]. As in Study 1, action discontinuities were associated with increases in coarse segmentation and spatial-temporal discontinuities were associated with decreases; however, neither effect was statistically significant [action discontinuities: mean regression weight = 1.75 ($SD = 4.93$), $t(1, 23) = 1.73$, $p = .096$; spatial-temporal discontinuities: mean regression weight = -1.34 ($SD = 3.48$), $t(1, 23) = -1.88$, $p = .073$].

The most striking observation that can be drawn from these data is that only discontinuity in action predicted increases in both coarse and fine segmentation. In fact, action discontinuity appears to be a requisite for the perception of a coarse event boundary in this film. However, the presence of continuity edits and spatial-temporal discontinuity edits led to small increases in the likelihood of perceiving a fine event boundary. Thus, the major conclusion to be drawn from this analysis is that action discontinuities have a large and robust effect on event segmentation, but visual or ST discontinuities do not. The two




Fig. 3. Mean proportion of event segmentation judgments for time bins with no cuts, continuity edit, spatial/temporal discontinuity edit, and action discontinuity edit for Study 2 (error bars reflect standard error).

datasets converge well (though Study 1b, with a smaller sample size, shows more variability). The differences between coarse and fine segmentation are suggestive and merit further investigation. Whereas edits and spatial-temporal discontinuities were not associated with increases in coarse segmentation, they were associated with small increases in fine segmentation. Perhaps if viewers were asked to segment at a yet finer grain, edits and ST discontinuities would come to the fore and the association between action discontinuities and segmentation would weaken.

These results are consistent with the perspective that continuity editing enables viewers to perceive a sense of cohesion across an editing boundary by maintaining a sense of continuity of action (Bordwell, 1985). Moreover, it is likely the case that continuity edits and edits with spatial-temporal discontinuities (but action continuity) require viewers to engage in processes that enable them to “bridge” across the edit. The fMRI data provided insights into the brain networks that might support these processes.

3. Study 2: Neural mechanisms associated with processing discontinuity

In Study 2, we tested the hypothesis that spatially localized brain mechanisms would be selectively engaged in processing editing discontinuities. Consider what may occur at an edit. First and foremost, the visual information changes dramatically, instantaneously, and discontinuously. This occurs at every edit. At a subset of these edits, visual discontinuity is accompanied by discontinuity of space and time. Finally, at a subset of the edits with ST discontinuities, there are also discontinuities in action. These edits constitute major scene changes, where one major unit of the film ends and another begins. As we described in the

Introduction, visual and ST discontinuities would be expected to produce *stimulus-driven increases* in processing that may contribute to bridging the discontinuities. These are the responses produced by the brain when it is presented with a salient change in the stimulus. Increased response to a changed stimulus may be important for bridging the change. For example, if a film cuts from a head-on close-up of an actor to a three-quarter view medium shot, stimulus-driven increases may help compute the higher level visual features necessary to match the identity of the character across shots. If this bridging is unnecessary at major scene changes and is effectively suppressed, the stimulus-driven increases should be reduced or eliminated at those points; this would be an instance of what we have termed *attention-driven regulation*.

A natural hypothesis is that visual discontinuity will produce stimulus-driven increases in visual processing regions. These would be expected to include the earliest cortical stages of visual processing in areas V1 and V2 and likely much wider activation throughout posterior occipital, temporal, and parietal cortex. Another natural hypothesis is that ST discontinuities would lead to stimulus-driven increases in areas known to respond selectively to the processing of new spatial locations. Two such areas are the “parahippocampal place area” (PPA) and the retrosplenial complex (RSC). Epstein (2008), in reviewing neuroimaging and neuropsychological studies of spatial navigation, argued that the PPA is selectively involved in the representation of specific places, and that the RSC is selectively involved in situating specific locations within a larger spatial context. Finally, it is possible that action discontinuities would also lead to stimulus-driven increases. The onset of a new action itself could activate new event representations, just as changes in visual stimuli activate visual representations or changes in location activate representations of space. A small number of studies provide evidence for selective responses when new actions are observed (Dinstein, Hasson, Rubin, & Heeger, 2007; Hamilton & Grafton, 2009). However, it is not clear whether these responses reflect stimulus-driven increases or attention-driven regulation; nor is it clear whether these responses reflect changes in the action person or in physical parameters associated with the action such as movement trajectory.

We hypothesized that the presence of action discontinuity would produce attention-driven down-regulation that would suppress stimulus-driven responses to visual discontinuities. Thus, we predicted a pattern that brain areas selectively engaged in bridging visual discontinuities would show increased activity in response to edits but decreased activity in response to the addition of action discontinuity. One reasonable place to expect such patterns is in high-level visual processing areas such as non-primary occipital cortex, inferior temporal cortex, and posterior parietal cortex. Regions involved in the extraction of complex meaningful patterns from visual input are likely to be helpful for bridging visual discontinuities—and potentially harmful when such bridging is not helpful.

3.1. Method

3.1.1. MRI scanning

Functional MRI images were acquired on a Siemens 3T Vision MRI Scanner (Erlangen, Germany) using a T2* weighted asymmetric spin-echo echo-planar sequence (slice

TR = 64 ms, TE = 25 ms) in 32 transverse slices (4.0 mm isotropic voxels) aligned with the anterior and posterior commissures. The images were pre-processed to correct for acquisition artifacts and participant motion, and aligned to a standard stereotaxic space (Talairach & Tournoux, 1988). During alignment, the images were resampled to 3.0 mm isotropic voxels. In the scanner, stimuli were back-projected onto a screen located at the head of the scanner and viewed through a mirror. (More details of the imaging methods are given in Zacks et al., 2010.)

