**Measuring Events: An Investigation into the Stability of Event Segmentation Patterns
Across Groups**

Karen Sasmita¹ and Khen. M. Swallow^{1*}

¹Department of Psychology, Cornell University

*Corresponding author: Khen Marie Swallow,
The Department of Psychology,
Cornell University, 211 Uris Hall, Ithaca, NY 14850, USA.
Email: kms424@cornell.edu

Abstract

People spontaneously divide everyday experience into smaller units (*event segmentation*). To measure event segmentation, studies typically ask participants to explicitly mark the boundary between events as they watch a movie (*segmentation task*). Their data may then be used to infer how others are likely to segment the same movie. However, significant variability in performance across individuals could undermine the ability to generalize across groups, especially as more research moves online. To address this concern, we used several widely employed and novel measures to quantify segmentation agreement across different sized groups ($n=2-32$) using data collected on different platforms and movie types (in-lab & commercial film vs. online & everyday activities). All agreement measures captured non-random and video-specific boundary identification with sample sizes as small as 2, though with notable between sample variability. As sample size increased, agreement values improved and eventually stabilized. Stabilization occurred at smaller sample sizes when measures reflected (1) agreement between two groups versus agreement between an individual and group, (2) boundary identification between small (*fine-grained*) rather than large (*coarse-grained*) events, and (3) segmentation of everyday activities online versus commercial film in-lab. These analyses inform the tailoring of sample sizes based on the comparison of interest, materials, and data collection platform. In addition to demonstrating the reliability of online as well as in-lab segmentation performance at moderate sample sizes, this study supports the use of these data as a means of inferring when everyday activities and commercial films are likely to be segmented.

Introduction

As part of everyday perception, continuous experience is segmented into distinct units (*events*). This cognitive process, event segmentation, involves identifying *event boundaries* or the moments when one event transitions to the next (Zacks, et al., 2001; Zacks et al., 2007). Event segmentation is typically studied by asking participants to watch a movie and simultaneously mark the boundaries between events (Newton, 1973; Newton & Engquist, 1976). In this paper, we examined the efficacy of measures that quantify segmentation task performance in capturing behavior that is stable and sensitive to the structure of the stimulus.

In a typical segmentation task, participants are asked to press a button whenever they believe that one natural and meaningful unit of activity has ended and another has begun (Newton, 1973; Newton & Engquist, 1976; Zacks et al., 2001). By design, participants are given limited directions and are told to rely on their own judgment when marking event boundaries. Despite the task's ambiguity, participants can effortlessly identify event boundaries at multiple levels of granularity (Kurby & Zacks, 2011; Newton, 1973). They also organize events based on the hierarchical structure of goal-directed activities: smaller *fine-grained* units (e.g., grinding coffee beans, pouring hot water) are grouped within larger *coarse-grained* units (e.g., making coffee; Hard et al., 2006; Hard et al., 2011; Zacks et al., 2001).

Segmentation task performance captures critical features of how the brain processes ongoing events. Previous fMRI studies have investigated the relationship between evoked neural responses during task-free movie viewing and boundaries defined through a segmentation task. These boundaries were either defined by the same participants following task-free viewing of the movie (Speer et al., 2003; Zacks, et al., 2001) or by a separate group of individuals (Baldassano et al., 2017; Ben-Yakov & Henson, 2018). Results showed that the time points in the movie

corresponding to event boundaries were associated with neural activity increases in regions involved in processing changes in the ongoing situation (Zacks, et al., 2010), motion processing (Speer et al., 2003; Zacks et al., 2001; Zacks et al., 2006) and memory encoding (Ben-Yakov & Henson, 2018). Event boundaries are also associated with transitions between stable patterns of activity within the posterior-medial network (Baldassano et al., 2017). This network of regions has been implicated in the construction of complex mental representations for remembering or imagining events (Ritchey & Cooper, 2020). Overall, these findings provide strong support for correspondence between behavioral and neural markers of event segmentation.

Event segmentation has been postulated to be an adaptive mechanism that optimizes attention and memory processing in part based on situational and perceptual cues (Zacks et al., 2007). Consistent with this notion, boundary placement typically correlates with mid-level perceptual changes in action, location, object interactions, or the actor's motion trajectory (Hard et al., 2006; Magliano et al., 2001; Newton et al., 1977; Swallow et al., 2018). Perceptual changes within the stimulus may therefore drive the systematicity of event segmentation across people. Indeed, although segmentation tasks are naturally subjective, people generally form good agreement with each other and are consistent in their own judgments about when event boundaries occur during a movie (Newton, 1973; Speer et al., 2003). When individual data are examined at the group level, it is possible to identify moments in a movie that most people are likely to report an event boundary (*normative boundaries*), and other moments that people are unlikely to report an event boundary. It is therefore reasonable for studies investigating the cognitive and neural implications of event segmentation to extrapolate individual segmentation behavior based on a separate group's segmentation data (e.g., Ben-Yakov & Henson, 2018; Faber et al., 2018; Swallow et al., 2009).

However, performance on the segmentation task can vary markedly from one person to the next. This variance could come from sources that do not necessarily reflect the underlying segmentation process, such as differing interpretations of ambiguous task instructions, erroneous button presses, or accidental misses. It could also be greater for online data collection, a practice that is increasing, which offers little control over either the context in which data collection occurs or the hardware that is being used. Additionally, there may be individual and group differences that could influence how events themselves are processed and segmented, though these factors have as yet received little attention in the literature (for exceptions, see Bläsing, 2015; Kurby & Zacks, 2011; Levine et al., 2017; Newberry et al., 2021; Sargent et al., 2013; Swallow & Wang, 2020).

Errors in task performance and the potential contributions of individual differences to segmentation (e.g., Speer et al., 2003) should reduce the degree to which boundaries identified by small groups are predictive of individual task performance. Yet, studies have often used segmentation data from a wide range of sample size ($N = 1$ to 41) and participant expertise to infer when other individuals are likely to segment (e.g. Ben-Yakov & Henson, 2018; Chen et al., 2017; Faber et al., 2018; Kosie & Baldwin, 2019; Swallow et al., 2009). Although studies over the last several decades have provided a wealth of data pointing to the segmentation task's effectiveness in capturing meaningful behavior, a formal examination of its sensitivity to the commonly perceived structure of naturalistic stimuli has been limited. In this study, we investigated various ways of quantifying segmentation performance, the sensitivity of these estimates to the structure within the stimulus itself, and its stability as the sample size increases.

In prior studies, segmentation performance at the group and individual levels has been characterized using a variety of methods. At the group level, segmentation performance has been

quantified by creating a *group time series* that captures the proportion of participants in a group that identified a boundary at each moment (often intervals of 5 s or less) in the movie (e.g., Massad et al., 1979, Zacks, et al., 2001). This time series can then be correlated with that from another group (as in, e.g., Hard et al., 2006) or used to identify intervals during which most participants press a button (normative boundaries; as in, e.g., Massad et al., 1979). In other studies, group performance has been evaluated by comparing the average number of button presses between groups (Bläsing, 2015) or the overlap of normative boundaries identified by two different groups (Massad et al., 1979). Individual performance has been assessed in multiple ways as well. In some cases, researchers analyzed the overlap of boundaries identified by an individual with normative boundaries defined by a separate group (Zalla et al., 2013), who were sometimes experts in the field (e.g., Levine et al., 2017). More commonly, individual segmentation performance has been quantified by correlating the group time series with the moments that an individual pressed the button (Kurby & Zacks, 2011; Newberry & Bailey, 2019; Swallow & Wang, 2020; Zacks, et al., 2006).

In this paper, we systematically examined the efficacy of new and existing measures of segmentation performance in characterizing behavior at the group and individual levels. Measures were selected based on their ability to capture the expected reproducibility of group level boundaries as well as the degree to which an individual's performance is consistent with group performance. Specifically, these measures quantified agreement in boundary placement within a group of individuals (*peakiness*), between two groups of individuals (*peak-to-peak distance*), and between an individual and a separate group (*surprise index*, adapted from Katori et al., 2018 ; and *agreement index*, Zacks et al., 2006). We asked how each of these agreement measures perform according to three criteria: 1) the ability to distinguish segmentation behavior

from noise; 2) the ability to distinguish segmentation of one movie from another; and 3) the sample size needed for the measure to stabilize and show little improvement with additional participants.

We analyzed two independent datasets: the first consisted of segmentation data collected in the lab with commercially made movies. The second consisted of segmentation data collected online with videos depicting everyday activities. With these datasets, we were able to assess the effectiveness of conducting event segmentation tasks online relative to in the lab. While lab-based experiments offer a more controlled environment, online experiments offer benefits such as the capacity to recruit from a more diverse population and efficiency in collecting data from a large number of participants (e.g., Birnbaum, 2004). These datasets also differed in whether they used edited but readily available commercial films or lab-produced videos of everyday activities. Although everyday-activity videos can be more realistic, they often lack the structure and richness of commercial films (Cutting et al., 2011).

Materials and Methods

This paper analyzed datasets collected for two other projects that will be (Sasmita & Swallow, in prep) or have already been published (Swallow & Wang, 2020). The projects differed in the type of video stimuli used, how the data were collected, and the sampled population. The first project (commercial-lab) examined segmentation of excerpts from commercially produced movies. These data were collected in the lab with undergraduate participants. The second project (everyday-online) utilized unedited videos of actors performing everyday activities produced by the lab. Data were collected online with participants recruited through Amazon's Mechanical Turk and CloudResearch (Swallow & Wang, 2020).

