Individualized event structure drives individual differences in whole-brain functional connectivity

Richard F. Betzel[1−4],* Sarah A. Cutts[1,3], Sarah Greenwell[1], and Olaf Sporns[1−4]
[1]*Department of Psychological and Brain Sciences,*
[2]*Cognitive Science Program,* [3]*Program in Neuroscience,*
[4]*Network Science Institute, Indiana University, Bloomington, IN 47405*
(Dated: March 12, 2021)

Resting-state functional connectivity is typically modeled as the correlation structure of whole-brain regional activity. It is studied widely, both to gain insight into the brain's intrinsic organization but also to develop markers sensitive to changes in an individual's cognitive, clinical, and developmental state. Despite this, the origins and drivers of functional connectivity, especially at the level of densely sampled individuals, remain elusive. Here, we leverage novel methodology to decompose functional connectivity into its precise framewise contributions. Using two dense sampling datasets, we investigate the origins of individualized functional connectivity, focusing specifically on the role of brain network "events" – short-lived and peaked patterns of high-amplitude cofluctuations. Here, we develop a statistical test to identify events in empirical recordings. We show that the patterns of cofluctuation expressed during events are repeated across multiple scans of the same individual and represent idiosyncratic variants of template patterns that are expressed at the group level. Lastly, we propose a simple model of functional connectivity based on event cofluctuations, demonstrating that group-averaged cofluctuations are suboptimal for explaining participant-specific connectivity. Our work complements recent studies implicating brief instants of high-amplitude cofluctuations as the primary drivers of static, whole-brain functional connectivity. Our work also extends those studies, demonstrating that cofluctuations during events are individualized, positing a dynamic basis for functional connectivity.

INTRODUCTION

Functional connectivity (FC) measures the temporal correlation of regional BOLD activity, often in the absence of explicit task instructions, i.e. in the "resting state" [1, 2]. Although usually estimated over an extended period of time and using all available data, a growing number of studies have shown that FC can be well approximated using relatively few observations, suggesting that FC may be driven by a temporally sparse process [3–7].

In parallel, a growing body of work has demonstrated that, like fingerprints, FC is unique to each individual and expresses features that reliably distinguish one brain from another [8–12]. These observations hold tremendous translational promise, and open up the possibility of designing personalized interventions [13] and developing increasingly potent connectivity-based biomarkers for cognition, development, and disease [14–16].

However, there remains a key open question: how does FC become individualized in the first place? One possibility is that, like FC itself, personalized information is encoded through time-varying connectivity patterns and distributed dynamically and sparsely throughout a scan session. Indeed, recent findings broadly support this hypothesis [17, 18]. In [19], for instance, we demonstrated that using a small subset of frames classified as "events" – brief and infrequent periods of high-amplitude cofluctuation – we could produce accurate reconstructions of FC while simultaneously rendering participants identifiable, amplifying their connectional fingerprints. In contrast, low-amplitude frames yielded poorer estimates of FC and contained little personalized information.

Although these observations support the hypothesis that personalized information is expressed selectively during high-amplitude frames, they also raise additional theoretical questions. For instance, do cofluctuation patterns during events repeat from one scan to another? If so, do they reflect a single repeating pattern or a repertoire of different patterns? Are these patterns shared across individuals but expressed in different proportions, thereby giving rise to individualized FC? Or does the individualization of FC arise from equally idiosyncratic patterns of high-amplitude cofluctuations? Addressing these questions is critical for linking patterns of brain connectivity with individual differences in behavior [20], and would help clarify the role of brain dynamics in shaping the individualization of FC [21], complementing other approaches that have focused on the collective influence of cortical expansion rates, post-natal experience, and genetics [22].

Here, we address these questions directly. Our approach leverages a recently-proposed method for decomposing FC into its framewise contributions, detecting events, and assessing the impact of events on time-averaged FC [19, 23–26]. We apply this framework to two independently acquired datasets: the Midnight Scan Club [11, 12] and the MyConnectome project [10, 27]. In agreement with our previous studies, we show that FC is accurately reconstructed from event data alone. Next, we focus on the properties of individual events, revealing

* rbetzel @ indiana.edu

that they repeat within and between scans of the same individual. We also show that event cofluctuations can be clustered across participants, revealing broad archetypes that are subtly yet systematically modified at the level of individuals. Finally, we construct a simple model of FC, demonstrating that FC can be predicted with a high level of accuracy using individualized event data, exclusively.

RESULTS

Edge time series as a mathematical precise link between brain dynamics and FC

In this paper, we analyze data from eight participants in the Midnight Scan Club, each scanned ten times (participants MSC08 and MSC09 were dropped due to data quality issues). We analyzed two versions of these data; one in which participants' brains were parcellated into $N = 333$ group-level parcels [28] and another in which parcels were defined on an individual basis, resulting in a different set parcels for each participant ($N = 612 \pm 28$) [29]. The primary analyses were carried out using the group-level parcels. We also analyzed data from the MyConnectome project, a study in which a single individual was scanned > 100 times [10, 27].

For each dataset, we transformed regional fMRI BOLD time series into cofluctuation or edge time series (ETS). Briefly, ETS are calculated as the element-wise product between pairs of z-scored regional (nodal) time series (Fig. 1a; see **Materials and Methods** for details). This operation results in a new time series – one for every node pair (edge) – whose elements index the direction and magnitude of instantaneous cofluctuations between the corresponding pair of brain regions. For instance, if the activity of region i and j deflect above (or below) their time-averaged means at the same instant, the value of the edge time series will be positive. On the other hand, if they deflect in opposite directions, then the edge time series returns a negative value. If one deflects and the other does not, then the value will be close to zero. The temporal mean of an edge time series is equal to the Pearson sample correlation coefficient, and therefore ETS is an exact decomposition of FC into its framewise contributions.

If we calculate edge time series for all pairs of regions, we obtain an edge×time matrix (Fig. 1b) whose temporal average yields a vector (Fig. 1c) that, when reshaped into the upper triangle elements of a node×node matrix, is exactly the FC matrix (Fig. 1d). The edge×time matrix can also be "sliced" temporally and the corresponding vector once again reshaped into the upper triangle elements of a node×node matrix, yielding an instantaneous estimate of whole-brain cofluctuations. These matrices vary in terms of their mean cofluctuations, which we summarize with a root sum of squares measure (Fig. 1e).

In our previous study, we showed that RSS values followed a heavy-tailed distribution, such that a small number of frames exhibited exceptionally high-amplitude RSS [19]. We also demonstrated that FC reconstructed using only these high-amplitude frames accurately recapitulated time-averaged FC, suggesting that FC weights are not driven equally by all frames, but by a select set of frames. We also demonstrated that these high-amplitude frames were underpinned by a principal mode of brain activity, emphasizing oppositional activation of default mode and control networks with sensorimotor and attentional networks. In this paper, however, high-amplitude frames were selected heuristically as the top $P\%$ by RSS value and, beyond the first mode of activity, we did not investigate other activity patterns that occur during events.

A statistical test for high-amplitude cofluctuation events

In previous work, we identified putative cofluctuation events as the top $P\%$ frames in terms of root sum squared (RSS) amplitude of cofluctuation weights. Although this heuristic is pragmatic – it is easy to implement and interpret – it has some unwanted characteristics. Notably, the parameter $P\%$ lacks statistical justification and, due to slow temporal fluctuations and serial correlations in the fMRI BOLD signal, can result in event samples that disproportionately represent only a small number of RSS peaks. Here, we present a simple statistical test to identify events that addresses both of these issues.

In essence, we identify high-amplitude frames by comparing the RSS time series estimated using real data with an ensemble of RSS time series generated under a null model. Here, as a null model we apply the circular shift operator independently, randomly, and bidirectionally to each region's time series, which exactly preserves its mean and variance (and its autocorrelation approximately). We then transform the shifted data into edge time series and estimate their RSS. This step is repeated 100 times yielding 100 sets of surrogate RSS time series, against which we compare the observed RSS data and identify sequences of frames whose RSS exceeds the null distribution (non-parametric permutation test at each frame; accepted false discovery rate fixed at $q = 0.05$; Figure. 2a). This entire procedure is repeated for every participant and every scan.

This procedure allows us to segment the time series into three categories: contiguous frames whose RSS is greater than expected, less than expected, or consistent with that of the null distribution. Rather than consider all frames, we select representative frames from each block for subsequent analysis. For segments whose RSS is greater than that of the null distribution or not significant, we extract peak cofluctuation pattern corresponding to the maximum RSS frame; for segments whose RSS is less than that of the null, we extract the pattern corresponding to the minimum RSS frame (trough).

To demonstrate that these categories of frames cap-

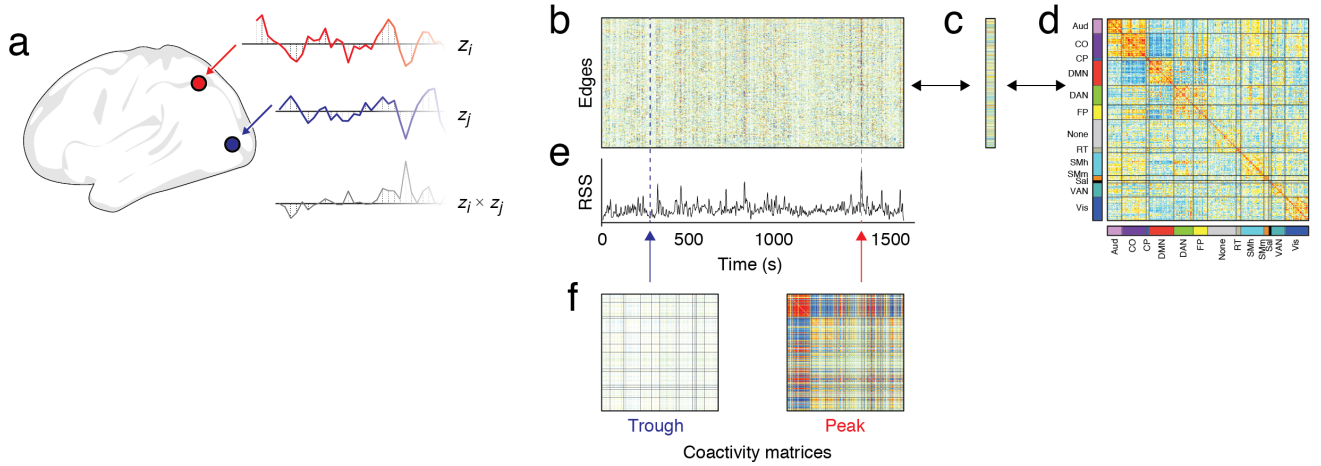

FIG. 1. **Edge time series.** (*a*) An edge time series is constructed for pairs of brain regions, i and j, by computing the elementwise product of their z-scored activities, z_i and z_j, respectively. The result is a new time series, z_{ij}, which indexes the framewise cofluctuations between i and j. (*b*) This procedure can be repeated for all pairs of regions, generating a matrix of edge time series. At each instant in time, a "slice" through this matrix yields a region-by-region cofluctuation matrix that can be modeled as a network. (*c*) At every moment in time we can calculate the root sum of squares (RSS) over all edge time series. The RSS time series is bursty, such that it takes on low values most of the time, but is punctuated by short, intermittent, high-amplitude bursts. (*d*) The temporal average over all edge time series yields a vector that corresponds to the upper triangle elements of a correlation matrix, i.e. functional connectivity (*e*). In this way, edge time series offer a means of tracking moment-to-moment fluctuations in network topology and links them to functional connectivity through an exact decomposition. In *f* we show examples of cofluctuations during a trough and peak (when RSS is small *versus* large).

ture distinct features of cofluctuations, we compare them along several different dimensions (ANOVA; for all comparisons $p < 10^{-15}$). First, we show that, as expected, high-amplitude frames express greater RSS values than low-amplitude and non-significant frames (Fig. 2b). On the other hand, scans contain fewer high-amplitude frames (Fig. 2c), form fewer contiguous segments (Fig. 2d), and exhibit shorter segment length (Fig. 2e) than low-amplitude frames. Consistent with our previous study, reconstructing FC using only high-amplitude frames results in a pattern of FC strongly correlated with the FC estimated using all frames, and greater in magnitude than that of the non-significant and low-amplitude frames (Fig. 2f). Lastly, we find that high-amplitude frames are almost never among those censored for excessive in-scanner motion (Fig. 2g); low-amplitude frames, on the other hand, were more likely to be associated with censored frames. This observation is consistent with our previous study, in which we reported a weak but consistent negative correlation between RSS and framewise displacement. These findings, in general, replicate using individualized parcels for participants in the Midnight Scan Club (Fig. S1) and with MyConnectome data (Fig. S2).

Taken together, these results suggest that the proposed statistical test segments frames into classes with distinct features. This test addresses two concerns associated with previous estimates of high-amplitude "event" frames. First, it defines high-amplitude frames according to a statistical criterion, rather than heuristically. Second, by extracting representative frames from each contiguous segment, we obtain a more heterogeneous sample of high-amplitude frames and avoid selecting multiple frames from around a single peak.

Peak cofluctuation patterns are repeated across scan sessions – troughs are not

In the previous section we presented a simple method for estimating statistically significant cofluctuation events and demonstrated that the peak and trough frames of these segments exhibit distinct spatiotemporal properties. Here, we investigate representative cofluctuation patterns extracted from blocks of high- and low-amplitude frames. We first compare the similarity of these cofluctuation patterns, first within-individuals and later between. Then, we present evidence that cofluctuation patterns expressed during the peaks of high-amplitude events recur across scans of the same individual and that these patterns exhibit subject-specificity.

First, we applied the statistical test to MSC scans (excluding MSC08 and MSC09 due to data quality issues). We found that each scan included 65.9 ± 9.2 and 72.2 ± 17.1 high- and low-amplitude segments, respectively (paired sample t-test; $p = 0.0017$; $t(79) = 3.24$). After additional quality control in which we excluded segments that included any motion-censored frames, the number of segments whose RSS was significantly greater than the null changed little (61.26 ± 14.7). However, the

number of segments with lower-than-expected RSS was reduced dramatically (52.5±23.9), reflecting the fact that those frames often coincide with periods of excessive in-scanner motion.

Next, we calculated the spatial similarity of motion-free cofluctuation patterns extracted during RSS peaks and troughs. We performed this analysis separately for each subject, resulting in eight similarity (correlation) matrices. We grouped these values based on whether similarity was measured between two peaks, two troughs, or a peak and trough co-fluctuation pattern. We found peak-peak similarity was significantly greater than trough-trough and peak-trough (t-test; $p < 10^{-15}$) (Fig. 3a,b). We see an identical effect in the MyConnectome data (Fig. S3a-c) and when parcels are individualized (Fig. S3d). These observations suggest that high-amplitude events encode subject-specific patterns of cofluctuations.