3.2. Results and discussion

To characterize the relations between editing and evoked brain activity, we used a two-stage statistical procedure. Statistical analyses were performed with in-house software (<http://www.nil.wustl.edu/~fidl/>). In the first stage, we fitted linear models of a form that paralleled the logistic regression models used to analyze the segmentation data. There were predictor variables coding the presence of an edit, the presence of a ST discontinuity, and the presence of an action discontinuity. Each of these was represented as a train of impulse responses convolved with a model hemodynamic response function (Boynton, Engel, Glover, & Heeger, 1996). We also included predictor variables controlling for effects not of theoretical interest: the presence of the movie, scan-to-scan differences in baseline response, and the linear trend within each scan. In the second stage, the regression weights for each effect of interest for each voxel were submitted to *t*-tests with participant as the random effect. The resulting maps of *t*-statistics were converted to *z*-statistics and corrected for multiple comparisons. Based on the Monte Carlo simulations of McAvoy, Ollinger, and Buckner (2001), we retained only voxels that exceeded a *z* threshold of 4.5 and were part of a cluster of two or more voxels exceeding this threshold. The resulting statistical maps are depicted in Figs. 4–6 for edits, spatial-temporal discontinuities, and action discontinuities, respectively. To summarize these maps, we identified local extrema within each map after smoothing with a 5 mm spherical kernel, and then iteratively combined extrema by averaging the coordinates of extrema that were closer than 30 mm. Each voxel was then assigned to a region based on the maximum or minimum to which it was closest. (Regions were formed separately for positive and negative *z*-statistics.) The locations of the largest *z*-statistic within each region, the *z*-value at that location, and the volume of the region are reported in Tables 1–3.

As can be seen in Fig. 4 and Table 1, the presence of a continuity edit was associated with large increases in early visual cortex bilaterally. Increases were also observed in lateral occipitotemporal cortex and the precuneus bilaterally, and in a small region of left medial frontal cortex. We observed decreases in primary somatomotor cortex, the supplementary motor area, the cingulate motor regions (Picard & Strick, 1996) the insula, and small regions of lateral inferior frontal cortex and lateral inferior temporal cortex. All decreases were bilateral.

As can be seen in Fig. 5 and Table 2, discontinuities in space or time in the movie were associated with increases in a pair of focal regions in the parahippocampal cortex, likely corresponding to the PPA (see Epstein, 2008). Increases were also observed in the superior parietal lobule bilaterally, in the insula bilaterally, and in the lateral-posterior temporal lobes bilaterally. We observed decreases bilaterally in early visual cortex, the precuneus and




Fig. 4. Regions of significant change at edits (orange-red reflects transient increases in activation, and blue indicates transient decreases in activation).




Fig. 5. Regions of significant change at spatial/temporal discontinuities (orange-red reflects transient increases in activation, and blue indicates transient decreases in activation).

posterior cingulate cortex, the lateral superior parietal lobule, and lateral premotor cortex. There were also decreases in small regions of the right posterior inferior parietal lobule and right lateral superior frontal cortex.

Finally, as can be seen in Fig. 6 and Table 3, discontinuities in action were associated with significant increases only in lateral occipital cortex, bilaterally. We observed decreases in the anterior superior parietal cortex bilaterally, and in the left superior temporal sulcus and left premotor cortex.




Fig. 6. Regions of significant change at action discontinuities (orange-red reflects transient increases in activation, and blue indicates transient decreases in activation).

Figs. 4–6 and Tables 1–3 provide evidence both for stimulus-driven increases and for attention-driven regulation, in particular for attention-driven reductions. The most salient responses detected were robust increases in activity in visual processing areas at continuity edits. This is quite likely stimulus driven—the large amount of novel visual information produces larger responses in the neurons processing those visual features. Continuity edits were also associated with increases in nearby areas at the juncture of the parietal, temporal, and occipital lobes. These regions have been shown to be activated for a wide range of tasks involving the analysis of human behavior, including processing biological and animate motion and understanding actors goals (Beauchamp, Lee, Haxby, & Martin, 2003; Frith, 2007; Grossman et al., 2000; Mar, Kelley, Heatherton, & Macrae, 2007; Martin & Weisberg, 2003). Similarly, the selective increases in parahippocampal cortex at ST discontinuities are likely a stimulus-driven response. This region has been found to be selectively active when observers look at pictures of buildings compared to pictures of other classes of objects (Epstein & Kanwisher, 1998), and when they experience a visual environment from an immersive first-person viewpoint compared to an overhead map-like view (Shelton & Gabrieli, 2002). In a previous study using a feature film, this region was selectively activated when viewers saw shots of landscapes or buildings (Hasson, Nir, Levy, Fuhrmann, & Malach, 2004). Other increases at discontinuities are harder to explain as stimulus-driven responses. These include the increases in lateral temporal cortex and the insula at ST discontinuities and the increases in lateral occipital cortex at action discontinuities.

