



# Human brain activity time-locked to perceptual event boundaries

Jeffrey M. Zacks<sup>1</sup>, Todd S. Braver<sup>1</sup>, Margaret A. Sheridan<sup>1</sup>, David I. Donaldson<sup>1</sup>, Abraham Z. Snyder<sup>2</sup>, John M. Ollinger<sup>2</sup>, Randy L. Buckner<sup>1,2,3</sup> and Marcus E. Raichle<sup>2,4</sup>

<sup>1</sup> Department of Psychology, Washington University, One Brookings Drive, Saint Louis, Missouri 63130, USA

<sup>2</sup> Department of Radiology, Washington University Medical School, One Brookings Drive, Saint Louis, Missouri 63130, USA

<sup>3</sup> Department of Anatomy and Neurobiology, Washington University Medical School, One Brookings Drive, Saint Louis, Missouri 63130, USA

<sup>4</sup> Department of Neurology, Washington University Medical School, One Brookings Drive, Saint Louis, Missouri 63130, USA

Temporal structure has a major role in human understanding of everyday events. Observers are able to segment ongoing activity into temporal parts and sub-parts that are reliable, meaningful and correlated with ecologically relevant features of the action. Here we present evidence that a network of brain regions is tuned to perceptually salient event boundaries, both during intentional event segmentation and during naive passive viewing of events. Activity within this network may provide a basis for parsing the temporally evolving environment into meaningful units.

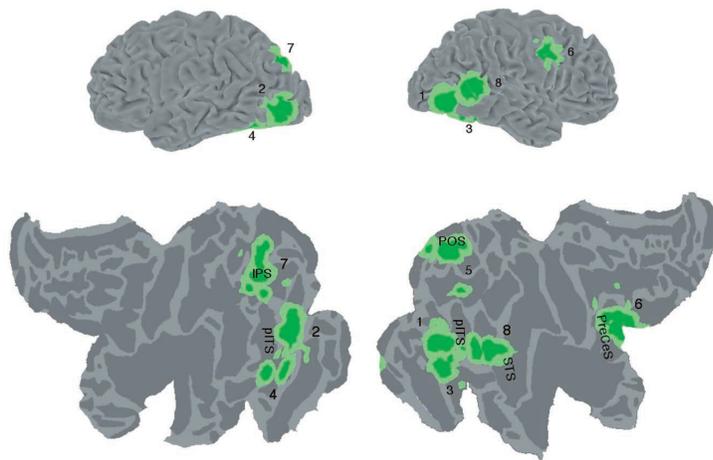
When one person observes another making a sandwich or changing a flat tire, knowledge about the parts of the activity can allow the observer to fill in missing information, predict what will happen next and plan an appropriate response. However, the information available to human perception is continuous, dynamic and unsegmented. How does the human perceptual system extract these parts, their beginnings and endings, and the relationships among them from the flux of sensory data?

We know that knowledge of event structure is pervasive and becomes increasingly elaborated throughout life: infants can subdivide continuous activity into discrete events<sup>1</sup>, children know the parts of common events such as going to school or to a birthday party<sup>2</sup>, and adults use knowledge of event structure in reading, remembering and planning<sup>3</sup>. Behavioral experiments indicate that observers can segment ongoing activity into temporal parts when asked to do so, that the resulting segmentation is reliable and systematically related to objective features of the stimulus<sup>4</sup>, and that it is predictive of later memory<sup>5,6</sup>. Event segmentation exhibits a hierarchical structure, with segments at a coarse temporal grain corresponding to groups of fine-grain segments<sup>7</sup>.

In short, human observers know about the parts of everyday activities, use this information extensively in cognition and can intentionally segment ongoing activity into parts when asked to do so. But the fact that they are able to segment activity by no means implies that they do so as a normal, spontaneous component of neural information processing. Moreover, the fact that observers are able to segment events when asked reveals nothing about how they do so. Is the human perceptual system sensitive to temporal segmentation during normal perception? If so, how does the brain accomplish this segmentation? These are two basic questions about perception, analogous to the questions of whether and how people ordinarily perceive the spatial parts of objects. To address these two questions, we measured local brain activity with functional magnetic resonance imaging (fMRI) while participants observed video depictions of everyday activities.

