

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

Richard F. Betzel^{a,b,c,d,*}, Sarah A. Cutts^{a,c}, Sarah Greenwell^a, Joshua Faskowitz^{a,c},
Olaf Sporns^{a,b,c,d}

^a Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405, United States

^b Program in Neuroscience, Indiana University, Bloomington, IN 47405, United States

^c Network Science Institute, Indiana University, Bloomington, IN 47405, United States

^d Cognitive Science Program, Indiana University, Bloomington, IN 47405, United States

A B S T R A C T

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 co-fluctuations. Here, we develop a statistical test to identify events in empirical recordings. We show that the patterns of co-fluctuation 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 co-fluctuations, demonstrating that group-averaged co-fluctuations are suboptimal for explaining participant-specific connectivity. Our work complements recent studies implicating brief instants of high-amplitude co-fluctuations as the primary drivers of static, whole-brain functional connectivity. Our work also extends those studies, demonstrating that co-fluctuations during events are individualized, positing a dynamic basis for functional connectivity.

1. 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” (Friston, 1994; Rogers et al., 2007). 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 (Tagliazucchi et al., 2012; Allan et al., 2015; Tagliazucchi et al., 2016; Petridou et al., 2013; Cifre et al., 2020).

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 (Finn et al., 2015; Demeter et al., 2020; Laumann et al., 2015; Gordon et al., 2017b; Gratton et al., 2018). These observations hold tremendous translational promise, and open up the possibility of designing personalized interventions (Gratton et al., 2019) and developing increasingly potent connectivity-based biomarkers for cognition, development, and disease (Rosenberg et al., 2016; Chan et al., 2014; Lynall et al., 2010).

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 con-

nectivity patterns and distributed dynamically and sparsely throughout a scan session. Indeed, recent findings broadly support this hypothesis (Liu et al., 2018; Fong et al., 2019). In (Esfahlani et al., 2020), for instance, we demonstrated that using a small subset of frames classified as “events” – brief and infrequent periods of high-amplitude co-fluctuation – we could produce accurate reconstructions of FC while simultaneously rendering participants identifiable, amplifying their connective 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 (Fig. 1). For instance, do co-fluctuation patterns during events repeat from one scan to another (Fig. 1b)? 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 (Fig. 1c)? Or does the individualization of FC arise from equally idiosyncratic patterns of high-amplitude co-fluctuations (Fig. 1d)? Addressing these questions is critical for linking patterns of brain connectivity with individual differences in behavior (Dubois and Adolphs, 2016), and would help clarify the role of brain dynamics in shaping the individualization of FC (Chen et al., 2015), complementing other approaches

* Corresponding author.

E-mail address: rbetzel@indiana.edu (R.F. Betzel).

<https://doi.org/10.1016/j.neuroimage.2022.118993>.

Received 31 October 2021; Received in revised form 25 December 2021; Accepted 10 February 2022

Available online 19 February 2022.

1053-8119/© 2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>)




Fig. 1. Competing hypotheses for how co-fluctuations contribute to the individualization of FC. (a) Edge time series decompose FC into its framewise contributions. (b) FC can be well-approximated from co-fluctuations expressed during high-amplitude “events”. Recurrences of event co-fluctuation patterns can be grouped into clusters or “states”. FC can then be approximated from state centroids and the relative frequency with which each state is visited. Why does FC differ between individuals? (c) One hypothesis (H1) is that the same states are shared across individuals and inter-individual differences are driven by differences in the frequency with which those shared states are visited. (d) Another hypothesis (H2) is that the states, themselves, are subject-specific. In this case, inter-individual differences in FC are driven by differences in the state centroids across subjects. (e) To adjudicate between these hypotheses, we can try to approximate FC matrices with centroids estimated from that same subject’s data (different scans) or from group-averaged data.

that have focused on the collective influence of cortical expansion rates, post-natal experience, and genetics (Mueller et al., 2013).