Participants

All participants provided informed consent and all procedures were approved by the Cornell Institutional Review Board.

Commercial-lab

Participants were recruited from the Cornell University community. Analyses reported in this paper focused on data from 64 participants (27 males, 37 females) between the ages of 18 and 41 ($M = 21.04$, $SD = 4.27$). Participants were compensated with course credit or \$15 for their time.

Everyday-online

Participants were recruited through Amazon's Mechanical Turk using CloudResearch. Data were collected from participants in India and in the US, but analyses in this paper only utilized data from participants in the US ($N = 72$; 41 males, 31 females; 19 – 58 years old. $M = 33.79$, $SD = 8.26$). Participants were compensated \$5 or \$7.50 for their time.

Experimental design

Commercial-lab

Video stimuli were constructed using excerpts from two commercial movies: *3 Iron* (Kidd, 2004) and *Corn Island* (Ovashili, 2014). These movies were selected because they depict distinct activities (e.g., cooking in the kitchen, sawing and hammering wood), have little to no dialogue, and are set in naturalistic settings.

For each movie, excerpts were created by first splitting scenes depicting one or more activities at natural breakpoints (e.g., an editor's cut, scene/ location changes, breaks in the narrative). The excerpts were then joined together to form one continuous video (*3 Iron* = 9.82 mins; *Corn Island* = 9.43 mins). Scenes in the final video were not necessarily contiguous in the original movie, but their order was preserved. The audio tracks for both videos were removed.

The videos were divided into ten, 1 minute long *clips* (the last clip was shorter than 1 minute). For half of the participants, the 10 clips from each movie were presented without interruption and in order (*continuous*). For the other half of the participants, the clips were presented in order but with a 3 second white noise display interruption in between clips (*interrupted*). Data from a third group were collected for the original project but were not included in the following analyses.

Everyday-online

Stimuli were videos of actors performing everyday activities set in the US: doing laundry (5.41 mins), making coffee (5.73 mins), doing the dishes (6.03 mins), and making a bed (4.55 mins). The original project also used recordings of the same four activities set in India; however, data from this condition were not included in the following analyses.

Videos were recorded using a GoPro Hero 4 silver edition (1920 x 1080 pixels, 29.98 fps, narrow field of view). The activities were performed by different actors and were filmed from several feet away. Therefore, in each video, an actor can be seen performing one activity within one room, with no changes in camera angle or scene cuts. Videos were presented with the audio tracks removed.

Segmentation task

All participants segmented every video that they watched. They were instructed to press the spacebar every time they believed one natural and meaningful unit of activity had ended and another had begun. Participants were instructed to mark the smallest (*fine grain*) and largest (*coarse grain*) units of activity change. They were told to press the spacebar as many times as necessary to mark all of the units in the video.

All experiment sessions began with a practice segmentation task. For *Commercial-lab*, a short practice video was constructed by extracting a 1.5-minute clip from the movie *3 Backyards* (Mendelsohn, 2010). For *Everyday-online*, a short practice video (1.65 mins) was recorded depicting an actor sitting outside, eating snacks, looking through books, and using his phone. To increase similarity across participants and reduce overlap between fine and coarse conditions, practice was repeated until participants reached a performance criterion. The criterion was based on mean unit durations reported in earlier literature (e.g., Zacks, et al. 2001) or on pilot segmentation data with the practice video (as described in Swallow & Wang, 2020): *commercial-lab*, mean duration of coarse = 11.25-30s (3-8 button presses), fine = 2.5-6s (15-36 button presses); *everyday-online*, mean duration of coarse = 16.5-49.5s (2-6 button presses), fine = 3.3-8.25s (12-30 button presses).

Participants in *commercial-lab* performed coarse and fine segmentation on one movie before moving on to segmenting the next movie. In contrast, participants in *everyday-online* segmented all videos in one grain first before moving on to segmenting the videos again in the other grain. Video and grain orders were fully counterbalanced in *commercial-lab*. Although most participants in *everyday-online* started with coarse grain segmentation due to a programming error, segmentation patterns between participants who started with fine segmentation and those who started with coarse segmentation were similar (Swallow & Wang, 2020).

Following segmentation of each video, participants in *commercial-lab* performed a free recall task, typing everything they remember happening in the video into a document while participants in *everyday-online* completed some questionnaires. The following analyses focused only on performance on the segmentation task.

Data Processing

Our analyses focused on various methods of quantifying group and individual level segmentation agreement. Measures were selected based on their use in prior research or their potential to provide additional insight into the consistency of segmentation patterns across individuals. We compared segmentation agreement along four measures: a summary of group level agreement (*peakiness*), a comparison of segmentation performance between two groups (*peak-to-peak distance*), and two measures that evaluate individual segmentation performance by comparing it to the group (*agreement index* and *surprise index*).

Creating datasets of different sample sizes

We used a bootstrapping procedure to estimate the distribution of each agreement measure when *sub-samples* of different sizes were used (Singh & Xie, 2008). For each sample, we randomly selected with replacement n individuals from the group of interest 100 times, where n ranged from 2 to 32, incremented by 2. This process was performed separately for each agreement measure.

Calculating measures of segmentation agreement

To allow for comparison of segmentation performance across different movies, segmentation data were restricted to the length of the shortest videos (commercial-lab = 566s, everyday-online = 273s).

Group level agreement. Previous studies often quantified group segmentation performance by calculating the proportion of participants that pressed a button at fixed intervals (*bins*) resulting in a *group time series* (Massad et al., 1979; Zalla et al., 2013). In such cases, the choice of bin-size was not typically motivated by the authors. However, it could have outsized effects on the results of an analysis: whereas small bin-sizes run the risk of preserving variance

from sources that do not operate at the group level (e.g., lapses in attention and motor speed), large bin-sizes may obscure the effects of brief, but relevant sources of variation (e.g., changes in the movie content). Binning the data in this way further distorts the continuous nature of events by treating a boundary as a binary event that is either present or absent during a discrete interval of time, rather than as an entity that could be probabilistically distributed over time. Finally, even if boundaries are punctate, binning them potentially treats two button presses that are far from each other as the same and two button presses that are near each other as different, depending on whether they fall within the same bin or not.

Rather than binning the data, we characterized group segmentation performance as the observed probability density of a button press over time (*group density*). Using the density function in base R (stats package; R Core Team, 2018), the sub-sample of participants' button press times were combined and every time point was smoothed using a gaussian kernel. The kernel's bandwidth was determined using a smoothing function that accounts for variations in the data (option 'sj' in the density function; Sheather & Jones, 1991). For each grain, the computed bandwidth was then multiplied by a value that generated a group time series with visually distinctive peaks and valleys for all sample sizes ('adj' parameter: coarse = 0.1, fine = 0.05; Figure 1A). *Normative boundaries* were then defined as the time points of the highest j peaks in the group density distribution, where j is the mean number of button presses generated by participants in the group. Whenever the number of observed peaks was lower than the average number of group button presses, j was defined as the maximum number of peaks identified in the group density estimate. This method provides straightforward criteria for defining normative boundaries while avoiding the need for what may be arbitrary decisions about the unit of time at which the data should be binned.

Figure 1. A) Example of density estimate with different adjustment bandwidths for small ($n = 2$; upper panel) and large ($n = 32$; lower panel) sample sizes. In all cases, the lower bandwidth (0.01) seems to capture individual button presses rather than the group's consensus button presses. For the large sample size (lower panel), the middle and higher bandwidth values do not seem to influence the shape of the peaks and valleys of the density estimate. However, for small sample size (upper panel), distinctive peaks and valleys are formed in the density estimate using the middle bandwidth (0.1 for coarse, 0.05 for fine). The highest bandwidth value (0.2 for coarse and 0.1 for fine) reduces the difference between the peaks and valleys, and even eliminates several peaks (arrows). Therefore, we chose the middle bandwidth (0.1 for coarse and 0.05 for fine) value for our density estimation for all sample sizes. B) Examples of growth (left) and decay (right) function

fits. Small dots represent the average agreement estimate for individual bootstrap iteration. Large dots represent the average agreement estimate from all bootstrap iterations in each sample size. Functions with the lowest BIC value were selected as the best fitting curve.

Peakiness. To quantify agreement among individuals in a single group, we introduce a new measure that we refer to as *peakiness*. This measure reflects the expectation that greater amounts of within-group agreement will lead to group density functions with higher ‘peaks’ and lower ‘valleys’ (i.e., they are less flat). For every sub-sample, we calculated the moment-to-moment change in amplitude of the group density function over the duration of the movie (rugosity) using the *ru*go function (seewave package; Sueur et al., 2008). This value was then scaled by dividing the observed rugosity by a rugosity value that captured rugosity when there is no agreement in boundary placement across individuals. This was defined as rugosity when the combined participants’ button presses were uniformly distributed over time and accounts for the varying number of button presses across sub-samples. A high peakiness value, therefore, suggests high within-group agreement on boundary placement.