Several of the decreases observed here are suggestive of attention-driven down-regulation. First, the fact that motor regions decreased at continuity edits is consistent with attention-driven down-regulation. Viewers of human activity consistently activate motor representations corresponding to the actions they observe (Iacoboni et al., 2005). Such activity may be temporarily suppressed in response to continuity edits, because these edits

Table 1
Regions of significant change at continuity edits

Atlas Coordinate of Peak			Location	Volume (cm ³)	Peak z-statistic
X	Y	Z			
4	-84	-3	Early visual cortex (BA 17/18)	30.13	8.19
2	-54	51	Precuneus (BA 7)	10.21	6.73
-40	-69	12	L. inferior temporal cortex (BA 19/37), lateral occipitotemporal cortex (BA 19/39)	14.45	6.81
14	-75	3	R. inferior temporal cortex (BA 19/37), lateral occipitotemporal cortex (BA 19/39)	25.14	7.13
22	-30	-3	R. medial temporal (BA 27/30)	3.62	6.52
-20	-30	-6	L. medial temporal (BA 27/30)	1.4	6.28
38	-21	18	R. insula	29.43	-7.56
-34	-21	12	L. insula	21.79	-7.26
40	-27	54	R. somatomotor cortex (BA 1/2/4)	16.66	-6.84
-40	-24	48	L. somatomotor cortex (BA 1/2/4)	13.47	-6.52
-38	0	12	L. frontal operculum (BA 22/44)	9.53	-6.83
10	-12	42	Supplementary motor area (BA 6)	13.28	-6.14
34	6	15	R. frontal operculum (BA 22/44)	3.94	-6.3
10	-54	-18	Brainstem	4.37	-6.17
-50	-57	48	L. inferior parietal lobule (BA 40)	1.67	-5.33
-28	-24	21	L. parietal white matter	3.27	-5.49

Table 2
Regions of significant change at spatiotemporal discontinuities

Atlas Coordinate of Peak			Location	Volume (cm ³)	Peak z-statistic
X	Y	Z			
-34	-72	48	L. superior parietal lobule (BA 7)	4.08	7.07
40	-63	45	R. superior parietal lobule (BA 7)	1.24	6.22
-20	-36	-15	L. parahippocampal cortex	1.51	6.23
20	-39	-15	R. parahippocampal cortex	1.05	5.54
38	-18	21	R. insula	1.19	5.68
-58	-54	-9	L. insula	1.78	6.29
-4	-36	36	Precuneus/posterior cingulate (BA 31)	1.16	5.23
8	-78	3	Early visual cortex (BA 17/18)	10.45	-6.84
4	-51	54	Precuneus (BA 7)	10.21	-6.79
38	-9	57	R. motor/premotor cortex (BA 4/6)	5.26	-6.81
44	-69	12	R. temporoparietal junction (BA 19/37/39)	5.37	-6.42
-44	-72	3	L. temporoparietal junction (BA 19/37/39)	1.67	-6.57
20	-87	27	R. cuneus (BA 19)	1.24	-5.81
4	42	39	Medial frontal gyrus (BA 8)	1.40	-5.28
-14	-51	48	L. parietal (BA 7/40)	2.11	-5.45

Table 3
Regions of significant change at action discontinuities

Atlas Coordinate of Peak			Location	Volume (cm ³)	Peak z-statistic
X	Y	Z			
28	-87	3	R. lateral occipital cortex (BA 18)	4.48	5.73
14	-60	24	L. lateral occipital cortex (BA 18)	0.78	5.22
-16	-66	54	L. inferior parietal lobule/superior parietal lobule (BA 7/40)	3.29	-5.66
-52	-72	0	L. posterior inferior temporal gyrus (BA 19/37)	1.70	-5.66
56	-42	36	R. supramarginal gyrus (BA 40)	1.11	-5.4
10	-57	57	R. superior parietal lobule (BA 7)	1.46	-5.51
-26	-12	54	L. premotor cortex (BA 4/6)	0.81	-5.07
-46	-42	60	L. supramarginal gyrus and adjacent parietal cortex (BA 40)	1.16	-5.37

disrupt the ability to map from visual input to motor representations. Second, the reduction in early visual cortex at ST discontinuities may reflect attention-driven suppression. Edit boundaries in general produce increases in these regions (see Fig. 4), and we have suggested that this is a stimulus-driven response. Such responses would be important for establishing connections between the current visual input and the visual information that preceded the edit. However, when the viewer detects a ST discontinuity, it may be adaptive to suppress this increased visual processing because the new visual input is likely not related to the pre-edit visual information. Finally, the reductions in lateral parietal cortex during action