The method we used was constrained by the fact that everyday events make for complex, dynamic stimuli. It was critical to leave intact as much as possible of the subjective experience of viewing continuous activity, unbroken by experimental trials or other externally imposed boundaries. However, we also wanted to isolate from the complex cascade of neural processing that goes on during such an experience those components that are specifically associated with the process of identifying event boundaries. Because event segmentation is subject to individual differences and is hard to characterize by normative criteria, it was desirable to mark event boundaries based on each participant's own judgments. To address these challenges, we adopted the following experimental approach. First, brain images were continuously acquired with fMRI while participants passively viewed movies without any specific task requirement or knowledge of the purpose of the study. Second, after passive viewing, each participant segmented the same stimuli, by tapping a button to mark segment boundaries during viewing, at a coarse and fine temporal grain (that is, segmented into large and small units, respectively). Brain activity was measured during active segmentation as well as during passive viewing. Third, each observer's own segment boundaries were applied to the corresponding moments in time from the passive viewing fMRI data, using event-related analyses to identify transient changes in brain activity time-locked to those segment boundaries (see Methods).

If segmentation is an ongoing component of normal perceptual processing, one should expect to see transient changes in neural activity correlated with perceptual event boundaries. We hypothesized that we would observe transient changes both during active segmentation and during passive viewing of events. Activity during active segmentation would be weak evidence for ongoing processing of perceptual event boundaries, but activity during passive viewing would strongly suggest that the brain tracks temporal event structure during normal perception. Based on the hypothesis that observers encode events hierarchically, grouping small event parts into larger events, we predicted that we



**Fig. 1.** Clusters of activity projected onto the Visible Man cortex<sup>10,26</sup>. Darker green indicates a 10-mm radius within the brain volume around each cluster peak; lighter green indicates a 15-mm radius. Top, lateral views of each hemisphere; bottom, flattened views. Numeric labels refer to the cluster numbers in Table 1. In the flattened images, prominent sulci near clusters are labeled (PreCeS, precentral sulcus; STS, superior temporal sulcus; POS, parieto-occipital sulcus; pITS, posterior inferior temporal sulcus; IPS, intraparietal sulcus). In the upper right image, cluster 5 is fully occluded by more lateral tissue. (For brain atlas and associated software, see <http://stp.wustl.edu>.)

would observe differences between the response to fine and coarse segment boundaries during passive viewing.

Both hypotheses were supported. The data revealed transient changes in neural activity at event boundaries, during both passive viewing and active segmentation. Throughout the activated network, responses were greater to coarse unit boundaries than to fine unit boundaries.

#### RESULTS

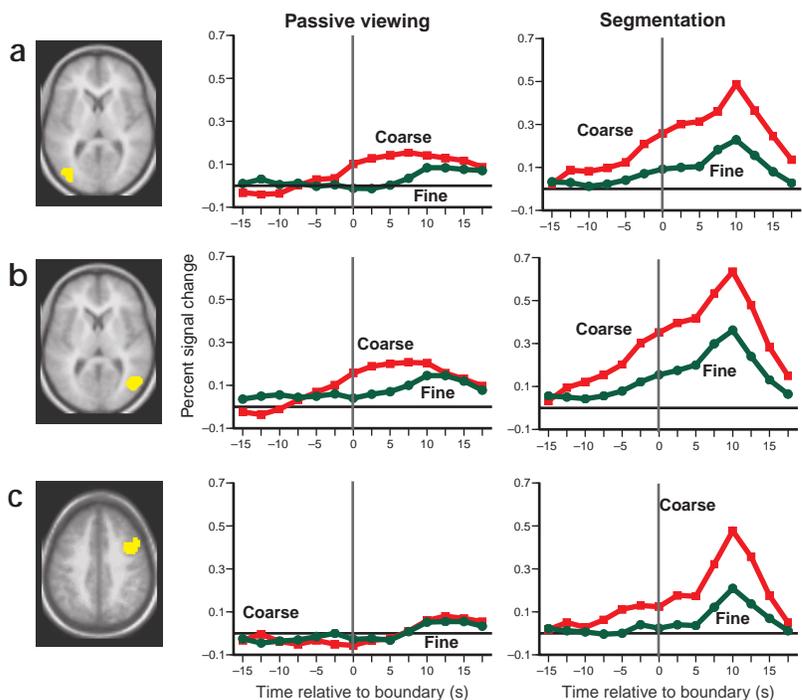
During passive viewing of everyday activities, transient changes occurred in a network of brain regions that included a larger bilateral region of contiguous activity in posterior cortex and a smaller region in right frontal cortex (Fig. 1). In posterior cortex, activity was most prominent at the occipital/temporal junction, Brodmann's areas 19/37 bilaterally (Fig. 2a and b). The location of this peak, particularly in the right hemisphere, corresponds closely to the area identified as the human MT (or V5) complex<sup>8–10</sup>. This area is involved in the processing of motion, and has been implicated in the understanding of biological motion and human action<sup>11,12</sup>. The right frontal activity was located in the precentral sulcus, Brodmann's area 6 (Fig. 2c). This region has been proposed as the