Peak-to-peak distance. To measure the consistency of normative boundaries identified across groups, we calculated the peak-to-peak distance. This measure is defined as the mean distance between one group’s normative boundaries to another group’s normative boundaries, with a value of 0 indicating perfect agreement.

To calculate peak-to-peak distance, we first obtained normative boundaries for one sub-sample (reference group). We repeated this process for a second sub-sample of the same size (matching group) and extracted the same number of normative boundaries as were identified for the reference group. Peak-to-peak distance was quantified as the mean distance between each peak for one group and the nearest peak for the other group.

Individual level agreement. The next two measures quantify agreement between boundaries identified by an individual and those identified by an independent group of participants. The first measure, agreement index, has been commonly used in prior research (Kurby & Zacks, 2011). We applied the second measure, the surprise index (Katori et al., 2018), to segmentation data for the first time to evaluate whether it provides a more stable index of individual-group agreement. To calculate agreement index, we transformed individual button press data into a time series indicating the presence or absence of a button press in 1s long bins (*individual timeseries*), as is standard for this measure. Calculation of surprise index was based on each individual's raw button press times.

Individual segmentation performance has also been quantified as a measure of individual accuracy in identifying normative boundaries (e.g., Zalla et al., 2013). We briefly explored this approach and present our findings in Appendix A. Importantly, quantifying individual performance using d' assumes that normative boundaries capture true signal in the stimuli (i.e., the presence of a boundary in specific bins) and that individuals generate button presses based on signals embedded in normally distributed noise. In contrast, the surprise index models the overlap between individual and group button presses probabilistically (using the Poisson distribution) without imposing assumptions about how individuals generate button presses.

Agreement index. This measure captures the similarity of segmentation patterns between an individual and an independent group of participants (Zacks, et al., 2006). We correlated every individual timeseries in a sub-sample with the group timeseries created by calculating the proportion of the remaining participants who pressed a button within each 1s bin (obs. r). The correlation was then scaled with the minimum and maximum possible correlations (min r and max r) for that individual using the following formula (Kurby & Zacks, 2011): $a_i = (\text{obs. } r - \text{min}$

$r)/(\max r - \min r)$. The agreement index therefore ranged from 0-1, with higher values indicating greater agreement.

Surprise index. This measure reflects the extent to which an individual's button presses are expected, given the normative boundaries identified by an independent group. It is adapted from work that estimates the probability of spike co-occurrence across two neurons using the Poisson distribution (Katori et al., 2018). We used the surprise index to capture the co-occurrence of an individual's button presses with a group's normative boundaries. For every participant, we first extracted their button press times (individual boundaries). Next, we identified normative boundaries for the rest of the group using the group density time series. We then calculated k as the number of times the individual's boundaries overlapped with (i.e., fell within 1 s of) a group normative boundary and *rarity* of such overlap as the summed probability of overlap occurring more than k -number of times. Surprise index was then calculated by $-\log_2$ transforming rarity (Katori et al., 2018).

Evaluating the Sensitivity and Specificity of Segmentation Agreement Measures

To evaluate each measure's sensitivity to the event boundaries participants are likely to identify in a movie, we calculated agreement values for: (1) segmentation data from one movie (*same*; as described above), (2) randomly generated segmentation data (*random*), and (3) segmentation data from different movies (*cross-movie*). Random data were generated for each participant by randomly sampling (without replacement) l times from a continuous uniform distribution with a minimum value of 0 and a maximum value equal to the video duration (in seconds), where l was the number of times the participant pressed the button. Random agreement was then calculated using a group of random data (in the case of peakiness) or by comparing two random datasets (in the case of peak-to-peak distance, agreement index, and surprise index).

Random agreement measures therefore represent agreement of completely random segmentation data. If a measure reliably captures information about how individuals or groups segment a video, then it should show better agreement for same movie comparisons than for random data, or cross-movie comparisons.

Statistical analyses

We fitted linear mixed effects models on each of the segmentation agreement measures with sample size (2 to 32, incremented by 2), segmentation grain (fine vs coarse), and measurement condition (same, random data, or cross-movie) as fixed effects and videos as random effects using the `lmer` function from the `lme4` package in R (Bates et al., 2015). Separate models were fitted for the commercial-lab and everyday-online datasets. We used planned contrasts to test whether same agreement values differed from random data and cross-movie values. For completeness, we also tested differences between random and cross-movie agreement values. Because we did not have specific predictions for this comparison, we only report findings that show clear and consistent patterns in all segmentation conditions.

To investigate how each agreement measure changes with increasing sample size, we examined polynomial contrasts for the effect of sample size using `emmeans` (Lenth et al., 2018). Segmentation agreement should improve with the increase in sample size. Therefore, we expected a significant positive linear contrast for peakiness, agreement index and surprise index and a significant negative linear contrast for peak-to-peak distance.

If segmentation agreement estimates stabilize with increasing sample size, there will be a point at which the rate of change in agreement values slows as sample size increases (the relationship will be nonlinear). Therefore, we expected significant negative quadratic effects for peakiness, agreement index, and surprise index (for which positive values are better) and a

positive quadratic effect for peak-to-peak distance (for which negative values are better).

Whenever a quadratic effect was significant, we estimated the sample size at which agreement stabilized. To do this, we first fit each agreement measure to several growth and decay functions (Figure 1B): linear, exponential, logarithmic, asymptotic regression, power curve, and yield loss using `lm`, `nls` (R core team, 2013), `drm` (`drc` package, Ritz et al., 2015) and `aomisc` (Onofri, 2020) in R. These functions were selected for their potential ability to fit the shape of the data and were not theoretically motivated. Next, we determined the best fitting curve as the fit with the lowest AIC and BIC values. In all but one case (see Appendix B), the best fitting curve was nonlinear. Finally, the elbow was defined as the point on the fitted curve that was furthest from a straight line connecting the start and end points of the curve. The elbow was then taken as the sample size at which the measure was stabilizing and including more participants would provide diminishing returns in estimating agreement. Asymptotic values for the best fitting functions are reported in Table 1.

Results

In this section, we report analyses for each agreement measure separately. For each measure, we first report analyses that assess the measure's sensitivity in distinguishing boundaries people are likely to identify during a movie from random noise (same vs random comparison) and from boundaries they are likely to identify for a different movie (same vs cross-movie comparison) at different sample sizes. Next, we describe how agreement values change as sample size increases. Should the measure of agreement stabilize, we report the sample size at which stabilization begins. Lastly, we report the presence of any differences between random and cross-movie comparison.

Peakiness

As a measure of within group segmentation agreement, our analyses confirmed that peakiness was sensitive to participants' agreement about when event boundaries occurred in a movie (Figure 2). Same peakiness values were overall higher than random peakiness values in the commercial-lab dataset, $t(12735) = 125.79$, 95% CI [4.52, 4.66] for coarse, $t(12735) = 160.66$, $p < .001$, 95% CI [5.79, 5.93] for fine and in the everyday-online dataset, $t(25528) = 62.45$, $p < .001$, 95% CI [15.17, 16.15] for coarse, $t(25528) = 26.41$, $p < .001$, 95% CI [6.13, 7.11] for fine. Pairwise comparisons revealed that peakiness began to differentiate real data from noise at sample size of 6 for coarse commercial-lab segmentation, $ts(12735) > 9.19$, $ps < .001$, smallest 95% CI [1.05, 1.63] for $N = 6$ to 32, 4 for fine commercial-lab segmentation, $ts(12735) > 2.38$, $ps < .017$, smallest 95% CIs [0.061, 0.63] for $N = 4$ to 32, 2 for coarse everyday-online segmentation, $ts(25528) > 10.67$, $ps < .001$, smallest 95% CIs [8.75, 12.69] for $N = 2$ to 32, and 2 for fine everyday-online segmentation $ts(25528) > 3.67$, $ps < .002$, smallest 95% CI [1.72, 5.65] for $N = 2$ to 32. Because peakiness provides a summary of segmentation agreement within one group and not a comparison between different groups, cross-movie peakiness values could not be calculated.

Increasing sample size increased peakiness in most cases. Significant linear contrasts indicated that same peakiness values increased with sample size in the commercial-lab data, $t(12735) = 81.44$, $p < .001$, 95% CI [302.35, 317.26] for coarse, $t(12735) = 88.43$, $p < .001$, 95% CI [328.92, 343.83] for fine, and for coarse segmentation in the everyday-online data, $t(25528) = 7.54$, $p < .001$, 95% CI [146.04, 248.62]. However, the linear contrast for fine segmentation in the everyday-online data showed a decreasing trend instead, $t(25528) = -3.08$, $p = .0021$, 95% CI [-131.68, -29.19], potentially reflecting greater variance across individuals in the online dataset. There was also evidence that in some cases, peakiness stabilized at larger sample sizes: the

negative quadratic effect of sample size on same peakiness values was significant for fine segmentation in commercial-lab, $t(12735) = 9.83, p < .00, 95\% \text{ CI } [-91.88, -61.31]$, and for coarse segmentation in everyday-online, $t(25528) = 2.74, p < .001, 95\% \text{ CI } [-251.99, -41.74]$, stabilizing at $N = 16$ and $N = 8$, respectively. However, we found no evidence for stabilization for coarse segmentation in commercial-lab as peakiness continued to increase with the increase in sample size, resulting in a significant positive quadratic contrast, $t(12735) = 4.27, p < .001, 95\% \text{ CI } [18.03, 48.59]$. Similarly, for fine segmentation in everyday-online, peakiness did not show evidence of stabilization as suggested by the non-significant quadratic contrast, $t(25528) = 0.60, p = .55, 95\% \text{ CI } [-72.98, 137.07]$.