Based on these observations, along with the fact that high-amplitude events are less likely to be impacted by motion, we calculated the spatial similarity of peak cofluctuation patterns for all pairs of detected events, for all scans, and for all subjects (Fig. 3c). We then compared these similarity values based on whether they came from the same or different participants (Fig. 3d). We found that within-individual similarity exceeded between-individual similarity (non-parametric permutation test, $p < 10^{-15}$; Fig. 3e,f).

Collectively, these results are in line with our previous study and suggest that low-amplitude cofluctuations contains little participant-specific information [19]. Rather, our findings support the hypothesis that high-amplitude cofluctuations contribute significantly more information about an individual than low-amplitude cofluctuations.

High-amplitude events can be divided into distinct communities based on their cofluctuation patterns

In the previous section we found that co-fluctuation patterns expressed during peaks of high-amplitude event segments are not related to motion and that they are repeatable across scans. This is in contrast to co-fluctuation patterns expressed during low-amplitude segments, which tend to coincide with excessive in-scanner motion and are dissimilar across scans and individuals, even within the same scan session. These observations motivated us to focus on high-amplitude events in yet greater detail. In this section, we test whether the cofluctuation patterns expressed during high-amplitude events are entirely subject-specific and not shared across individuals or whether they belong to a general archetype that is fine-tuned to single participants.

To test whether this is the case, we aggregated across participants all cofluctuation patterns that occurred during event peaks, calculated the similarity matrix of those patterns, and clustered this matrix using a variant of multi-resolution consensus clustering (modularity maximization with a uniform null model [30, 31]; see **Materials and Methods** for details). The results of this analysis yielded two large communities (clusters) along with many very small communities. We found that every participant was represented in the two large communities (labeled 1 and 2 in Fig. 4a) and that instances of those communities appeared in every scan session, accounting for 54.1% and 19.0% of all event peaks, respectively. Every participant was also represented in the two next-largest communities, although they appeared infrequently across scan sessions and collectively accounted for only 8.6% of all event peaks. Accordingly, we aggregated the smaller communities to form a third larger community (labeled 3 in Fig. 4a). Most subsequent analyses will focus on communities 1 and 2 unless otherwise noted. For completeness, we analyze community 3 in greater detail in the **Supplementary Material** (see Fig. S4 and Fig. S5).

To better understand why certain cofluctuation patterns were grouped together, we examined group-representative centroids for each community by calculating the mean cofluctuation pattern of all frames assigned to that community (Fig. 4b,c). We found that community 1 reflected a topology that expressed strong and anticorrelated cofluctuations mostly between cingulo-opercular and default mode networks (Fig. 4b). Community 2, on the other hand, expressed strong anticorrelations between the default mode with dorsal attention and fronto-parietal networks (Fig. 4b). See Fig. S6 for topographic depiction of systems on cortical surface.

To better understand how brain *activity* drives the cofluctuation patterns described above, we extracted activation profiles (regional BOLD activity, rather than cofluctuations) during event peaks, grouped them by community, and performed principal components analysis. The first principal component for each community represents the mode of fMRI BOLD activity that tended to occur during frames assigned to that community. We found the first principal components of communities 1 and 2 (accounting for 24% and 26% of variance, respectively) to be uncorrelated ($r = -0.025$; $p = 0.65$; Fig. 4d). The first principal component for community 1 (Fig. 4e) exhibited significant activation of the default mode and inactivation of cingulo-opercular, visual, and somatomotor networks (spin test; false discovery rate fixed at 5%; $p_{adj} = 5.3 \times 10^{-4}$; Fig. 4f). The first principal component for community 2 (Fig. 4h) exhibited significant activation of dorsal attention and fronto-parietal networks and inactivation of default mode, ventral attention, and visual networks (spin test; false discovery rate fixed at 5%; $p_{adj} = 0.02$; Fig. 4f). We show individual-level principal components in the **Supplementary Material** (Fig. S7) and derive similar modes of activity using an alternative procedure (Fig. S8). Note that the PCA analysis results in modes of activity; in all cases, the sign of PCs can be flipped and result in the same pattern of co-activity.

In the **Supplementary Material** we perform a similar analysis of Midnight Scan Club data parcels fit to each

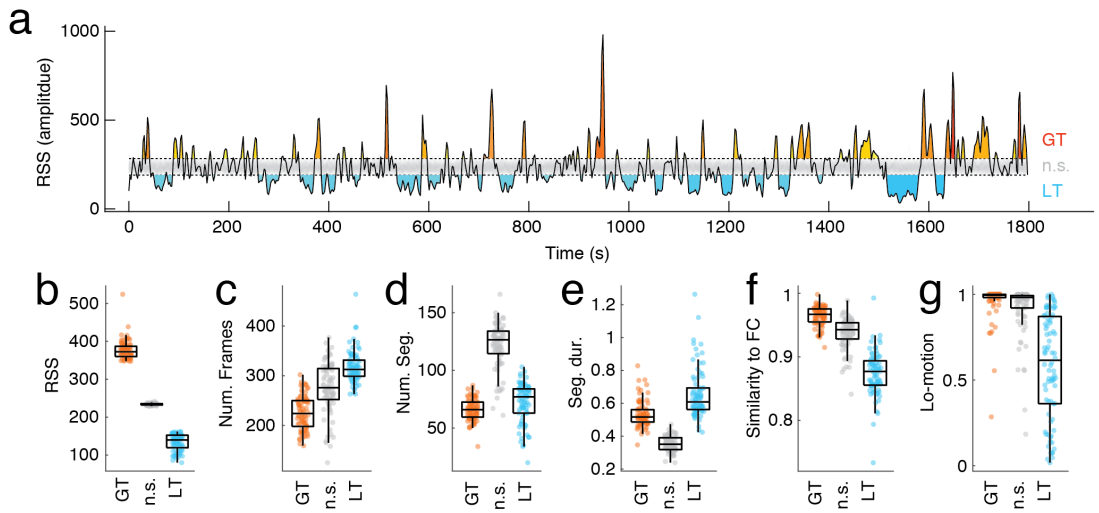

FIG. 2. **A statistical test for network-wide events.** (*a*) We generated edge time series and computed the RSS time series. We compared this time series to a null RSS time series estimated from edge time series that had been generated after circularly shifting the original fMRI BOLD time series. At each point in time, we calculated the probability that the observed RSS value exceeded the null distribution, controlled for multiple comparisons, and identified sequences of frames that exceeded the null distribution. This allowed us to categorize time points into three classes: those whose RSS was greater than null (GT), those whose RSS was significantly less than the null (LT), and those that were in between (n.s.). For subsequent analysis, we extracted a single representative cofluctuation pattern for every contiguous sequence of frames that was greater/less than the null distribution. This pattern corresponded to the frame with the maximum/minimum RSS. In panels *b*-*g*, we separate the frames of each scan into these three classes and compare their features to one another. Each point represents the mean value over all frames assigned to a given class. The features that are compared are: *b* mean RSS, *c*, the number of frames assigned to a given class, *d* the number of contiguous sequences of each class, *e* the mean duration of sequences (log_{10} transformed), *f* the similarity of time-average FC with FC reconstructed using only frames assigned to each class, and *g* the fraction of frames censored for high levels of in-scanner motion.

participant individually. We show that, when we examine activity during RSS peaks defined from the group level parcellations, the activation patterns using individualized parcels closely match those reported here (Fig. S9). Even when we perform event detection separately using the individualized parcels, we find similar modes of activity (mean similarity of $r = 0.89$; $p < 10^{-11}$; Fig. S10). Note that because the number of parcels differ across individuals, this comparison was carried out at a system level.

Again, we use MyConnectome data as a replication dataset, finding analogous communities (Fig. S11). We also take advantage of the fact that MyConnectome data includes many more samples from an individual than the Midnight Scan Club (84 scans versus 10), to identify several communities not evident in the group analysis of the Might Scan Club data. We also find evidence of interdigitated communities with similar system-level profiles but drawing on different regions from within those systems (Fig. S12).

These observations build on our previous study, which focused on a single pattern of cofluctuation during high-amplitude frames. Here, we use data-driven methods to show that high-amplitude cofluctuation is not monolithic and can be divided into meaningful sub-patterns, each driven by a distinct mode of brain activity.

Event communities are individualized

In the previous section, we showed that cofluctuation patterns expressed at the peaks of high-amplitude events could be grouped into meaningful communities. Within each community, are these patterns individualized or are they shared across participants? If they are individualized, which node pairs contribute the most?

To address this question, we analyzed communities 1 and 2 separately and in greater detail. Although collectively each community is cohesive (similarity is greater among cofluctuation patterns assigned to the same community than to other communities; see Fig. S13), we also found evidence that the similarity between cofluctuation patterns are stronger still when they come from the same participant (t-test comparing within- and between-individual similarity; $p < 10^{-15}$; Fig. 5a-d). Further, we identified the pairs of brain regions whose cofluctuations were most variable across individuals by computing the standard deviation of edge weights across participant centroids (Fig. 5e-f). For the cofluctuation pattern expressed by community 1, we found that the most variable

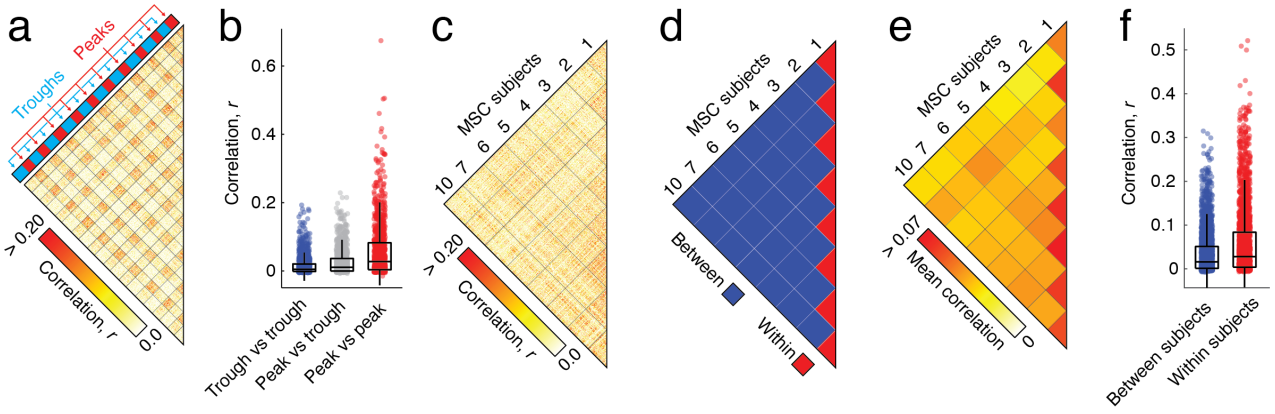

FIG. 3. **Intra-/Inter-individual similarity of cofluctuation patterns.** (*a*) For each individual separately, we aggregated all low-motion cofluctuation patterns during peaks and troughs. We then computed the similarity between cofluctuation patterns (data from participant MSC06 is shown here as an example). (*b*) Boxplot showing similarity values broken down by trough *versus* trough, peak *versus* trough, and peak *versus* peak. Because only cofluctuation at peaks exhibited similarity across scans, we focused on these patterns only, discarding cofluctuation that occurs during troughs and focusing on comparisons of participants to one another (*c*). We found that cofluctuation at peaks tended to be more similar within participants than between. We show the raw similarity matrix in *d* and the averaged values in *e*. (*f*) Boxplot of similarity values broken down according to whether they fell within or between participants.

edges linked the cingulo-opercular network to the dorsal attention, fronto-parietal, and ventral attention networks (Fig. 5e). In the case of community 2, the most variable edges were linked to default mode, dorsal attention, and fronto-parietal.

These observations suggest that high-amplitude cofluctuation patterns reflect a topology that is broadly shared across individuals but is systematically and individually altered.

Functional connectivity is modeled accurately only when using individual-specific co-fluctuation patterns

Here and in [19], we argued that network dynamics are bursty such that short-lived, high-amplitude, cofluctuation "events" contribute significantly to static, time-averaged FC. Here, we test this hypothesis directly by modeling participant- and scan-specific FC as a function of the cofluctuation patterns generated in the previous section.

Motivated by previous studies showing that FC is driven by high-amplitude frames [3, 4, 6, 7, 19], our model assumes that FC is driven exclusively by high-amplitude events and that frames with low RSS make no contribution (Fig. 6a). Our model generates predictions of participant- and scan-specific FC. It works by calculating, for that scan, the fraction of time points assigned to communities 1, 2, and 3, and uses those values to weight the centroids of each community (Fig. 6b). In also counts the fraction of time during the scan spent in a low-amplitude frame, using that value to weight a matrix of zeros (corresponding to the hypothesis that low-amplitude frames make a negligible contribution).

Although the fraction of time spent in a given community is fixed, the centroids can be defined in several different ways. To predict the FC of any given participant and scan, we can estimate centroids: (a) using group all available data from all participants and scans combined, (b) using group data from all participants and scans excluding the scan we are trying to predict, (c) using all available data from the participant whose FC we are trying to predict, (d) using all available data from that participant but excluding data from the scan we are trying to predict, and (e) using data from a different participant. In all cases end, the model generates a $[N \times N]$ matrix of predicted connection weights.

In general, we found heterogeneity across models in terms of their performance ($ANOVA$; $p < 10^{-15}$; Fig. 6c). The model in which FC is predicted using data from different participants performed the worse ($r = 0.52 \pm 0.05$), followed by the two models in which centroids were estimated using group data ($r = 0.69 \pm 0.05$ and $r = 0.70 \pm 0.05$). However, the two participant-specific models outperformed all others by a wide margin ($r = 0.86 \pm 0.05$ and $r = 0.83 \pm 0.06$; paired sample t-test of mean average performance on both participant-specific models *versus* average performance on other three models; $p < 10^{-15}$). In the **Supplementary Material** we perform a sensitivity analysis to identify which of the three communities drive these effects. For all models, we find that model performance suffers the most by removing community 1, which alone accounts for 54.1% of all events (Fig. S14).