Table 4
Regions selectively activated during bridging at edits

Atlas Coordinate of Center of Mass			Location	Volume (cm ³)
X	Y	Z		
-1	-71	44	Cuneus/precuneus (BA 7/18/19)	22.06
-45	-70	7	L. inferior temporal cortex (BA 19/37), lateral occipitotemporal cortex (BA 19/39)	7.45
45	-65	8	R. inferior temporal cortex (BA 19/37), lateral occipitotemporal cortex (BA 19/39)	2.57
-16	-61	-6	L. inferior temporal cortex (BA 19/37)	1.97
16	-60	-6	R. inferior temporal cortex (BA 19/37),	1.08
-25	-59	51	L. superior parietal lobule (BA 7)	0.89
25	-6	58	R. precentral sulcus (BA 6)	0.57
-55	-38	30	L. temporoparietal junction (BA 19/39)	0.49
-36	-8	47	L. precentral sulcus (BA 6)	0.43
25	-59	57	R. superior parietal lobule (BA 7)	0.24
-43	-59	-25	L. fusiform gyrus (BA 36)	0.22
-1	-1	-8	Ventricle	0.22
-24	0	56	L. middle frontal gyrus (BA 6)	0.14

discontinuities could reflect attention-driven suppression. This region has been argued to selectively represent the goals of action. For example, Hamilton and Grafton (2009) used a repetition suppression paradigm in which the goal or location of a reaching movement was repeated. The left supramarginal gyrus showed reduced activity when goals were repeated. Action discontinuities are major scene breaks, at which it may be adaptive to transiently suppress goal processing while the parameters of the new scene are established.

More speculatively, we note that a number of the posterior regions that increased in activity at continuity edits corresponded to regions that increase in activity at event boundaries (Zacks et al., 2001, 2006). These included early visual cortex, the precuneus, and the lateral inferior parietal lobule. Zacks et al. (2010) found that these regions responded at event boundaries in *The Red Balloon*, and that these responses were significant after controlling for the presence of edits. We propose that these responses could reflect in part attentional up-regulation. Such responses have also been observed in these regions (excepting early visual cortex) at event boundaries in narrative texts (Speer et al., 2007; Whitney et al., 2009). One possibility is that perceptual processing pathways are up-regulated as a consequence of the mechanisms associated with detecting an event boundary (Zacks et al., 2007).

In sum, these analyses provide evidence for both stimulus-driven increases and attention-driven regulation. Continuity editing theory makes a very specific prediction about the contributions of these two mechanisms to film comprehension. To process a continuity edit, the viewer must bridge the visual discontinuity to maintain continuity of space and time. However, when a major scene break occurs at an action discontinuity, there is no need to perform this bridging. This suggests that brain regions that are particularly important for bridging in continuity edits should show a distinctive pattern: They should show relative increases at cuts, but relative *decreases* when action discontinuities are present. In other

words, such regions should increase in activity selectively for continuity edits, but these increases should be attenuated or action discontinuities. We performed an additional analysis to test specifically for this pattern.

The final analysis was conducted to identify brain activity that is selectively involved in processing continuity edit. To identify brain regions specifically engaged in bridging visual discontinuities at edits, we performed a conjunction analysis using the z -statistic maps described in the previous section. The conjunction included only those voxels that increased in response to edits with a z -statistic >2.5 and decreased in response to action discontinuities with a z -statistic <-2.5 . This corresponds to a voxelwise Bonferroni-corrected false-positive rate of 0.00015. The regions that satisfied these criteria are shown in Fig. 7 and listed in Table 4. They were predominantly non-primary visual cortex, inferotemporal cortex, and parietal cortex. As noted previously, the parietal regions included regions that have been associated with processing the higher order features of human movement and actors' goals. Also included were small regions in the precentral sulcus. All activations were bilateral.

We interpret this pattern as indicating that edits push the viewer to engage a set of higher order visually driven processes, in order to maintain a consistent representation of an ongoing event in the face of dramatic low-level visual discontinuity. This includes matching objects from shot to shot, which may stimulate the inferotemporal cortex; processing biological and animate motion to match characters and their actions, which may stimulate the posterior superior temporal sulcus and adjacent regions; and matching the motor actions of characters, which may stimulate motor areas in the precentral sulcus. These are all stimulus-driven effects. If an edit is recognized as a scene change, observers may transiently inhibit much of this processing, waiting until the new scene is established to engage in it. This would be an instance of attention-driven down-regulation.




Fig. 7. Regions selectively activated during bridging at edits.

4. General discussion

To understand a film, viewers must understand and build a mental model of the intended event structure, and filmmakers must construct an editing sequence that supports this endeavor. At its most basic level, editing must indicate when viewers are intended to perceive a continuity or discontinuity of the event sequence across an editing boundary. As specified by classical views of continuity editing (Bordwell, 1985), filmmakers must manage the flow of visual content across the boundary so that it supports the perception of spatial-temporal and action continuities across boundaries. To our knowledge, there are few if any studies that have directly studied how viewers cognitively, behaviorally, or neurologically respond to continuity editing (for exceptions, see Magliano et al., 2001; Zacks et al., 2009).

4.1. Classical editing supports psychological mechanisms for segmentation activity on action

The behavioral data from the current study clearly indicate that discontinuity of action is the strongest predictor of the perception of event boundaries at editing locations. This qualifies the finding from Magliano et al. (2001) who compared the relative contribution of shifts in spatial regions and time on the perception of event boundaries. That study used several feature length films, which afforded an exploration of the relative impact of shifts in narrative space and temporal ellipsis on event segmentation in narrative film. Magliano et al. (2001) found that a shift in narrative regions alone did not engender the perception of an event boundary, primarily because filmmakers can shift from one region to another in order to create a sense of temporal concurrence of events occurring in the two locations. Magliano et al. (1996) showed that doing so supports predictions that the events happening in the two different locations will eventually converge (e.g., the cars engaged in the chase will run into the pane of glass carried by the hapless workers as they cross the street). However, Magliano et al. (2001) found that temporal ellipsis by itself increased the likelihood of perceiving an event boundary. Magliano et al. (2001) did not code for action continuity across editing boundaries and the results of the present study suggest that the effects of temporal ellipsis on event segmentation may have at least in part been driven by a co-occurrence with discontinuities in action.