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 (Esfahlani et al., 2020; Faskowitz et al., 2020; Jo et al., 2020a; Sporns et al., 2020; Greenwell et al., 2021; Esfahlani et al., 2021; Liu et al., 2021; Pope et al., 2021). We apply this framework to two independently acquired datasets: the Midnight Scan Club (Gordon et al., 2017b; Gratton et al., 2018) and the MyConnectome project (Laumann et al., 2015; Poldrack et al., 2015). 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 that they repeat within and between scans of the same individual. We also show that event co-fluctuations 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.

2. Results

2.1. Edge time series as a mathematically 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 (Gordon et al., 2016) and another in which parcels were defined on an individual basis, resulting in a different set of parcels for each participant ($N = 612 \pm 28$) (Wig et al., 2014). 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 (Laumann et al., 2015; Poldrack et al., 2015).

For each dataset, we transformed regional fMRI BOLD time series into co-fluctuation or edge time series (ETS). Briefly, ETS are calculated as the element-wise product between pairs of z-scored regional (nodal) time series (Fig. 2a; 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 co-fluctuations 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 $edge \times time$ matrix (Fig. 2b) whose temporal average yields a vector (Fig. 2c) that, when reshaped into the upper triangle elements of a




Fig. 2. 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 confuctuations 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 confuctuation 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 confuctuations during a trough and peak (when RSS is small versus large).

$node \times node$ matrix, is exactly the FC matrix (Fig. 2d). The $edge \times time$ matrix can also be “sliced” temporally and the corresponding vector once again reshaped into the upper triangle elements of a $node \times node$ matrix, yielding an instantaneous estimate of whole-brain confuctuations. These matrices vary in terms of their mean confuctuations, which we summarize with a root sum of squares measure (Fig. 2e).

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 (Esfahlani et al., 2020). 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.

2.2. A statistical test for high-amplitude confuctuation events

In previous work, we identified putative confuctuation events as the top $P\%$ frames in terms of root sum squared (RSS) amplitude of confuctuation 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 $q = 0.05$; Fig. 3a). 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 confuctuation 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 capture distinct features of confuctuations, 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. 3b). On the other hand, the number of high-amplitude frames in a scan is smaller than the number of low-amplitude frames (Fig. 3c). Additionally, high-amplitude frames form fewer contiguous segments than low-amplitude frames (Fig. 3d), and, when they do, those segments tend to be of shorter duration than contiguous segments of low-amplitude frames (Fig. 3e). 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. 3f). We note, however, that here the gap in correlation between the high- and low-amplitude is narrower than in previous studies (Esfahlani et al., 2020). This is due to differences in the total number of frames used to reconstruct FC and how those frames were selected. In fact, when controlling for the number of frames, we find that the gap widens (see Fig. S1). This relationship holds over a range of sample sizes, from 20 to 260 frames. Lastly, we find that high-amplitude frames are almost never among those censored for excessive in-scanner motion (Fig. 3g); low-amplitude frames, on the other hand, were more likely to be associated with censored frames, but were also more variable, a result that can be attributed, at least in part, to stable inter-individual differences in motion (Bolton et al., 2020; Zeng et al., 2014) (Fig. S2). This observation is consistent with our previous study, in which we reported a weak but consistent negative correlation between RSS and framewise displacement (Esfahlani et al., 2020). These findings, in general, replicate using individualized parcels for partici-




Fig. 3. 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 co-fluctuation 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 (log10 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.

pants in the Midnight Scan Club (Fig. S3) and with MyConnectome data (Fig. S4).

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.

2.3. Peak co-fluctuation patterns are repeated across scan sessions – troughs are not

In the previous section we presented a simple method for estimating statistically significant co-fluctuation events and demonstrated that the peak and trough frames of these segments exhibit distinct spatiotemporal properties. Here, we investigate representative co-fluctuation patterns extracted from blocks of high- and low-amplitude frames. We first compare the similarity of these co-fluctuation patterns, first within-individuals and later between. Then, we present evidence that co-fluctuation 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; see Materials and Methods for more details). 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 co-fluctuation patterns extracted during RSS peaks and troughs (periods when the RSS was significantly greater or less than the null model; labeled G.T. and L.T. in Fig. 3). 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. 4a,b). We see an identical effect in the MyConnectome data (Fig. S5a–c) and when parcels are individualized (Fig. S5d). These observations suggest that high-amplitude events encode subject-specific patterns of co-fluctuations.