These results suggest that the cofluctuation patterns

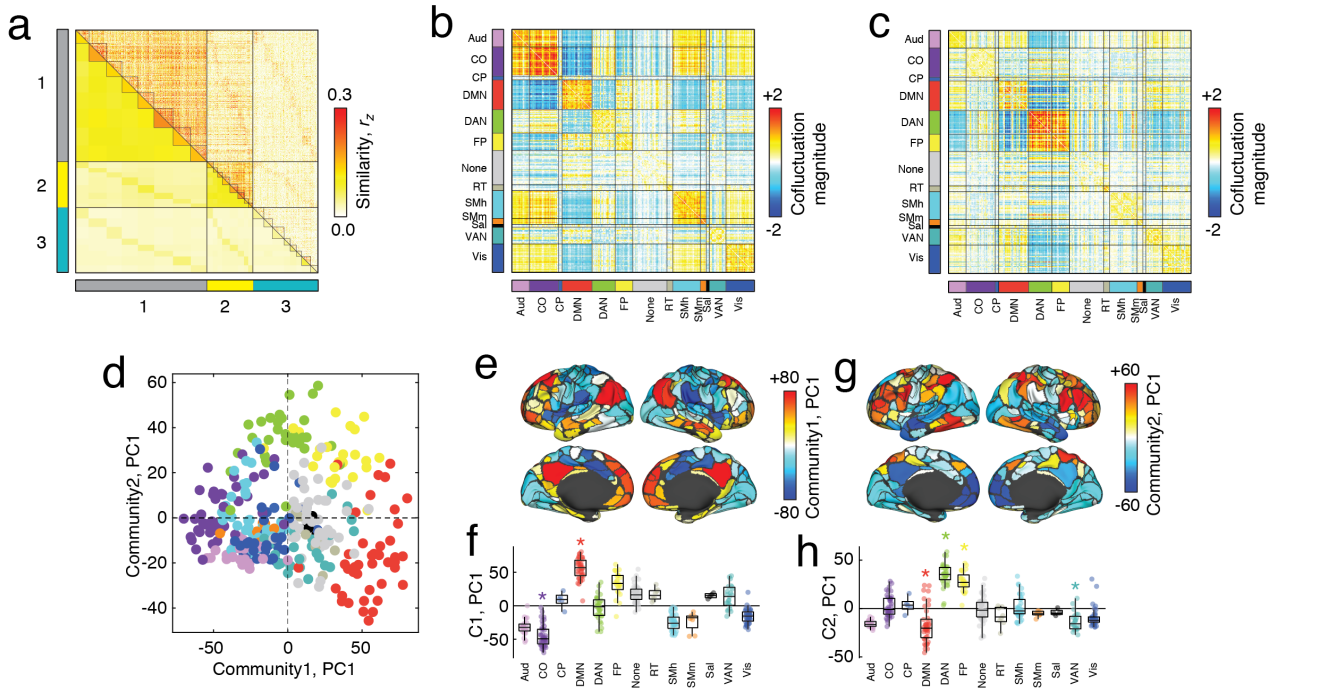

FIG. 4. **Clustering peak cofluctuation.** (*a*) We clustered the cofluctuation similarity matrix using a multi-scale consensus clustering algorithm, resulting in two large communities (1 and 2) and a third set of much smaller communities, grouped together here to form community 3. Here, black lines divide communities from one another and, internally, participants from one another. The mean cofluctuation pattern for communities 1 and 2 are shown in *b* and *c*. To understand activity that underpins each community, we pooled together corresponding activity time series separately for communities 1 and 2 and performed principal component analysis on each set, returning the primary mode of activity (PC1). (*d*) Scatterplot of PC1 for community 1 and community 2. Colors denote brain systems. (*e*) Topographic depiction of PC1 for community 1. (*f*) PC1 grouped according to brain system. Asterisks indicate $p < p_{adj}$ (FDR fixed at $q = 0.05$). Panels *g* and *h* show corresponding plots for community 2.

expressed at the peaks of high-amplitude events can, on their own, explain a significant fraction of participant-specific variance in time-averaged FC. More importantly, our results reaffirm that cofluctuation patterns during event peaks are participant specific; even when predicting a held-out scan, when centroids are estimated using participant-specific data, we find marked improvement in model performance.

DISCUSSION

Here, we extended our recent analyses of edge time series and putative high-amplitude cofluctuation events [19, 26]. We proposed a simple null model that allowed us to identify frames whose amplitude was significantly greater or less than chance. We then analyzed the cofluctuation patterns expressed during these frames, discovering that across scans of the same individual, the cofluctuation patterns expressed during frames with lower-than-expected amplitude were dissimilar. In contrast, we found that the cofluctuation during high-amplitude frames was repeatable within an individual and dissimilar between individuals. We then clustered patterns of cofluctuation expressed during high-amplitude frames, identifying a small number of cofluctuation patterns that were shared across individuals. These patterns, however, were altered subtly yet systematically, so that they could be used to reliably distinguish participants from one another. Finally, we tested the hypothesis that FC could be predicted exclusively from the co-fluctuation patterns expressed during events, and constructed a model that generated estimates of FC given a set of cofluctuation community centroids and the frequency that those centroids are expressed by that individual. We found that the model performed well only when the centroids were estimated using data from the participants whose FC we were aiming to predict.

High-amplitude cofluctuations can be partitioned into different communities based on their topology

Intrinsic or resting-state functional connectivity reflects the coupling of spontaneous activity between distant brain regions [1, 32]. It is often used to construct a

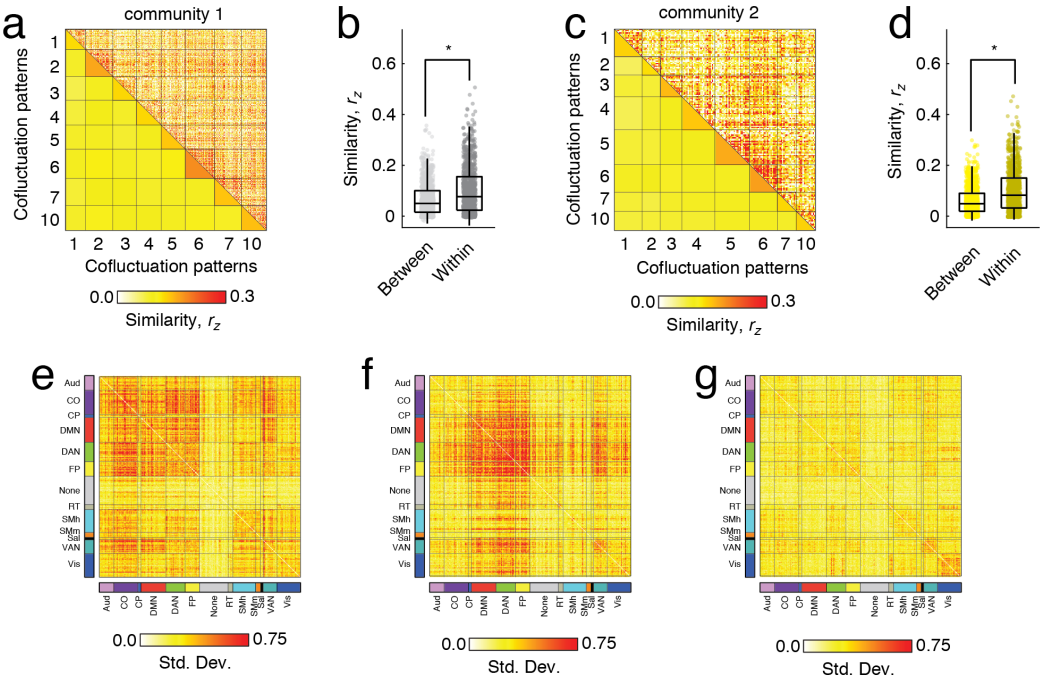

FIG. 5. **Personalization of cofluctuation patterns.** We separately extracted cofluctuation patterns for communities 1 and 2 and computed the pairwise similarity. Similarity matrices are shown in panels *a* and *c*. In both cases, we found that within-individual similarity was statistically greater than between-individual similarity. Boxplots of similarity scores are depicted in panels *b* and *d*. Asterisks indicate $p < p_{adj}$ based on permutation test (FDR fixed at $q = 0.05$). For each community, we calculated mean cofluctuation patterns for each participant and, across participants, computed the standard deviation of each node-pairs cofluctuation magnitude. The patterns of variability for communities 1, 2, and 3 are depicted in panels *e*, *f*, and *g*, respectively.

graphical representation of the brain to be analyzed using tools from network science [33, 34]. Although inter-individual differences in FC have been linked to an individual's clinical [16], cognitive [35], and developmental state [15], the dynamic origins of individualized FC remain unknown.

Recently, we presented a method for decomposing FC into its framewise (instantaneous) contributions [19, 23]. Our work, in agreement with other recent studies [3, 4, 6, 7], demonstrated that all frames do not contribute equally to FC – rather only a small number of high-amplitude frames – "events" – when averaged together, are necessary for explaining a high proportion of variance in FC. In that study, however, we only examined the mean pattern of high-amplitude cofluctuations and did not investigate individual events nor did we characterize variation in the co-fluctuation patterns across events.

Here, we address these issues using two dense phenotyping datasets. Leveraging a statistical test for identifying high-amplitude frames, we show that "events" are *not* monolithic and comprise distinct patterns of cofluctuations. Using a data-driven clustering method, we find evidence of two patterns of cofluctuation that are conserved across all participants and scans. These patterns emphasize opposed activation of default mode regions with cingulo-opercular and sensorimotor systems (community 1) and with dorsal attention and fronto-parietal systems (community 2). We also find evidence of smaller communities corresponding to less frequent events involving only a fraction of participants. In the main text we grouped these patterns into a single community, but find that the two largest (accounting for 8.6% of event peaks) involve cofluctuations of sensorimotor systems, which we explore in the **Supplementary Material**.

The two largest communities have interesting properties. For instance, the default mode is cohesive (densely interconnected, internally) in both, but is selectively decoupled from distinct sets of brain systems associated with processing sensorimotor information [36] and coordinating flexible, goal-directed behavior [37]. In neither community does the default mode couple strongly to other systems. Rather, in these high-amplitude states it maintains relative autonomy, in agreement with studies that have examined its "hubness" using the participation coefficient [38] – a graph-theoretic measure that describes the extent to which a node's connections are distributed across or concentrated within communities [39].

Another interesting feature of these communities is their possible relationship to network states obtained by clustering sliding-window estimates of time-varying FC

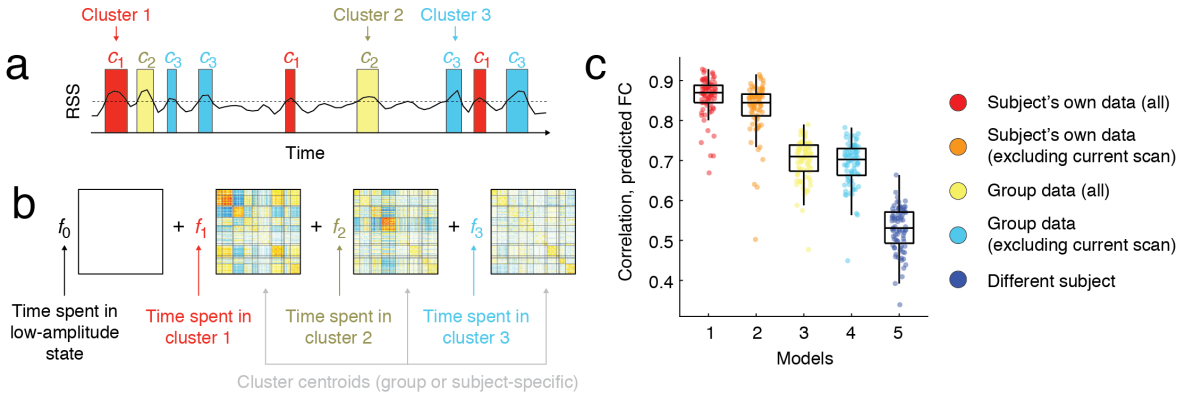

FIG. 6. **Modeling participant- and scan-specific FC with cofluctuation patterns.** (*a*) We hypothesized that FC is driven by brief high-amplitude patterns of cofluctuation. A scan session can be abstracted, then, as periods of time where cofluctuation is close to zero punctuated by periods of time where cofluctuation corresponds to one of the cofluctuation communities. (*b*) We modeled FC as a linear combination of cofluctuation patterns corresponding to communities 1, 2, and 3 as well as a blank state where cofluctuation was treated as zeros (low-amplitude frames). We varied the cofluctuation patterns so that they came either from the participant whose FC we were modeling and were estimated using all data or data excluding the current scan. We tested a variant of the model where the cofluctuation patterns were estimated at the group level and therefore contained information from all participants. Again, we tested versions of this model where cofluctuation patterns were estimated using all available data or all available data excluding the current scan. Lastly, we tested a version of the model where data from a different participant was used to predict FC. (*c*) Performance of the five models. Overall, models that included participant-specific information outperformed other models.

[40] or co-activation patterns (CAPs) [41]. We note, however, that these relationships may be superficial, as there are several key methodological differences. Here, we analyze edge time series, a parameter-free method for estimating instantaneous cofluctuations between regional activity (localized to individual frames) and whose sum is precisely time-averaged FC. In contrast, sliding-window estimates of time-varying FC require users to specify a window duration and overlap fraction (the number of frames shared by successive estimates of FC) and, due to the sliding window, lead to temporally blurred connectivity estimates that cannot be precisely localized in time [42, 43]. In fact, we speculate that brief high-amplitude events may be present in sliding-window estimates of time-varying FC, but because they evolve over timescales much shorter than that of the typical window duration, are effectively obscured due to blurring. CAPs, on the other hand, identifies whole-brain cofluctuation patterns in frames when a seed voxel's activity exceeds some threshold, parameters selected by the user. Critically, because edge time series are a decomposition of FC, events can be viewed as the "atoms" or "building blocks" of FC.

Collectively, our findings suggest that cofluctuation patterns expressed during putative network-wide events are variable but can be described in terms of two principal patterns. These findings extend our previous study [19] and open up opportunities for future studies to investigate inter-individual differences in these patterns as well as the smaller and less frequent patterns.

High-amplitude cofluctuation patterns are individualized and drive time-averaged FC

One of the questions we aimed to address was whether the individualization of FC occurred because: *a*) high-amplitude cofluctuation patterns are shared across individuals but expressed in different subject-specific proportions or *b*) cofluctuation patterns expressed during high-amplitude events are inherently subject-specific. The answer to this question is important, as it speaks to the origins of individual differences in FC [22, 44], has implications for brain-behavior studies [14], and also informs our understanding of time-varying FC [45].

Here, we addressed these questions by aggregating and clustering high-amplitude cofluctuation patterns from across all participants. This analysis returned two large communities in which every participant and scan were represented, indicating that, to some extent, patterns of high-amplitude cofluctuations are indeed shared across individuals. However, when we examined these communities in greater detail, we found that within communities there existed more cohesive sub-communities corresponding to individual participants.

To better adjudicate between hypotheses, we constructed a simple model to predict an individual's scan-specific pattern of FC. Motivated by previous studies [3, 4, 6, 7], this model assumed that FC is driven exclusively by high-amplitude cofluctuations and that all other time points made negligible contributions to FC. We then replaced FC during high-amplitude frames with the centroid of the community to which those frames were

assigned. We found that, when centroids were generated using data pooled from across all participants we could explain only ≈50% of the variance in functional connection weights. However, when we estimated centroids using data from the same participant whose FC we were predicting, the model exhibited a significant increase in performance, accounting for ≈75% of variance in connection weights.

