

Synchronous high-amplitude co-fluctuations of functional brain networks during movie-watching

Jacob C. Tanner^{1,2}, Joshua Faskowitz^{3,4}, Lisa Byrge³, Daniel

P. Kennedy^{1,3,4}, Olaf Sporns^{1,3–5}, and Richard F. Betzel^{1,3–5*}

¹*Cognitive Science Program, ²School of Informatics, Computing, and Engineering,*

³*Department of Psychological and Brain Sciences, ⁴Program in Neuroscience,*

⁵*Network Science Institute, Indiana University, Bloomington, IN 47405*

(Dated: June 30, 2022)

Recent studies have shown that functional connectivity can be decomposed into its exact framewise contributions, revealing short-lived, infrequent, and high-amplitude time points referred to as “events.” Although events contribute disproportionately to the time-averaged connectivity pattern, improve identifiability and brain-behavior associations, and differences in their expression have been linked to endogenous hormonal fluctuations and autism, their origins remain unclear. Here, we address this question using two independently-acquired imaging datasets in which participants passively watched movies. We find that events synchronize across individuals and based on the level of synchronization, can be categorized into three distinct classes: those that synchronize at the boundaries between movies, those that synchronize during movies, and those that do not synchronize at all. We find that boundary events, compared to the other categories, exhibit greater amplitude, distinct co-fluctuation patterns, and temporal propagation. We show that underlying boundary events is a specific mode of co-fluctuation involving the activation of control and salience systems alongside the deactivation of visual systems. Finally, we find a strong positive relationship between the similarity of time-locked co-fluctuation patterns and the propensity for those time-locked frames to involve synchronous events. In other words, the inter-subject similarity of brain networks during passive movie-watching is strongly related both to the moment in the movie when they occur, and the likelihood that this moment evoked an event. Collectively, our results suggest that the spatiotemporal properties of events during passive movie-watching are non-random and locked to time-varying stimuli.

INTRODUCTION

The human brain is fundamentally a complex network comprised of anatomically connected neural elements [1–3]. This physical network constrains dynamical interactions between brain regions, inducing statistical dependencies in the activity of distant brain regions, i.e. functional connectivity (FC) [4–6]. A growing number of studies have focused on characterizing the architectural features of FC [7, 8] and linking inter-individual differences to cognition [9–11], disease [12, 13], and development [14, 15].

Recent methodological advances have made it possible to precisely decompose FC into its framewise contributions [16, 17]. This “edge-centric” approach yields time-varying estimates of the co-fluctuation magnitude and valence for every pair of brain regions (edge). Previous studies have shown that, collectively, edges exhibit bursty behavior, such that long periods of quiescence are punctuated by brief, high-amplitude “events” in which many edges simultaneously and strongly co-fluctuate with one another [16, 18–24]. The whole-brain co-fluctuation patterns expressed during events are closely related to the static (time-averaged) FC [16], improve subject identification and brain-behavior correlations [16], are individualized [19], can be linked to endogenous hormone fluc-

tuations [18], clinical status [25], and memory processes [26], and are shaped by the underlying anatomical connectivity [27].

However, the principles that determine the timing of events are undisclosed. The first “edge-centric” study showed that the whole-brain co-fluctuation amplitude is correlated during movie-watching but not at rest [16], suggesting that audiovisual stimuli may induce correlations in the timing of events. However, the spatiotemporal structure of events in naturalistic stimuli was not investigated further. In contrast, several recent studies have called into question whether events have any neurocognitive underpinnings at all, suggesting that events are mathematical consequences of systems with modular and stationary correlation structure [28, 29].

Here, we seek to uncover potential drivers of event timing, and analyze their inter-subject synchronization during movie-watching. As in [16], we find that the timing of events synchronizes across participants watching the same movie. We replicate this finding using two independently-acquired, -processed, and -parcellated datasets, representing eight separate movies (each movie consisting of a small number of movie scenes, or trailers). We analyze events further, grouping them into three distinct categories: events that synchronize at the boundaries between movies, those that synchronize within the body of a movie, and those that are asynchronous. We focus further on boundary events, showing that they exhibit distinct overall co-fluctuation amplitudes, co-fluctuation pattern, and temporal structure. We show

* rbetzel@indiana.edu

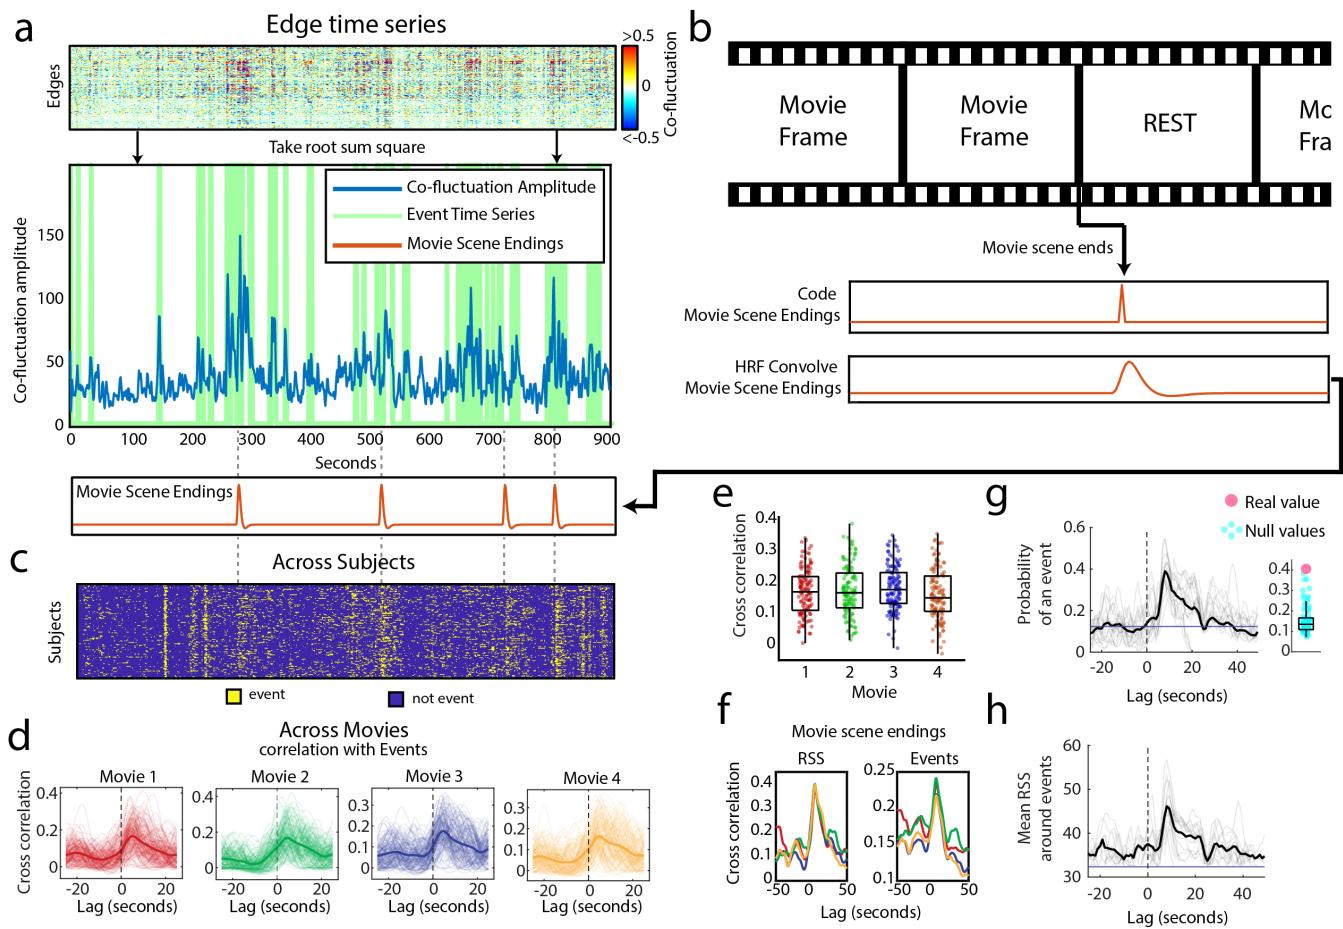


FIG. 1. Events occurs at the end of movie scenes. (a) We compute edge time series (top), take the root sum square (RSS) to derive the co-fluctuation amplitude time series. Some of the peaks in co-fluctuation amplitude are determined to be events using a previously published statistical method [30]. The green line represents the binary event time series created from this method. (b) Schematic showing that movie scene endings were manually coded by two individuals authors (JT and RB), and then convolved with the hemodynamic response function (HRF). [Stills from movie scenes were removed for biorxiv submission, but will be in the published version.] Example of HRF-convolved movie scene endings (orange) plotted below the relevant co-fluctuation amplitude time series. Note how many of the events from the event time series line up with movie scene endings. (c) Here, we “stack” the event time series across all subjects in a representative scan to illustrate that subjects tend to have events at the same moment in time, and that the moments when subjects synchronize their events often correspond to movie scene endings. This result holds across choice of FDR. (d) Here, we show this alignment statistically by correlating each subject’s event time series with the HRF-convolved movie scene endings while circularly shifting (referred to in figure as “Lag”) the movie scene endings in either direction. We plot this separately for each of four movies. The bold lines represent the mean across movie scene endings. (e) Boxplots showing the distribution of cross correlation values for the same time lag. (f) Plot showing the time-lagged correlation between the indices of movie scene endings, and mean RSS across subjects (first plot) per scan. Plot showing the time-lagged correlation between the indices of movie scene endings, and the number of events per time point (second plot), with one line for each movie. (g) Plot showing the probability of an event occurring near a movie scene ending (movie scene endings centered on zero). Black line is the probability over all scans, and the grey lines are the probability per scan. (h) Plot showing the mean co-fluctuation amplitude (RSS) near movie scene endings. Black line is the mean across all scans, and the grey lines are the mean per scan.

that boundary events are also underpinned by a distinct pattern of activity that involves visual processing, attentional, and control systems. Finally, we find a strong positive relationship between the similarity of time-locked co-fluctuation patterns and the propensity for those time-locked frames to involve synchronous events. In other words, we find that the inter-subject similarity of brain

networks during passive movie-watching is strongly related both to the moment in the movie when they occur, and the likelihood that this moment evoked an event. This work sets the stage for future studies to take advantage of the shared structure of events during naturalistic stimuli, and to explore other relevant features in such stimuli that elicit events.