locus of the human analog of the frontal eye field (FEF) of the monkey, and is active during shifts of spatial attention as well as eye movements<sup>13</sup>. (Within this putative FEF region, a distinction has been made between an anterior subregion associated with shifting attention, and a posterior subregion associated with eye movements<sup>14</sup>. The present activity corresponds more closely to the anterior subregion.) The overall spatial distribution of activity also resembles that observed for processing of faces during memory encoding and retrieval<sup>15</sup>.

Observers had no knowledge of the segmentation task during passive viewing. Thus, this brain activity reflects spontaneous neural processing that has the same temporal dynamics as intentional event segmentation.

Activity within these regions was explored by averaging over voxels within each region and conducting analyses of variance (ANOVAs) on these average time courses (see Methods). All areas that showed transient changes in activity during passive viewing also showed reliable changes during the active segmentation runs (min  $F_{13,195} = 12.41$ ,  $p < 0.001$ ). In general, the time course of activity during segmentation (Fig. 2, right) was characterized by a larger late increase in activity than that observed during passive viewing (Fig. 2, left). During both passive viewing and active segmentation, transient changes gen-

**Fig. 2.** Time courses of focal brain activity in a subset of activated locations. Data from the strongest left posterior focus (cluster 2, Talairach coordinates  $-37, -81, 3$ ) (a), strongest right posterior focus (cluster 1, Talairach coordinates  $46, -69, 3$ ) (b) and the right frontal focus (cluster 6, Talairach coordinates  $40, 9, 39$ ) (c). (The homologous left-hemisphere precentral location showed no evidence of time-locked activity, even with a more lenient criterion of  $Z > 3.0$ .) In each panel, the left image shows the extent of the cluster, superimposed on an averaged anatomical image for the 16 participants. The graphs to the right show the time course of activity of that location in the passive viewing and active segmentation conditions, respectively. Green lines plot activity correlated with fine unit boundaries; red lines plot activity correlated with coarse unit boundaries. Vertical lines indicate the frame at which a segment boundary was placed during the intentional segmentation scans, on which the time courses were aligned for estimation using the general linear model (see Methods).



**Table 1. Cortical areas showing transient changes in activity during passive viewing of perceptual event boundaries.**

Cluster	Hemi-sphere	x	y	z	Mean z	BA	Gyrus/sulcus	Passive viewing, coarse – fine difference	Intentional segmentation, coarse – fine difference
1	R	46	-69	3	10.5	19	posterior inferior temporal sulcus	*	*
2	L	-37	-81	3	8.39	19	posterior inferior temporal sulcus	*	*
3	R	40	-54	-15	7.23	37	fusiform gyrus		*
4	L	-37	-57	-15	7.1	37	fusiform gyrus	*	*
5	R	19	-66	24	6.8	31	precuneus		*
6	R	40	9	39	6.78	6	precentral sulcus		*
7	L	-19	-78	36	6.36	19	cuneus		*
8	R	58	-48	12	5.88	39	superior temporal sulcus		*