These observations suggest that incorrectly ascribing group-level features to an individual participant distorts our prediction of their FC, favoring the hypothesis that high-amplitude cofluctuation patterns are individualized. This endorsement of hypothesis *b* comes with some caveats, however. Although FC is the product of individualized cofluctuation patterns, those patterns appear to be variants of archetypal patterns, i.e. the two large communities discussed in the previous section. These observations align with other studies showing that the individualization of FC is generally a subtle modulation of features that are evident in group-level data, from brain systems [46–48] to regional FC fingerprints [8].

Our results have implications for studies of brain-behavior correlations as well as state-based analyses of time-varying FC. We show that inter-individual differences in FC are largely shaped by differences in high-amplitude cofluctuation patterns. However, our sensitivity analysis (Fig. S14) demonstrated that of the three communities we considered, one contributed disproportionately to the individualization of FC relative to the other two. This suggests that, rather than linking inter-individual differences in FC across individuals, it may be more profitable to directly investigate specific community centroids, e.g. those that drive individual variation, potentially leading to improvements in brain-behavior correlations.

Our results also have implications for studies of network states in time-varying FC [40, 43]. In general, these studies cluster time points together based on the similarity of networks to one another. To facilitate ease of comparison across individuals, this step is usually performed using concatenated data from many participants or conditions. Different metrics can be calculated from these partitions, e.g. cumulative time a participant spends in any cluster, transition matrices, etc., and linked to behavioral and clinical phenotypes. Our results suggest that, although methodologically convenient, clustering time points together and treating them as recurrences of the same "network state" likely obscures meaningful participant-level variation.

Future directions

The results presented here raise important questions that should be investigated in future research. First, because events contribute disproportionately to the organization of time-averaged FC and because they appear to be drivers of individualization, they should be the target of future studies. We investigated events in two dense sampling studies and in a total of nine brains. Although these data allowed us to investigate the extent to which cofluctuation patterns during events are shared *versus* individualized, the small number of participants precludes the possibility of investigating behavioral, cognitive, or disease correlates of events [49]. Future studies should investigate communities of high-amplitude events in larger datasets.

Relatedly, our study examined event structure exclusively during task-free resting-state conditions. We demonstrated that time-averaged FC could be well-approximated using only high-amplitude frames and individualized estimates of cofluctuations during those frames. What happens to events when participants are asked to perform tasks in the scanner? Previous studies have demonstrated that tasks systematically modulates patterns of FC [35]. Do these changes reflect different patterns of cofluctuations during events? Are they the same patterns as rest but in different proportions? And directly related to the aims of this study, are task events similarly personalized or can they be used to strengthen brain-behavior associations [50]?

Here, we focused on the contributions of high-amplitude "events" to patterns of time-averaged FC. The simple model we proposed even goes so far as to consider contributions from all other frames as negligible. Is this really the case? What biases might we reinforce by focusing on high-amplitude frames? High-amplitude cofluctuations make proportionally bigger contributions to time-averaged FC than low-amplitude cofluctuations. This statement is non-controversial; edge time series are a mathematically precise "temporal unwrapping" of the Pearson correlation into its framewise contributions, the average of which is simply FC [19, 23–26]. For this reason, it makes sense to focus on frames where many edges simultaneously make big contributions – those same frames necessarily will, on average, make bigger contributions to FC than, say, frames where only a few edges exhibit high-amplitude edge time series. However, this does not rule out the possibility that frames outside of high-RSS events, which are more numerous, make contributions that outweigh or match those of high-amplitude frames. Additionally, in focusing on *global* high-amplitude events, we may miss out on events involving small brain systems, which will fail to meet statistical criteria for significance due to their size.

Finally, a key overarching and open question concerns the origins of high-amplitude cofluctuations. In our previous study, we demonstrated that movie-watching leads to synchronization of events across participants [19], suggesting that their timing can be modulated selectively, in that case by sensory input. But what about rest? In that previous study, we found no differences in event amplitude between rest and movie-watching, suggesting that spontaneous events are just as large as those driven by sensory input. A couple recent studies help us speculate on the origins of events. One possibility is that

events help preserve brain circuit function in the absence of use. In [51], the authors demonstrated that disuse of motor circuits by casting participants' arms leads to increases in high-amplitude pulses and that manually censoring these pulses reduces FC magnitude [52]. Another possibility is that burst activation of distinct systems, including default mode networks, co-occurs with the recall or previously-observed stimuli [53]. In both cases, these activation patterns are likely constrained by the underlying anatomical connectivity and reflect groups of mutually connected brain regions [54]. A final possibility is that events are truly stochastic and are mere bi-products of modular correlated time series, where the activation of one element within a module implies the activation of the others. In any case, future experimental studies – possibly invasive studies that allow for more targeted and temporally resolved recording – should be directed to investigate their origins.

Limitations

The conclusions of this study are limited in several ways. Notably, we analyze fMRI data, which affords whole-brain coverage but at a spatial resolution of millimeters and a temporal resolution of, at best tenths of a second. Moreover, fMRI BOLD samples a slow signal slowly, and is only indirectly related to population activity. However, many studies have demonstrated a correspondence between BOLD fluctuations and FC estimated from fMRI with other modalities, including local field potentials [55], intracranial EEG [56], and optically recorded calcium imaging signals [57], positing a neural basis for functional connectivity estimated from the fMRI BOLD signal. Future studies could apply methods similar to those used here to other imaging modalities.

A second limitation concerns data quality, in-scanner motion, and other artifacts, which are known to impact estimates of FC and can produce burst-like behavior in fMRI time series [58, 59]. Here, we adopted a conservative approach and discarded putative events that occurred near censored frames. Notably, this procedure impacted low-amplitude frames to a greater extent than high-amplitude frames, suggesting that high-amplitude frames, in addition to contributing disproportionately to FC, are also less likely to be contaminated by artifacts. Nonetheless, how to adequately address motion-related issues in the analysis of FC remains an ongoing and disputed topic [60–62].

A final limitation concerns the use of PCA to extract "modes" of activity underlying communities. While we find that each community's first component explains ≈25% variance, we do not investigate the remaining components, which likely include other meaningful modes of activity. Future studies should further investigate the link between brain activity and connectivity. Because edge time series is a mathematically exact decomposition of FC into its time-varying contributions, it represents a useful framework for doing so.

Conclusion

In conclusion, we find that FC can be explained using a small number of high-amplitude frames. These frames can be clustered into a small number of communities corresponding to archetypal patterns of cofluctuation, broadly shared across individuals. However, these patterns undergo refinement at the level of individual participants, yielding reliable and individualized cofluctuation patterns. Finally, we show that participants' FC is more accurately predicted using participant-specific estimates of their high-amplitude cofluctuation patterns compared to group-level estimates. Our study discloses high-amplitude, network-wide cofluctuations as dynamical drivers of individualized FC and introduces methodology for exploring their role in cognition, development, and disease in future studies.

MATERIALS AND METHODS

Datasets

Midnight Scan Club

The description of the Midnight Scan Club dataset acquisition, pre-processing, and network modeling is described in detail in [11]. Here, we provide a high-level overview. Data were collected from ten healthy, right-handed, young adult participants (5 females; age: 24-34). Participants were recruited from the Washington University community. Informed consent was obtained from all participants. The study was approved by the Washington University School of Medicine Human Studies Committee and Institutional Review Board. This dataset was previously reported in [11, 12] and is publicly available at `https://openneuro.org/datasets/ds000224/versions/00002`. Imaging for each participant was performed on a Siemens TRIO 3T MRI scanner over the course of 12 sessions conducted on separate days, each beginning at midnight. In total, four T1-weighted images, four T2-weighted images, and 5 hours of resting-state BOLD fMRI were collected from each participant. For further details regarding data acquisition parameters, see [11].

High-resolution structural MRI data were averaged together, and the average T1 images were used to generate hand-edited cortical surfaces using Freesurfer [63]. The resulting surfaces were registered into fs_LR_32k surface space as described in [64]. Separately, an average native T1-to-Talaraich [65] volumetric atlas transform was calculated. That transform was applied to the fs_LR_32k surfaces to put them into Talaraich volumetric space.

Volumetric fMRI pre-processing included including slice-timing correction, frame-to-frame alignment to cor-

rect for motion, intensity normalization to mode 1000, registration to the T2 image (which was registered to the high-resolution T1 anatomical image, which in turn had been previously registered to the template space), and distortion correction [11]. Registration, atlas transformation, resampling to 3 mm isotropic resolution, and distortion correction were all combined and applied in a single transformation step [66]. Subsequent steps were all completed on the atlas transformed and resampled data.

Several connectivity-specific steps were included (see [67]): (1) demeaning and de-trending of the data, (2) nuisance regression of signals from white matter, cerebrospinal fluid, and the global signal, (3) removal of high motion frames (with framewise displacement (FD) > 0.2 mm; see [11]) and their interpolation using power-spectral matched data, and (4) bandpass filtering (0.009 Hz to 0.08 Hz). Functional data were sampled to the cortical surface and smoothed (Gaussian kernel, $\sigma = 2.55$ mm) with 2-D geodesic smoothing.

The following steps were also undertaken to reduce contributions from non-neuronal sources [67, 68]. First, motion-contaminated frames were flagged. Two participants (MSC03 and MSC10) had high-frequency artifacts in the motion estimates calculated in the phase encode (anterior-posterior) direction. Motion estimate time courses were filtered in this direction to retain effects occurring below 0.1 Hz. Motion contaminated volumes were then identified by frame-by-frame displacement (FD, described in [69]), calculated as the sum of absolute values of the differentials of the 3 translational motion parameters (including one filtered parameter) and 3 rotational motion parameters. Frames with FD > 0.2 mm were flagged as motion-contaminated. Across all participants, these masks censored 28%±18% (range: 6% – 67%) of the data; on average, participants retained 5929±1508 volumes (range: 2733 – 7667). Note that in this paradigm, even the worst participant retained almost two hours of data.

Time courses were extracted from $N = 333$ cortical regions using a common (group) functional parcellation [28]. We also analyze time courses estimated from using individualized parcellations (see [70] for details). Both group and individualized time series were used for FC estimation and edge time series generation.

MyConnectome dataset

All data and cortical surface files are freely available and were obtained from the *MyConnectome Project*'s data-sharing webpage (`http://myconnectome.org/wp/data-sharing/`). Specifically, we studied pre-processed parcel fMRI time series for scan sessions 14–104. Details of the pre-processing procedure have been described elsewhere [10, 27]. Each session consisted of 518 time points during which the average fMRI BOLD signal was measured for $N = 630$ parcels or regions of interest (ROIs). With a TR of 1.16 s, the analyzed segment of each session was approximately 10 minutes long.

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 $\mathbf{z}_i = [z_i(1), \ldots, z_i(T)]$, where T is the number of samples. The Pearson correlation is then calculated as:

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

In other words, the correlation is equal to the temporal average of two regions' cofluctuation.

Edge time series

Recently, we proposed a method for decomposing FC into its framewise contributions. This is accomplished by simply omitting the averaging step in computing Pearson's correlation. This omission results in a new time series:

$$r_{ij}(t) = z_i(t) \cdot z_j(t). \tag{2}$$

where the value of $r_{ij}(t)$ indexes the instantaneous cofluctuation between regions i and j at time t. When regions i and j both deflect from their mean in the same direction $r_{ij} > 0$, when they deflect in opposite direction, $r_{ij} < 0$, and when one (or both) of their activities is near their mean then $r_{ij} \approx 0$. Importantly, the mean of this time series is exactly equal to the FC between regions i and j, and therefore we can think of $r_{ij}(t)$ as the instantaneous contribution of frame t to the overall FC.

If we consider the set of cofluctuation between all pairs of regions $\{ij\}$ at time t, we can arrange those elements into a node-by-node connectivity matrix and analyze it as a network. We can also calculate the total amplitude of cofluctuation between all node pairs as their root sum square, $RSS(t) = \sqrt{\sum_{i,j>i} r_{ij}(t)^2}$.

Multiresolution consensus clustering

We used a variation of modularity maximization [71] to group cofluctuation patterns into clusters or "communities". Briefly, modularity maximization is a community detection method for partitioning relational data,

e.g. networks, into non-overlapping communities. This is accomplished by optimizing a modularity quality function:

$$Q(\gamma) = \sum_{rs}[S_{rs} - \gamma P_{rs}]\delta(z_r, z_s). \quad (3)$$

In this expression, S_{rs} is the similarity of cofluctuation patterns r and s; P_{rs} is the level of similarity expected by chance; $\delta(z_r, z_s)$ is the Kronecker delta function and is equal to 1 when the community assignments of patterns r and s, denoted as z_r and z_s, are identical, and 0 otherwise. The structural resolution parameter, γ, controls the importance of S_{rs} relative to P_{rs} and, in effect, can be tuned to recover smaller or larger communities.

Here, we use modularity maximization to obtain a representative set of multi-resolution communities [30]. That is, a partition of cofluctuation patterns into communities that takes into account how strongly coupled patterns are to one another at different scales, from finescale partitions of patterns into many small communities to coarse partitions of patterns into a few large communities. To do this, we sample communities at various scales by changing the value of γ. Specifically, we sample 10,000 different values of γ based from the distribution of S_{rs} values. At each value, we use a generalized version of the Louvain algorithm to optimize the corresponding $Q(\gamma)$ [31].

This procedure results in 10000 estimates of communities at a range of scales. We transform these estimates into a probabilistic co-assignment matrix, whose element T_{rs} is equal to the fraction of the 10000 partitions in which nodes r and s were assigned to the same community. From this matrix, we construct a new modularity:

$$Q^c = \sum_{rs}[T_{rs} - P^c_{rs}]\delta(z^c_r, z^c_s). \quad (4)$$

In this expression, P^c_{rs} is the expected co-assignment of r and s to the same community and can be estimated as the mean value of the elements in the empirical co-assignment matrix. z^c_r and z^c_s are the consensus community assignment of patterns r and s, respectively.

Optimizing Q^c tends to return partitions that are, as a group, more similar to one another and, possibly, identical. If this is the case, the algorithm ends and the resulting partition is accepted as the representative consensus partition. However, if there is any variability in the output so that the algorithm does not arrive at the same solution each run, the a new co-assignment matrix is estimated and its modularity optimized. This procedure repeats until the algorithm converges.

Predictive model of FC

In the main text we described a procedure for modeling FC in terms of cofluctuation community centroids. In this section, we provide more details of how the model works.