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 co-fluctuation patterns for all pairs of detected events, for all scans, and for all subjects (Fig. 4c). We then compared these similarity values based on whether they came from the same or different participants (Fig. 4d). We found that within-individual similarity exceeded between-individual similarity (non-parametric permutation test, $p < 10^{-15}$; Fig. 4f).

Collectively, these results are in line with our previous study and suggest that low-amplitude co-fluctuations contains little participant-specific information (Esfahlani et al., 2020). Rather, our findings support the hypothesis that high-amplitude co-fluctuations contribute significantly more information about an individual than low-amplitude co-fluctuations.

2.4. High-amplitude events can be divided into distinct communities based on their co-fluctuation 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 dis-




Fig. 4. Intra-/Inter-individual similarity of co-fluctuation patterns. (a) For each individual separately, we aggregated all low-motion co-fluctuation patterns during peaks and troughs. We then computed the similarity between co-fluctuation 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 co-fluctuation at peaks exhibited similarity across scans, we focused on these patterns only, discarding co-fluctuation that occurs during troughs and focusing on comparisons of participants to one another. We found that co-fluctuation at peaks tended to be more similar within participants than between. We show the raw similarity matrix in c and the averaged values in d. (e) Boxplot of similarity values broken down according to whether they fell within or between participants.

similar 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 co-fluctuation 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 co-fluctuation patterns that occurred during event peaks, calculated the similarity matrix of those patterns, and clustered this matrix using a variant of multiresolution consensus clustering (modularity maximization with a uniform null model (Jeub et al., 2018; Jutla et al., 2011); 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. 5a) and that instances of those communities appeared in every scan session, accounting for 54.1% and 19.0% of all event peaks, respectively (see Fig. S6). 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. 5a). 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. S7 and Fig. S8). Note that we also repeated this clustering analysis for each participant individually and found that subject-level partitions of events were highly similar to partitions estimated with the group-aggregated data (mean $\pm$ standard deviation adjusted Rand index across subjects of 0.81 ± 0.10 ; $p < 10^{-4}$, permutation test in which each subject's community labels were randomly shuffled).

To better understand why certain co-fluctuation patterns were grouped together, we examined group-representative centroids for each community by calculating the mean co-fluctuation pattern of all frames assigned to that community (Fig. 5b,c). We found that community 1 reflected a topology that expressed strong and anticorrelated co-fluctuations mostly between cingulo-opercular and default mode networks (Fig. 5b). Community 2, on the other hand, expressed strong anticorrelations between the default mode with dorsal attention and fronto-parietal networks (Fig. 5b). See Fig. S9 for topographic depiction of systems on cortical surface.

We next wanted to better understand how brain activity drives the co-fluctuation patterns described above. To do this, we extracted activity profiles (regional BOLD activity, rather than co-fluctuations) dur-

ing event peaks, grouped them by community, and performed principal components analysis (PCA). 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 25% and 26% of variance, respectively, with a sharp drop-off for increasing component numbers, Fig. S10) to be uncorrelated ($r = -0.025$; $p = 0.65$; Fig. 5d). The first principal component for community 1 (Fig. 5e) exhibited significant activation of the default mode and inactivation of cingulo-opercular, visual, and somatomotor networks (distance-preserving permutation test of node order, i.e. spin test (Alexander-Bloch et al., 2018; Markello and Misic, 2021; Váša et al., 2018); false discovery rate fixed at 5%; $p_{adj} = 5.3 \times 10^{-4}$; Fig. 5f). The first principal component for community 2 (Fig. 5h) exhibited significant activation of default mode, ventral attention, and visual networks (spin test; false discovery rate fixed at 5%; $p_{adj} = 0.02$; Fig. 5f). We show individual-level principal components in the Supplementary Material (Fig. S11) and derive similar modes of activity using an alternative procedure (Fig. S12). 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 participant individually. We show that the first principal component of brain activity during high-amplitude co-fluctuations is similar irrespective of whether we use group-level or individualized parcels (Fig. S13). 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. S14). 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. S15). 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 Midnight 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. S16).