Each row describes a cluster of voxels that, during passive viewing, showed a reliable main effect of time point in a voxel-wise ANOVA (see Methods). For each cluster, the peak location is given in Talairach coordinates<sup>22</sup>, accompanied by the mean Z statistic and location in terms of Brodmann's areas and sulcal/gyral locus. The final two columns denote whether the difference between time courses for fine and coarse segment boundaries were statistically reliable for the passive viewing and intentional segmentation conditions, respectively. (This was calculated as a test of the time-by-grain interaction, critical  $F_{13,195} = 2.22$ , which corresponds to a statistical threshold of  $p < 0.01$ , uncorrected.) The cluster-identification algorithm finds multiple peaks within large contiguous regions of activated voxels (for example, clusters 1 and 3, 2 and 4). Such large contiguous groups of voxels may reflect true functional units, but can also arise from the spatial blurring introduced by pre-processing and averaging across participants.

erally were greater for coarse unit segment boundaries (red lines) than for fine unit boundaries (green lines), though this difference was not always statistically reliable in the passive viewing condition (Table 1).

Visual inspection of the time courses suggested that activity in some areas began before the event boundary. To quantify this pattern, we conducted analyses including only the six scan frames immediately before (but not including) the event boundary. The most strongly activated posterior clusters and the sole frontal cluster were examined. ANOVAs were conducted on the data from the first six time points of each 'trial.' Separate analyses were done for the coarse-grain and fine-grain response in each cluster during active segmentation and passive viewing. In the active segmentation runs, both posterior locations increased reliably during coarse segmentation (left,  $F_{5,75} = 5.3$ ,  $p < 0.001$ ; right,  $F_{5,75} = 12.2$ ,  $p < 0.001$ ) and fine segmentation (left,  $F_{5,75} = 3.94$ ,  $p = 0.003$ ; right,  $F_{5,75} = 4.98$ ,  $p = 0.001$ ). For the passive viewing runs, the same trend was observed during coarse segmentation, and was reliable for the right hemisphere (left,  $F_{5,75} = 2.00$ ,  $p = 0.088$ , NS; right,  $F_{5,75} = 4.63$ ,  $p = 0.001$ ). For fine-grain boundaries in the passive viewing condition, there was no evidence of pre-boundary activity in the posterior regions (left,  $F_{5,75} = 0.487$ ,  $p = 0.785$ , NS; right,  $F_{5,75} = 0.332$ ,  $p = 0.892$ , NS). The frontal region had a reliable change only in the coarse-grain active segmentation runs ( $F_{5,75} = 2.729$ ,  $p = 0.026$ ; every other  $F \leq 1.34$ ).

Analyses of the behavioral data confirmed that, during the active segmentation runs, participants modulated the grain of their segmentation in accord with the instructions. Fine units had a mean length across participants of 11.3 s (range of participant means, 4.26 s to 23.2 s). Coarse units had a mean length across participants of 31.2 s (range of participant means, 18.3 to 48.0 s). The ratio of coarse unit length to fine unit length had a mean across participants of 3.09 (range, 1.21 to 5.36).

The behavioral data also indicated that participants' coarse unit boundaries tended to align with a subset of their fine unit boundaries. Alignment was analyzed using methods described previously<sup>7</sup>. Coarse unit boundaries were, on average, 1.54 s from their nearest fine unit boundary, 4.20 s closer than would be predicted if they were independent,  $t_{63} = 9.04$ ,  $p < 0.001$ .

(Similar results were obtained from a discrete-time method.) These results are consistent with the view that participants encoded the activity in terms of hierarchical relationships between parts and sub-parts.