In our previous work [19], we claimed that FC is driven by high-amplitude frames. One way to test whether this is the case is to "zero out" all low-amplitude and non-significant frames and to compute FC as the sum of whatever cofluctuation patterns are expressed at high-amplitude frames. Here, we take this claim one step further and state that the cofluctuation patterns expressed during high-amplitude frames are recurrences of one of three template patterns, which we obtained from the community detection analysis.

This intuition can be formalized by the following model:

$$\mathbf{FC}_{subject,scan} = f_1\mathbf{c}_1 + f_2\mathbf{c}_2 + f_3\mathbf{c}_3 + f_0\mathbf{0}. \quad (5)$$

In this expression, f_1, f_2, and f_3 are the fractions of all low-motion frames in which a participant expresses communities 1, 2, and 3, respectively. The parameter f_0 is the fraction of frames in which a participant is in a low-amplitude or non-significant state. The values of these parameters come from the results of the multiresolution consensus clustering analysis. The other parameters $\mathbf{c}_1$, $\mathbf{c}_2$, and $\mathbf{c}_3$ represent the average pattern of cofluctuation for of the three communities. The final parameter, $\mathbf{0}$ is a node-by-node matrix where all elements are zero.

We test five versions of this model that vary in terms of how $\mathbf{c}_1$, $\mathbf{c}_2$, and $\mathbf{c}_3$ are defined. In the first model, we define centroids using *all* available data, including data from other participants and from the scan whose FC we are trying to predict. The second model is similar, but estimates centroids excluding data from the current scan. Models three and four are personalized estimate centroids using data from participant whose scan is being predicted. Like models one and two, we test variants of models three and four that use all data and all data excluding the scan being predicted. Lastly, we tested a model in which centroids are estimated using data from participants other than the participant whose FC we are trying to predict.

In general, the model does not yield admissible correlation matrices. Accordingly, we transform each matrix (translation/rotation/scaling) to match the nearest admissible matrix. In all cases, we measure model fitness as the correlation of upper triangle elements in the true FC with those of the predicted FC matrix.

DATA AVAILABILITY

Midnight Scan Club raw data and derivatives are available here: `https://openfmri.org/dataset/ds000224/` and in processed, parcellated form here `https://www.dropbox.com/sh/tb694nmpu2lbpnc/AABKU_Mew7hyjtAC4ObzGVaKa?dl=0`. Processed and parcellated MyConnectome data is available here: `http://myconnectome.org/wp/data-sharing/`.

AUTHOR CONTRIBUTIONS

RFB designed study, analyzed data, wrote code, generated figures, and wrote initial draft of manuscript. SC, SG, and OS designed study, revised manuscript, and wrote final version of manuscript.

ACKNOWLEDGMENTS

We thank Evan Gordon for sharing individualized parcellations for MSC participants. This material is based upon work supported by the National Science Foundation under Grant No. 076059-00003C (RFB, OS).

[1] K. J. Friston, Human brain mapping **2**, 56 (1994).
[2] B. P. Rogers, V. L. Morgan, A. T. Newton, and J. C. Gore, Magnetic resonance imaging **25**, 1347 (2007).
[3] E. Tagliazucchi, P. Balenzuela, D. Fraiman, and D. R. Chialvo, Frontiers in physiology **3**, 15 (2012).
[4] T. W. Allan, S. T. Francis, C. Caballero-Gaudes, P. G. Morris, E. B. Liddle, P. F. Liddle, M. J. Brookes, and P. A. Gowland, PloS one **10**, e0124577 (2015).
[5] E. Tagliazucchi, M. Siniatchkin, H. Laufs, and D. R. Chialvo, Frontiers in neuroscience **10**, 381 (2016).
[6] N. Petridou, C. C. Gaudes, I. L. Dryden, S. T. Francis, and P. A. Gowland, Human brain mapping **34**, 1319 (2013).
[7] I. Cifre, M. Zarepour, S. Horovitz, S. Cannas, and D. Chialvo, Papers in Physics **12**, 120003 (2020).
[8] E. S. Finn, X. Shen, D. Scheinost, M. D. Rosenberg, J. Huang, M. M. Chun, X. Papademetris, and R. T. Constable, Nature neuroscience **18**, 1664 (2015).
[9] D. V. Demeter, L. E. Engelhardt, R. Mallett, E. M. Gordon, T. Nugiel, K. P. Harden, E. M. Tucker-Drob, J. A. Lewis-Peacock, and J. A. Church, Iscience **23**, 100801 (2020).
[10] T. O. Laumann, E. M. Gordon, B. Adeyemo, A. Z. Snyder, S. J. Joo, M.-Y. Chen, A. W. Gilmore, K. B. McDermott, S. M. Nelson, N. U. Dosenbach, *et al.*, Neuron **87**, 657 (2015).
[11] E. M. Gordon, T. O. Laumann, A. W. Gilmore, D. J. Newbold, D. J. Greene, J. J. Berg, M. Ortega, C. Hoyt-Drazen, C. Gratton, H. Sun, *et al.*, Neuron **95**, 791 (2017).
[12] C. Gratton, T. O. Laumann, A. N. Nielsen, D. J. Greene, E. M. Gordon, A. W. Gilmore, S. M. Nelson, R. S. Coalson, A. Z. Snyder, B. L. Schlaggar, *et al.*, Neuron **98**, 439 (2018).
[13] C. Gratton, B. T. Kraus, D. J. Greene, E. M. Gordon, T. O. Laumann, S. M. Nelson, N. U. Dosenbach, and S. E. Petersen, Biological psychiatry (2019).
[14] M. D. Rosenberg, E. S. Finn, D. Scheinost, X. Papademetris, X. Shen, R. T. Constable, and M. M. Chun, Nature neuroscience **19**, 165 (2016).
[15] M. Y. Chan, D. C. Park, N. K. Savalia, S. E. Petersen, and G. S. Wig, Proceedings of the National Academy of Sciences **111**, E4997 (2014).
[16] M.-E. Lynall, D. S. Bassett, R. Kerwin, P. J. McKenna, M. Kitzbichler, U. Muller, and E. Bullmore, Journal of Neuroscience **30**, 9477 (2010).
[17] J. Liu, X. Liao, M. Xia, and Y. He, Human brain mapping **39**, 902 (2018).
[18] A. H. C. Fong, K. Yoo, M. D. Rosenberg, S. Zhang, C.-S. R. Li, D. Scheinost, R. T. Constable, and M. M. Chun, NeuroImage **188**, 14 (2019).
[19] F. Z. Esfahlani, Y. Jo, J. Faskowitz, L. Byrge, D. Kennedy, O. Sporns, and R. Betzel, Proceedings of the National Academy of Sciences (2020).
[20] J. Dubois and R. Adolphs, Trends in cognitive sciences **20**, 425 (2016).
[21] J. E. Chen, C. Chang, M. D. Greicius, and G. H. Glover, Neuroimage **111**, 476 (2015).
[22] S. Mueller, D. Wang, M. D. Fox, B. T. Yeo, J. Sepulcre, M. R. Sabuncu, R. Shafee, J. Lu, and H. Liu, Neuron **77**, 586 (2013).
[23] J. Faskowitz, F. Z. Esfahlani, Y. Jo, O. Sporns, and R. F. Betzel, Nature neuroscience **23**, 1644 (2020).
[24] Y. Jo, F. Z. Esfahlani, J. Faskowitz, E. Chumin, O. Sporns, and R. Betzel, bioRxiv (2020).
[25] Y. Jo, J. Faskowitz, F. Z. Esfahlani, O. Sporns, and R. F. Betzel, bioRxiv (2020).
[26] O. Sporns, J. Faskowitz, S. Teixera, and R. Betzel, bioRxiv (2020).
[27] R. A. Poldrack, T. O. Laumann, O. Koyejo, B. Gregory, A. Hover, M.-Y. Chen, K. J. Gorgolewski, J. Luci, S. J. Joo, R. L. Boyd, *et al.*, Nature communications **6**, 1 (2015).
[28] E. M. Gordon, T. O. Laumann, B. Adeyemo, J. F. Huckins, W. M. Kelley, and S. E. Petersen, Cerebral cortex **26**, 288 (2016).
[29] G. S. Wig, T. O. Laumann, A. L. Cohen, J. D. Power, S. M. Nelson, M. F. Glasser, F. M. Miezin, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen, Cerebral cortex **24**, 2036 (2014).
[30] L. G. Jeub, O. Sporns, and S. Fortunato, Scientific reports **8**, 1 (2018).
[31] I. S. Jutla, L. G. Jeub, and P. J. Mucha, URL http://netwiki. amath. unc. edu/GenLouvain (2011).
[32] B. Horwitz, Neuroimage **19**, 466 (2003).
[33] E. Bullmore and O. Sporns, Nature reviews neuroscience **10**, 186 (2009).
[34] M. Rubinov and O. Sporns, Neuroimage **52**, 1059 (2010).
[35] J. R. Cohen and M. D'Esposito, Journal of Neuroscience **36**, 12083 (2016).
[36] B. Biswal, F. Zerrin Yetkin, V. M. Haughton, and J. S. Hyde, Magnetic resonance in medicine **34**, 537 (1995).
[37] N. U. Dosenbach, D. A. Fair, F. M. Miezin, A. L. Cohen, K. K. Wenger, R. A. Dosenbach, M. D. Fox, A. Z. Snyder, J. L. Vincent, M. E. Raichle, *et al.*, Proceedings of the National Academy of Sciences **104**, 11073 (2007).
[38] J. D. Power, B. L. Schlaggar, C. N. Lessov-Schlaggar, and S. E. Petersen, Neuron **79**, 798 (2013).
[39] R. Guimera and L. A. N. Amaral, nature **433**, 895 (2005).
[40] E. A. Allen, E. Damaraju, S. M. Plis, E. B. Erhardt, T. Eichele, and V. D. Calhoun, Cerebral cortex **24**, 663 (2014).
[41] X. Liu and J. H. Duyn, Proceedings of the National Academy of Sciences **110**, 4392 (2013).
[42] R. Hindriks, M. H. Adhikari, Y. Murayama, M. Ganzetti, D. Mantini, N. K. Logothetis, and G. Deco, Neuroimage

127, 242 (2016).
[43] S. Shakil, C.-H. Lee, and S. D. Keilholz, Neuroimage **133**, 111 (2016).
[44] B. A. Seitzman, C. Gratton, T. O. Laumann, E. M. Gordon, B. Adeyemo, A. Dworetsky, B. T. Kraus, A. W. Gilmore, J. J. Berg, M. Ortega, *et al.*, Proceedings of the National Academy of Sciences **116**, 22851 (2019).
[45] M. G. Preti, T. A. Bolton, and D. Van De Ville, Neuroimage **160**, 41 (2017).
[46] K. M. Anderson, T. Ge, R. Kong, L. M. Patrick, R. N. Spreng, M. R. Sabuncu, B. T. Yeo, and A. J. Holmes, Proceedings of the National Academy of Sciences **118** (2021).
[47] R. Kong, J. Li, C. Orban, M. R. Sabuncu, H. Liu, A. Schaefer, N. Sun, X.-N. Zuo, A. J. Holmes, S. B. Eickhoff, *et al.*, Cerebral cortex **29**, 2533 (2019).
[48] B. T. Kraus, D. Perez, Z. Ladwig, B. A. Seitzman, A. Dworetsky, S. E. Petersen, and C. Gratton, NeuroImage , 117743 (2020).
[49] S. Marek, B. Tervo-Clemmens, F. J. Calabro, D. F. Montez, B. P. Kay, A. S. Hatoum, M. R. Donohue, W. Foran, R. L. Miller, E. Feczko, *et al.*, BioRxiv (2020).
[50] A. S. Greene, S. Gao, D. Scheinost, and R. T. Constable, Nature communications **9**, 1 (2018).
[51] D. J. Newbold, T. O. Laumann, C. R. Hoyt, J. M. Hampton, D. F. Montez, R. V. Raut, M. Ortega, A. Mitra, A. N. Nielsen, D. B. Miller, *et al.*, Neuron **107**, 580 (2020).
[52] D. J. Newbold, E. M. Gordon, T. O. Laumann, D. F. Montez, N. A. Seider, S. J. Gross, A. Zheng, A. N. Nielsen, C. R. Hoyt, J. M. Hampton, *et al.*, bioRxiv (2020).
[53] C. Higgins, Y. Liu, D. Vidaurre, Z. Kurth-Nelson, R. Dolan, T. Behrens, and M. Woolrich, Neuron (2020).
[54] A. Avena-Koenigsberger, B. Misic, and O. Sporns, Nature Reviews Neuroscience **19**, 17 (2018).
[55] N. K. Logothetis, J. Pauls, M. Augath, T. Trinath, and A. Oeltermann, nature **412**, 150 (2001).
[56] R. F. Betzel, J. D. Medaglia, A. E. Kahn, J. Soffer, D. R. Schonhaut, and D. S. Bassett, Nature biomedical engineering **3**, 902 (2019).
[57] E. M. Lake, X. Ge, X. Shen, P. Herman, F. Hyder, J. A. Cardin, M. J. Higley, D. Scheinost, X. Papademetris, M. C. Crair, *et al.*, Nature Methods **17**, 1262 (2020).
[58] J. D. Power, M. Plitt, S. J. Gotts, P. Kundu, V. Voon, P. A. Bandettini, and A. Martin, Proceedings of the National Academy of Sciences **115**, E2105 (2018).
[59] C. J. Lynch, B. M. Silver, M. J. Dubin, A. Martin, H. U. Voss, R. M. Jones, and J. D. Power, Nature communications **11**, 1 (2020).
[60] M. F. Glasser, T. S. Coalson, J. D. Bijsterbosch, S. J. Harrison, M. P. Harms, A. Anticevic, D. C. Van Essen, and S. M. Smith, Neuroimage **181**, 692 (2018).
[61] M. F. Glasser, T. S. Coalson, J. D. Bijsterbosch, S. J. Harrison, M. P. Harms, A. Anticevic, D. C. Van Essen, and S. M. Smith, Neuroimage **197**, 435 (2019).
[62] J. D. Power, Neuroimage **197**, 650 (2019).
[63] A. M. Dale, B. Fischl, and M. I. Sereno, Neuroimage **9**, 179 (1999).
[64] M. F. Glasser, S. N. Sotiropoulos, J. A. Wilson, T. S. Coalson, B. Fischl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni, *et al.*, Neuroimage **80**, 105 (2013).
[65] J. Talairach, An approach to cerebral imaging (1988).
[66] S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, *et al.*, Neuroimage **23**, S208 (2004).
[67] J. D. Power, A. Mitra, T. O. Laumann, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen, Neuroimage **84**, 320 (2014).
[68] R. Ciric, D. H. Wolf, J. D. Power, D. R. Roalf, G. L. Baum, K. Ruparel, R. T. Shinohara, M. A. Elliott, S. B. Eickhoff, C. Davatzikos, *et al.*, Neuroimage **154**, 174 (2017).
[69] J. D. Power, K. A. Barnes, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen, Neuroimage **59**, 2142 (2012).
[70] E. M. Gordon, T. O. Laumann, B. Adeyemo, and S. E. Petersen, Cerebral cortex **27**, 386 (2017).
[71] M. E. Newman and M. Girvan, Physical review E **69**, 026113 (2004).