These findings suggest that peakiness can be used to quantify within group agreement about boundary placement. Further, while it is sensitive to this information at relatively small sample sizes of 2 to 6 participants (depending on grain and dataset) it did not appear to consistently stabilize as sample size increased. When it did, stabilization occurred with 16 or

fewer participants.

Figure 2. $\log_{10}$ transformed peakiness values over increasing sample sizes for: A) commercial-lab and B) everyday-online datasets. Small shapes depict individual values calculated from a single bootstrap iteration. Larger shapes depict the average value from all bootstrapping iterations. One low peakiness value and seven high peakiness values for coarse everyday activity segmentation were excluded from the plot due to the y-axis limit. Error bars represent 95% confidence interval and carets (^) represent the elbows. Due to $\log_{10}$ scaling a peakiness value of 1 (when observed rugosity = rugosity for uniformly distributed button presses) will correspond to 0 on this graph.

Table 1

Asymptotic values of the best fitting nonlinear functions (see Appendix B) for measures showing stabilization with increasing sample sizes.

Measure	Commercial-lab		Everyday-online	
	Coarse	Fine	Coarse	Fine
Peakiness	-	22.68	22.95	-
Peak-to-peak distance	3.57	1.32	10.87	1.19
Agreement Index	0.50	0.63	0.54	0.69
Surprise Index	9.74	11.08	7.07	6.44

Note: Peak-to-peak distance values are in seconds.

Peak-to-peak distance

Peak-to-peak distance was sensitive to boundaries identified in specific movies, distinguishing group boundaries from the same movie from group boundaries identified using random data and other movies (Figure 3). Same peak-to-peak distance was overall lower than random peak-to-peak distance in the commercial-lab dataset, $t(19103) = 84.14, p < .001$, 95% CI [2.98, 3.12] for coarse, $t(19103) = 20.83, p < .001$, 95% CI [0.68, 0.83] for fine, and in the everyday-online dataset, $t(38301) = 26.59, p < .001$, 95% CI [2.49, 2.89] for coarse and $t(38301) = 3.72, p < .001$, 95% CI [0.18, 0.57] for fine. Same peak-to-peak distance was also overall lower than cross-movie comparisons in commercial-lab, $t(19103) = 93.77, p < .001$, 95% CI [3.33, 3.47] for coarse, $t(19103) = 20.16, p < .001$, 95% CI [0.66, 0.80] for fine and in the everyday-online dataset, $t(38301) = 52.53, p < .001$, 95% CI [5.12, 5.51] for coarse, $t(38301) = 7.22, p < .001$, 95% CI [0.53, 0.93] for fine. Pairwise comparisons revealed that these differences were evident starting at sample size 2 for coarse commercial-lab segmentation, $ts(19103) > 14.48, ps < .001$, smallest 95% CIs [1.82, 2.38] for $N = 2$ to 32, and fine

commercial-lab segmentation, $t(19103) > 3.84$, $p < .001$, smallest 95% CIs [0.27, 0.84] for $N = 2$ to 32, and for coarse segmentation in everyday-online, $t(38301) > 5.17$, $ps < .001$, smallest 95% CIs [1.30, 2.88], for $N = 2$ to 32. However, same peak-to-peak distances values for fine segmentation in everyday-online were not significantly lower than random, $t(38301) < 1.07$, $ps > .28$, largest 95% CIs [-.036, 1.23] for $N = 2$ to 32, or cross-movie, $t(38301) < 1.96$, $ps > .05$, largest 95% CIs [-.000053, 1.59] for $N = 2$ to 32 peak-to-peak distance at all sample sizes. This set of findings may again reflect the high variance in fine peak-to-peak distance observed in the everyday-online dataset, similar to peakiness values for fine everyday-online segmentation.

Peak-to-peak distance decreased as sample size increased and eventually stabilized. Significant linear contrasts indicated that peak-to-peak distance decreased with increasing sample size in the commercial-lab dataset: $t(19103) = 56.98$, $p < .001$, 95% CI [-222.79, -207.97] for coarse, $t(19103) = 17.53$, $p < .001$, 95% CI [-73.84, -59.02] for fine, and in the everyday-online dataset: $t(38301) = 13.68$, $p < .001$ 95% CI [-165.03, -123.66] for coarse, $t(38301) = 3.09$, $p = .0011$, 95% CI [-53.33, -11.96] for fine. Significant positive quadratic fits in commercial-lab coarse, $t(19103) = 38.16$, $p < .001$ 95% CI [280.47, 310.84] and fine $t(19103) = 11.57$, $p < .001$, 95% CI [74.48, 104.85] segmentation and in everyday-online coarse, $t(38301) = 9.27$, $p < .001$, 95% CI [158.18, 242.95] and fine, $t(38301) = 2.48$, $p = .013$, 95% CI [11.15, 95.92] segmentation were consistent with stabilization in this measure. Subsequent analyses suggested that peak-to-peak distance started to stabilize at sample size 10 for coarse and fine segmentation in the commercial-lab dataset and 8 for coarse segmentation and 6 for fine segmentation in the everyday-online dataset. Asymptotic values for peak-to-peak distance differed greatly across grains and data sets, ranging from 1.19 s to 10.87 s (Table 1 & Figure 3).

Overall, cross-movie and random peak-to-peak distance did not consistently differ. In the commercial-lab dataset, while cross-movie peak-to-peak distance was higher than random peak-to-peak distance for coarse segmentation, $t(19103) = 9.63, p < .001$, 95% CI [0.28, 0.42], cross-movie and random peak-to-peak distance were not different for fine segmentation, $t(19103) = .667, p = .503$, 95% CI [-0.095, 0.047]. In addition, although cross-movie peak-to-peak distance was higher than random for coarse, $t(38301) = 25.94, p < .001$, 95% CI [2.43, 2.82] and fine, $t(38301) = 3.50, p < .001$, 95% CI [0.16, 0.55] segmentation in the everyday-online dataset, pairwise comparison revealed that cross-movie and random peak-to-peak distances for fine everyday-online datasets did not differ significantly at any sample size, all $ts(38301) < 1.01, ps > .31$, largest 95% CIs [-0.39, 1.20].

These findings suggest that peak-to-peak distance can be used to quantify two group's agreement about boundary placement within a specific movie. Information about normative group boundaries captured by peak-to-peak distance was present in samples as small as 2 for all segmentation conditions except for fine segmentation in everyday-online. Further, as sample size increased, peak-to-peak distance stabilized with 10 or fewer participants.

Figure 3. $\log_{10}$ transformed peak-to-peak distance over increasing sample sizes for segmentation of: A) commercial-lab and B) everyday-online. Small shapes depict individual values calculated from a single bootstrap iteration. Larger shapes depict the average value from all bootstrapping iterations. The minimum and maximum values of the y-axis for each plot is adjusted between grains to better capture the degree of change in peak-to-peak distance values, but the ranges are kept consistent. Twelve high peak-to-peak distance values for coarse commercial-lab and five high peak-to-peak distance values for coarse everyday-online were excluded from the plot due to the limits set for the y-axes. Error bars represent 95% confidence intervals and carets (^) represent the elbows.

Agreement Index

As a measure of how well an individual agrees with a separate group of observers, our analyses confirmed that agreement index was sensitive to the movie being segmented (Figure 4).

It distinguished segmentation of the same movie from random button presses and segmentation of a different movie in both the commercial-lab dataset and the everyday-online dataset. Same agreement index values were overall higher than those for random data in the commercial-lab dataset, $t(19103) = 134.77, p < .001, 95\% \text{ CI} = [0.155, 0.160]$ for coarse segmentation, $t(19103) = 103.04, p < .001, 95\% \text{ CI} = [0.118, 0.122]$ for fine segmentation, and in the everyday-online dataset, $t(38301) = 181.20, p < .001, 95\% \text{ CI} = [0.188, 0.192]$ for coarse segmentation, $t(38301) = 153.42, p < .001, 95\% \text{ CI} = [0.159, 0.163]$ for fine segmentation. Same agreement index values were also higher than cross-movie agreement index values in the commercial-lab dataset, $t(19103) = 195.79, p < .001, 95\% \text{ CI} [0.227, 0.231]$ for coarse segmentation, $t(19103) = 174.11, p < .001, 95\% \text{ CI} [0.201, 0.206]$ for fine segmentation, and in the everyday-online dataset, $t(38301) = 212.40, p < .001, 95\% \text{ CI} [0.221, 0.225]$ for coarse segmentation, $t(38301) = 177.36, p < .001, 95\% \text{ CI} [0.184, 0.188]$ for fine segmentation. Pairwise comparison revealed that agreement index distinguished same movie data from both random data and cross-movie data starting at sample size of 2 for the commercial-lab dataset, $ts(19103) > 12.09, ps < .001$, smallest 95% CIs $[0.047, 0.066]$ for coarse segmentation, $ts(19103) > 8.53, ps < .001$, smallest 95% CIs $[0.031, 0.049]$ for fine segmentation, and for the everyday-online dataset, $ts(38301) > 11.15, ps < .001$, smallest 95% CIs $[0.039, 0.055]$ for coarse segmentation, $ts(38301) > 15.67, ps < .001$, smallest 95% CIs $[0.058, 0.074]$ for fine segmentation.