#### DISCUSSION

These data can be summarized in terms of four points regarding the role of the neocortex in human event perception. First, a network of cortical areas showed task-independent transient changes in activity correlated with perceptual segmentation. The fact that these changes occurred during passive viewing indicates that they are a concomitant component of normal event perception. Second, the response throughout this network was stronger for coarse-grain event boundaries than for fine-grain boundaries, indicating that its activity was modulated by the hierarchical structure of the activity. Third, areas in close proximity to the human MT complex and FEF may be heavily involved in perceptual event segmentation. Finally, response in at least part of this network began substantially before the event segment boundary, suggesting a buildup of information that anticipated the observer's explicit report of an event boundary.

The current study also makes a methodological contribution: it demonstrates the possibility of conducting functional neuroimaging experiments in which participants define the 'trial' structure of the experiment based on a continuous behavior recorded after the imaging data are collected. Self-paced designs in general offer some methodological advantages (L. Maccotta, J.M. Zacks & R.L. Buckner, *Ann. Meeting Cogn. Neurosci. Soc.*, 54, 2001) and have been applied successfully in studies of binocular rivalry<sup>16-18</sup>. However, in this case, the definition of trials was based on behavior that took place after the functional imaging data to be analyzed, allowing the measurement of brain activity uncorrupted by an interfering task. The present results indicate that such designs allow for the measurement and quantitative analysis of the hemodynamic response during an ecologically valid perceptual experience, without previous assumptions about the shape or timing of that response.

The correspondence we observed among three of the group functional activation foci and regions corresponding to the



human MT complex and FEF is provocative, but it should not be taken as definitive. The functional localization of both areas varies across individuals<sup>19,20</sup>. Establishing the precise role of MT complex and FEF in event segmentation will require functional localization of these regions in individual observers who also complete the segmentation protocol.

The finding that this network responded more strongly to coarse segment boundaries suggests that it was modulated by the top-down influence of internal event representations. Such representations are proposed to be involved in integrating information about the physical structure of activity with information about goals, plans and causes<sup>3,21</sup>. One possibility is that physical correlates of object and actor motion lead to the detection of fine unit boundaries, whereas cognitive representations of events determine how fine units group together into large units. Based on this view, the pre-boundary activity observed here can be taken as evidence that event segmentation reflects active monitoring rather than simple reactive processing. Similarly, the increased responsiveness of the network to coarse boundaries may reflect feedback in the system from activation of internal knowledge structures that represent the relationships between large and small event segments, including part-subpart and goal-subgoal relationships. That is, information about goals, plans and causes may be integrated with physical cues to direct ongoing perceptual processing. Consistent with this hypothesis, the behavioral data replicated previous findings that large event boundaries align with a subset of the small event boundaries<sup>7</sup>, indicating that participants encoded the segmentation structure in terms of hierarchical part relationships.

However, the data are consistent with another possibility. The observed transient changes in cortical activity may reflect bottom-up processing of physical features of the activity that correlate with natural event segmentation. Based on this view, the possible involvement of the MT complex can be taken as evidence that features of the motion signal are important for conscious event segmentation, and the activity of the right frontal region can be interpreted as due to shifts of gaze or visual attention due to object or actor motion. This alternative hypothesis can be tested by measuring or controlling motion information in the stimuli, and by measuring or controlling eye movements during viewing.

Whether activity in this network is strongly influenced by top-down modulation from cognitive representations of events or is primarily driven by low-level perceptual processing, the present results strongly indicate that a network of cortical regions tracks observers' temporal segmentation of ongoing action. Moreover, the neural signals associated with event segmentation seem to be a concomitant component of ongoing perception, rather than being tied to a particular task.

## METHODS

**Participants.** Sixteen participants (18–50 years old, 10 female) participated in the experiment. (Three additional participants were discarded, one due to equipment malfunction, one due to behavioral noncompliance and one due to MRI artifact.) All participants gave written informed consent for the experimental procedures, which were approved by the Washington University Medical School Human Studies Committee.