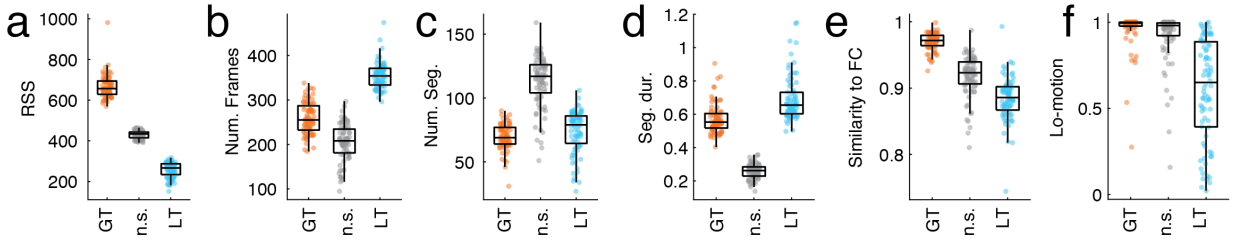

FIG. S1. **Statistical comparison of high- and low-amplitude frames using individualized parcels and Midnight Scan Club data.** IIn the main text, we described a statistical procedure for partitioning frames into three categories based on how their RSS values compared to those of a null model. We applied that model to data from the Midnight Scan Club in which participants' brains were parcellated into individualized regions and compared the three categories across multiple features, including their RSS amplitude, number of frames associated with each category, the number of contiguous frames of the same category (we call these "sequences"), sequence duration, the similarity of FC reconstructed using frames of each category to the static FC matrix, and the fraction of frames of a given category that were censored due to motion/data quality issues.

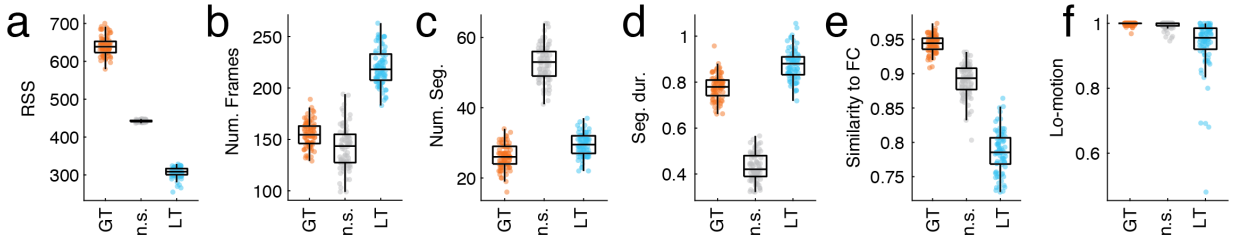

FIG. S2. **Statistical comparison of high- and low-amplitude frames using MyConnectome dataset.** In the main text, we described a statistical procedure for partitioning frames into three categories based on how their RSS values compared to those of a null model. We applied that model to data from the MyConnectome Project and compared the three categories across multiple features, including their RSS amplitude, number of frames associated with each category, the number of contiguous frames of the same category (we call these "sequences"), sequence duration, the similarity of FC reconstructed using frames of each category to the static FC matrix, and the fraction of frames of a given category that were censored due to motion/data quality issues.

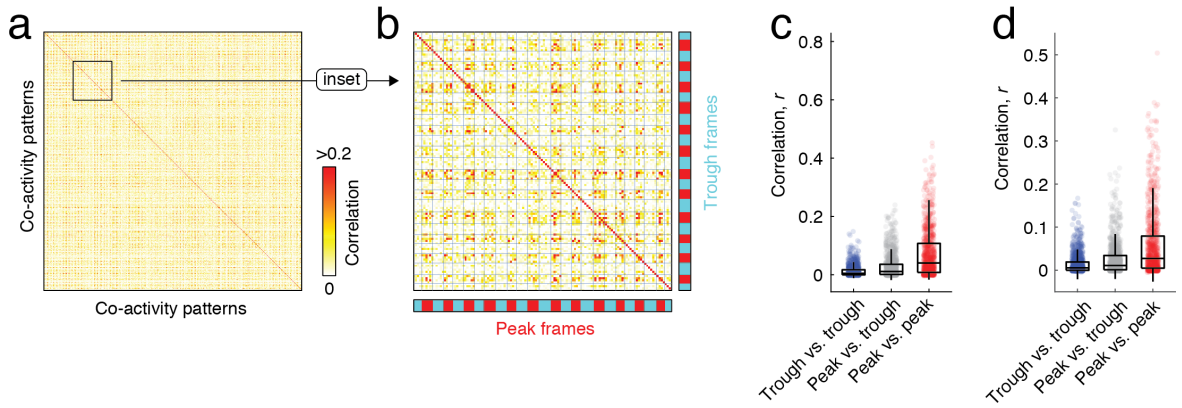

FIG. S3. **Cross-scan similarity of cofluctuation patterns in MyConnectome data and Midnight Scan Club data with individualized parcels.** In the main text we demonstrated that peak cofluctuation patterns were repeated across scans, while trough cofluctuation were not. Here, we replicate this finding using data from 84 resting-state scans from the MyConnectome data. (*a*) Similarity matrix of cofluctuation patterns, including both peaks and troughs. (*b*) We show a smaller section of the complete similarity matrix in greater detail to highlight the relationship between peak and trough similarity. (*c*) Boxplot of similarity scores, aggregated based on whether the similarity was computed between two trough patterns, a trough and a peak pattern, or a peak and a peak pattern. We also show that we obtain a similar effect when using individualized parcellations of participant's brains in the Midnight Scan Club. As in panel *c*, panel *d* depicts similarity scores, aggregated based on whether the similarity was computed between two trough patterns, a trough and a peak pattern, or a peak and a peak pattern.

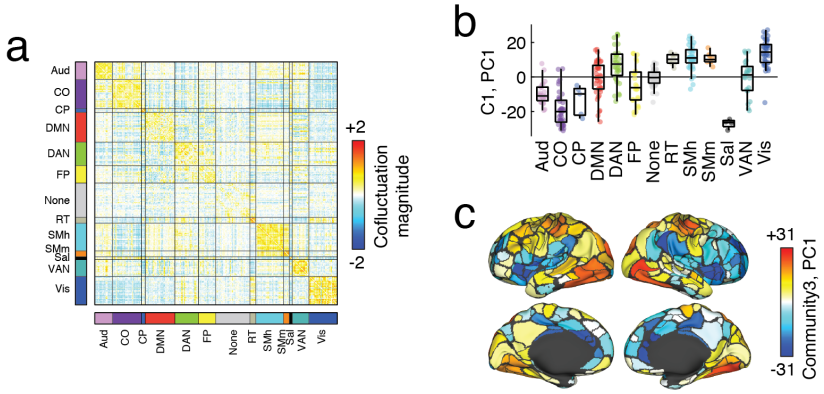

FIG. S4. **Summary of community 3 as an aggregated community.** In the main text, we described the two large, cohesive communities that resulted from clustering peak cofluctuation patterns. The clustering algorithm also generated a large number of smaller communities. Only a fraction of participants were represented in each of these communities (there were no cases where all participants were represented). For simplicity, we grouped the smaller communities and treated them as if they had the same community label (community 3). Here, we summarize some features of that community. (*a*) cofluctuation pattern. We applied PCA to the fMRI BOLD time series associated with each peak cofluctuation pattern. In panel *b*, we show the first component divided into brain systems. (*c*) The same component projected onto the cortical surface.

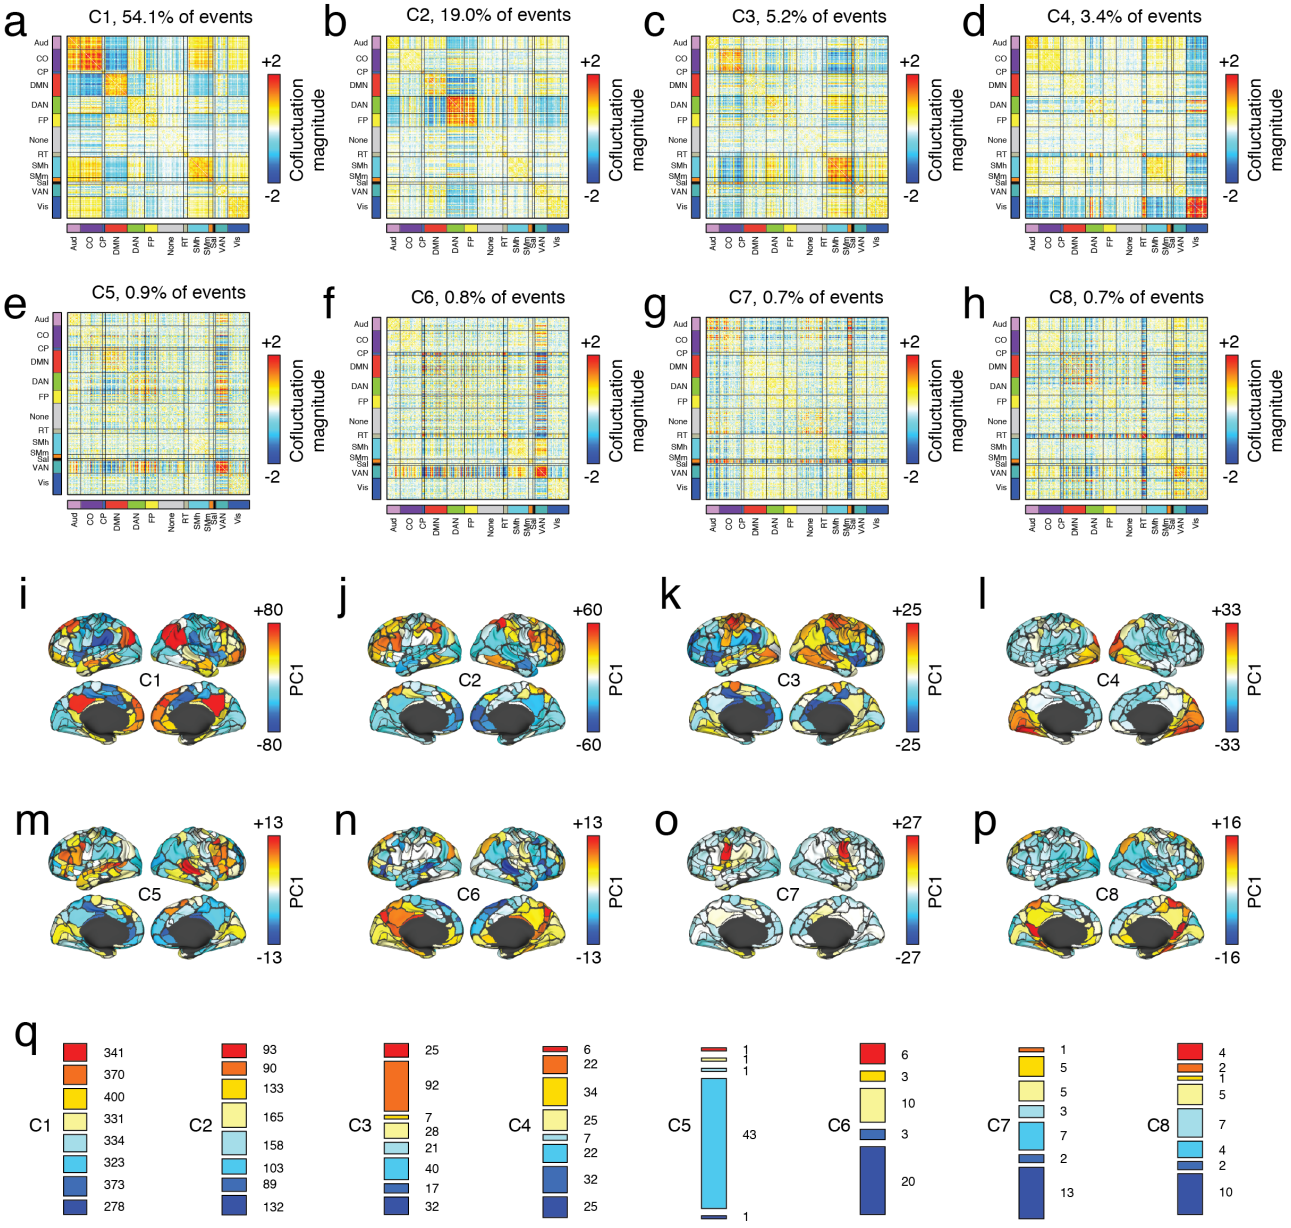

FIG. S5. **Summary of the individual communities that comprise community 3.** In the main text, we described the two large, cohesive communities that resulted from clustering peak cofluctuation patterns. The clustering algorithm also generated a large number of smaller communities. In the main text we grouped these smaller communities into a single community labeled '3'. Here, we examine those smaller communities in greater detail along with communities 1 and 2, for completeness. In panels *a* -*h*, we show cofluctuation patterns for eight communities that included data from multiple participants and accounted for at least 0.5% of all detected events. Above each matrix, we show the fraction of all events accounted for by that community. In panels *i*-*p*, we show the activity patterns that underpin each community, plotting the first principal component of corresponding activity on the surface of the brain. In panel *q*, we show participant composition of each community. Here, different colors correspond to different participants. The height of each bar indicates the relative number of events assigned to a given community from a given participant.

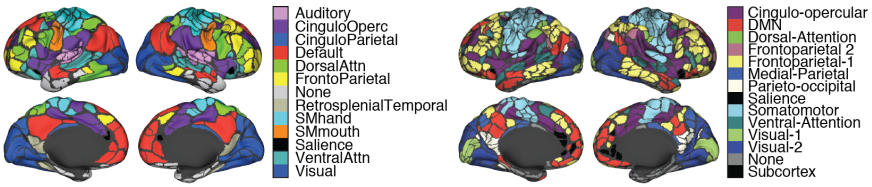

FIG. S6. **Brain systems projected onto cortical surface.** (*left*) Gordon atlas. (*right*) MyConnectome atlas.