Agreement index also improved and eventually stabilized with increasing sample size. Positive, significant linear contrasts indicated that same agreement index increased with increasing sample size in the commercial-lab dataset, $t(19103) = 83.41, p < .001, 95\% \text{ CI} [9.94, 10.42]$ for coarse segmentation, $t(19103) = 85.04, p < .001, 95\% \text{ CI} [10.14, 10.62]$ for fine segmentation and the everyday-online dataset, $t(38301) = 104.50, p < .001, 95\% \text{ CI} [11.24,$

11.67] for coarse segmentation, $t(38301) = 73.90, p < .001$, 95% CI [7.88, 8.32] for fine segmentation. This effect of increasing sample size on the agreement index decreased as sample size increased, as evidenced by the significant negative quadratic fit in the commercial-lab dataset, $t(19103) = 40.59, p < .001$, 95% CI [-10.65, -9.66] for coarse segmentation, $t(19103) = 47.30, p < .001$, 95% CI [-12.33, -11.34] for fine segmentation, and in the everyday-online dataset, $t(38301) = 48.00, p < .001$, 95% CI [-11.22, -10.34] for coarse segmentation, $t(38301) = 44.44, p < .001$, 95% CI [-10.42, -9.54] for fine segmentation. Subsequent analyses indicated that the agreement index started to stabilize at sample size 14 for coarse and fine segmentation in the commercial-lab dataset, 12 for coarse segmentation in the everyday-online dataset and 10 for fine segmentation in the everyday-online dataset. Asymptotic values for the agreement index differed across grains, but were comparable across data sets (Table 1 & Figure 4). These values were also comparable to those reported in other work (e.g., Kurby & Zacks, 2011).

Unlike the other segmentation agreement measures, the agreement index differentiated cross-movie comparisons from comparisons to random data. The cross-movie agreement index was overall lower than the random agreement index for all segmentation grains in the commercial-lab dataset, $t(19103) = 61.02, p < .001$, 95% CI [0.069, 0.074] for coarse segmentation, $t(19103) = 71.07, p < .001$, 95% CI [0.081, 0.085] for fine segmentation and in the everyday-online dataset, $t(38301) = 31.20, p < .001$, 95% CI [0.031, 0.035] for coarse segmentation, $t(38301) = 23.94, p < .001$, 95% CI [0.023, 0.027] for fine segmentation. Pairwise comparisons revealed that this difference was evident with a sample size as small as 6, $ts(19103) > 2.75, ps < .006$, smallest 95% CIs [0.0037, 0.022] for $N = 6$ to 32 for coarse segmentation and 4, $ts(19103) > 2.49, ps < .013$, smallest 95% CIs [0.0025, 0.021] for $N = 4$ to 32 for fine segmentation in the commercial-lab data. For the everyday-online dataset, the difference started

at sample size 12 for coarse segmentation, $t_s(38301) > 4.01$, $p_s < .001$, smallest 95% CIs [0.0086, 0.025] for $N = 12$ to 32, and sample size 6 for fine segmentation, $t_s(38301) > 2.58$, $p_s < .001$, smallest 95% CIs [0.0026, 0.019] for $N = 6$ to 32. Lower cross-movie agreement values could reflect incompatible information about event structure in different movies. Consistent with this possibility, the group timeseries from different movies on the original sample were weakly correlated (range of $r(272)s = .00 - .16$, $p_s > .007$, Appendix C).

These findings suggest that agreement index can be used to quantify how well the boundaries identified by an individual agree with those identified by a separate group. This information was captured by agreement index at sample size as small as 2 for all segmentation conditions. Further, as sample size increased, agreement index eventually stabilized and this stabilization occurred with 14 participants or less. Lastly, agreement index also seems to be sensitive enough to differentiate random segmentation data from segmentation on different movies. This property was not observed in the other agreement measures tested in this study.

Figure 4. Agreement index over increasing sample size for segmentation in: A) commercial-lab and B) everyday-online. Small shapes depict individual values calculated from a single bootstrap iteration. Larger shapes depict the average value from all bootstrapping iterations. Error bars represent 95% confidence interval and carets (^) represent the elbows.

Surprise Index

Overall, surprise index was sensitive to whether individual segmentation was compared to normative boundaries from groups that were segmenting the same movie (Figure 5). The surprise index distinguished segmentation of the same movie from random button presses and segmentation of a different movie in both the commercial-lab dataset and the everyday-online dataset. Same surprise index was higher than random surprise index in the commercial-lab

datasets, $t(19103) = 96.78, p < .001$ 95% CI [3.36, 3.50] for coarse segmentation, $t(19103) = 188.96, p < .001$, 95% CI [6.64, 6.77] for fine segmentation, and in the everyday-online datasets, $t(38301) = 113.33, p < .001$, 95% CI [2.51, 2.59] for coarse segmentation, $t(38301) = 180.11, p < .001$, 95% CI [4.01, 4.10] for fine segmentation. Surprise index was also overall higher for same than cross-movie group comparisons in the commercial-lab dataset, $t(19103) = 100.61, p < .001$, 95% CI [3.50, 3.63] for coarse segmentation, $t(19103) = 188.24, p < .001$, 95% CI [6.61, 6.75] for fine segmentation, and in the everyday-online dataset, $t(38301) = 100.56, p < .001$, 95% CI [2.22, 2.31] for coarse segmentation, $t(38301) = 172.68, p < .001$, 95% CI [3.84, 3.93] for fine segmentation. Pairwise comparisons revealed that differences between same movie comparisons and both random data and cross-movie comparisons started at a sample size of 2 in the commercial-lab datasets, all $ts(19103) > 9.38, ps < .001$, smallest 95% CI [1.05, 1.61] for coarse segmentation, all $ts(19103) > 16.37, ps < .001$, smallest 95% CI [2.04, 2.60] for fine segmentation, and in the everyday-online dataset, all $ts(38301) > 13.41, ps < .001$, smallest 95% CI [1.03, 1.38] for coarse segmentation, all $ts(38301) > 20.86, ps < .001$, smallest 95% CI [1.70, 2.05] for fine segmentation.

Surprise index improved and eventually stabilized with increasing sample size, as indicated by significant positive linear contrasts for the commercial-lab dataset, $t(19103) = 42.47, p < .001$, 95% CI [149.96, 164.47] for coarse, $t(19103) = 77.65, p < .001$, 95% CI [280.16, 294.67] for fine and the everyday-online dataset, $t(38301) = 33.30, p < .001$, 95% CI [73.56, 82.76] for coarse, $t(38301) = 41.50, p < .001$, 95% CI [92.79, 101.99] for fine. The positive effect of sample size on surprise index grew smaller with larger sample sizes, as evidenced by the significant negative quadratic fit in the commercial-lab dataset, $t(19103) = 5.76, p < .001$, 95% CI [-58.56, -28.82] for coarse, $t(19103) = 27.83, p < .001$, 95% CI [-226.00, -196.27] for

fine, as well as in the everyday-online dataset, $t(38301) = 6.50, p < .001, 95\% \text{ CI } [-40.70, -21.85]$ for coarse, $t(38301) = 17.73, p < .001, 95\% \text{ CI } [-94.68, -75.83]$ for fine. Additional analyses indicated that surprise index began to stabilize at sample size 16 for coarse segmentation in commercial-lab and everyday-online, 14 for fine segmentation in commercial-lab and 10 for fine segmentation in everyday-online. Asymptotic surprise index values differed markedly across conditions (Table 1; Figure 5).

Surprise index did not consistently differ between cross-movie and random data comparisons. Cross-movie surprise index was lower than random surprise index for coarse, $t(19103) = 3.83, p < .001, 95\% \text{ CI } [0.07, 0.21]$ but not fine, $t(19103) = .71, p = .47, 95\% \text{ CI } [-0.095, 0.044]$ segmentation in the commercial-lab dataset. For fine and coarse segmentation in the everyday-online dataset, cross-movie surprise index was higher than random surprise index, $t(38301) = 12.78, p < .001, 95\% \text{ CI } [0.24, 0.33]$ for coarse, $t(38301) = 7.43, p < .001, 95\% \text{ CI } [0.12, 0.21]$.

These findings suggest that surprise index can be used to quantify agreement between an individual's boundary placement with normative group boundaries within a specific movie. This information was captured by surprise index at a sample size as small as 2 for all segmentation conditions. Further, as sample size increased, surprise index eventually stabilized with 16 or fewer participants.

Figure 5. $\log_{10}$ transformed surprise index over increasing sample size for segmentation in: A) commercial-lab and B) everyday-online. Small shapes depict individual values calculated from a single bootstrap iteration. Larger shapes depict the average value from all bootstrapping iterations. Error bars represent 95% confidence intervals and carets (^) represent the elbow.