The significant effect of the presence of continuity edits on fine grain segmentation merits particular comment because at first blush it appears to conflict with the results of Zacks et al. (2009). In that analyses, once situational changes were taken into account, edit boundaries were found to be associated with *decreases* in the likelihood of an event boundary. However, the situational changes coded by Zacks et al. included several types of change not considered in theoretical accounts of continuity editing, and certainly factors outside of continuity editing affect the perceived event structures for narrative film (Magliano et al., 2001; Zacks et al., 2009). For example, Zacks et al. (2009) found character changes were strongly correlated both with editing boundaries and with event boundaries. In the current analysis, it is possible that the increase in event boundary identification at continuity edits was due to changes to other situational dimensions (e.g., the introduction of new characters) rather than to visual discontinuities. Thus, the major conclusions to be drawn from this analysis are that

action discontinuities have a large and robust effect on event segmentation, but visual or ST discontinuities do not.

New actions were frequently perceived as new events much more strongly than continuity edits or ST discontinuities. This finding could lead one to conclude that space and time have little to do with why continuity editing works—that action dominates viewers' segmentation. However, our view is that space and time likely *are* critical for how continuity editing functions, because they play a key role in defining discontinuities in action (Bordwell, 1985). In classical editing, space and time are managed to make the structure of the action clear to the viewer. If spatial or temporal continuities at editing boundaries are mismanaged or intentionally left ambiguous, it will be more challenging to understand the action. As an example, adherence to the 180° rule, a central aspect of managing space across editing boundaries, is critical for creating a sense of continuity in action (Bordwell & Thompson, 2006). For example, if one crossed the 180° axis across an editing boundary, the image will flip direction on the frame and movement will appear to occur in the opposite direction than it occurred prior to the edit (e.g., action moving from right to left across the frames preceding the editing boundary will appear to occur from left to right after the boundary).

In the present study, ST discontinuities were associated with small increases in fine segmentation, but small *decreases* in coarse segmentation. We suspect that this reflects the effective management of space and time in the service of depicting action. When a shift in space and time occurs without a shift in action, filmmakers may provide explicit cues (e.g., continuity of music, dialog, and other sounds) to the viewer that lead the viewer to bridge the discontinuity. The fMRI data provide some clues to the networks that may support these bridging processes; we will return to this possibility shortly.

Although these results demonstrate that continuity editing works as intended, it is important to note that there is debate as to whether following these conventions is necessary for effective narration in film (e.g., Germeys & d'Ydewalle, 2007; Wurtzel, 1983). For example, Germeys and d'Ydewalle (2007) showed that early perceptual processes are not dramatically disrupted when editing rules are violated (e.g., a violation of the 180° rule, which would flip the images on the frame), as would be expected by film theorists who emphasize the importance of continuity editing in film. Intentionally breaking the rules of continuity editing may be artistically desirable and may not disrupt early perceptual processing of a film (Germeys & d'Ydewalle, 2007). However, there may be costs for violating these conventions in terms of comprehension and memory (see Graesser, Kassler, Kreuz, and McLain-Allen [1998] and Ohtsuka and Brewer [1992] for findings supporting this claim in the context of text).

The further exploration of continuity editing on event perception may provide valuable clues to how visual perception operates in the face of eye movements, blinks, and occlusions. This is particularly the case when editing conventions roughly correspond to the perceptual flow of information as we interact in the real world, which does occur in a subset of types of continuity edits (Schwan & Ildirar, 2010). Additionally, the present analysis did not take into account the variety of techniques that comprise continuity editing (see Bordwell & Thompson, 2003), but rather focused on different levels of continuity that can

arise from adherence to these techniques. A study that would explore the processing implication of the rules of continuity editing would require the use of variety of films that contain a rich set of editing locations that both adhere to and violate these rules, which is not the case with the *Red Balloon*.

4.2. Neural processing of discontinuities reflects their hierarchical organization

The fMRI data showed that different brain regions responded during different levels of editing discontinuity. We interpreted these changes as reflecting either stimulus-driven or attention-driven processing in response to these changes. Specifically, the different types of cuts reflect different changes in perceptual and conceptual information in the frames surrounding the edit boundary. Some of the changes require the viewer to process new content and map it onto the content of the frames preceding the edit (at both perceptual and conceptual levels), but doing so may direct the viewers' attention from other levels of processing (e.g., visual mapping may direct attention away from conceptual mapping).

Continuity edits reflect a situation in which there is a disruption in the visual flow of information, but viewers need to establish how the frames across the boundary are related. That is, given that the frames around continuity edits most often contain overlapping content (e.g., one frame shows a medium shot of a character, and the frame across the cut shows a closeup of his face), viewers likely have to engage in mapping processes to establish how the perceptually salient features across the frames are related. The fMRI data support this interpretation and revealed a transient increase in activation in networks that are associated with early visual processing and the precuneus, which we interpreted as being reflective of stimulus-driven processing. Additionally, there were transient increases in extrastriate regions associated with the processing of rigid motion and biological motion (Blake & Shiffrar, 2007). Interestingly, although the behavioral data indicated that viewers maintained a sense of event continuity at continuity edits, we found bilateral transient decreases in the somatomotor cortex and the insula. We interpreted the changes in the somatomotor cortex as reflecting attention-driven down-regulation. Specifically, mapping the visual features across the editing boundary may have momentarily directed the viewers' attention away from mapping action continuities across the editing boundary.