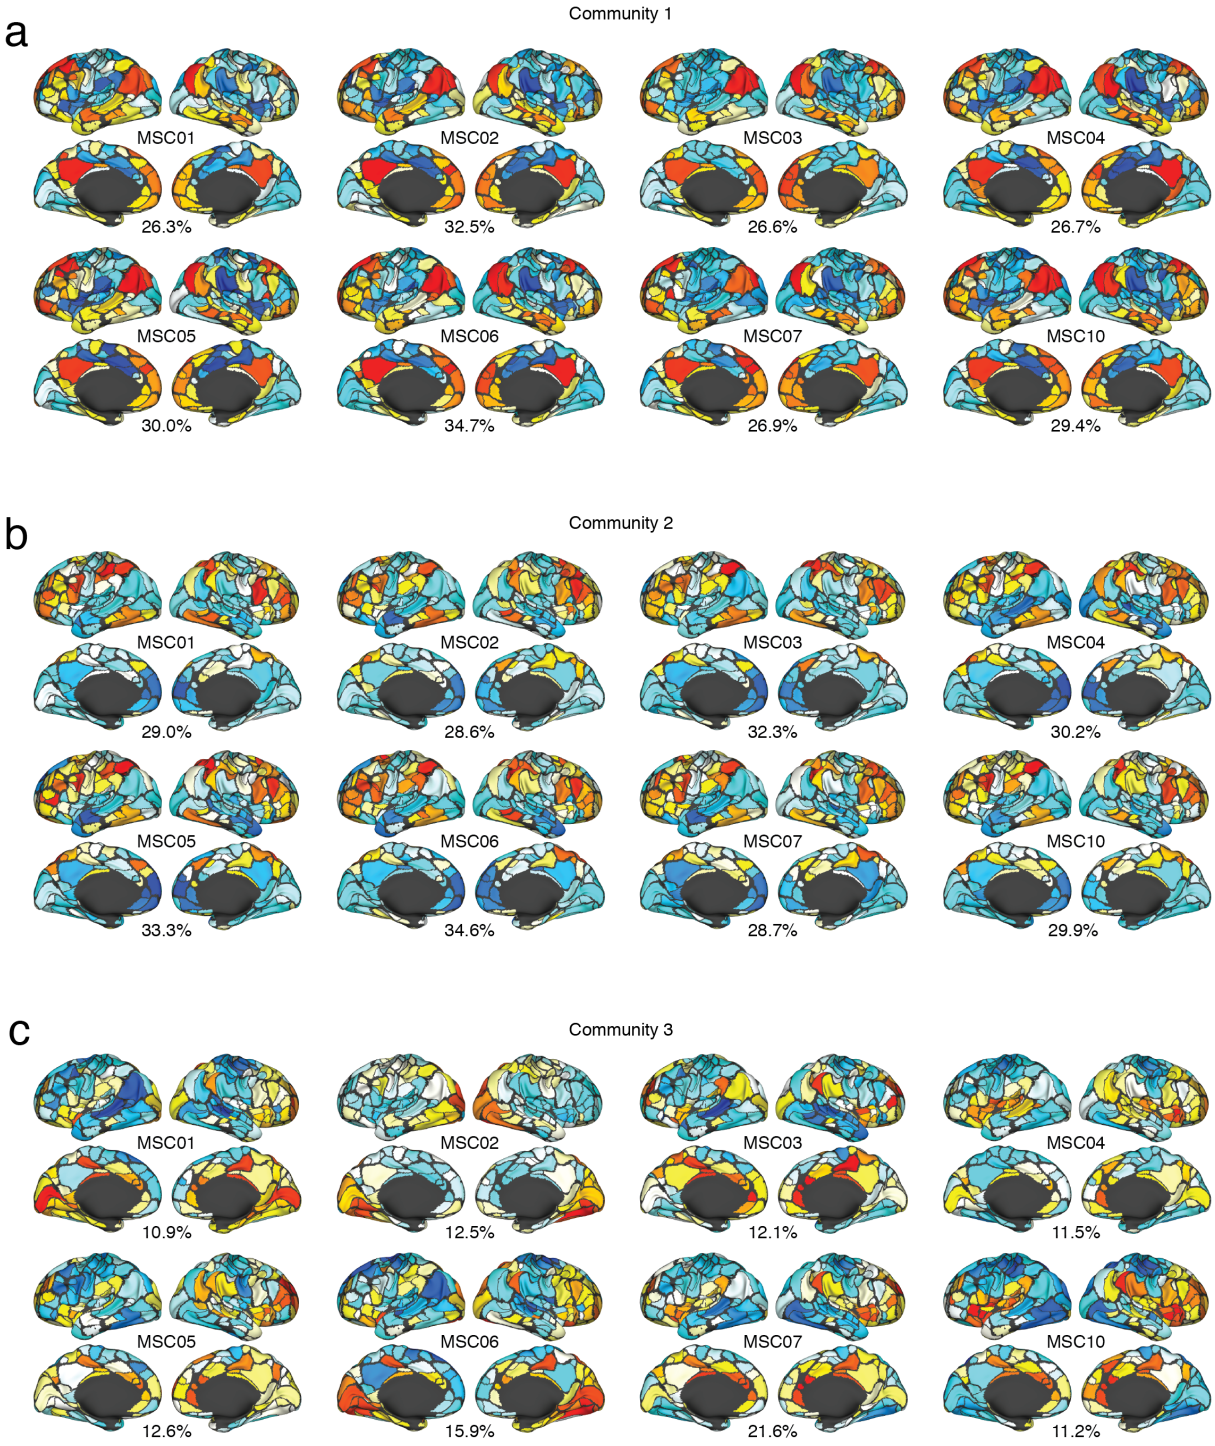

FIG. S7. **Participant-level PCA analysis of peak activity patterns.** In the main text we, applied PCA to activity patterns corresponding to peak cofluctuation. Those analyses were performed at the group level. Here, we repeat this analysis at the level of individual participants. Panels *a*-*c* show PC1 for each of the three communities. The text below each surface plot indicates the percent variance explained by that PC.

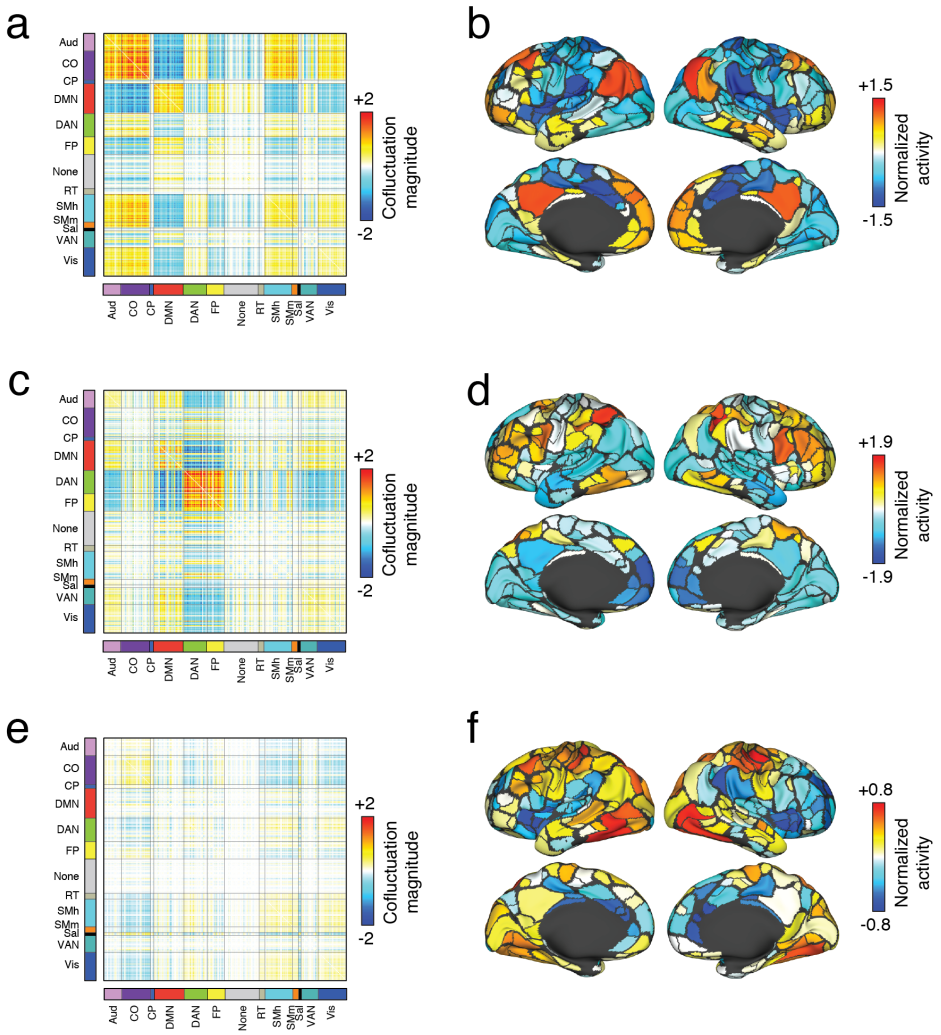

FIG. S8. **Alternative method for estimating cofluctuation and activity patterns.** In the main text we identified three recurring patterns of cofluctuation and used a PCA-based method to identify their corresponding modes of activity. However, the cofluctuation between nodes i and j must satisfy for all $\{i, j\}$ the following relationship: $C_{ij} = x_i x_j$, where x_i is the activity level of region i. In general, PCA will generate modes of activity that do not satisfy this relationship. To estimate this pattern, we used a greedy algorithm to identify the vector $\mathbf{x} = [x_1, \ldots, x_N]$ whose cofluctuation matrix $\mathbf{C} = \{C_{ij}\}$ minimizes the distance (root sum of squares) between the community-representative centroid. In panels *a*, *c*, and *e*, we show the cofluctuation matrices for communities 1, 2, and 3. Panels *b*, *d*, and *f* depict the activity pattern used to estimate cofluctuation.

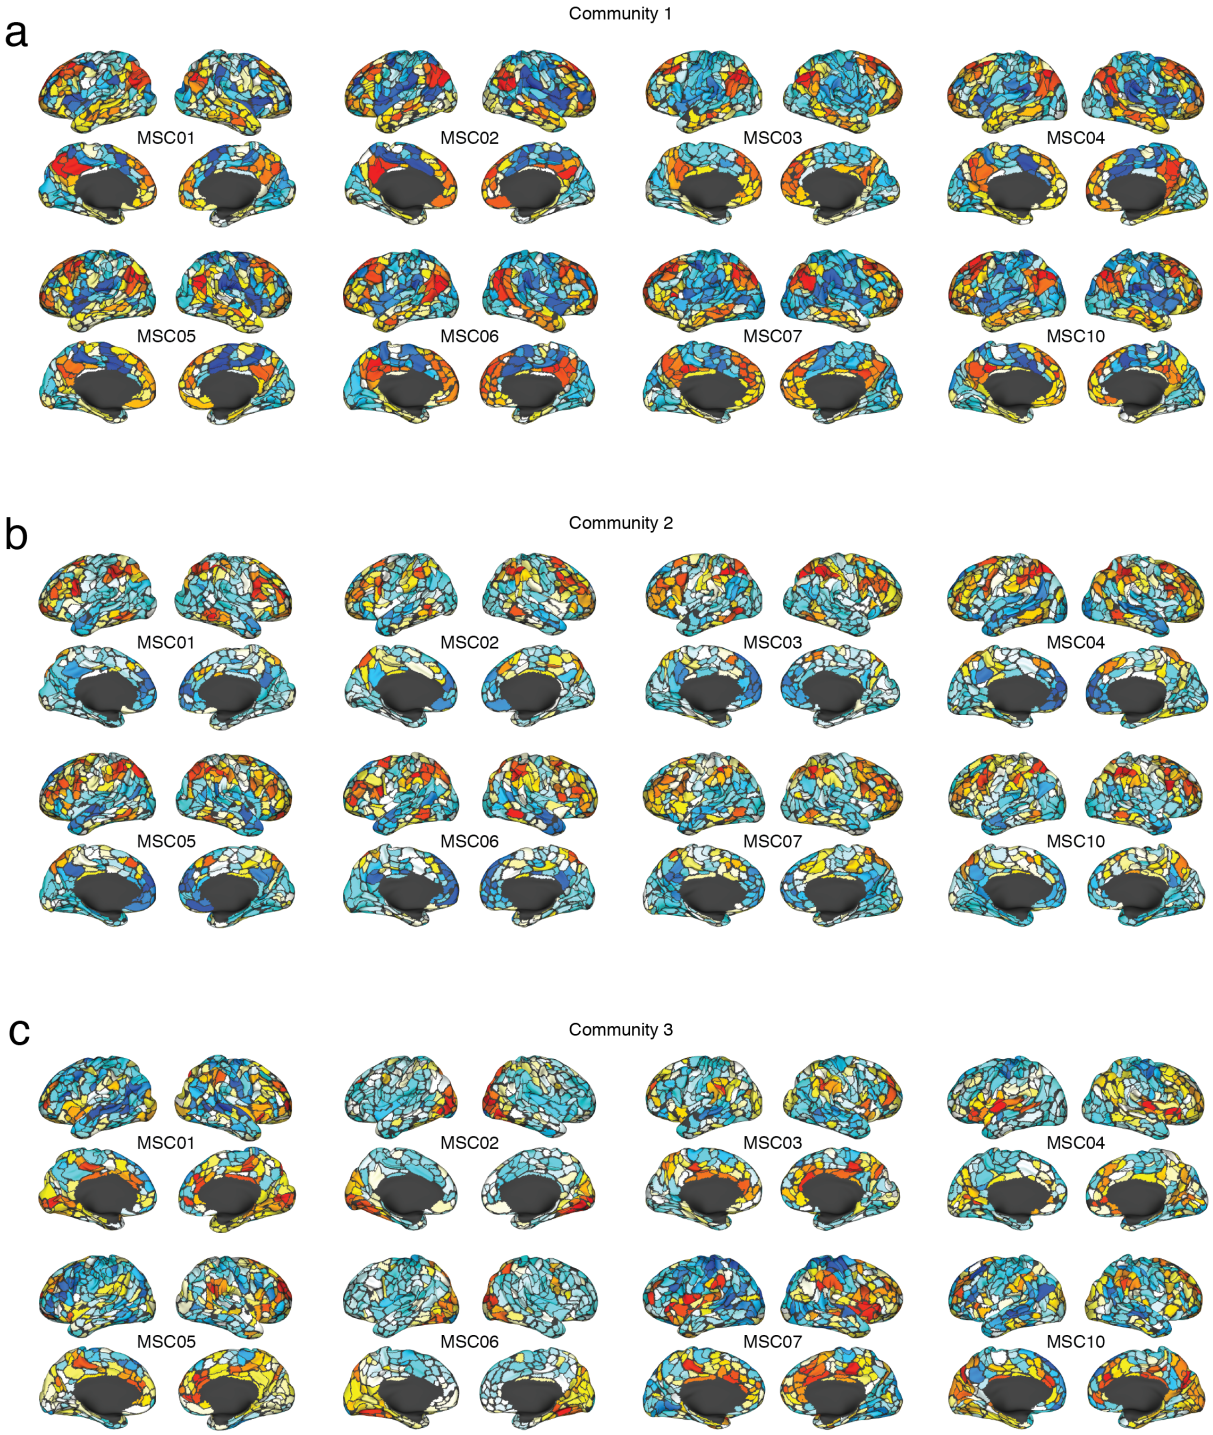

FIG. S9. **participant-level PCA analysis of peak activity patterns using individualized parcels.** In the main text, we applied PCA to activity patterns corresponding to peak cofluctuation. Those analyses were performed at the group level using a common set of $N = 333$ brain regions. Here, we repeat this analysis at the level of individual participants using parcels fit individually to each participant. Panels *a*-*c* show PC1 for each of the three communities.