**Task and stimuli.** Participants watched four movies of everyday activities during fMRI scanning. The activities were 'making a bed' (316 s), 'doing the dishes' (258 s), 'fertilizing a houseplant' (120 s) and 'ironing a shirt' (298 s). Each activity was performed by a single actor, and filmed from a fixed location at approximate standing head height. Cinematic effects such as cuts, pans and zooms were not used, to avoid injecting

external segmentation cues<sup>7</sup>. Each stimulus was presented three times. During the first presentation, participants were instructed simply to watch the movie and try to learn as much about it as possible. During the second and third viewings, participants were instructed to segment the activity into natural and meaningful units by pressing a handheld button<sup>5,7</sup>. During one viewing, participants were asked to press the buttons to mark off the behavior of the person they would be seeing into the "largest units that seem natural and meaningful" (coarse units). On the other viewing, they were instructed to mark off the smallest natural and meaningful units (fine units). Order of temporal grain (coarse versus fine) was counterbalanced across participants.

**Imaging.** Scanning was done on a 1.5 T Siemens Vision MRI scanner (Siemens, Erlangen, Germany). Structural images were acquired using a sagittal MP-RAGE T1-weighted sequence with 2 mm (isotropic) resolution. Functional images were acquired using a T2\*-weighted asymmetric spin-echo echo-planar sequence (T.E. Conturo *et al. Soc. Neurosci. Abstr.* 22, 7, 1996). Sixteen slices, 8 mm in thickness, with an in-plane resolution of  $3.75 \times 3.75$  mm, were acquired every 2.5 s. Functional runs included four frames to allow the T2\* signal to stabilize, followed by a variable number of frames equal to the length of the movie (see above). (For the 'making a bed' stimulus, scanning was performed only during the first 310 s because of a limit on scanner acquisition time per run.) Before analysis, the functional data were pre-processed to remove participant motion and imaging artifacts. They were then warped to a standard stereotaxic space<sup>22</sup> by aligning each participant's MP-RAGE to an MP-RAGE atlas target image based on 12 normal adults, using a 12-parameter affine transformation. Timing offsets between slices were corrected using sinc interpolation. Slice intensity differences from contiguous interleaved slice acquisition were removed using suitably chosen scale factors. The data were spatially smoothed with a Gaussian kernel (full width at half maximum, 6.0 mm).

**Voxel-wise analysis.** The pre-processed and stereotaxically normalized fMRI data were subjected to a series of event-related analyses based on modeling typical responses for each individual using the general linear model and making generalizations across individuals using random effect ANOVAs<sup>23</sup>. For each observer, variations in brain activity were assessed based on the location of segment boundaries in their second and third viewings. A 'trial' was defined as a 35-s periresponse interval centered at the location within the movie of a fine or coarse segment boundary (that is, 17.5 s before and after the segment boundary). (This window size was chosen to permit measurement of the bulk of the hemodynamic response for changes occurring up to several seconds after a segment boundary, allowing for the lag of the hemodynamic response, and an interval of equal length was adopted for detecting activity before a boundary.) For each observer, the BOLD response during fine and coarse segment boundaries during passive viewing was estimated using the general linear model<sup>24</sup>, yielding two 35-s time courses per voxel per viewing per observer. Each time course can be thought of as an 'average' response for one observer for one segment boundary type (fine or coarse). However, the time courses are model fits that take into account overlaps between nearby segment boundaries, rather than simple arithmetic means. (The model also included covariates that removed any linear trend and low-frequency noise; cutoff, 0.009 Hz.) The data were analyzed across subjects with a two-factor ANOVA. One factor was time, with 14 levels corresponding to each estimated time point. The other factor was grain, with levels of 'coarse' and 'fine.' The main effect of time identified voxels with time courses that deviated reliably from the baseline estimated by the general linear model (that is, the rest state), whereas the task-by-time interaction identified regions with time courses that differed across the two grains. Brain regions were identified that showed a reliable transient change in activity correlated with the segment boundaries (that is, a main effect of time) by converting the resulting *F* statistics to *Z* statistics and identifying clusters of at least five contiguous voxels with *Z* statistics that exceeded the threshold, subject to the constraint that clusters could be no closer than 20 mm. A threshold of  $Z = 5.69$  was chosen; this corresponds to a  $p = 0.05$  probability of falsely identifying one or more activated clusters, as adjusted by the Bonferroni correction for multiple comparisons and the Box correction for correlated predictor variables<sup>25</sup>.