Editing boundaries with spatial or temporal discontinuities require viewers to understand that the action is continuing despite the fact that the camera is shifting to a new location (or between locations in instances where the shots are shifting between concurrent actions occurring at different locations). The fMRI data showed a transient bilateral increase in activation in the parahippocampal cortex, which we interpreted as a stimulus-driven response to the change in spatial location. This change in location likely provides a cue that there is no need to map visual features across the frame, and consequently we found transient decreases in some of the same areas of the early visual cortex that may have been involved in mapping visual features across continuity edits. We interpreted this as reflective of an attention-driven down-regulation in visual mapping across these edit boundaries.

Action discontinuities primarily occurred at breaks between scenes, which require viewers to shift from updating a mental model for the prior scene to building a new mental model for the new scene (Gernsbacher, 1990; Zwaan & Radvansky, 1998). However, we found surprisingly little transient increase or decrease associated with action discontinuities and none involved areas associated with high-level visual processing, as hypothesized. Nonetheless, the transient decreases are most amenable to interpretation in terms of attention-driven down-regulation. Specifically, we found reductions in the lateral parietal cortex, which has been argued to selectively represent the goals of action (Hamilton & Grafton, 2009). The suppression of goal processing as new scenes are established may aid in shifting from one structure in the mental model to another.

4.3. Bridging low-level discontinuities is associated with specialized neural mechanisms

At continuity edits viewers are faced with massive visual discontinuity but maintain continuity of action. We hypothesized that this would be associated with selective neural processing to bridge the low-level discontinuities. We found evidence for such selective processing in the non-primary visual cortex, inferotemporal cortex, and parietal cortex, all bilaterally. This may reflect higher order, visually driven processing that is requisite for maintaining a consistent representation of an ongoing event in the face of visual discontinuities across the boundary. Transient increases in the inferotemporal cortex may reflect processing that supports matching objects from shot to shot. Transient increases in the posterior superior temporal sulcus and the precentral sulcus may reflect processes that allow viewers to match characters and their actions. These particular interpretations are admittedly speculative, but these results provide strong support for a broader conclusion: Perceptual processing appears to be regulated in response to cues given by continuity editing. This provides empirical support for the prescriptions of classical continuity editing (Bordwell, 1985; Bordwell & Thompson, 2006).

4.4. A cognitive science approach to film editing

The present study shows the value of an inclusive cognitive science approach to film. Cognitively oriented theories of viewer or reader response that have been proposed in literary or film studies may warrant empirical testing (Kreuz & MacNealy, 1996). Cognitive scientists interested in comprehension and affective response have attempted to test some of these proposals empirically or by detailed literature review (e.g., Bortolussi & Dixon, 2003; Brewer & Ohtsuka, 1988; Tan, 1996). Similarly, vision scientists have investigated the effects of cuts over decades (e.g., Carroll & Bever, 1976; Cutting, DeLong, & Nothelfer, 2010; Hochberg & Brooks, 2006; Kraft, 1986; Schwan, Garsoffky, & Hesse, 2000; Smith & Henderson, 2008). The present study clearly shows that classical theories of continuity editing are amenable to psychological testing and that related theories of event comprehension in cognitive science, namely event segmentation theory (e.g., Zacks et al., 2007) and the event-indexing models (Zwaan & Radvansky, 1998), provide a theoretical framework for doing so.

Notes

1. The proportion of situational shifts identified by Zacks et al. (2009) that occurred in time bins with no cut, continuity edit, ST discontinuity edit, and action discontinuity edit:

	New Characters	Character–Character Interactions	Character–Object Interactions	Time	Space	Causal Breaks	Goal
Continuity edit	0.25	0.22	0.20	0.04	0.22	0.02	0.28
ST discontinuity	0.28	0.13	0.07	0.70	0.60	0.20	0.16
Action discontinuity	0.10	0.04	0.02	0.26	0.16	0.17	0.03
No cut	0.38	0.61	0.72	0.00	0.02	0.61	0.53

2. It is important to note that there are conventions for breaking the 180° rule and still maintaining situational continuity across an editing boundary and, as such, adherence to this rule is not strictly necessary.

Acknowledgments

The project described was supported by grant RO1MH70674 from the National Institute of Mental Health and grant R01AG031150 from the National Institute of Aging awarded to the second author and the Center for the Interdisciplinary Study of Language and Literacy at Northern Illinois University.