Discussion

This study systematically explored various segmentation agreement measures, each quantifying performance at the group (*peakiness* and *peak-to-peak distance*) or individual (*agreement index* and *surprise index*) level. The results confirmed previous findings on the segmentation task's utility in measuring meaningful behavior. They demonstrated that multiple agreement measures can capture group and individual segmentation patterns with relatively few

participants. This observation is true for data collected using structured commercial films in the lab and unedited videos of everyday activities online. The results also validated the use of group segmentation data to infer when other people are likely to segment, especially when large sample sizes are used. They provide insight into potential limits to the amount of agreement that can be observed at the group level and inform the selection of sample sizes and measures for future research.

In almost all cases, we found that segmentation agreement calculated using performance on one movie (*same*) faired significantly better than segmentation agreement calculated using randomly generated data (*random*). This finding shows that the explored agreement measures are able to distinguish real segmentation data from random button presses. Further, *same* agreement values were also significantly better than segmentation agreement calculated by comparing performance on two different movies (*cross-movie*). This result indicates that the agreement measures are sensitive to responses driven by specific movies. Importantly, such non-random and movie-specific segmentation behavior was observed with sample sizes of 6 individuals or less (2 in most cases). Collectively, these findings highlight the efficacy of using the segmentation task to capture genuine and meaningful segmentation of naturalistic events.

One advantage of using the segmentation task to study event perception is its use of naturalistic stimuli, such as movies, to approximate the complex and continuous nature of everyday experience. However, unlike presenting simple stimuli such as images or tones in a discrete and temporally controlled manner, the mapping between stimuli and responses in a segmentation task is undefined. Further, although boundary identification is associated with changes in the features of a movie (e.g., action or location changes; Hard et al., 2006; Magliano et al., 2001; Newton et al., 1977; Swallow et al., 2018), events are still fundamentally a

manifestation of the mind. Therefore, establishing a ground truth to compare segmentation performance with is not only challenging but potentially inappropriate. This study showed that, despite no “correct” behavior in a segmentation task, segmentation performance is far from random and is not generalizable across different movies. Therefore, findings from this study should increase confidence in the use of the segmentation task and naturalistic stimuli to study event perception, its cognitive implications, and its neural correlates (Sonkusare et al., 2019).

Importantly, we found that the reliability of segmentation task measures extends beyond the confines of controlled laboratory settings. In this study, we performed all analyses on two separate datasets: segmentation of commercially made movies performed in the lab (*commercial-lab*) and segmentation of everyday activities performed online (*everyday-online*). Despite differences in the type of videos used, how the data were collected, and the sampled population, on average, group and individual level agreement measures were comparable across both datasets. This finding provides further support on two main fronts. First, it affirms that the segmentation task can capture meaningful information about how people segment everyday activities or events depicted in commercial movies. Second, it demonstrates that the segmentation task is a reliable paradigm for assessing segmentation behavior in a diverse population and on a platform where there is minimal experimental control. This information is valuable considering the increased interest in online data collection methods for studying human cognition over the recent years (Stewart et al., 2017) and the heightened reliance on web-based alternatives due to global pressures (e.g., the Covid-19 pandemic).

Unsurprisingly, in all measures and conditions, we found that normative boundaries differed across samples of participants, particularly for small samples. This finding calls for caution when using data from a small group of participants as a basis for inferring when

individuals are likely to segment an event. For small sample sizes, normative boundaries could shift by several seconds (on average) from one sample to another (Figure 2). Shifts in boundary location remained as large as 10 s (Table 1) even at the largest sample sizes that we examined. Therefore, future research should take into account the limited precision of normative boundaries and ensure that they are acquired from groups that are large enough to maximize their expected reproducibility across groups.

Despite inherent variability in segmentation performance across individuals, in this study agreement started to stabilize with moderate sample sizes (16 participants or less). Although adding participants beyond this stabilization point resulted in diminishing gains in agreement, studies investigating potentially small differences in segmentation agreement may need to aim for larger sample sizes. Further, the sample sizes at which segmentation agreement started to stabilize differed between individual and group level agreement measures. Thus, studies interested in estimating the similarity between individual and group performance would require larger sample sizes than studies interested in identifying normative boundaries. However, the inherent variability of segmentation performance poses a limit on estimates of group and individual levels of agreement (see Table 1), and future studies should always be cautious of this limit.

This study introduced a few novel methods to quantify different elements typical to a segmentation task. First, it introduced a new systematic method for defining normative group boundaries that preserve the continuous nature of segmentation behavior and account for subtle individual variation in when participants pressed a button to report an event boundary. Second, this study introduced new ways to quantify individual and group-level agreement: the *surprise index* and *peakiness*. The surprise index compared individual segmentation performance to the

group's normative boundary placement, whereas peakiness summarizes within-group segmentation agreement. Of note, in our analysis, peakiness is the measure in which we did not find consistent evidence for stabilization over increasing sample size. In commercial-lab coarse segmentation, peakiness continued to increase with increases in sample size, up to the largest sample tested in this study. This finding further highlights the considerable variation embedded within individual performance that can affect overall group agreement, particularly when agreement is examined at the temporal resolution used in this study.

Of all the agreement measures explored in this paper, the most commonly used method, the *agreement index*, also reliably differentiated agreement on randomly generated data from agreement on segmentation of different movies. Specifically, cross-movie agreement index was significantly worse than random agreement index in all segmentation grains and conditions. This finding potentially reflects the agreement index's added sensitivity to capturing responses to incompatible signals coming from different movies. Consistent with this reasoning, we found weak correlations between different movies' group boundary placement over time. Our findings also suggest that this sensitivity is unique to the agreement index. However, verification of this conjecture is beyond the scope of this paper. Future studies could validate this finding by more closely examining the relationship between boundary placement and features of the movies.

This study does have several limitations. Although the overall pattern of results between commercial-lab and everyday-online datasets were comparable, there were notable differences. First, we found greater variance within the everyday-online dataset that mainly affected group-level agreement (peakiness and peak-to-peak distance) for fine segmentation. Additionally, while analyses on individual to group agreement measures show the same effects in both datasets, more variability was observed across the online data sub-samples. These differences are difficult

to interpret. Greater variability in the everyday-online data could reflect common concerns with online task performance, such as differences in: levels of attentional engagement, equipment used to perform the task, or in the sample population. It could also reflect the use of lab-produced videos of everyday activities in the everyday-online dataset. Unlike commercially produced movies, everyday activity videos are not edited and do not include features (e.g., cuts and scene changes) that may systematically facilitate boundary identification (Cutting, 2019; Magliano et al., 2020; Magliano & Zacks, 2011). However, because everyday activity videos were segmented online, we could not disentangle online performance effects from video type. Future work is necessary to address this question. Finally, although our data sets were relatively large for this type of research ($N = 64$, and $N = 72$), our analyses and the conclusions that can be drawn from them are limited by sample size, the population that they were drawn from, and by the degree to which participants in our data were representative of the populations from which they were drawn (an assumption of bootstrapping procedures for estimating sampling distributions).

Conclusion

This study demonstrates that group-to-group or individual-to-group segmentation agreement measures reflect behavior that is meaningful and sensitive to the structure of the movie, and that this is true for segmentation data collected online as well as in the lab. When estimating segmentation agreement, large samples may not show large advantages over moderately sized samples, and even small samples may be sufficient to detect non-random task performance. However, the inherent variability of segmentation performance across participants may limit the certainty with which normative boundaries can be identified. Given the extensive influence of event segmentation on other cognitive processes such as attention and memory, this

research can inform and supplement other studies on event segmentation and its cognitive implications.

Acknowledgements

The authors would like to thank Melissa Yap, Ariella Ratner, Sofie Prieto, & Joel Cruz for assistance in data collection for the in-lab dataset and Taylor Beauvais, Clara Lee, Isabella Poulous, and Hyung Sop Shin for assistance in conducting the online study. The authors declare no conflicts of interest. This research was financially supported by the Cornell Center for Social Sciences (CCSS 297) and Cornell College of Arts and Sciences.

Declarations

Funding

This research was financially supported by the Cornell Center for Social Sciences (CCSS 297) and Cornell College of Arts and Sciences.

Competing interests

The authors have no conflicts of interest to declare.

Ethics Approval

All procedures involved in collecting and analyzing data from human subjects were conducted in accordance with the ethical standards of the 1964 Declaration of Helsinki and its

subsequent amendments. Procedures were reviewed and approved by the Cornell Institutional Review Board.

Consent to Participate

Informed consent was obtained from all individual participants who contributed data to this study.

Data and Code Availability

Data and analytical code used for this study are publicly available in the Github repository, at the following link: <https://github.com/ksasmita/esMethods>.

Author Contributions

K. Sasmita – study concept and design, data collection, data analysis design, code development, writing first draft of manuscript, and manuscript revisions and editing.