We also conducted a set of control analyses in which functional data from other participants, scanned during rest, were analyzed as if those participants had marked event boundaries at the same times as the participants in the current experiment (that is, the behavioral data from the current experiment were applied to functional data from resting participants). In 10 replications of the full experimental protocol with different random assignments of functional data to behavioral segmentation data, no false positive clusters were observed.

**Region-wise analysis.** Within the activated voxels, peaks of activity were identified by an automatic algorithm that locates local maxima subject to the constraint that no two peaks can be within 20 mm, and that each peak must be within a group of 5 or more contiguous activated voxels. Clusters were identified as all voxels within 12 mm of each peak with Z scores above 3.0. Time courses were averaged across the voxels in each activated cluster for each participant and segment boundary type. Time courses were calculated for the passive viewing data from which the threshold had been defined, and also for the data from the intentional segmentation runs. For the segmentation runs, only functional data from the coarse-grain segmentation runs were used to calculate time courses for coarse-grain segment boundaries, and only functional data from the fine-grain segmentation runs were used to calculate time courses for fine-grain segment boundaries. These average time courses were submitted to ANOVAs of the same form as the original voxel-wise ANOVA. Passive viewing and intentional segmentation were analyzed separately. The main effect of time in the passive viewing data was not analyzed, because this effect was used to define the clusters. The main effect of grain and the interaction between grain and time scale, and the main effect of time in the intentional segmentation data, are fully independent of this effect.

**Controlling for sampling and attentional set.** We conducted two control analyses to rule out potential artifactual explanations of the differences between the time courses for coarse and fine segment boundaries. First, there were differences in the mean number of observations that went into coarse unit and fine unit time course estimates. We controlled for this by sampling the fine unit boundaries: for each movie for each participant, a number of fine unit boundaries was randomly chosen to be modeled, equal to the number of coarse-unit boundaries that participant produced. Results were similar to those of the original analyses. Second, there may have been effects due to differences in attentional set during the coarse unit and fine unit active segmentation runs. (Of course, for the passive viewing analyses, attentional set is not a potential confound.) We controlled for these by performing an analysis based only on the fine-grain segmentation data. Fine unit boundaries were grouped into those that occurred near a coarse unit boundary and those that did not. Comparing these two classes of fine unit boundaries gave rise to the same patterns as did comparing coarse and fine boundaries. In sum, the control analyses argue against attributing the observed differences between coarse and fine segment boundaries to artifacts.

**Analysis of hierarchical structure in segmentation.** The alignment of each observer's coarse and fine segment boundaries was analyzed for the presence of hierarchical structure as described previously<sup>7</sup>. The analysis begins with the continuous locations of the fine segment boundaries, and calculates the mean distance from all points during the movie to the nearest segment boundary, which is given by the following formula.

$$\frac{f_i^2}{2} + \frac{\sum_{l=1}^{i-1} \left[ \frac{f_{i+1} - f_l}{2} \right]^2}{f_{\text{Fine}}}$$

Here, *Fine* is the number of fine unit boundaries and  $f_i$  is the location of the *i*-th fine unit boundary. If coarse segmentation is statistically independent of fine segmentation, the mean distance from coarse segment boundaries will tend toward the mean distance from all points. To the

extent that the observed mean distance is smaller, coarse unit boundaries are aligned with fine unit boundaries. (Similar results were obtained with a discrete-time method, also described previously<sup>7</sup>.)

#### ACKNOWLEDGEMENTS

This research was supported in part by grants from the Mallinckrodt Institute of Radiology and the McDonnell Center for Higher Brain Function. The authors thank E. Akbudak, T.E. Conturo and D.C. Van Essen for their assistance, and B. Tversky for her comments.