References

- Alaerts, K., Swinnen, S. P., & Wenderoth, N. (2009). Is the human primary motor cortex activated by muscular or direction-dependent features of observed movements? *Cortex*, *45*, 1148–1155.
- Beauchamp, M. S., Lee, K. E., Haxby, J. V., & Martin, A. (2003). fMRI responses to video and point-light displays of moving humans and manipulable objects. *Journal of Cognitive Neuroscience*, *15*, 991–1001.
- Blake, R., & Shiffrar, M. (2007). Perception of human motion. *Annual Review of Psychology*, *58*, 47–73.
- Bordwell, D. (1985). *Narration in the fiction film*. Madison: University of Wisconsin Press.
- Bordwell, D. (2002). Intensified continuity: Visual style in contemporary American film. *Film Quarterly*, *55*, 16–28.
- Bordwell, D., & Thompson, K. (2006). *Film art: An introduction*. New York: McGraw-Hill.
- Bortolussi, M., & Dixon, P. (2003). *Psychonarratology: Foundations for the empirical study of literary response*. Cambridge, UK: Cambridge University Press.
- Boynton, G. M., Engel, S. A., Glover, G. H., & Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience*, *16*, 4207–4221.
- Brewer, W. F., & Ohtsuka, K. (1988). Story structure, characterization, just world organization, and reader affect in American and Hungarian short stories. *Poetics*, *17*, 395–414.
- Carroll, J. M., & Bever, T. G. (1976). Segmentation in cinema perception. *Science*, *191*, 1053–1055.
- Cutting, J. E. (2005). Perceiving scenes in film and in the world. In J. D. Anderson & B. F. Anderson (Eds.), *Moving image theory: Ecological considerations* (pp. 9–27). Carbondale, IL: Southern Illinois University Press.

- Cutting, J. E., DeLong, J. E., & Nothelfer, C. E. (2010). Attention and the evolution of Hollywood film. *Psychological Science*, *21*, 432–439.
- Dancyger, K. (2007). *The technique of film and video editing: History, theory, and practice* (4th ed.). Burlington, MA: Focal Press.
- Dinstein, I., Hasson, U., Rubin, N., & Heeger, D. J. (2007). Brain areas selective for both observed and executed movements. *Journal of Neurophysiology*, *98*, 1415–1427.
- Epstein, R. A. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends in Cognitive Sciences*, *12*, 388–396.
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*, 598–601.
- Fleming, V. (1939). *The wizard of Oz*. Loew's Incorporated.
- Frith, C. (2007). The social brain? *Philosophical Transactions of the Royal Society B-Biological Science*, *362*, 671–678.
- Germeys, F., & d'Ydewalle, G. (2007). The psychology of film: Perceiving beyond the cut. *Psychological Research*, *71*, 458–466.
- Gernsbacher, M. A. (1990). *Language comprehension as structure building*. Hillsdale, NJ: Erlbaum.
- Graesser, A. C., Kassler, M. A., Kreuz, R. J., & McLain-Allen, B. (1998). Verification of statements about story worlds that deviate from normal conceptions of time: What is true about *Einstein's Dreams*? *Cognitive Psychology*, *35*, 246–301.
- Grill-Spector, K., & Malach, R. (2001). fMR-adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychologica*, *107*, 293–321.
- Grossman, E. D., Donnelly, M., Price, R., Pickens, D., Morgan, V., Neighbor, G., & Blake, R. (2000). Brain areas involved in perception of biological motion. *Journal of Cognitive Neuroscience*, *12*, 711–720.
- Hamilton, A. F. D. C., & Grafton, S. T. (2009). Repetition suppression for performed hand gestures revealed by fMRI. *Human Brain Mapping*, *30*, 2898–2906.
- Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science*, *303*, 1634–1640.
- Hochberg, J., & Brooks, V. (1978). *The perception of motion pictures, (Vol. 10)*. New York: Academic Press.
- Hochberg, J., & Brooks, V. (2006). Film cutting and visual momentum. In M. A. Peterson, B. Gillam, & H. A. Sedgwick (Eds.), *In the mind's eye: Julian Hochberg on the perception of pictures, films, and the world* (pp. 206–228). New York: Oxford University Press.
- Iacoboni, M., Molnar-Szakacs, I., Gallese, V., Buccino, G., Mazziotta, J., & Rizzolatti, G. (2005). Grasping the intentions of others with one's own mirror neuron system. *PLoS Biology*, *3*, 1–7.
- Kraft, R. N. (1986). The role of cutting in the evaluation and retention of film. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *12*, 155–163.
- Kreuz, R. J., & MacNealy, M. S. (Eds.) (1996). *Empirical approaches to literature and aesthetics* (pp. 201–220). Norwood, NJ: Ablex.
- Kurby, C. A., & Zacks, J. M. (2008). Segmentation in the perception and memory of events. *Trends in Cognitive Sciences*, *12*, 72–79.
- Lamorisse, A. (1956). *The red balloon [Motion picture]*. Films Montsouris.
- Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of Macaque visual cortex. *Journal of Neurophysiology*, *77*, 24–42.
- Magliano, J. P., Dijkstra, K., & Zwaan, R. (1996). Generating predictive inferences while viewing a movie. *Discourse Processes*, *22*, 199–224.
- Magliano, J. P., Miller, J., & Zwaan, R. A. (2001). Indexing space and time in film understanding. *Applied Cognitive Psychology*, *15*(5), 533–545.
- Magliano, J. P., Taylor, H. A., & Kim, H. J. (2005). When goals collide: Monitoring the goals of multiple characters. *Memory & Cognition*, *33*(8), 1357–1367.