K. Swallow – study concept and design, data collection, data analysis design, manuscript revisions and editing.

References

- Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., & Norman, K. A. (2017). Discovering Event Structure in Continuous Narrative Perception and Memory. *Neuron*, 95(3), 709-721.e5. <https://doi.org/10.1016/j.neuron.2017.06.041>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using **lme4**. *Journal of Statistical Software*, 67(1). <https://doi.org/10.18637/jss.v067.i01>
- Ben-Yakov, A., & Henson, R. N. (2018). The Hippocampal Film Editor: Sensitivity and Specificity to Event Boundaries in Continuous Experience. *Journal of Neuroscience*, 38(47), 10057–10068. <https://doi.org/10.1523/JNEUROSCI.0524-18.2018>
- Birnbaum, M. H. (2004). Human Research and Data Collection via the Internet. *Annual Review of Psychology*, 55(1), 803–832. <https://doi.org/10.1146/annurev.psych.55.090902.141601>
- Bläsing, B. E. (2015). Segmentation of dance movement: Effects of expertise, visual familiarity, motor experience and music. *Frontiers in Psychology*, 5. <https://doi.org/10.3389/fpsyg.2014.01500>
- Chen, J., Leong, Y. C., Honey, C. J., Yong, C. H., Norman, K. A., & Hasson, U. (2017). Shared memories reveal shared structure in neural activity across individuals. *Nature Neuroscience*, 20(1), 115–125. <https://doi.org/10.1038/nn.4450>
- Cutting, J. E., DeLong, J. E., & Brunick, K. L. (2011). Visual activity in Hollywood film: 1935 to 2005 and beyond. *Psychology of Aesthetics, Creativity, and the Arts*, 5(2), 115–125. <https://doi.org/10.1037/a0020995>
- Faber, M., Radvansky, G. A., & D’Mello, S. K. (2018). Driven to distraction: A lack of change gives rise to mind wandering. *Cognition*, 173, 133–137. <https://doi.org/10.1016/j.cognition.2018.01.007>

- Hard, Bridgette M., Tversky, B., & Lang, D. S. (2006). Making sense of abstract events: Building event schemas. *Memory & Cognition*, 34(6), 1221–1235.
<https://doi.org/10.3758/BF03193267>
- Hard, Bridgette M., Recchia, G., & Tversky, B. (2011). The shape of action. *Journal of Experimental Psychology: General*, 140(4), 586–604. <https://doi.org/10.1037/a0024310>
- Katori, K., Manabe, H., Nakashima, A., Dunfu, E., Sasaki, T., Ikegaya, Y., & Takeuchi, H. (2018). Sharp wave-associated activity patterns of cortical neurons in the mouse piriform cortex. *European Journal of Neuroscience*, 48(10), 3246–3254.
<https://doi.org/10.1111/ejn.14099>
- Kosie, J. E., & Baldwin, D. (2019). Attention rapidly reorganizes to naturally occurring structure in a novel activity sequence. *Cognition*, 182, 31–44.
<https://doi.org/10.1016/j.cognition.2018.09.004>
- Kurby, C. A., & Zacks, J. M. (2011). Age differences in the perception of hierarchical structure in events. *Memory & Cognition*, 39(1), 75–91. <https://doi.org/10.3758/s13421-010-0027-2>
- Levine, D., Hirsh-Pasek, K., Pace, A., & Michnick Golinkoff, R. (2017). A goal bias in action: The boundaries adults perceive in events align with sites of actor intent. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(6), 916–927.
<https://doi.org/10.1037/xlm0000364>
- Magliano, J. P., Miller, J., & Zwaan, R. A. (2001). Indexing space and time in film understanding. *Applied Cognitive Psychology*, 15(5), 533–545.
<https://doi.org/10.1002/acp.724>

- Massad, C. M., Hubbard, M., & Newton, D. (1979). Selective perception of events. *Journal of Experimental Social Psychology*, 15(6), 513–532. [https://doi.org/10.1016/0022-1031\(79\)90049-0](https://doi.org/10.1016/0022-1031(79)90049-0)
- Newberry, K. M., & Bailey, H. R. (2019). Does semantic knowledge influence event segmentation and recall of text? *Memory & Cognition*. <https://doi.org/10.3758/s13421-019-00926-4>
- Newberry, K. M., Feller, D. P., & Bailey, H. R. (2021). Influences of domain knowledge on segmentation and memory. *Memory & Cognition*. <https://doi.org/10.3758/s13421-020-01118-1>
- Newton, D. (1973). Attribution and the unit of perception of ongoing behavior. *Journal of Personality and Social Psychology*, 28(1), 28–38. <https://doi.org/10.1037/h0035584>
- Newton, D., & Engquist, G. (1976). The perceptual organization of ongoing behavior. *Journal of Experimental Social Psychology*, 12(5), 436–450. [https://doi.org/10.1016/0022-1031\(76\)90076-7](https://doi.org/10.1016/0022-1031(76)90076-7)
- Newton, D., Engquist, G., & Bois, J. (1977). *The Objective Basis of Behavior Units*. 16.
- Ritchey, M., & Cooper, R. A. (2020). Deconstructing the Posterior Medial Episodic Network. *Trends in Cognitive Sciences*, 24(6), 451–465. <https://doi.org/10.1016/j.tics.2020.03.006>
- Ritz, C., Baty, F., Streibig, J. C., & Gerhard, D. (2015). Dose-Response Analysis Using R. *PLOS ONE*, 10(12), e0146021. <https://doi.org/10.1371/journal.pone.0146021>
- Sargent, J. Q., Zacks, J. M., Hambrick, D. Z., Zacks, R. T., Kurby, C. A., Bailey, H. R., Eisenberg, M. L., & Beck, T. M. (2013). Event segmentation ability uniquely predicts event memory. *Cognition*, 129(2), 241–255. <https://doi.org/10.1016/j.cognition.2013.07.002>

- Sheather, S. J., & Jones, M. C. (1991). A Reliable Data-Based Bandwidth Selection Method for Kernel Density Estimation. *Journal of the Royal Statistical Society. Series B (Methodological)*, 53(3), 683–690.
- 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. <https://doi.org/10.3758/CABN.3.4.335>
- Sueur, J., Aubin, T., & Simonis, C. (2008). Seewave, a Free Modular Tool for Sound Analysis and Synthesis. *Bioacoustics*, 18(2), 213–226.
<https://doi.org/10.1080/09524622.2008.9753600>
- Swallow, K. M., Kemp, J. T., & Candan Simsek, A. (2018). The role of perspective in event segmentation. *Cognition*, 177, 249–262. <https://doi.org/10.1016/j.cognition.2018.04.019>
- Swallow, K. M., & Wang, Q. (2020). Culture influences how people divide continuous sensory experience into events. *Cognition*, 205, 104450.
<https://doi.org/10.1016/j.cognition.2020.104450>
- Swallow, K. M., Zacks, J. M., & Abrams, R. A. (2009). Event boundaries in perception affect memory encoding and updating. *Journal of Experimental Psychology: General*, 138(2), 236–257. <https://doi.org/10.1037/a0015631>
- 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(6), 651–655. <https://doi.org/10.1038/88486>
- Zacks, J. M., Speer, N. K., & Reynolds, J. R. (2009). Segmentation in reading and film comprehension. *Journal of Experimental Psychology. General*, 138(2), 307–327.
<https://doi.org/10.1037/a0015305>

- 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.
<https://doi.org/10.1037/0033-2909.133.2.273>
- Zacks, J. M., Speer, N. K., Vettel, J. M., & Jacoby, L. L. (2006). Event understanding and memory in healthy aging and dementia of the Alzheimer type. *Psychology and Aging*, 21(3), 466–482. <https://doi.org/10.1037/0882-7974.21.3.466>
- Zacks, J. M., Swallow, K. M., Vettel, J. M., & McAvoy, M. P. (2006). Visual motion and the neural correlates of event perception. *Brain Research*, 1076(1), 150–162.
<https://doi.org/10.1016/j.brainres.2005.12.122>
- Zacks, J. M., Tversky, B., & Iyer, G. (2001). Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology: General*, 130(1), 29–58.
<https://doi.org/10.1037/0096-3445.130.1.29>
- Zalla, T., Labruyère, N., & Georgieff, N. (2013). Perceiving Goals and Actions in Individuals with Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders*, 43(10), 2353–2365. <https://doi.org/10.1007/s10803-013-1784-0>

Appendix A

Aside from the agreement index and surprise index, individual segmentation agreement with a group has also been quantified as the proportion of individual button presses that overlap with a normative boundary (Zalla et al., 2013). In this section, we present limited analyses on this quantification method. For clarity, we refer to this method as *predictive value*. We also include analyses on *detection accuracy*, which incorporates information about the rate at which normative boundaries occurred in the video. We applied the same set of analyses on these measures as the main text, investigating: each measure's ability to distinguish segmentation of one movie from 1) random data and 2) segmentation of a different movie, and 3) the stability of each segmentation measure as sample size increases. Because of the exploratory nature of these measures, we only performed analyses on the commercial-lab dataset.

Both predictive value and detection accuracy involve comparisons between boundaries identified by an individual and boundaries identified by a separate group (*normative boundaries*). Because analyses performed in previous studies binned the data, we defined normative boundaries for these analyses by first calculating the proportion of individuals who pressed a button at every 1s interval (bin) of the movie rather than using group density (see Methods in the main paper). We then extracted the timepoints containing the m highest button press rates. Here, m was defined as the average number of button presses made by the group (Massad et al., 1979). In contrast, Zalla et al. (2013) quantified normative boundaries as the 1s interval time bins that were identified as boundaries by 70% or more of the participants in a separate group of 10 participants.