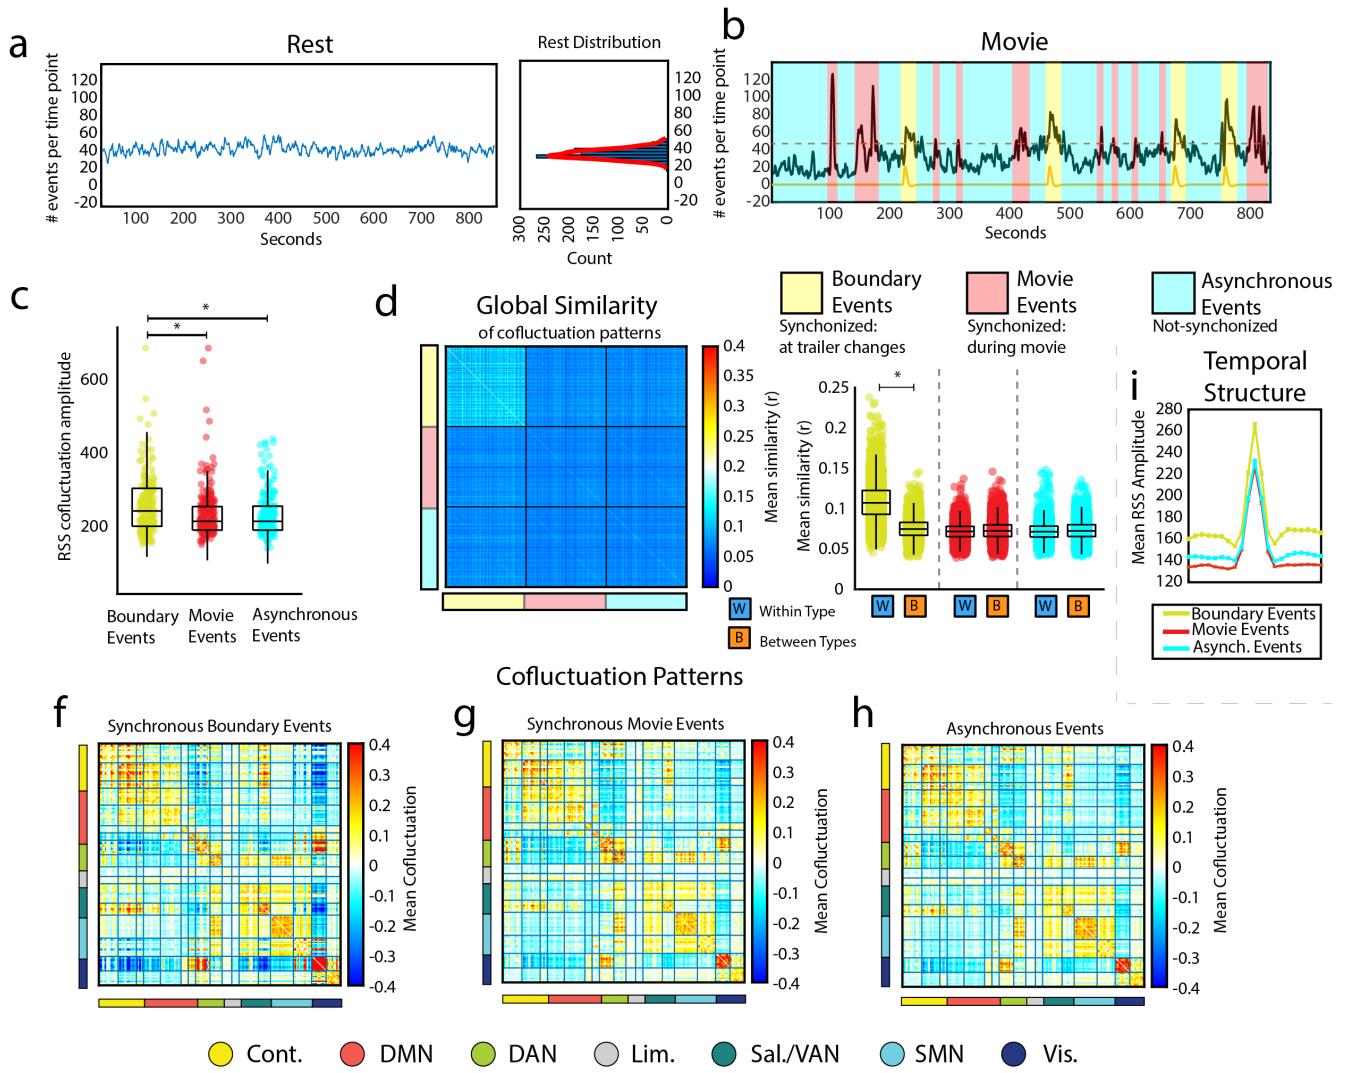


FIG. 2. Boundary events exhibit distinct spatiotemporal characteristics. In the main text we described methods for detecting events and partitioning frames into categories based on when they occurred during the movie and the number of subjects exhibiting temporally coincident events. To detect coincident (synchronous) events, we compared event counts at each movie frame with a null distribution estimated from resting-state data. (a) Resting-state event count distribution (averaged across time) and event counts per frame across time. (b) Event counts for movie data, with frames labeled according to whether events were synchronous and occurred at movie boundaries (boundary event) or within a movie (movie event) or were not synchronized across individuals (asynchronous). Note that given the size of this cohort ($n = 129$), at least one participant exhibited an event at every frame. In principle, there could exist frames where no participant exhibited an event, necessitating a fourth category of time point (non-events). (c) Co-fluctuation amplitudes grouped by event type. (d) Global similarity of whole-brain co-fluctuation patterns grouped by event type (within each type, subject order is identical; mean across scans; see methods for details). (e) Given each of the three event types, its self similarity and between-type similarity. Panels f-h depict mean co-fluctuation pattern for each of the three event types. (i) Typical temporal profile (co-fluctuation amplitude) for each event type (all events locked to their respective peaks).

RESULTS

Previous studies have identified brief, high-amplitude co-fluctuations in resting-state fMRI [16, 27, 30]. The origins of these “events” are unclear. In this section, we address this question directly using naturalistic stimuli (movie-watching) data from the Human Connectome Project [31] and a second, independently-acquired

dataset [32]. In the following sections, we describe results of these analyses. Specifically, we show that events synchronize across subjects at certain parts of the movie. We provide a tripartite classification scheme for the events that occur while subjects watch movies, as follows. Synchronous boundary events occur at the end of movie scenes (and beginning of rest blocks). Synchronous movie events occur during the movie, and asynchronous events

do not synchronize across subjects. We show that boundary events exhibit reproducible and distinct spatiotemporal characteristics, as well as a unique activation pattern, and finally we observe a strong positive relationship between the similarity of time-locked co-fluctuation patterns and the propensity for those time-locked frames to involve synchronous events.

Events occur at the end of movie scenes

We first aimed to determine whether events synchronize during movie watching and, if so, when in the movies these synchronous events occurred. To investigate this, we estimated edge time series for every subject in every movie, calculated the co-fluctuation amplitude (root sum square of all edges' at every time point; RSS), and used a previously-described algorithm to detect event frames whose RSS exceeded that of a null distribution [30]. This resulted in an “event time series” (Fig. 1a). Two of the authors (JT and RB) without reference to event time series, then manually coded and hemodynamically convolved a time series to indicate the frames at which each movie scene ended and rest blocks began (Fig. 1b; see Methods for details on the structure of this naturalistic data, and our coding protocol). Authors were blinded to brain data during coding procedure. We then “stacked” event time series and discovered an alignment between frames when many subjects had events and the indices of movie scene endings (Fig. 1c).

To confirm this observation statistically, we then performed several complementary tests. First, we correlated this “event time series” with the time series of hemodynamically convolved indices of movie scene endings (Fig. 1d). We found that event time series were correlated with the movie-scene endings (Fig. 1e; mean correlation per movie $r = 0.17 \pm 0.07$, $r = 0.18 \pm 0.08$, $r = 0.18 \pm 0.07$, $r = 0.16 \pm 0.08$ respectively). When we averaged event time series across subjects to obtain a group-level and pseudo-continuous estimate of event time series, we found that the fraction of subjects exhibiting an event at any given instance was maximally correlated with movie-scene endings at a lag of 5 seconds (Fig. 1f; Correlation at lag: $r = 0.35$, $r = 0.39$, $r = 0.39$, $r = 0.40$ respectively, all p-values $p < 10^{-15}$). This effect is also evident in the raw and unthresholded co-fluctuation amplitude (Fig. 1f; Correlation with mean RSS across subjects at lag: $r = 0.40$, $r = 0.45$, $r = 0.45$, $r = 0.48$ respectively, all p-values $p < 10^{-15}$). Finally, we found that these effects resolved to an increased probability of having an event at a lag of 5 seconds from movie-scene endings (Fig. 1g; probability of having an event $Pr(E) = 0.41$; compared with probability at other nearby time points; two-sample t -test $p < 10^{-15}$).

Importantly, this result was replicated in an independently collected, processed, and parcellated data set. Instead of movie scenes, this data set presented participants with movie trailers. JT and RB again manually coded

and hemodynamically convolved movie trailer endings, and compared these with the stacked event time series (Fig. S1a). We then computed the number of events within a window of 10 seconds on either side of each movie trailer ending and compared this with a null model where this window was circularly shifted 100 times. We found that there were more events near movie trailer endings than should be expected by chance (Fig. S1b,c; p -values for each movie: $p = 4.40 \times 10^{-3}$, $p = 1.25 \times 10^{-5}$, $p = 2.30 \times 10^{-2}$, $p = 1.26 \times 10^{-2}$). Note that the 10 second window was selected to include the peak lag of 5 seconds, but also be broad enough to account for inter-individual differences in timing.

Taken together, these results suggest that the ending of movies scenes (or trailers) correspond to time points when events are likely to occur in many subjects.

Events synchronize within movie scenes

In the previous section, we demonstrated that the end of movie scenes coincide with high-amplitude co-fluctuations and are often categorized as events. It is unclear, however, if synchronized co-fluctuations also occur *within* individual movies. Here, we develop a statistical test to determine whether there are other points in time, specifically during the movies, when events occur coincidentally across subjects

To address this question, we calculated the group-averaged event time series by summing across participant-level event time series. Each frame of this group time series indicated the number of subjects exhibiting an event at that instant. To identify temporally coincident events, we compared the observed group-averaged event time series with a null distribution estimated using resting-state scans rather than movie-watching (Fig. 2a). In the resting data, because subjects' fMRI BOLD timecourses are not temporally locked to a movie stimulus, any observed synchronization is due to chance fluctuations. We also created a null model by circularly shifting event time series while subjects watched movies in order to maintain event number and relative timing while breaking alignment with the movie stimulus, but we ultimately chose the rest null model after determining it was more conservative (Fig. S6; two-sample t -test $p < 10^{-15}$).

Comparing the observed event time series with the null distribution of subjects at rest, we found evidence of synchronous events occurring, as expected, at movie boundaries (we refer to these as “synchronous boundary events” or simply “boundary events”; but also within movie scenes (we refer to these as “synchronous movie events” or “movie events” (we controlled for multiple comparisons by fixing the false discovery rate to $q = 0.05$, $p_{adj} = 0.03$). We also defined a third category of time point – those for which the number of observed events across subjects was consistent with or less than that of the null distribution (referred to as “asynchronous events”; Fig. 2b).

Boundary events exhibit reproducible and distinct spatiotemporal characteristics

Given this tripartite classification scheme, we asked whether the three event types exhibited distinguishable characteristics. First, we tested whether boundary events exhibited dissimilar amplitudes than the other types. We found that, on average, boundary events had greater RSS than both movie and asynchronous events (Fig. 2c; two-sample t -test $p < 10^{-15}$).

Next, we asked whether the whole-brain co-fluctuation patterns expressed during events were dissimilar across event types. To do this, we calculated the similarity (bivariate linear correlation) between all pairs of detected events. Then, we averaged these scores by subject and event type. For every pair of subjects, this yielded three similarity scores – one per each event type (Fig. 2d). For a given event type (indicated by the blocks shown Fig. 2d), we compared within- and between-event type similarity. Intuitively, this corresponds to comparing elements in the diagonal blocks with the off-diagonal blocks. We found that “boundary events” were more similar to other boundary events than to other event types (Fig. 2e; two-sample t -test $p < 10^{-15}$), suggesting that boundary events represent a distinct category of events, dissimilar from the other two, which are largely indistinguishable from one another.

Interestingly, we also found strong subject-level effects across event types. That is, irrespective of which types of events are being compared, subjects in most scans maintained a high level of identifiability (Fig. S4; synchronous movie events did not pass identifiability tests for two scans). To confirm this, we directly compared the within-subject similarity to between subject similarity and found that it was statistically significant for nearly every block in the similarity matrix (two-sample t -test $p < 8.66 \times 10^{-4}$). More importantly, we found that identifiability is actually higher in boundary events than any other event type (Fig. S4; two-sample t -test $p < 10^{-15}$).

We then examined the co-fluctuation patterns in greater detail. For each event type, we calculated its mean co-fluctuation pattern across all subjects and scans (Fig. 2f-h). In general, we found that the mean co-fluctuation patterns shared a relatively high baseline level of similarity (mean similarity $r = 0.91 \pm 0.07$). However, we observed subtle yet systematic deviations between event types. Notably, boundary events displayed stronger co-fluctuation within many systems, specifically interactions of the central visual system with salience and control networks (spin test, false discovery rate fixed at $q = 0.05$, $p_{adj} = 1.73 \times 10^{-4}$; Fig. S3g).

We then examined the local temporal structure around events. Briefly, this involved temporally aligning every instance of all event types to their respective peaks (Fig. 2i). We found, in agreement with the previous RSS analysis, that at their peaks, event types were stratified based on their amplitudes, with boundary events exhibit-

ing greater RSS than the other event types. Interestingly, we also found off-peak effects. That is, the frames before and after peaks, exhibited a similar relationship, suggesting that the frames immediately before and after boundary events also exhibited greater RSS than other event types (searched 10 frames on either side of event peak, two-sample t -test Bonferroni corrected $p_{adj} = 4.76 \times 10^{-5}$ for all 21 frames.).

Critically, all of these effects were directly replicated in a second dataset (Fig. S2). In addition, we found that the mean co-fluctuation pattern for each event type also replicated. We also found that time-averaged FC is mostly driven by synchronous movie events in the Human Connectome Project data, and that this effect is likely due to synchronous movie events containing more frames per scan than other event types (Fig. S5; two-sample t -test $p < 10^{-15}$) but this result did not replicate in the Indiana University data set (Fig. S5e-g).

Collectively, these results suggest that timing of events and their synchronicity across individuals shapes their amplitude, co-fluctuation pattern, and temporal evolution.

Activation patterns during boundary events

To this point, we have analyzed co-fluctuations, which are defined at the level of edges (pairs of brain regions). While the calculation of edge-time series from activations is straightforward, obtaining the activation pattern underpinning a given co-fluctuation matrix is not. This is because every co-fluctuation matrix could have been generated by two different patterns of activity that differ in their sign at every node. For example, the product of two positively-valued nodal activations yields a positive co-fluctuation. However, had the activation amplitude been identical but negatively-valued (a deactivation) we would still obtain an identical co-fluctuation. Here, we investigate the activation patterns that underlie co-fluctuations during events, focusing specifically on their configuration at movie boundaries.