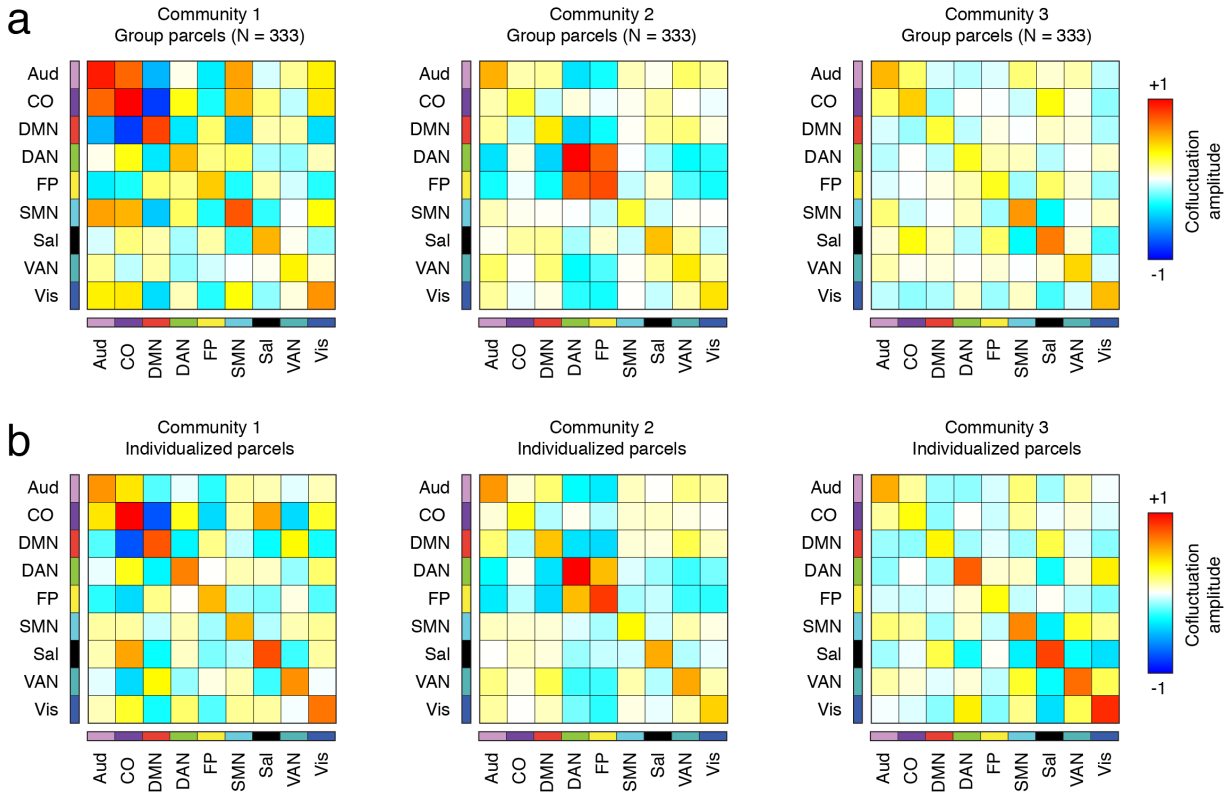

FIG. S10. **Similar system-level centroids using both group and individualized parcels.** In the main text, we described three distinct cofluctuation patterns that were broadly shared across participants. Those patterns were estimated using a common set of $N = 333$ parcels. We also repeated those analyses using individualized parcels, which differed in number across participants. This variability made it impossible to compare cofluctuation patterns between participants as well as with the cofluctuation patterns estimated using the common set of parcels. To circumvent this issue, we estimated centroids at a system level by averaging the cofluctuation magnitude within and between nine systems that were expressed both at the group and individual level: Auditory, cingulo-opercular, default mode, dorsal attention, fronto-parietal, somatomotor, salience, ventral attention, and visual networks. Note, that with the group parcels the somatomotor network contains the somatmotor-hand and somatomotor-mouth networks. In the case of the individualized parcels, the somatomotor network contains the somatmotor-hand, somatomotor-foot, and somatomotor-mouth networks. Similarly, the visual network contains both lateral and primary visual networks for the individualized parcels. In general, we found a high level of correspondence between the two parcellations. The similarity values (Pearson correlation) of centroids for communities 1, 2, and 3 between datasets were $r = 0.85$, $r = 0.97$, and $r = 0.84$. In contrast, the similarity between centroids of $r = 0.63 \pm 0.07$. In a and b we show centroids estimated using the group and individualized parcels.

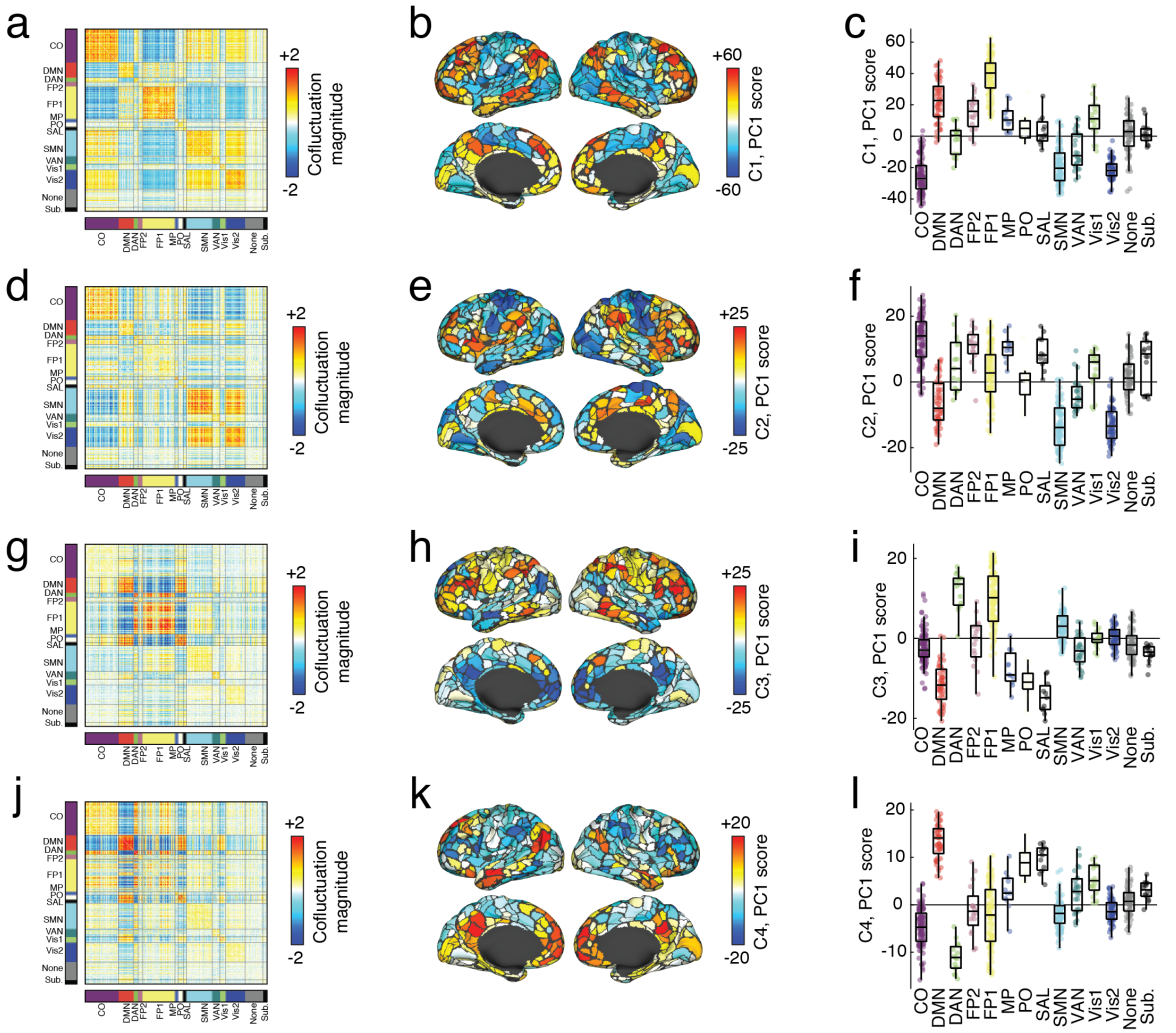

FIG. S11. **Results of multi-resolution consensus cluster applied to peak cofluctuation in MyConnectome dataset.** In the main text we extracted peak cofluctuation patterns in the Midnight Scan Club dataset and applied a clustering algorithm to group patterns into communities. Here, we apply an identical analysis to MyConnectome data. The algorithm identified four large communities, which we summarize here. In panels *a*, *d*, *g*, and *j*, we show the mean cofluctuation for each community. Note that community 1 and community 4 in the MyConnectome data is nearly identical to community 1 in the Midnight Scan Club. Similarly, communities 2 and 3 bear strong resemblance to community 2 in the Midnight Scan Club analysis. Panels *b*, *e*, *h*, and *k* show the first principal component (PC1) of activity during each peak. Panels *c*, *f*, *i*, and *l* show boxplots of PCs sorted by system.

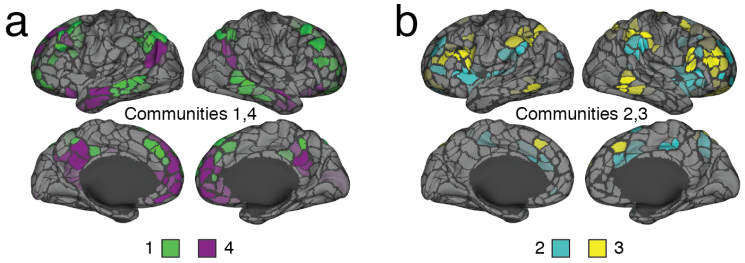

FIG. S12. **Interdigitated activity during peaks.** We clustered cofluctuation during peaks and found that the four main communities we reported could be grouped into pairs that involved similar brain systems. Here, we show these pairs for interdigitated sub-divisions of canonical systems. To do this, we calculated PC1 for each of the four communities. We then normalized PC1, dividing every element by $\sqrt{\sum_i v_i.^2)}$, where v_i is the ith element (corresponding to region i) of PC1. This normalization ensures that the eigenvector has magnitude equal to unity. We then paired communities 1 with 4 and 2 with 3 and, for each region, we calculated which of the two normalized eigenvectors had a larger value. Here, we visualize the winners of the top quartile (25%) of regions. In a we show communities 1 and 4 and in b we show communities 2 and 3.

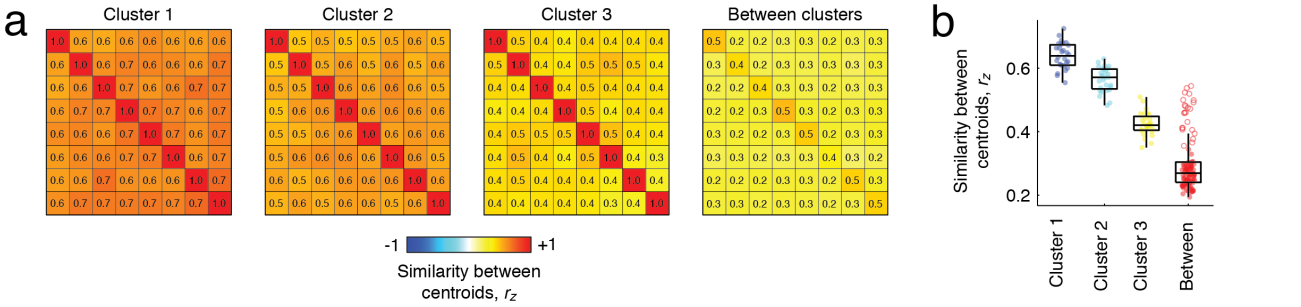

FIG. S13. **Within and between community similarity of centroids.** In the main text we clustered cofluctuation peaks into three communities. Here, we calculated each participant's centroid (mean cofluctuation pattern for each community) using only those peaks assigned to that community. (a) The similarity (Pearson correlation) of participants' centroids to one another for communities 1, 2, and 3. We also show the similarity of centroids from different communities (*right*). (b) Similarity scores grouped by community. Note that in the "between" category, some points are solid and others outlined. The outlined points represent similarity scores from the same participant across different centroids.

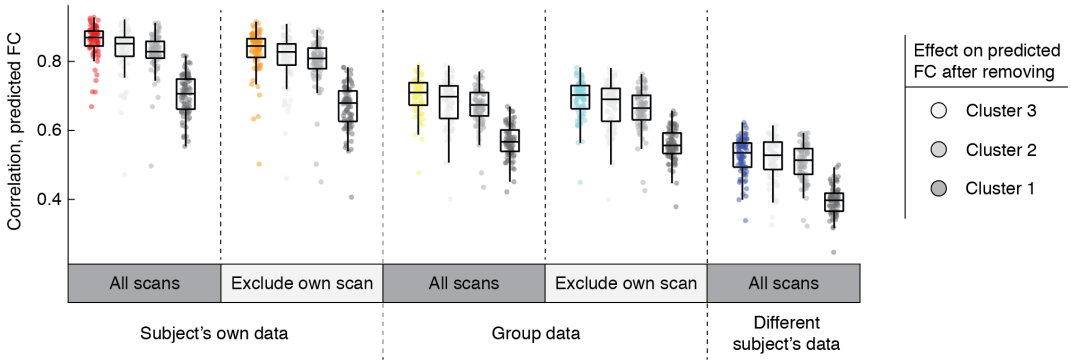

FIG. S14. **Sensitivity analysis using predictive model.** In the main text we predicted participant's FC using three community centroids estimated several different ways. Here, we perform a sensitivity analysis in which we selectively exclude each of the three communities so that it does not contribute to the prediction. For instance, when we remove data from community 1 we then only use data from communities 2 and 3 to make a prediction.