Calculation of predictive value

This measure quantifies the likelihood that boundaries identified by an individual overlapped with a normative boundary (Zalla et al., 2013). To compute the predictive value, we first counted the number of times an individual pressed a button during a bin that also contained a normative boundary (*hits*) and the number of times they pressed a button during a bin that did not include a normative boundary (*false alarms*). Next, we calculated the proportion of the individual's button presses that were hits (*precision*) versus false alarms (*false discovery rate*) by dividing these counts by the total number of individual button presses. The predictive value was then defined as $z(\text{precision}) - z(\text{false discovery rate})$. It is important to note that, unlike sensitivity, predictive value does not consider the total number of normative boundaries, ignoring both *correct rejections* and normative boundaries that individuals failed to identify (*misses*). To address this issue, we also analyzed detection accuracy.

Calculation of detection accuracy

This measure quantifies how well individuals detect the same boundaries as the group. First, the number of hits and false alarms were calculated in the same manner as for predictive value. However, for detection accuracy, the hit rate was calculated as the number of hits divided by the number of normative boundaries. The false alarm rate was calculated as the number of false alarms divided by the number of non-boundary time bins. Similar to predictive value, detection accuracy was calculated as $z(\text{hit rate}) - z(\text{false alarm rate})$. Unlike predictive value, detection accuracy accounts for the number of boundaries identified by both the individual as well as the likelihood that a bin contained a normative boundary.

Statistical analyses results

We performed the same analyses as in the main text separately for predictive value and detection accuracy. Overall, predictive value (Figure A1-A) and detection accuracy (Figure A1-

B) were sensitive to the structure of the movie. Both measures were able to distinguish segmentation of one movie from random button presses and segmentation of different movies. Same predictive value was overall higher than random for coarse, $t(18932.41) = 172.26, p < .001$, 95% CI [1.57, 1.60] and fine, $t(18932.41) = 122.88, p < .001$, 95% CI [1.10, 1.13] segmentation. Same predictive value was also overall higher than and cross-movie predictive value for coarse, $t(18932.41) = 161.82, p < .001$, 95% CI [1.46, 1.50] and fine, $t(18932.41) = 123.96, p < .001$, 95% CI [1.11, 1.14] segmentation. Further, same detection accuracy was higher than random for coarse, $t(18932) = 187.19, p < .001$, 95% CI [0.80, 0.82] and fine, $t(18932) = 152.95, p < .001$, 95% CI [0.64, 0.66] segmentation. Same detection accuracy was also higher than cross-movie detection accuracy for coarse, $t(18932) = 179.79, p < .001$, 95% CI [0.77, 0.78] and fine, $t(18932) = 155.13, p < .001$, 95% CI [0.65, 0.67] segmentation. Pairwise comparisons revealed that differences between same movie comparisons and both random data and cross-movie comparisons started at sample size of 2 for predictive value, all $ts(18932.41) > 18.50, ps < .001$, smallest 95% CI [.066, 0.81] for coarse segmentation, all $ts(18932.41) > 12.05, ps < .001$, smallest 95% CI [0.37, 0.51] for fine segmentation, and for detection accuracy, all $ts(18932) > 18.32, ps < .001$, smallest 95% CI [0.31, 0.38] for coarse segmentation, all $ts(18932) > 14.81, ps < .001$, smallest 95% CI [0.22, 0.29] for fine segmentation.

Both predictive value and detection accuracy improved and eventually stabilized with the increased sample size. Same predictive value increased as sample size increased for coarse, $t(18932.41) = 34.49, p < .001$, 95% CI [31.68, 35.50] and fine, $t(18932.41) = 29.34, p < .001$, 95% CI [25.97, 29.68] segmentation. Same detection accuracy also increased as sample size increased for coarse, $t(18932) = 36.71, p < .001$, 95% CI [15.91, 17.70] and fine, $t(18932) = 35.72, p < .001$, 95% CI [15.05, 16.79] segmentation. The rate of increase in predictive value

decreased as sample size increased, as evidenced by the significant negative quadratic contrast for coarse, $t(18932.41) = 13.14, p < .001, 95\% \text{ CI } [-30.31, -22.44]$ and fine, $t(18932.41) = 15.85, p < .001, 95\% \text{ CI } [-34.62, -27.00]$ segmentation. Similarly, the rate of increase in detection accuracy also decreased with increasing sample size, as evidenced by the significant negative quadratic contrast for coarse, $t(18932) = 13.84, p < .001, 95\% \text{ CI } [-14.91, -11.21]$ and fine, $t(18932) = 19.07, p < .001, 95\% \text{ CI } [-19.22, -15.63]$ segmentation. Predictive value started to stabilize at sample size 12 for coarse and fine segmentation. Detection accuracy started to stabilize at sample size 14 for coarse and fine segmentation.

Cross-movie predictive value and detection accuracy did not consistently differ from random predictive value and detection accuracy. For predictive value, cross-movie was higher than random for coarse segmentation, $t(18932.41) = 11.26, p < .001, 95\% \text{ CI } [0.085, 0.12]$, but cross-movie was not different than random for fine segmentation, $t(18932.41) = 1.07, p = .28, 95\% \text{ CI } [-0.028, 0.0080]$. For detection accuracy, cross movie was higher than random for coarse segmentation, $t(18932) = 8.30, p < .001, 95\% \text{ CI } [0.027, 0.044]$, but cross movie was lower than random for fine segmentation, $t(18932) = 2.19, p = .029, 95\% \text{ CI } [-0.018, -0.00096]$.

Figure A1. Individual segmentation performance over increasing sample size calculated as: A) predictive value and B) detection accuracy. Small shapes depict individual values calculated from a single bootstrap iteration. Larger shapes depict the average value from all bootstrapping iterations. Error bars represent 95% confidence interval and carets (^) represent the elbows.

Conclusion

Overall, similar to the agreement index and surprise index, predictive value and detection accuracy sensitivity are able to capture segmentation performance that is distinguishable from random button presses and segmentation on different movies. Additionally, predictive value and detection accuracy also improved with increased sample size and eventually stabilized.

Stabilization in both predictive value and detection accuracy also occurred at similar sample sizes compared to the surprise index and agreement index. These findings suggest comparable

performance between agreement index, surprise index, predictive value, and detection accuracy in quantifying individual level segmentation performance.

Appendix B

Table A1

Linear and non-linear model fits (BIC values) for change in various agreement values for various measures over increasing sample size.

Measure	Function	Commercial-lab		Everyday-online	
		Coarse	Fine	Coarse	Fine
Peakiness	Linear	-	12589.82	59571.45	-
	Exponential	-	12715.42	59571.81	-
	Logarithmic	-	12918.82	59566.83	-
	Asymptotic	-	12531.16*	59576.03	-
	Power curve	-	12686.83	59567.19	-
	Yield loss	-	13476.48	59565.84*	-
Peak-to-peak distance	Linear	13637.26	4698.12	45611.13	19457.01
	Exponential	13265.46	4201.39	45606.75	19424.59
	Logarithmic	12712.42	3125.50	45563.75	19282.91
	Asymptotic	12315.88*	2317.42*	45546.68	19174.01
	Power curve	12366.13	2476.85	45557.59	19236.61
	Yield loss	12697.54	2922.57	45537.09*	19159.27*
Agreement index	Linear	-6592.99	-7296.07	-11344.15	-12733.11
	Exponential	-6437.69	-7128.15	-11041.12	-12620.51
	Logarithmic	-7257.05	-8478.62	-12430.95	-13823.93
	Asymptotic	-7300.43	-8743.75	-12468.02	-14127.58
	Power curve	-7085.64	-8216.03	-12150.16	-13672.72
	Yield loss	-7336.21*	-8812.12*	-12516.55*	-14181.85*
Surprise index	Linear	13198.20	15740.91	25174.77	29739.64
	Exponential	13227.42	15856.87	25185.30	29759.92
	Logarithmic	13285.25	15580.96	25191.83	29636.57*
	Asymptotic	13189.83*	15548.75*	25162.00*	29642.75
	Power curve	13226.17	15584.85	25177.86	29647.39
	Yield loss	13397.86	15577.51	25327.81	29642.92

Note. Polynomial contrast for coarse commercial-lab and fine everyday-online peakiness did not suggest evidence of stabilization (see main text), therefore, model fits were not performed for these conditions.

*Lowest BIC values.

Appendix C

Table C1.

Test statistics for correlation between group button press timeseries on all possible video pairs for commercial-lab and everyday-online.

Dataset (df)	Video 1	Video 2	Coarse		Fine	
			<i>r</i>	<i>p</i>	<i>r</i>	<i>p</i>
Commercial-lab (565)	3Iron	Corn Island	0.0011	.98	-0.0053	.90
Everyday-online (272)	Bed	Coffee	0.075	.22	0.0031	.96
		Dishes	0.080	.18	0.15	.016
		Laundry	0.12	.039	0.14	.023
	Coffee	Dishes	0.16	.007	0.13	.033
		Laundry	.031	.61	0.15	.015
	Dishes	Laundry	.063	.30	0.12	.052

References

- Massad, C. M., Hubbard, M., & Newtonson, D. (1979). Selective perception of events. *Journal of Experimental Social Psychology*, 15(6), 513–532. [https://doi.org/10.1016/0022-1031\(79\)90049-0](https://doi.org/10.1016/0022-1031(79)90049-0)
- Zalla, T., Labruière, N., & Georgieff, N. (2013). Perceiving Goals and Actions in Individuals with Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders*, 43(10), 2353–2365. <https://doi.org/10.1007/s10803-013-1784-0>