First, we identified brain regions whose activity significantly increased or decreased during boundary events compared to all other time points (Fig. 3a,b; as many as 227 scans significant per region, Bonferroni corrected for number of scans tested $p_{adj} = 8.77 \times 10^{-5}$). Interestingly, we found that these correlations were highly system-specific, with increased activation among control b and salience/ventral attention nodes b and decreased activation in central visual (spin test; Bonferroni corrected $p_{adj} = 2.90 \times 10^{-3}$; Fig 3b & Fig S7). We also expanded our activation analysis to include subcortical and cerebellar regions of interest [33–36]. We found decreased activation in thalamic regions that are associated, based on their functional connectivity, with cortical control networks (Fig. 3c,e). We also found changes in activation patterns among cerebellar regions, specifically those associated with default mode, control, and dorsal

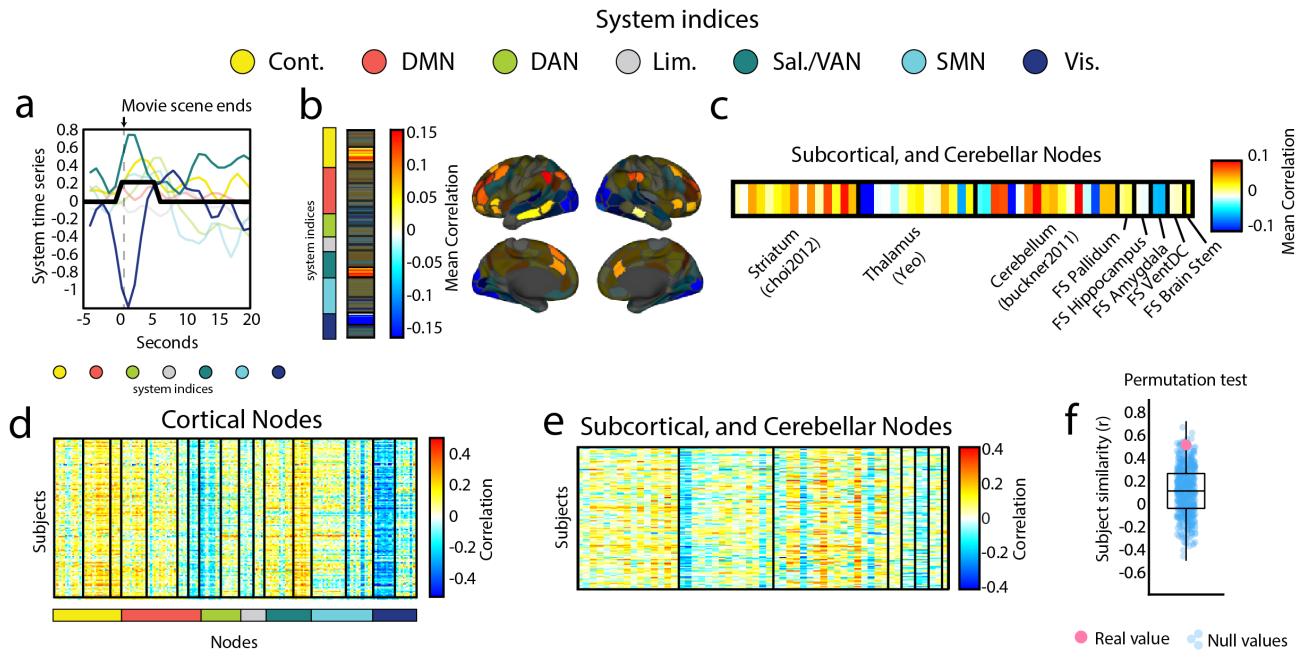


FIG. 3. Boundary events correspond to activation of control, and salience systems, and deactivation of visual systems. (a) Plot of mean system activity for a window of time. The black line represents an index of boundary events. Notice the decrease in visual systems activation and an increase in salience and control. (b) Mean correlation values (across subjects) per region organized according to brain system (indices of brain system are color coded on the left of the plot). A space-preserving spin test was used to test if any of these brain systems tended to have a higher or lower concentration of correlation values than expected by chance. Control b, salience b, and central visual systems passed this test ($p < 2.90 \times 10^{-3}$). Systems that did not pass this test are shown shaded in grey. (c) Mean correlation values (across subjects) per subcortical and cerebellar region. Regions are organized mainly into parcellations for striatum, thalamus, and cerebellum. (d & e) Plot of the correlation values per subject and node (cortical, and subcortical respectively). (f) The pattern of correlation values across nodes is significantly similar across subjects, when compared with a time-shifted null model where we circularly shifted the boundary event indices before correlating them with the nodal time series ($p = 0.03$).

attention networks (Fig. 3c,e; as many as 151 time series significant per region; Bonferroni corrected for number of scans tested $p_{adj} = 8.77 \times 10^{-5}$).

Importantly, we found that these correlation patterns were highly stable across individuals, both at the level of cortex as well as subcortex and cerebellum (Fig. 3d,e; mean correlation $r = 0.51 \pm 0.18$, compared against correlation patterns found when using randomly permuted boundary event indices; $p < 0.05$).

In summary, these results suggest that, although co-fluctuation patterns could, in principle, arise from a degenerate set of activity patterns, in the case of boundary events, they are underpinned by a single mode of activation and deactivation, involving a specific constellation of brain systems and regions.

Time-locked movie events exhibit distinct co-fluctuation patterns

To this point we have shown that movie boundaries tend to elicit high-amplitude co-fluctuations that are synchronous across individuals and exhibit distinct spa-

tiotemporal characteristics that distinguish them from other types of events. On one hand, perhaps this is to be expected; movie boundaries correspond to periods during the movie with similar audiovisual features, e.g. black or darkened screen and the absence of auditory stimuli. However, we have largely overlooked synchronous events that occur within the context of a single movie segment—the so-called “synchronous movie events.” In this section, we take advantage of the repeated presentation of the same movie scene across multiple scans to characterize the stable properties of synchronous and time-locked events across that occur within the same movie.

To address this question, we first extracted frames corresponding to the same movie scene presented across four independent scans (a total of 516 viewings; Fig. 4a). Using the previously-described algorithm for detecting synchronous movie events, we identified and focused on three peaks (frames) in the event time series (Fig. 4b). These peaks also corresponded to points in time when subjects' co-fluctuation patterns were highly similar (Fig. 4k) and the 129 individuals who watched the movie scene four times reliably exhibited events upon each viewing (Fig. 4c).

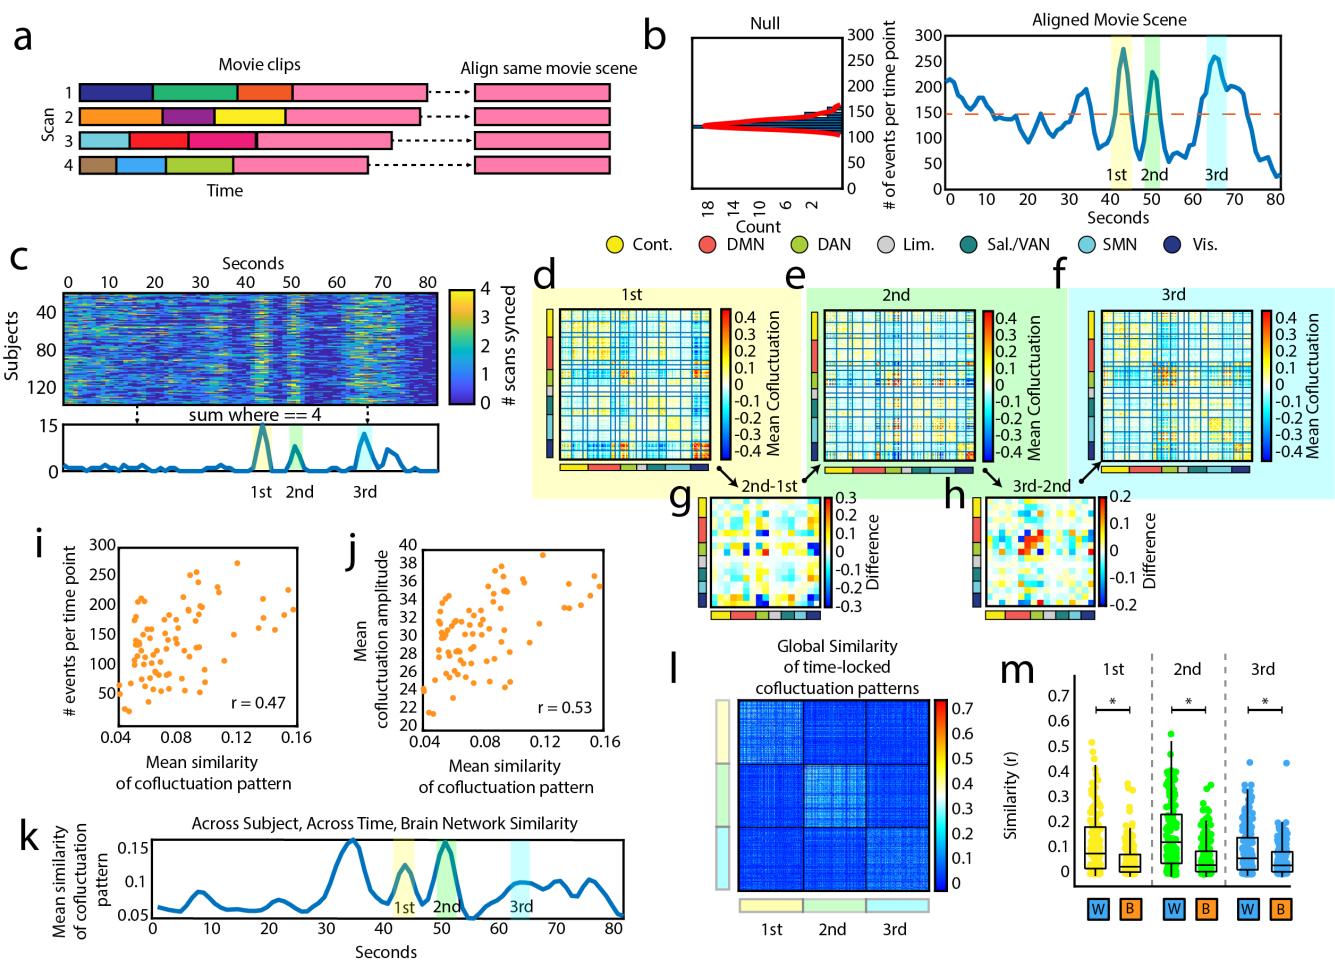


FIG. 4. Time-locked movie events have similar co-fluctuation patterns. (a) Schematic showing that all four scans ended with the same movie clip, and how we temporally aligned the time series data from the four different presentations of this clip. (b) Histogram on the left shows the null distribution of subjects at rest. The plot on the right shows the number of events per time point in blue, with the dashed line indicating when this number is significantly greater than chance. We highlighted and numbered three peaks that we will focus on in further analyses. (c) This plot shows the 129 subjects who watched this movie scene four times. The values in this matrix show how many times a given subject had an event during each frame, where four is the maximum number possible. Three areas stand out in this matrix, and as the plot below the matrix shows, these areas correspond to frames where as many as 15 subjects had an event every time they watched this movie scene. (d,e & f) These three matrices show the mean co-fluctuation patterns (across subjects) at the frames highlighted in other plots. (g) This matrix shows the main differences between the 1st peak and the 2nd peak in their system by system interactions. (h) This matrix shows the main differences between the 2nd peak and the 3rd peak in their system by system interactions. (i) Plot showing the positive linear relationship between the number of events per time point, and the mean similarity of co-fluctuation patterns (across subjects). (j) Plot showing the positive linear relationship between the mean co-fluctuation amplitude (across subjects), and the mean similarity of co-fluctuation patterns (across subjects). (k) Mean similarity of co-fluctuation patterns (across subjects) plotted across time. Three frames/peaks are highlighted. (l) Matrix showing the similarity of co-fluctuation patterns for the three time-locked frames. (m) Boxplots showing the data in the previous figure divided into similarity within and between the highlighted peaks (1st, 2nd, & 3rd). Each peak is significantly more similar to itself than to the other peaks.

Next, we examined the co-fluctuation patterns during these three peaks (Fig. 4d,e,f). Interestingly, we found that these co-fluctuation patterns were dissimilar to one another, driven by system-specific differences involving the dorsal attention and temporo-parietal networks (Fig. 4g,h). This pattern is evident across subjects (Fig. 4l), where co-fluctuation patterns could be reliably distinguished from one another based on which peak they

corresponded to (Fig. 4l,m; $p < 10^{-15}$).

Upon reflection, this pattern – wherein peaks were dissociable from one another based on co-fluctuation patterns and event status – represents the tail of a more general phenomenon in which inter-subject similarity is positively associated with both event times and co-fluctuation amplitude ($r = 0.47$, $p = 8.81 \times 10^{-6}$; $r = 0.53$, $p = 2.64 \times 10^{-7}$, respectively; Fig. 4i,j).

In summary, these results posit a link between the similarity of co-fluctuation patterns across subjects and the propensity for those patterns to occur synchronously across individuals. Further, it reveals repeatable, fine-scale, and spatiotemporally dissociable event structure within movie scenes.

DISCUSSION

Previous studies have shown that spontaneous resting-state fMRI data exhibited “events” – brief, high-amplitude, and network-level co-fluctuations. The drivers of events remain unclear. Here, we investigate this question using data acquired while subjects passively viewed movie scenes (or trailers) in the scanner. We find evidence for three types of events: boundary events, movie events, and asynchronous events. We show that boundary events are likely to occur across subjects during the boundaries between sequential movie scenes (or trailers). Further, we show that these boundary events are of greater amplitude than non-boundary events, exhibit similar co-fluctuation patterns across subjects, and follow distinct temporal trajectories. Interestingly, boundary events also carry the most subject-specific information of any event type. At the nodal level, boundary events are underpinned by a distinct pattern of activity, corresponding to deactivation of visual areas and the activation of control and ventral attention networks. In addition to showing that boundary events are distinct from other event types in many of their features, we show that movie events (synchronous events that occur during the movie) share features that are time-locked to the movies that subjects are watching. That is, when a movie scene elicits an event in many subjects, the whole-brain co-fluctuation patterns of these subjects tend to be similar. Additionally, we found that this specific pattern is a feature of a more general phenomenon wherein inter-subject synchrony in events is associated with higher inter-subject similarity in co-fluctuation patterns.

Movie boundaries reliably elicit high-amplitude co-fluctuations

Events explain a large fraction of variance in static FC and can improve subject identification and brain-behavior correlations [16, 37]. They are individualized [30], are correlated with quotidian variability in hormones across the menstrual cycle [18], can be used to distinguish healthy controls from individuals with autism [25], and can arise spontaneously in networks whose underlying anatomical structure is modular [27]. However, the timing of events – why they occur when they do – is not understood. Resting-state datasets—the focus of most previous edge time series papers—are poorly suited for addressing this question; the unconstrained nature of rest and the absence of any temporal correlations across

participants makes it difficult to establish possible event drivers.

Here, we leverage naturalistic movie-watching data [38–40] to begin addressing this question. We find evidence that, while events occur at all times during movies, they reliably and synchronously occur at the end of movie scenes. This observation is consistent with several recent studies, which link activations and intersubject synchronization at the boundaries of movies to memory [41, 42], attention processes [43], and general cognition [44].

Why events reliably occur at boundaries remains unclear. One possibility is that boundary events serve as brain-based markers for the discretization of experience [45, 46]. That is, that boundary events signify the ending of one subjectively and perceptually defined segment and the beginning of another [47–49]. Although appealing, directly testing this hypothesis necessitates the collection of additional data to confirm that movie boundaries align with participants’ subjectively defined segment boundaries [50].