- Magliano, J. P., Zwaan, R. A., & Graesser, A. C. (1999). The role of situational continuity in narrative understanding. In H. van Oostendorp & S. R. Goldman (Eds.), *The construction of mental representations during reading* (pp. 219–245). Mahwah, NJ: Lawrence Erlbaum Associates.
- Mar, R. A., Kelley, W. M., Heatherton, T. F., & Macrae, C. (2007). Detecting agency from the biological motion of veridical vs. animated agents. *Social Cognitive and Affective Neuroscience*, 2, 199–205.
- Martin, A., & Weisberg, J. (2003). Neural foundations for understanding social and mechanical concepts. *Cognitive Neuropsychology*, 20, 575–587.
- Maunsell, J., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, 29, 317–322.
- McAvoy, M., Ollinger, J. M., & Buckner, R. L. (2001). Cluster size thresholds for assessment of significant activation in fMRI. *NeuroImage*, 13, S198.
- Murch, W. (2001). *In the blink of an eye: A perspective on film editing* (2nd ed.). Los Angeles: Silman-James Press.
- Newton, D. (1973). Attribution and the unit of perception of ongoing behavior. *Journal of Personality and Social Psychology*, 28, 28–38.
- Newton, D., & Engquist, G. (1976). The perceptual organization of ongoing behavior. *Journal of Experimental Social Psychology*, 12, 436–450.
- Ohtsuka, K., & Brewer, W. F. (1992). Discourse organization in the comprehension of temporal order in narrative events. *Discourse Processes*, 15, 317–336.
- Picard, N., & Strick, P. L. (1996). Motor areas of the medial wall: A review of their location and functional activation. *Cerebral Cortex*, 6, 342–353.
- Reynolds, J., Pasternak, T., & Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron*, 26, 703–714.
- Schwan, S., Garsoffky, B., & Hesse, F. W. (2000). Do film cuts facilitate the perceptual and cognitive organization of activity sequences? *Memory & Cognition*, 28, 214–223.
- Schwan, S., & Ildirar, S. (2010). Watching film for the first time: How adult viewers interpret perceptual discontinuities in film. *Psychological Science*, 21, 1–7.
- Shelton, A. L., & Gabrieli, J. D. E. (2002). Neural correlates of encoding space from route and survey perspectives. *Journal of Neuroscience*, 22, 2711–2717.
- Smith, T. J., & Henderson, J. M. (2008). Edit blindness: The relationship between attention and global change in dynamic scenes. *Journal of Eye Movement Research*, 2, 1–17.
- Speer, N. K., Reynolds, J. R., Swallow, K. M., & Zacks, J. M. (2009). Reading stories activates neural representations of perceptual and motor experiences. *Psychological Science*, 20, 989–999.
- Speer, N. K., Reynolds, J. R., & Zacks, J. M. (2007). Human brain activity time-locked to narrative event boundaries. *Psychological Science*, 18(5), 449–455.
- Speer, N. K., Swallow, K. M., & Zacks, J. M. (2003). Activation of human motion processing areas during event perception. *Cognitive, Affective & Behavioral Neuroscience*, 3(4), 335–345.
- Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain*. Stuttgart, Germany: Thieme.
- Tan, E. S. (1996). *Emotion and the structure of narrative film*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Trabasso, T., & Suh, S. (1993). Understanding text: Achieving explanatory coherence through online inferences and mental operations in working memory. *Discourse Processes*, 16, 3–34.
- Trabasso, T., van den Broek, P., & Suh, S. (1989). Logical necessity and transitivity of causal relations in the representation of stories. *Discourse Processes*, 12, 1–25.
- Troje, N., Westhoff, C., & Lavrov, M. (2005). Person identification from biological motion: Effects of structural and kinematic cues. *Perception and Psychophysics*, 67, 667–675.
- Whitney, C., Huber, W., Klann, J., Weis, S., Krach, S., & Kircher, T. (2009). Neural correlates of narrative shifts during auditory story comprehension. *NeuroImage*, 47, 360–366.
- Wurtzel, A. (1983). *Television production*. New York: McGraw-Hill.

- Zacks, J. M., Braver, T. S., Sheridan, M. A., Donaldson, D. I., Snyder, A. Z., Ollinger, J.M., Buckner, R. L., & Raichle, M. E. (2001). Human brain activity time-locked to perceptual event boundaries. *Nature Neuroscience*, 4, 651–655.
- Zacks, J. M., Speer, N. K., & Reynolds, J. R. (2009). Segmentation in reading and film comprehension. *Journal of Experimental Psychology: General*, 138, 307–327.
- Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: A mind/brain perspective. *Psychological Bulletin*, 133(2), 273–293.
- Zacks, J. M., Speer, N. K., Swallow, K. M., & Maley, C. J. (2010). The brain's cutting-room floor: Segmentation of narrative cinema. *Frontiers in Human Neuroscience*, 4(168), 1–15.
- Zacks, J. M., & Swallow, K. M. (2007). Event segmentation. *Current Directions in Psychological Science*, 16(5), 80–84.
- Zacks, J. M., Swallow, K. M., Vettel, J. M., & McAvoy, M. P. (2006). Visual movement and the neural correlates of event perception. *Brain Research*, 1076, 150–163.
- Zwaan, R. A., Langston, M. C., & Graesser, A. C. (1995). The construction of situation models in narrative comprehension: An event-indexing model. *Psychological Science*, 6(5), 292–297.
- Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2), 162–185.