RECEIVED 20 MARCH; ACCEPTED 23 APRIL 2001

- Wynn, K. Infants' individuation and enumeration of actions. *Psychol. Sci.* 7, 164–169 (1996).
- Nelson, K. & Gruendel, J. in *Event Knowledge: Structure and Function in Development* (ed. Nelson, K.) 21–46 (Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1986).
- Zacks, J. M. & Tversky, B. Event structure in perception and conception. *Psychol. Bull.* 127, 3–21 (2001).
- Newtonson, D., Engquist, G. & Bois, J. The objective basis of behavior units. *J. Pers. Soc. Psychol.* 35, 847–862 (1977).
- Newtonson, D. & Engquist, G. The perceptual organization of ongoing behavior. *J. Exp. Soc. Psychol.* 12, 436–450 (1976).
- Avrahami, J. & Kareev, Y. The emergence of events. *Cognition* 53, 239–261 (1994).
- Zacks, J. M., Tversky, B. & Iyer, G. Perceiving, remembering, and communicating structure in events. *J. Exp. Psychol. Gen.* 130, 29–58 (2001).
- Watson, J. D. *et al.* Area V5 of the human brain: evidence from a combined study using positron emission tomography and magnetic resonance imaging. *Cereb. Cortex* 3, 79–94 (1993).
- Tootell, R. B. & Taylor, J. B. Anatomical evidence for MT and additional cortical visual areas in humans. *Cereb. Cortex* 5, 39–55 (1995).
- Van Essen, D. C. & Drury, H. A. Structural and functional analyses of human cerebral cortex using a surface-based atlas. *J. Neurosci.* 17, 7079–7102 (1997).
- Bonda, E., Petrides, M., Ostry, D. & Evans, A. Specific involvement of human parietal systems and the amygdala in the perception of biological motion. *J. Neurosci.* 16, 3737–3744 (1996).
- Decety, J. & Grezes, J. Neural mechanisms subserving the perception of human actions. *Trends Cogn. Sci.* 3, 172–178 (1999).
- Corbetta, M. *et al.* A common network of functional areas for attention and eye movements. *Neuron* 21, 761–773 (1998).
- Corbetta, M. Frontoparietal cortical networks for directing attention and the eye to visual locations: identical, independent, or overlapping neural systems? *Proc. Natl. Acad. Sci. USA* 95, 831–838 (1998).
- McDermott, K. B., Buckner, R. L., Petersen, S. E., Kelley, W. M. & Sanders, A. L. Set- and code-specific activation in frontal cortex: an fMRI study of encoding and retrieval of faces and words. *J. Cogn. Neurosci.* 11, 631–640 (1999).
- Lumer, E. D., Friston, K. J. & Rees, G. Neural correlates of perceptual rivalry in the human brain. *Science* 280, 1930–1934 (1998).
- Lumer, E. D. & Rees, G. Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proc. Natl. Acad. Sci. USA* 96, 1669–1673 (1999).
- Tong, F., Nakayama, K., Vaughan, J. T. & Kanwisher, N. Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron* 21, 753–759 (1998).
- Dumoulin, S. O. *et al.* A new anatomical landmark for reliable identification of human area V5/MT: a quantitative analysis of sulcal patterning. *Cereb. Cortex* 10, 454–463 (2000).
- Paus, T. Location and function of the human frontal eye-field: a selective review. *Neuropsychologia* 34, 475–483 (1996).
- Grafman, J. Similarities and distinctions among current models of prefrontal cortical functions. *Ann. NY Acad. Sci.* 769, 337–368 (1995).
- Talairach, J. & Tournoux, P. *Co-planar Stereotaxic Atlas of the Human Brain* (Thieme, Stuttgart, 1988).
- Miezin, F. M., Maccotta, L., Ollinger, J. M., Petersen, S. E. & Buckner, R. L. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *Neuroimage* 11, 735–759 (2000).
- Friston, K. A., Holmes, A., Worsley, K. & Poline, J. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 2, 189–210 (1995).
- Ollinger, J. M. & McAvoy, M. P. in *Functional Mapping of the Human Brain* (eds. Mesulam, M.-M. & Bandettini, P.) (Academic, San Antonio, Texas, 2000).
- Van Essen, D. C., Drury, H. A., Joshi, S. & Miller, M. I. Functional and structural mapping of human cerebral cortex: solutions are in the surfaces. *Proc. Natl. Acad. Sci. USA* 95, 788–795 (1998).