Another possibility is that synchronous events arise from perturbations to a dynamical system. Many studies have shown that the temporal evolution of brain activity can be modeled as a dynamical system. If these systems are broadly similar across individuals, we might expect that driving them with the same time-varying inputs (audiovisual stimuli) will yield similar response profiles. This partially explains intersubject correlations during naturalistic paradigms. We speculate that the ends of movies, which are often typified by lower levels of luminance and sound volume, represent particularly evocative or arousing stimuli so that the shared response is brain-wide and initiates an event. Indeed, previous studies have shown that state of arousal is a powerful modulator of brain activity [51, 52] and connectivity [53, 54], suggesting that it may also play a role in event induction and should be investigated further in future studies.

We note that these hypotheses – experience segmentation and perturbations to dynamical systems – represent only two possible explanations for boundary events and are, themselves, not mutually exclusive. Additionally, we note that while movie scene boundaries are factors that explain some synchronous high-amplitude co-fluctuations, we also observe synchronous events at other points within movies. The underlying mechanisms support “movie events” may be similar to boundary events, although stimuli during the movie are much richer and complex [55]. Future studies should focus on linking this second category of events with other features, e.g. movie annotations or eye-tracking data as a proxy for arousal.

Event timing is not fully explained by sampling variability

Several studies have suggested that events arise not from any neurocognitively meaningful mechanism, but simply from sampling variability around a fixed correla-

tion structure [28, 29]. That is, given a static FC matrix, nodal and edge time series must have a specific set of properties, including events. This is an important concern, and one that echoes other ongoing and unresolved debates in the network and cognitive neuroscience literature [56–58].

Our results speak directly to this controversy. Specifically, if events arise only due to stochastic fluctuations, then their occurrences should be independent and uncorrelated across individuals. However, we find that this is not the case and that, in line with previous studies [16], the timing of events within a scan session is correlated across individuals. These events cannot obviously be attributed to “rest blocks” that are interspersed between movie scenes, as we also find evidence of correlated events within movies.

These observations contribute to and enrich the ongoing dialogue surrounding exactly what features of brain networks are stationary *versus* non-stationary. Future studies – both empirical and *in silico* – should continue to investigate these and related questions.

Future Directions

There are several ways that the results of this study could be extended in the future. For instance, we examined naturalistic stimuli (passive movie-watching). We found robust evidence of inter-subject correlations, specifically the timing of events, both during and at the end of movie scenes. While it is clear that subjects synchronize events, the precise causes remain unclear. One way to gain insight into these causes involves leveraging the movie annotations that accompany the imaging data. That is, to link the timing of events, synchronized or otherwise, to the timing of particularly salient features (e.g. objects and actions) occurring in the movie [59]. Additionally, future studies with more specified and explicit task design could also be performed to adjudicate between competing hypotheses about the origins of events. For example, online experience sampling during acquisition could be used to better understand the subjective experiences that co-occur with events [60].

Another way to extend these results would be to examine multi-modal recordings made at, potentially, faster timescales, e.g. scalp or intracranial EEG or MEG [61–63]. Our work here, and related studies, have focused almost exclusively on fMRI data. However, fMRI has a number of known limitations, most notably it exhibits relatively poor temporal resolution, which may obscure rapidly occurring neural processes, e.g. fine-grained temporal structure that occurs just prior to or following an event [64–66]. Future studies should investigate inter-synchrony of events using imaging modalities that are acquired at faster rates.

Here, we focus on empirical recordings. However, recent studies have shown that events can be observed in synthetic data generated by dynamic models [27]. This

opens up several avenues for future work, including artificially “stimulating” *in silico* brains to induce events and, if we have two independent simulations, whether simultaneous stimulation leads them to synchronize [67]. The simulation-based analyses also make it possible to investigate and test related hypotheses using data from non-human subjects, e.g. macaque [68].

One of our key findings was that events are temporally coincident during movie-watching and especially near the boundaries between movie trailers. However, participants consistently exhibit events throughout the scan, even if the timing is not coincident with other individuals. What are the origins of these “asynchronous events?” One possibility is that they reflect an intrinsic mode of rest that is decoupled from the ongoing task/movie stimuli [69]. Future studies, likely with dedicated experimentation, are needed to clarify the distinction between asynchronous and synchronous events.

Another noteworthy point concerns how events were defined. Here, a time point was considered an event if global – i.e. whole-brain – fluctuations exceeded a statistical criterion. We found that these global events reliably occurred during movie boundaries, which correspond to large shifts in movie content. Future studies should investigate locally-defined events – e.g. at the level of brain systems – which may be linked to more subtle variation in movie features.

Finally, while we find evidence of synchronous events, we note that synchronization does not imply that all subjects experienced an event at the same moment (just that a sufficiently large number did). An important open question is whether those subjects that did not experience an event are distinguished from their peers along other dimensions as well, e.g. cognitive, clinical, demographic profiles. For example, perhaps subjects that experienced an event have better recall of related movie scenes than subjects that did not.

Limitations

There are several limitations of the current study. One notable issue is the use of fMRI data to infer brain activity. Functional imaging data is fundamentally an indirect (and slow) measure of the hemodynamic response to population-level activity. Follow-up studies should aim to reproduce these results using more direct assessments of activity, e.g. intracranial EEG.

Additionally, there are several limitations associated with the movie data themselves. The first concerns the subjectivity with which individuals experience the movie stimuli. While the stimuli are identical across individuals, their experiences may not be, e.g. participants may attend to different movie features at the same time point or draw on their personal history, leading to dissimilar experience. Without subjective post-scan reports or incorporation of additional data modalities, e.g. eye-tracking, it remains challenging to uncover the drivers of synchro-

nized events.

Furthermore, it is important to clarify that while previous studies have primarily focused on events during the resting-state, we focus here on events detected during passive movie-watching. While the statistical approach used to detect periods of high-amplitude co-fluctuations is identical, it remains unclear whether resting and movie events share common neurocognitive underpinnings. Future studies should investigate their differences and commonalities in more detail.

Lastly, there remain challenges related to data processing and analysis. Specifically, the procedures for disambiguating activations (first-order effects) from co-fluctuations (second-order effects) during movie-watching are not well defined. In blocked or event-related task design, regressors can be constructed to orthogonalize time series with respect to activations [70], making it possible to effectively remove activations as confounds prior to estimating connectivity or co-fluctuations. However, with movie-watching data, stimuli are presented in a naturalistic way, meaning that many stimuli overlap and co-occur, e.g. the presence of a face is often correlated with speech, making it difficult to control for all possible stimuli. Moreover, the stimuli are presented in a continuous stream, similarly making it difficult to construct appropriate regressors. This remains an active and contentious area of research. Future methodological studies will address this potential issue. One particular concern unique to our findings is the presence of “rest blocks” in the HCP dataset. An alternative interpretation of our results is that because the transition from a “movie state” to “rest state” is accompanied by a short pause, it is the pause itself that induces a non-stationarity in the activations, yielding a boundary event.

Conclusion

In conclusion, this study represents an analysis of the synchronization of high-amplitude co-fluctuations, or “events”, while participants are watching the same movie in two separately collected, processed and parcellated data sets. Our analysis suggests that events during movie watching can be divided into three categories based upon their synchronization. Synchronous boundary events occur at the end of movie scenes, or trailers, and are distinct from the other event types in overall co-fluctuation pattern, amplitude and temporal structure, as well as carrying the most subject-specific information. Synchronous movie events occur during the movie and their time-locked co-fluctuation patterns are highly similar across subjects. In fact, we found a general relationship between the degree of event synchronization across subjects and the across subject similarity in co-fluctuation pattern. These two types of synchronous events suggests that events can be elicited by environmental stimuli. Asynchronous events also occur

throughout the movie, suggesting an individualized element to the inducement of events.

MATERIALS AND METHODS

Human Connectome Project Data

The Human Connectome Project (HCP) 7T dataset [31] consists of structural magnetic resonance imaging (T1w), resting state functional magnetic resonance imaging (rsfMRI) data, movie watching functional magnetic resonance imaging (mwfMRI) from 184 adult subjects. These subjects are a subset of a larger cohort of approximately 1200 subjects additionally scanned at 3T. Subjects’ 7T fMRI data were collected during four scan sessions over the course of two or three days at the Center for Magnetic Resonance Research at the University of Minnesota. Subjects’ 3T T1w data collected at Washington University in St. Louis. The study was approved by the Washington University Institutional Review Board and informed consent was obtained from all subjects.

Demographics

We analyzed MRI data collected from $N_s = 129$ subjects (77 female, 52 male), after excluding subjects with poor quality data. Our exclusion criteria was as follows: where each spike is defined as relative framewise displacement of at least 0.25 mm, we excluded subjects who fulfill at least 1 of the following criteria: greater than 15% of time points spike, average framewise displacement greater than 0.2 mm; contains any spikes larger than 5mm. Following this filter, subjects who contained all four scans were retained. At the time of their first scan, the average subject age was 29.36 ± 3.36 years, with a range from 22 – 36. 70 of these subjects were monozygotic twins, 57 were non-monozygotic twins, and 2 were not twins.

MRI acquisition and processing

A comprehensive description of the imaging parameters and image preprocessing can be found in [71] and in HCP’s online documentation (<https://www.humanconnectome.org/study/hcp-young-adult/document/1200-subjects-data-release>). T1w were collected on a 3T Siemens Connectome Skyra scanner with a 32-channel head coil. Subjects underwent two T1-weighted structural scans, which were averaged for each subject (TR = 2400 ms, TE = 2.14 ms, flip angle = 8° , 0.7 mm isotropic voxel resolution). fMRI were collected on a 7T Siemens Magnetom scanner with a 32-channel head coil. All 7T fMRI data was acquired with a gradient-echo planar imaging sequence (TR = 1000 ms, TE = 22.2 ms, flip angle = 45° , 1.6 mm isotropic voxel

resolution, multi-band factor = 5, image acceleration factor = 2, partial Fourier sample = 7/8, echo spacing = 0.64 ms, bandwidth = 1924 Hz/Px). Four resting state data runs were collected, each lasting 15 minutes (frames = 900), with eyes open and instructions to fixate on a cross. Four movie watching data runs were collected, each lasting approximately 15 minutes (frames = 921, 918, 915, 901), with subjects passively viewing visual and audio presentations of movie scenes. Movies consisted of both freely available independent films covered by Creative Commons licensing and Hollywood movies prepared for analysis [72]. For both resting state and movie watching data, two runs were acquired with posterior-to-anterior phase encoding direction and two runs were acquired with anterior-to-posterior phase encoding direction.

Structural and functional images were minimally pre-processed according to the description provided in [71]. 7T fMRI images were downloaded after correction and re-processing announced by the HCP consortium in April, 2018. Briefly, T1w images were aligned to MNI space before undergoing FreeSurfer's (version 5.3) cortical reconstruction workflow. fMRI images were corrected for gradient distortion, susceptibility distortion, and motion, and then aligned to the corresponding T1w with one spline interpolation step. This volume was further corrected for intensity bias and normalized to a mean of 10000. This volume was then projected to the 2mm *32k_fs_LR* mesh, excluding outliers, and aligned to a common space using a multi-modal surface registration [73]. The resultant CIFTI file for each HCP subject used in this study followed the file naming pattern: **_Atlas_MSMA11_hp2000_clean.dtseries.nii*. Resting state and moving watching fMRI images were nuisance regressed in the same manner. Each minimally pre-processed fMRI was linearly detrended, band-pass filtered (0.008-0.25 Hz), confound regressed and standardized using Nilearn's `signal.clean` function, which removes confounds orthogonally to the temporal filters. The confound regression strategy included six motion estimates, mean signal from a white matter, cerebrospinal fluid, and whole brain mask, derivatives of these previous nine regressors, and squares of these 18 terms. Spike regressors were not applied. Following these preprocessing operations, the mean signal was taken at each time frame for each node, as defined by the Schaefer 400 parcellation [74] in *32k_fs_LR* space.

Indiana University Data

Demographics

We analyzed MRI data collected from $N_s = 29$ subjects (5 female, 24 male; 25 were right-handed). This cohort was male-dominant, as subjects were intended to serve as controls for a study in autism spectrum disorder, which is more common in men than women. At the time

of their first scan, the average subject age was 24.9 ± 4.7 years.

MRI acquisition and processing

MRI images were acquired using a 3T whole-body MRI system (Magnetom Tim Trio, Siemens Medical Solutions, Natick, MA) with a 32-channel head receive array. Both raw and prescan-normalized images were acquired; raw images were used at all preprocessing stages and in all analyses unless specifically noted. During functional scans, T2*-weighted multiband echo planar imaging (EPI) data were acquired using the following parameters: TR/TE = 813/28 ms; 1200 vol; flip angle = 60°; 3.4 mm isotropic voxels; 42 slices acquired with interleaved order covering the whole brain; multi-band acceleration factor of 3. Preceding the first functional scan, gradient-echo EPI images were acquired in opposite phase-encoding directions (10 images each with P-A and A-P phase encoding) with identical geometry to the EPI data (TR/TE = 1175/39.2 ms, flip angle = 60°) to be used to generate a fieldmap to correct EPI distortions, similar to the approach used by the Human Connectome Project [75]. High-resolution T1-weighted images of the whole brain (MPRAGE, 0.7 mm isotropic voxel size; TR/TE/TI = 2499/2.3/1000 ms) were acquired as anatomical references.

All functional data were processed according to an in-house pipeline using FEAT (v6.00) and MELODIC (v3.14) within FSL (v. 5.0.9; FMRIB's Software Library, www.fmrrib.ox.ac.uk/fsl), Advanced Normalization Tools (ANTs; v2.1.0) [76], and Matlab_R2014b. This pipeline was identical to the **GLM + MGTR** procedure described in [32].

In more detail, individual anatomical images were bias-corrected and skull-stripped using ANTs, and segmented into gray matter, white matter, and CSF partial volume estimates using FSL FAST. A midspace template was constructed using ANTs' `buildtemplateparallel` and subsequently skull-stripped. Composite (affine and diffeomorphic) transforms warping each individual anatomical image to this midspace template, and warping the midspace template to the Montreal Neurological Institute MNI152 1mm reference template, were obtained using ANTs.

For each functional run, the first five volumes (≈ 4 seconds) were discarded to minimize magnetization equilibration effects. Framewise displacement traces for this raw (trimmed) data were computed using `fsl_motion_outliers`. Following [32, 77], we performed FIX followed by mean cortical signal regression. This procedure included rigid-body motion correction, fieldmap-based geometric distortion correction, and non-brain removal (but not slice-timing correction due to fast TR [75]). Preprocessing included weak highpass temporal filtering (>2000 s FWHM) to remove slow drifts [75] and no spatial smoothing. Off-resonance geometric distor-

tions in EPI data were corrected using a fieldmap derived from two gradient-echo EPI images collected in opposite phase-encoding directions (posterior-anterior and anterior-posterior) using FSL topup.

We then used FSL-FIX [78] to regress out independent components classified as noise using a classifier trained on independent but similar data and validated on hand-classified functional runs. The residuals were regarded as “cleaned” data. Finally, we regressed out the mean cortical signal (mean BOLD signal across gray matter partial volume estimate obtained from FSL FAST). All analyses were carried out on these data, which were registered to subjects’ skull-stripped T1-weighted anatomical imaging using Boundary-Based Registration (BBR) with *epi_reg* within FSL. Subjects’ functional images were then transformed to the MNI152 reference in a single step, using ANTS to apply a concatenation of the affine transformation matrix with the composite (affine + diffeomorphic) transforms between a subject’s anatomical image, the midspace template, and the MNI152 reference. Prior to network analysis, we extracted mean regional time series from regions of interest defined as sub-divisions of the 17-system parcellation reported in [79] and used previously [80–82]. Wakefulness during movie and rest scans was monitored in real-time using an eye tracking camera (Eyelink 1000).

Naturalistic stimuli and coding boundaries

Human Connectome Project Data

Movies consisted of concatenated movie scenes with 20 second blocks of rest between them. The movies scenes were sourced from both freely available independent films covered by Creative Commons licensing and Hollywood movies prepared for analysis [72].

Both JCT and RFB manually coded the boundaries between the end of movie scenes and the beginning of rest blocks and confirmed consistency across codings. Additionally, JCT confirmed that these codings aligned with the end of movie scenes and the beginning of rest blocks using the RGB values from a digitized version of the movies. The movie screen during rest blocks, on average, contains more black pixels (See Fig 1b for example). By indexing the number of black pixels and overlaying the coded boundaries, JCT found an alignment between the beginning of rest blocks (where the number of black pixels increased) and coded boundaries.

Indiana University Data

All movies were obtained from Vimeo (<https://vimeo.com>). They were selected based on multiple criteria. First, to ensure that movie trailers represented novel stimuli, we excluded any movie that had a wide theatrical release. Secondly, we excluded movies with potentially

objectionable content including nudity, swearing, drug use, etc. Lastly, we excluded movies with intentionally startling events that could lead to excessive in-scanner movement.

Each trailer lasted between 45 and 285 seconds (approximately 1 to 5 minutes). Each movie scan comprised between four and six trailers with genres that included documentaries, dramas, comedies, sports, mystery, and adventure (See Table. S1 for more details). Both JCT and RFB manually coded the boundaries between trailers and confirmed consistency across codings.

Functional connectivity

Functional connectivity (FC) measures the statistical dependence between the activity of distinct neural elements. In the modeling of macroscale brain networks with fMRI data, this usually means computing the Pearson correlation of brain regions’ activity time series. To calculate FC for regions i and j , then, we first standardize their time series and represent them as z-scores. We denote the z-scored time series of region i as $\mathbf{z}_i = [z_i(1), \dots, z_i(T)]$, where T is the number of samples. The Pearson correlation is then calculated as:

$$r_{ij} = \frac{1}{T-1} \sum_{t=1}^T z_i(t) \cdot z_j(t). \quad (1)$$

In other words, the correlation is equal to the temporal average of two regions’ co-fluctuation.

Edge time series

We analyzed edge time series data. Edge time series can be viewed as a temporal decomposition of a correlation (functional connection) into its framewise contributions. Note that Pearson correlation is calculated as $r_{x,y} = \frac{1}{T-1} \sum_t z_x(t) \cdot z_y(t)$, where T is the number of samples and $z_x(t) = \frac{x - \mu_x}{\sigma_x}$ is the z-scored transformation of the time series $x = [x(1), \dots, x(T)]$. If we omit the summation in our calculation of $r_{x,y}$, we obtain a time series $r_{x,y}(t) = z_x(t) \cdot z_y(t)$, whose elements index the instantaneous co-fluctuation between variates x and y . Here, we estimated the edge time series for all pairs of brain regions $\{i, j\}$.

We analyzed edge time series using two distinct approaches.

Edge time series amplitude (RSS)

First, we calculated the amplitude at each frame as the root sum of squared co-fluctuations. That is, the amplitude at time t was given by:

Co-fluctuation Matrix Differences

$$RSS(t) = \sqrt{\sum_{i,j>i} r_{ij}(t)^2} \quad (2)$$

Whole-brain co-fluctuation patterns

The section way in which we analyzed edge time series was by considering the whole-brain co-fluctuation pattern at a given time point. That is, we focused on the co-fluctuation matrices, $C(t) \in \mathbb{R}^{N \times N}$. The element $\{i, j\}$ in this matrix denoted the co-fluctuation magnitude between regions i and j at time t , i.e. $r_{ij}(t)$. Although co-fluctuation matrices are not correlation matrices, they can nonetheless be analyzed using similar network science tools.

Global Similarity Matrix

We assessed the similarity of events between subjects and types (boundary, movie, and asynchronous). The procedure for doing so involved two steps. First, for a given pair of subjects, r and s , with n_r and n_s events, respectively, we calculated the $R \in \mathbb{R}^{[n_r \times n_s]}$ similarity matrix between all pairs of events. The elements of this matrix could be further aggregated and averaged based on the category to which events were assigned, reducing R to an asymmetric, $[3 \times 3]$ matrix (rows and columns correspond to the three event types).

This first step was repeated for all pairs of subjects, each time yielding an analogous $[3 \times 3]$ matrix. At its completion this, second step yielded a 387×387 similarity matrix, whose rows and columns could be reordered by event types. We show this matrix in Fig. 2d.

In the main text, we described differences in the co-fluctuation matrices associated with different event types. To assess these differences, we performed the following set of calculations. For a given pair of event types, we identified all instances of each type and calculated, at the level of edges, the differences in their co-fluctuation magnitude and the mean sign of those differences. We then retained the sign of these differences. The result is a $n \times n$ matrix for each subject whose elements are $\{-1, 1\}$. Note that this calculation was performed within subjects.

Next, we examined whether, across subjects, edges exhibited concordant behavior using a paired-samples t -test. That is, we assessed whether, across subjects, an edge's co-fluctuation magnitude under one event type was consistently greater than or less than that of another event type. This procedure was repeated for every pair of event types.

Author Contributions

JCT, OS, and RFB conceived of experiments, designed analyses, wrote the initial versions of the manuscript, and edited the manuscript. JCT performed all analyses. JF, LB, and DK contributed data. All authors helped edit the manuscript.

Acknowledgements

Data were provided [in part] by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. This material is based on work supported by NSF Grant IIS-2023985 (RFB, OS), NSF Grant 1735095, Interdisciplinary Training in Complex Networks (JCT), and Indiana University Office of the Vice President for Research Emerging Area of Research Initiative, Learning: Brains, Machines and Children (RFB). This research was supported in part by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute.

[1] Ed Bullmore and Olaf Sporns, “Complex brain networks: graph theoretical analysis of structural and functional systems,” *Nature reviews neuroscience* **10**, 186–198 (2009).

[2] Hae-Jeong Park and Karl Friston, “Structural and functional brain networks: from connections to cognition,” *Science* **342** (2013).

[3] Danielle S Bassett and Olaf Sporns, “Network neuroscience,” *Nature neuroscience* **20**, 353–364 (2017).

[4] Barry Horwitz, “The elusive concept of brain connectivity,” *Neuroimage* **19**, 466–470 (2003).

[5] Karl J Friston, “Functional and effective connectivity in neuroimaging: a synthesis,” *Human brain mapping* **2**, 56–78 (1994).

[6] Andrew T Reid, Drew B Headley, Ravi D Mill, Ruben Sanchez-Romero, Lucina Q Uddin, Daniele Marinazzo, Daniel J Lurie, Pedro A Valdés-Sosa, Stephen José Hanson, Bharat B Biswal, *et al.*, “Advancing functional connectivity research from association to causation,” *Nature neuroscience* **22**, 1751–1760 (2019).

[7] Jonathan D Power, Alexander L Cohen, Steven M Nelson, Gagan S Wig, Kelly Anne Barnes, Jessica A Church, Alecia C Vogel, Timothy O Laumann, Fran M Miezin, Bradley L Schlaggar, *et al.*, “Functional network organization of the human brain,” *Neuron* **72**, 665–678 (2011).

[8] Jonathan D Power, Bradley L Schlaggar, Christina N Lessov-Schlaggar, and Steven E Petersen, “Evidence for hubs in human functional brain networks,” *Neuron* **79**, 798–813 (2013).

[9] Monica D Rosenberg, Emily S Finn, Dustin Scheinost, Xenophon Papademetris, Xilin Shen, R Todd Constable, and Marvin M Chun, “A neuromarker of sustained attention from whole-brain functional connectivity,” *Nature neuroscience* **19**, 165–171 (2016).

[10] Emily W Avery, Kwangsun Yoo, Monica D Rosenberg, Abigail S Greene, Siyuan Gao, Duk L Na, Dustin Scheinost, Todd R Constable, and Marvin M Chun, “Distributed patterns of functional connectivity predict working memory performance in novel healthy and memory-impaired individuals,” *Journal of cognitive neuroscience* **32**, 241–255 (2020).

[11] Lucy S King, M Catalina Camacho, David F Montez, Kathryn L Humphreys, and Ian H Gotlib, “Naturalistic language input is associated with resting-state functional connectivity in infancy,” *Journal of Neuroscience* **41**, 424–434 (2021).

[12] Alex Fornito, Andrew Zalesky, and Michael Breakspear, “The connectomics of brain disorders,” *Nature Reviews Neuroscience* **16**, 159–172 (2015).

[13] Ye Tian, Andrew Zalesky, Chad Bousman, Ian Everall, and Christos Pantelis, “Insula functional connectivity in schizophrenia: subregions, gradients, and symptoms,” *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* **4**, 399–408 (2019).

[14] Shi Gu, Theodore D Satterthwaite, John D Medaglia, Muzhi Yang, Raquel E Gur, Ruben C Gur, and Danielle S Bassett, “Emergence of system roles in normative neurodevelopment,” *Proceedings of the National Academy of Sciences* **112**, 13681–13686 (2015).

[15] Micaela Y Chan, Denise C Park, Neil K Savalia, Steven E Petersen, and Gagan S Wig, “Decreased segregation of brain systems across the healthy adult lifespan,” *Proceedings of the National Academy of Sciences* **111**, E4997–E5006 (2014).

[16] Farnaz Zamani Esfahlani, Younghun Jo, Joshua Farkowitz, Lisa Byrge, Daniel Kennedy, Olaf Sporns, and Richard Betzel, “High-amplitude co-fluctuations in cortical activity drive functional connectivity,” *Proceedings of the National Academy of Sciences* (2020).

[17] Joshua Farkowitz, Farnaz Zamani Esfahlani, Younghun Jo, Olaf Sporns, and Richard F Betzel, “Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture,” *Nature neuroscience* **23**, 1644–1654 (2020).

[18] Sarah Greenwell, Joshua Farkowitz, Laura Pritschet, Tyler Santander, Emily G Jacobs, and Richard F Betzel, “High-amplitude network co-fluctuations linked to variation in hormone concentrations over menstrual cycle,” *bioRxiv* (2021).

[19] Richard F Betzel, Sarah A Cutts, Sarah Greenwell, Joshua Farkowitz, and Olaf Sporns, “Individualized event structure drives individual differences in whole-brain functional connectivity,” *NeuroImage* **252**, 118993 (2022).

[20] Richard Betzel, Sarah Cutts, Jaco Tanner, Sarah Greenwell, Thomas Varley, Joshua Farkowitz, and Olaf Sporns, “Hierarchical organization of spontaneous co-fluctuations in densely-sampled individuals using fmri,” *bioRxiv* (2022).

[21] Natalia Petridou, César Caballero Gaudes, Ian L Dryden, Susan T Francis, and Penny A Gowland, “Periods of rest in fmri contain individual spontaneous events which are related to slowly fluctuating spontaneous activity,” *Human brain mapping* **34**, 1319–1329 (2013).

[22] Enzo Tagliazucchi, Pablo Balenzuela, Daniel Fraiman, and Dante R Chialvo, “Criticality in large-scale brain fmri dynamics unveiled by a novel point process analysis,” *Frontiers in physiology* **3**, 15 (2012).

[23] I Cifre, M Zarepour, SG Horovitz, SA Cannas, and DR Chialvo, “Further results on why a point process is effective for estimating correlation between brain regions,” *Papers in Physics* **12**, 120003–120003 (2020).

[24] Xiao Liu and Jeff H Duyn, “Time-varying functional network information extracted from brief instances of spontaneous brain activity,” *Proceedings of the National Academy of Sciences* **110**, 4392–4397 (2013).

[25] Farnaz Zamani Esfahlani, Lisa Byrge, Jacob Tanner, Olaf Sporns, Daniel Kennedy, and Richard Betzel, “Edge-centric analysis of time-varying functional brain networks with applications in autism spectrum disorder,” *bioRxiv* (2021).

[26] Rose A Cooper, Kyle A Kurkela, Simon W Davis, and Maureen Ritchey, “Mapping the organization and dynamics of the posterior medial network during movie watching,” *NeuroImage* **236**, 118075 (2021).

[27] Maria Pope, Makoto Fukushima, Richard F Betzel, and Olaf Sporns, “Modular origins of high-amplitude cofluctuations in fine-scale functional connectivity dynamics,” *Proceedings of the National Academy of Sciences* **118** (2021).

[28] Leonardo Novelli and Adeel Razi, “A mathematical perspective on edge-centric functional connectivity,” *arXiv preprint arXiv:2106.10631* (2021).

[29] Z Ladwig, BA Seitzman, A Dworetzky, Y Yu, B Adeyemo, DM Smith, SE Petersen, and C Gratton, “Bold cofluctuation ‘events’ are predicted from static functional connectivity,” (2022).

[30] Richard Betzel, Sarah Cutts, Sarah Greenwell, and Olaf Sporns, “Individualized event structure drives individual differences in whole-brain functional connectivity,” *bioRxiv* (2021).

[31] David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, Wu-Minn HCP Consortium, *et al.*, “The wu-minn human connectome project: an overview,” *NeuroImage* **80**, 62–79 (2013).

[32] Lisa Byrge and Daniel P Kennedy, “Identifying and characterizing systematic temporally-lagged bold artifacts,” *NeuroImage* **171**, 376–392 (2018).

[33] Eun Young Choi, BT Thomas Yeo, and Randy L Buckner, “The organization of the human striatum estimated by intrinsic functional connectivity,” *Journal of neuro-*

physiology **108**, 2242–2263 (2012).

[34] BT Thomas Yeo, Fenna M Krienen, Jorge Sepulcre, Mert R Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L Roffman, Jordan W Smoller, Lilla Zöllei, Jonathan R Polimeni, *et al.*, “The organization of the human cerebral cortex estimated by intrinsic functional connectivity,” *Journal of neurophysiology* (2011).

[35] Randy L Buckner, Fenna M Krienen, Angela Castellanos, Julio C Diaz, and BT Thomas Yeo, “The organization of the human cerebellum estimated by intrinsic functional connectivity,” *Journal of neurophysiology* **106**, 2322–2345 (2011).

[36] Bruce Fischl, “Freesurfer,” *Neuroimage* **62**, 774–781 (2012).

[37] Sarah A Cutts, Joshua Faskowitz, Richard F Betzel, and Olaf Sporns, “Uncovering individual differences in fine-scale dynamics of functional connectivity,” *Cerebral cortex* (2022).

[38] Saurabh Sonkusare, Michael Breakspear, and Christine Guo, “Naturalistic stimuli in neuroscience: critically acclaimed,” *Trends in cognitive sciences* **23**, 699–714 (2019).

[39] Emily S Finn, “Is it time to put rest to rest?” *Trends in cognitive sciences* **25**, 1021–1032 (2021).

[40] Eleanor A Maguire, “Studying the freely-behaving brain with fmri,” *Neuroimage* **62**, 1170–1176 (2012).

[41] Hongmi Lee and Janice Chen, “A generalized cortical activity pattern at internally-generated mental context boundaries during unguided narrative recall,” *bioRxiv* (2021).

[42] Avital Hahamy, Haim Dubossarsky, and Timothy EJ Behrens, “The human brain reactivates context-specific past information at event boundaries of naturalistic experiences,” *bioRxiv* (2022).

[43] Gidon Levakov, Olaf Sporns, and Galia Avidan, “Fine-scale dynamics of functional connectivity in the face processing network during movie watching,” *bioRxiv* (2022).

[44] Reza Rajimehr, Haoran Xu, Asa Farahani, Simon Kornblith, John Duncan, and Robert Desimone, “Functional architecture of cerebral cortex during naturalistic movie-watching,” *bioRxiv* (2022).

[45] George A Miller, “The magical number seven, plus or minus two: Some limits on our capacity for processing information.” *Psychological review* **63**, 81 (1956).

[46] Mark J Landau, Jamie Arndt, Trevor J Swanson, and Michael N Bultmann, “Why life speeds up: Chunking and the passage of autobiographical time,” *Self and Identity* **17**, 294–309 (2018).

[47] Christopher A Kurby and Jeffrey M Zacks, “Segmentation in the perception and memory of events,” *Trends in cognitive sciences* **12**, 72–79 (2008).

[48] Jesse Q Sargent, Jeffrey M Zacks, David Z Hambrick, Rose T Zacks, Christopher A Kurby, Heather R Bailey, Michelle L Eisenberg, and Taylor M Beck, “Event segmentation ability uniquely predicts event memory,” *Cognition* **129**, 241–255 (2013).

[49] Yeon Soon Shin and Sarah DuBrow, “Structuring memory through inference-based event segmentation,” *Topics in Cognitive Science* **13**, 106–127 (2021).

[50] Anna Jafarpour, Elizabeth A Buffalo, Robert T Knight, and Anne GE Collins, “Event segmentation reveals working memory forgetting rate,” *Iscience* **25**, 103902 (2022).

[51] Catie Chang, David A Leopold, Marieke Louise Schölvink, Hendrik Mandelkow, Dante Picchioni, Xiao Liu, Q Ye Frank, Janita N Turchi, and Jeff H Duy, “Tracking brain arousal fluctuations with fmri,” *Proceedings of the National Academy of Sciences* **113**, 4518–4523 (2016).

[52] Javier Gonzalez-Castillo, Isabel Fernandez, Daniel Handwerker, and Peter Bandettini, “Contribution of ultra-slow fmri fluctuations in the 4th ventricle to our understanding of sleep,” *bioRxiv* (2022).

[53] Elena A Allen, Eswar Damaraju, Sergey M Plis, Erik B Erhardt, Tom Eichele, and Vince D Calhoun, “Tracking whole-brain connectivity dynamics in the resting state,” *Cerebral cortex* **24**, 663–676 (2014).

[54] Wha Jin Lee, Jesse A Brown, Hye Ryun Kim, Renaud La Joie, Hanna Cho, Chul Hyoung Lyoo, Gil D Rabinovici, Joon-Kyung Seong, William W Seeley, Alzheimer’s Disease Neuroimaging Initiative, *et al.*, “Regional α/β -tau interactions promote onset and acceleration of alzheimer’s disease tau spreading,” *Neuron* (2022).

[55] Uri Hasson, Rafael Malach, and David J Heeger, “Reliability of cortical activity during natural stimulation,” *Trends in cognitive sciences* **14**, 40–48 (2010).

[56] Raphael Liegeois, Timothy O Laumann, Abraham Z Snyder, Juan Zhou, and BT Thomas Yeo, “Interpreting temporal fluctuations in resting-state functional connectivity mri,” *Neuroimage* **163**, 437–455 (2017).

[57] Timothy O Laumann, Abraham Z Snyder, Anish Mitra, Evan M Gordon, Caterina Gratton, Babatunde Adeyemo, Adrian W Gilmore, Steven M Nelson, Jeff J Berg, Deanna J Greene, *et al.*, “On the stability of bold fmri correlations,” *Cerebral cortex* **27**, 4719–4732 (2017).

[58] Daniel J Lurie, Daniel Kessler, Danielle S Bassett, Richard F Betzel, Michael Breakspear, Shella Kheilholz, Aaron Kucyi, Raphaël Liégeois, Martin A Lindquist, Anthony Randal McIntosh, *et al.*, “Questions and controversies in the study of time-varying functional connectivity in resting fmri,” *Network Neuroscience* **4**, 30–69 (2020).

[59] Michael Hanke, Florian J Baumgartner, Pierre Ibe, Falko R Kaule, Stefan Pollmann, Oliver Speck, Wolf Zinke, and Jörg Stadler, “A high-resolution 7-tesla fmri dataset from complex natural stimulation with an audio movie,” *Scientific data* **1**, 1–18 (2014).

[60] Aaron Kucyi, Michael Esterman, James Capella, Allison Green, Mai Uchida, Joseph Biederman, John DE Gabrieli, Eve M Valera, and Susan Whitfield-Gabrieli, “Prediction of stimulus-independent and task-unrelated thought from functional brain networks,” *Nature communications* **12**, 1–17 (2021).

[61] Richard F Betzel, John D Medaglia, Ari E Kahn, Jonathan Soffer, Daniel R Schonhaut, and Danielle S Bassett, “Structural, geometric and genetic factors predict interregional brain connectivity patterns probed by electrocorticography,” *Nature biomedical engineering* **3**, 902–916 (2019).

[62] Brittany H Scheid, Arian Ashourvan, Jennifer Stiso, Kathryn A Davis, Fadi Mikhail, Fabio Pasqualetti, Brian Litt, and Danielle S Bassett, “Time-evolving controllability of effective connectivity networks during seizure progression,” *Proceedings of the National Academy of Sciences* **118** (2021).

[63] Catherine J Chu, Mark A Kramer, Jay Pathmanathan, Matt T Bianchi, M Brandon Westover, Lauren Wizon, and Sydney S Cash, “Emergence of stable functional networks in long-term human electroencephalography,” *Journal of Neuroscience* **32**, 2703–2713 (2012).

[64] Lior Fisch, Eran Privman, Michal Ramot, Michal Harel, Yuval Nir, Svetlana Kipervasser, Fani Andelman, Miri Y Neufeld, Uri Kramer, Itzhak Fried, *et al.*, “Neural “ignition”: enhanced activation linked to perceptual awareness in human ventral stream visual cortex,” *Neuron* **64**, 562–574 (2009).

[65] Stephan Krohn, Nina von Schwanenflug, Leonhard Waschke, Amy Romanello, Martin Gell, Douglas D Garrett, and Carsten Finke, “A spatiotemporal complexity architecture of human brain activity,” *bioRxiv* (2021).

[66] Anzar Abbas, Michaël Belloy, Amrit Kashyap, Jacob Billings, Maysam Nezafati, Eric H Schumacher, and Shella Keilholz, “Quasi-periodic patterns contribute to functional connectivity in the brain,” *Neuroimage* **191**, 193–204 (2019).

[67] Petra Ritter, Michael Schirner, Anthony R McIntosh, and Viktor K Jirsa, “The virtual brain integrates computational modeling and multimodal neuroimaging,” *Brain connectivity* **3**, 121–145 (2013).

[68] Jill X O'Reilly, Paula L Croxson, Saad Jbabdi, Jerome Sallet, MaryAnn P Noonan, Rogier B Mars, Philip GF Browning, Charles RE Wilson, Anna S Mitchell, Karla L Miller, *et al.*, “Causal effect of disconnection lesions on interhemispheric functional connectivity in rhesus monkeys,” *Proceedings of the National Academy of Sciences* **110**, 13982–13987 (2013).

[69] Michael W Cole, Danielle S Bassett, Jonathan D Power, Todd S Braver, and Steven E Petersen, “Intrinsic and task-evoked network architectures of the human brain,” *Neuron* **83**, 238–251 (2014).

[70] Michael W Cole, Takuya Ito, Douglas Schultz, Ravi Mill, Richard Chen, and Carrissa Cocuzza, “Task activations produce spurious but systematic inflation of task functional connectivity estimates,” *Neuroimage* **189**, 1–18 (2019).

[71] Matthew F Glasser, Stamatios N Sotiropoulos, J Anthony Wilson, Timothy S Coalson, Bruce Fischl, Jesper L Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R Polimeni, *et al.*, “The minimal preprocessing pipelines for the human connectome project,” *Neuroimage* **80**, 105–124 (2013).

[72] James E Cutting, Kaitlin L Brunick, and Ayse Canban, “Perceiving event dynamics and parsing hollywood films.” *Journal of experimental psychology: human perception and performance* **38**, 1476 (2012).

[73] Emma C Robinson, Saad Jbabdi, Matthew F Glasser, Jesper Andersson, Gregory C Burgess, Michael P Harms, Stephen M Smith, David C Van Essen, and Mark Jenkinson, “Msm: a new flexible framework for multimodal surface matching,” *Neuroimage* **100**, 414–426 (2014).

[74] Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and BT Thomas Yeo, “Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri,” *Cerebral cortex* **28**, 3095–3114 (2018).

[75] Stephen M Smith, Christian F Beckmann, Jesper Andersson, Edward J Auerbach, Janine Bijsterbosch, Gwenaëlle Douaud, Eugene Duff, David A Feinberg, Ludovica Griffanti, Michael P Harms, *et al.*, “Resting-state fmri in the human connectome project,” *Neuroimage* **80**, 144–168 (2013).

[76] Brian B Avants, Nicholas J Tustison, Gang Song, Philip A Cook, Arno Klein, and James C Gee, “A reproducible evaluation of ants similarity metric performance in brain image registration,” *Neuroimage* **54**, 2033–2044 (2011).

[77] Gregory C Burgess, Sridhar Kandala, Dan Nolan, Timothy O Laumann, Jonathan D Power, Babatunde Adeyemo, Michael P Harms, Steven E Petersen, and Deanna M Barch, “Evaluation of denoising strategies to address motion-correlated artifacts in resting-state functional magnetic resonance imaging data from the human connectome project,” *Brain connectivity* **6**, 669–680 (2016).

[78] Gholamreza Salimi-Khorshidi, Gwenaëlle Douaud, Christian F Beckmann, Matthew F Glasser, Ludovica Griffanti, and Stephen M Smith, “Automatic denoising of functional mri data: combining independent component analysis and hierarchical fusion of classifiers,” *Neuroimage* **90**, 449–468 (2014).

[79] BT Thomas Yeo, Fenna M Krienen, Jorge Sepulcre, Mert R Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L Roffman, Jordan W Smoller, Lilla Zöllei, Jonathan R Polimeni, *et al.*, “The organization of the human cerebral cortex estimated by intrinsic functional connectivity,” *Journal of neurophysiology* **106**, 1125–1165 (2011).

[80] Richard F Betzel, Lisa Byrge, Ye He, Joaquín Goñi, Xi-Nian Zuo, and Olaf Sporns, “Changes in structural and functional connectivity among resting-state networks across the human lifespan,” *Neuroimage* **102**, 345–357 (2014).

[81] Lisa Byrge, Julien Dubois, J Michael Tyszka, Ralph Adolphs, and Daniel P Kennedy, “Idiosyncratic brain activation patterns are associated with poor social comprehension in autism,” *Journal of Neuroscience* **35**, 5837–5850 (2015).

[82] Richard F Betzel, Bratislav Mišić, Ye He, Jeffrey Rumschlag, Xi-Nian Zuo, and Olaf Sporns, “Functional brain modules reconfigure at multiple scales across the human lifespan,” *arXiv preprint arXiv:1510.08045* (2015).

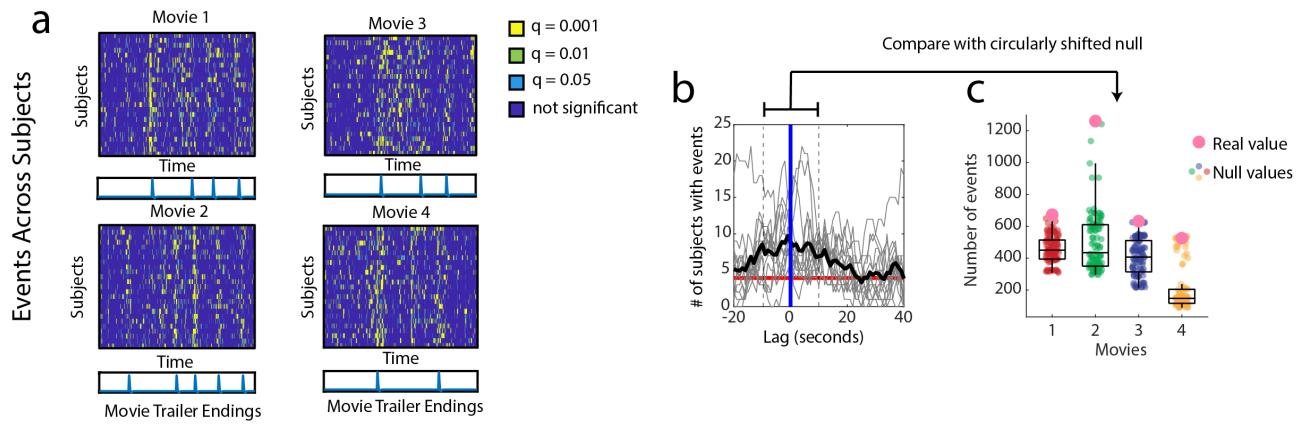


FIG. S1. **Replication of synchronous boundary events in Indiana University data set.** (a) Each panel represents two things: 1) the stacked event time series of all subjects, and 2) the hrf-convolved movie trailer endings. Each of the four figures is representative of a different scan/movie-presentation for the Indiana University data set. (b) Plot showing the number of subjects with an event for each trailer ending (for every scan). Bold line indicates the mean across all trailer endings. (c) Using a window of 10 seconds on either side of all movie trailer endings, we compared the number of events within this window to a null model where this window is circularly shifted 100 times.

| Scan | Title | Genre | Runtime |
|------|-------------------|---|---------|
| 1 | Man Up and Go | documentary/emotional | 4m20s |
| 1 | The First 70 | documentary | 3m |
| 1 | Fixation | documentary/adventure | 1m42s |
| 1 | The Living | drama | 2m |
| 1 | SAMSARA | documentary/“unparalleled sensory experience” | 1m35s |
| 1 | Blood Brother | documentary | 2m20s |
| 2 | Birdmen | documentary/adventure | 3m59s |
| 2 | Groomed | drama | 1m30s |
| 2 | Cold | outdoor/sports | 2m |
| 2 | Sleepwalkers | drama | 2m |
| 2 | A Kind of Show | comedy | 1m |
| 3 | Geofish | documentary/adventure | 4m40s |
| 3 | The Debut | outdoor/sports | 3m23s |
| 3 | Dreams of a Life | documentary/mystery | 2m10s |
| 3 | The Front Man | documentary | 2m30s |
| 3 | This Is Vanity | drama | 1m |
| 4 | Planetary | documentary | 4m30s |
| 4 | Sign Painters | documentary | 2m50s |
| 4 | Florida Man | documentary/drama | 2m |
| 4 | The Sleeping Bear | drama | 3m40 |

TABLE S1. Movies included in each movie scan for the Indiana University dataset.

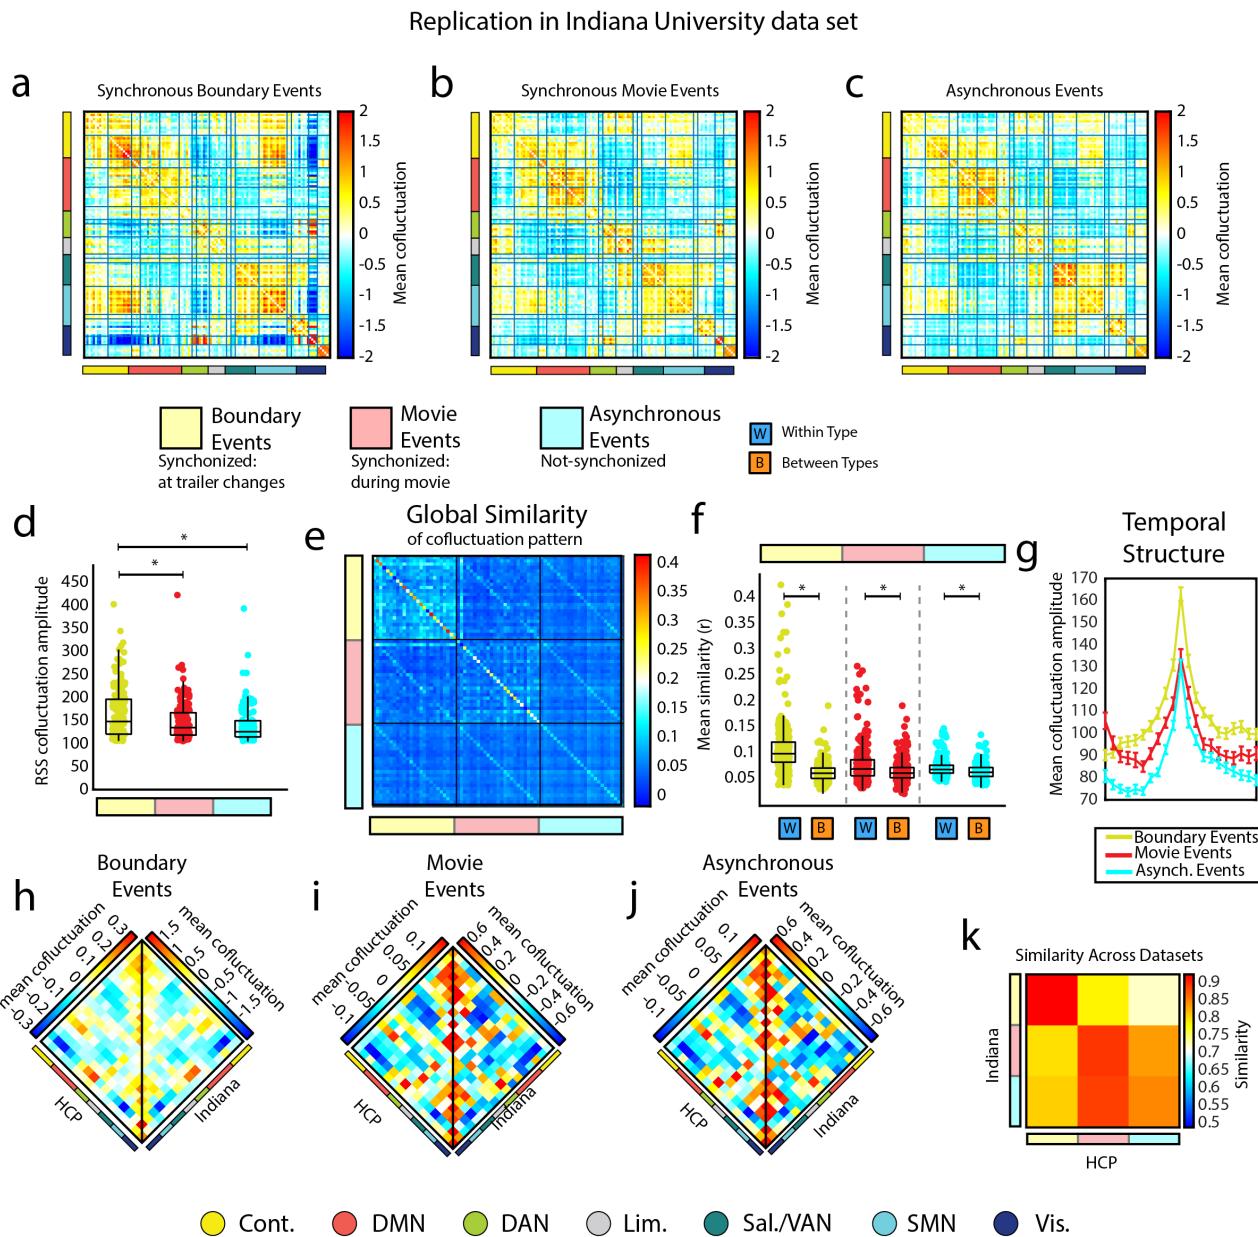


FIG. S2. Three types of events: replication in Indiana University data set. (a-c) Mean co-fluctuation matrices representing the mean of all events for a given event type. (d) Boxplots of the distribution of co-fluctuation amplitudes per event type. (e) Global similarity within and between different event types (global similarity is defined in the methods section). (f) Boxplots showing the data from the previous figure divided by similarity values within the same type and similarity values between different types. (g) Plot showing the mean temporal trajectory of each event type. (h-j) System by system mean co-fluctuation patterns across boundary events, movie events, and asynchronous events respectively for both data sets (HCP: left; Indiana: Right). (k) Across data set similarity of system by system mean co-fluctuation matrices for each event type.

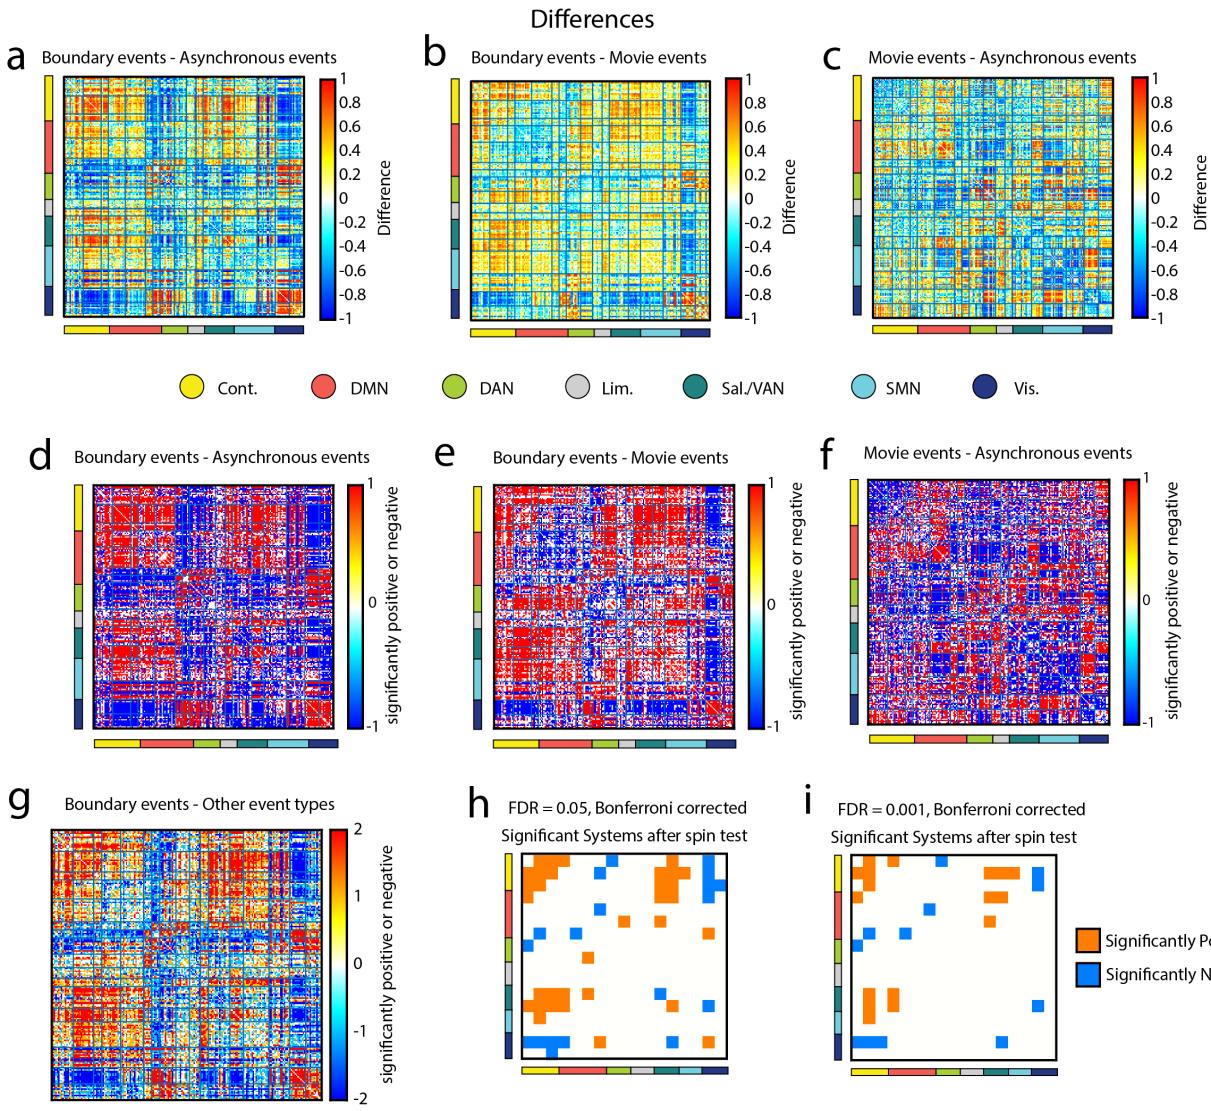


FIG. S3. Differences between the Event Types All differences described here are computed using the method detailed in the methods section. (a) Difference between boundary events and asynchronous events. (b) Difference between boundary events and movie events. (c) Difference between movie events and asynchronous events. (d) Significantly different edges between boundary events and asynchronous events each colored according to whether the difference was significantly positive or significantly negative. (e) Significantly different edges between boundary events and movie events each colored according to whether the difference was significantly positive or significantly negative. (f) Significantly different edges between movie events and asynchronous events each colored according to whether the difference was significantly positive or significantly negative. (g) Significantly different edges between boundary events and both asynchronous events as well as movie events each colored according to whether the difference was significantly positive or significantly negative. (h) Significant system by system edges after running a space-preserving null model (spin test) to test if the significant edges in the previous figure are more concentrated in system by system blocks than we should expect by chance. Here we controlled for multiple comparisons by fixing the false discovery rate to $q = 0.05$. (i) Same test as previous figure with a false discovery rate set to $q = 0.001$.

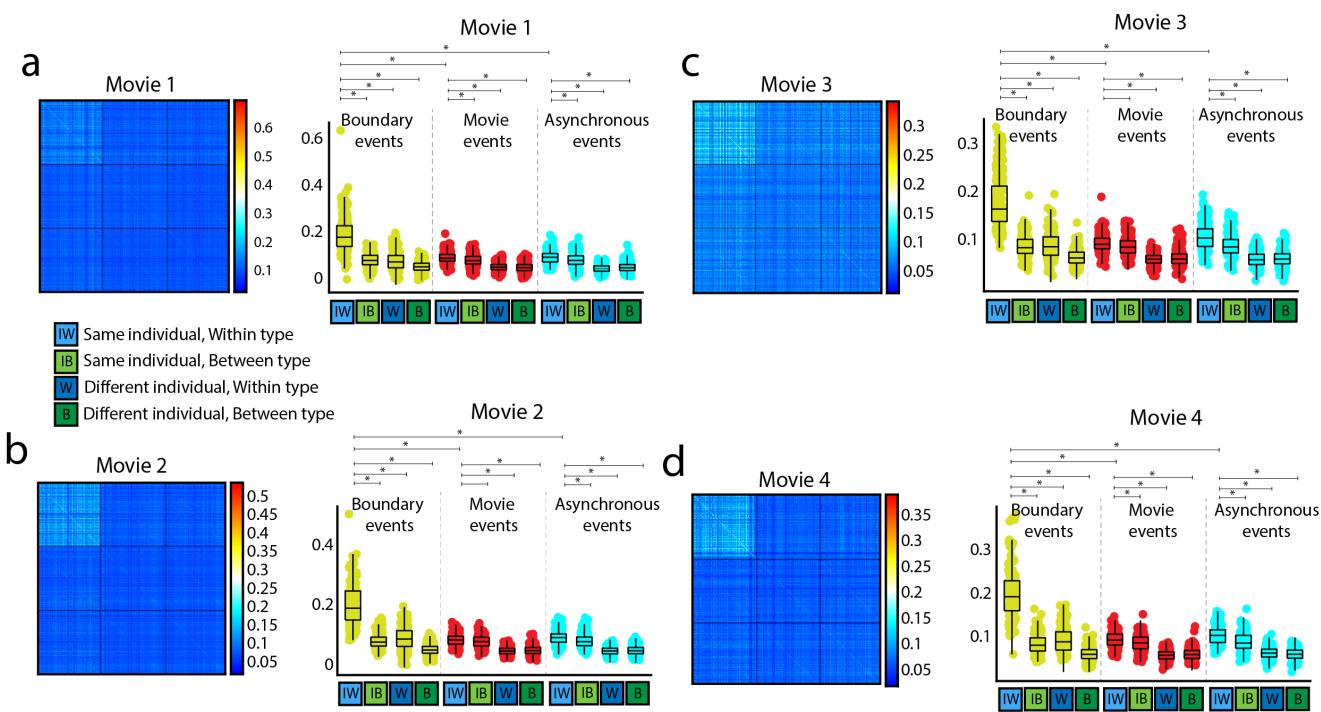


FIG. S4. Identifiability of the event types. (a-d) Each of the following figures show two plots: a similarity matrix and boxplots of these values divided into four categories per event type (same individual-within type, same individual-between type, different individual-within type, and different individual-between type). All p-values below $p < 8.66 \times 10^{-4}$, p-values for boundary events $p < 10^{-15}$.

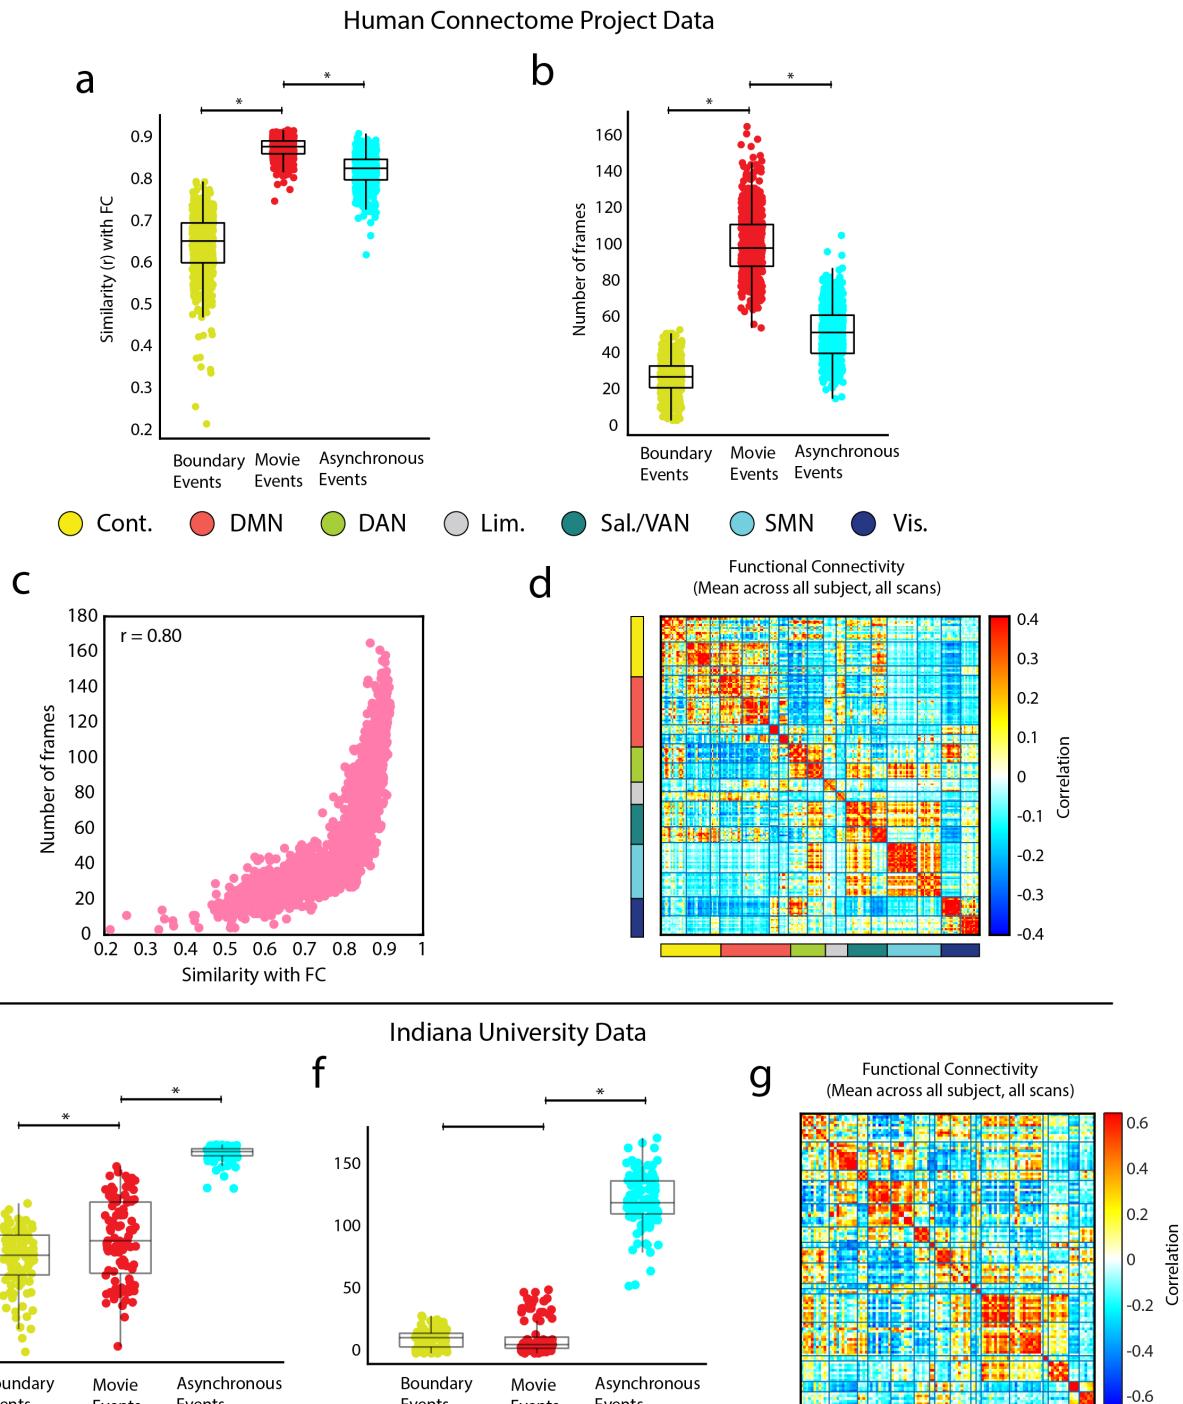


FIG. S5. Movie events are the most similar to functional connectivity in HCP data. (a) Boxplots showing the distribution of correlations between an event type (mean across the participants scan) and FC for the same scan. Movie events are the most similar to FC (all p -values below $p < 10^{-15}$). (b) Boxplots showing the number of frames assigned to each event type per scan/participant (all p -values below $p < 10^{-15}$). (c) Scatter plot showing the relationship between similarity with FC and number of frames. Similarity with FC appears to be at least partially driven by the number of frames used to compute the mean cofluctuation pattern for an event type. (d) Mean functional connectivity across all scans/participants. (e) Boxplots showing the distribution of correlations between an event type (mean across the participants scan) and FC for the same scan. (f) Boxplots showing the number of frames assigned to each event type per scan/participant. (g) Mean functional connectivity across all scans/participants.

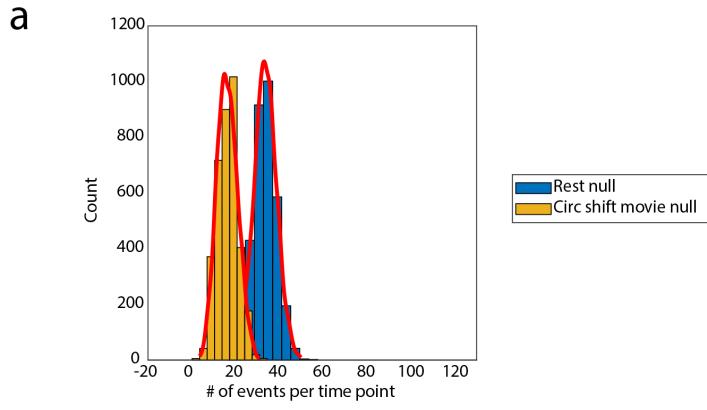


FIG. S6. Rest-based null model is more conservative than circularly shifted movie event time series. (a) Plot of two null distributions. The rest null was computed as the number of events per time point while subjects were at rest. The circ shift movie null was computed as the number of events per time point with the event time series circularly shifted to maintain the number and relative timing of events in the actual data, but break the alignment.

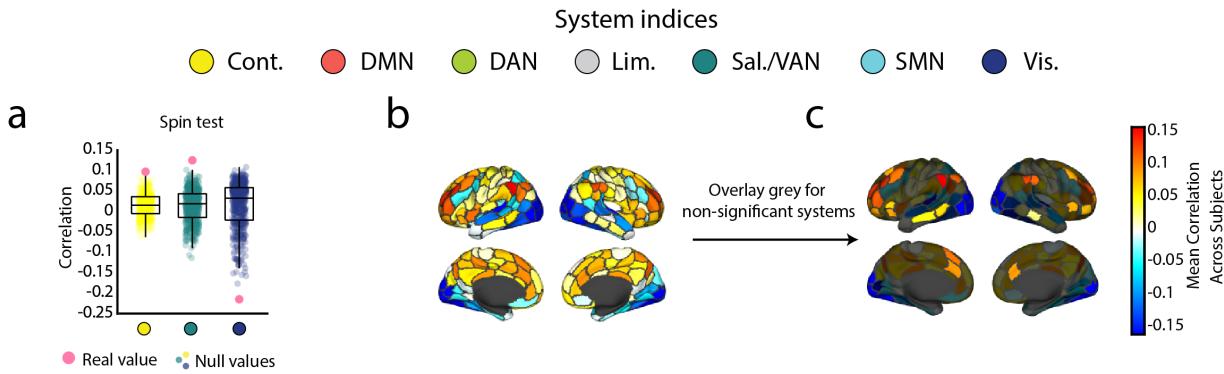


FIG. S7. Spin test shows that correlations are significantly concentrated in control b, salience b, and central visual systems. (a) We found that in three systems (control b, salience b, and central visual) the mean correlation value between nodal activations in that system and boundary events was significantly higher than a null distribution created using a space-preserving permutations test (spin test). (b) All correlations plotted to the brains surface. (c) Systems that were not significant were overlaid with grey.