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




Fig. 5. Clustering peak co-fluctuation. (a) We clustered the co-fluctuation 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 co-fluctuation 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. Analogous information about community 3 can be found in Fig. S7.

2.5. Event communities are individualized

In the previous section, we showed that co-fluctuation 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? Which node pairs are most variable between individuals and, therefore, more likely to be useful for subject fingerprinting and identifiability?

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 co-fluctuation patterns assigned to the same community than to other communities; see Fig. S17 and Fig. S18), we also found evidence that the similarity between co-fluctuation patterns is stronger still when they come from the same participant (t-test comparing within- and between-individual similarity; $p < 10^{-15}$; Fig. 6a-d). Further, we identified the pairs of brain regions whose co-fluctuations were most variable across individuals by computing the standard deviation of edge weights across participant centroids (Fig. 6e-f). Our rationale for doing so was that pairs of regions whose co-fluctuation amplitude was variable are also among those most likely to drive individualization (Amico and Goñi, 2018). For the co-fluctuation pattern expressed by community 1, we found that the most variable edges linked the cingulo-opercular network to the dorsal attention, fronto-parietal, and ventral attention networks (Fig. 6e). In the case of community 2, the most variable edges were linked to default mode, dorsal attention, and fronto-parietal.

Note that in Fig. S19 and Fig. S20, we further explored the individualization of high-amplitude co-fluctuation patterns. Briefly, we followed (Seitzman et al., 2019) and computed the similarity (correlation) of regional co-fluctuation patterns estimated at the group and subject levels for both communities 1 and 2 (Fig. S19a,b). This procedure yielded a similarity map for each subject and community (Fig. S19c-f), which we analyzed further. Specifically, we calculated the mean similarity of regions in each putative brain system (Gordon et al., 2016) and demon-

strated that, for community 1 regions in the ventral attention system was more dissimilar from the group than expected (space-preserving permutation test, 10000 repetitions, $p < 10^{-4}$; Fig. S19f). For community 2, regions in the somatomotor-hand system, along with those that lack a clear system assignment, were more dissimilar. Interestingly, regions in the default mode and cingulo-opercular systems (community 1) and default mode, fronto-parietal, and dorsal attention systems (community 2) were more similar than expected. These systems all participated in the strongest co-fluctuations and largely typified each cluster. Finally, we investigate whether subject deviations from the group are evident in individual scan sessions. To test this, we repeated the above procedure but using scan-resolved estimates of communities (see Fig. S20 for an example from MSC06 and community 1). In general, we found that similarity maps for any given subject were highly repeatable across scan sessions and dissimilar between individuals (t-test of mean within- and between-subject similarity; $p < 10^{-15}$; Fig. S19g).

These observations suggest that high-amplitude co-fluctuation patterns reflect a topology that is broadly shared across individuals but is systematically and individually refined. We note, however, that the origins of this individualization is likely multi-factorial and that true differences in co-fluctuation patterns likely depend on idiosyncratic anatomical differences as well as choice of parcellation (Bijsterbosch et al., 2020).

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

Previous studies have shown that time-averaged FC is individualized and can serve as a fingerprint of an individual (Gratton et al., 2018; Gordon et al., 2017a; Finn et al., 2015). Here and in (Esfahani et al., 2020), we showed that time-averaged FC can also be approximated using only a small number of high-amplitude frames (events) and that the co-fluctuation patterns expressed during those events can be grouped into a small number of clusters or communities. How do these dif-




Fig. 6. 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.

ferent communities produce individualized patterns of FC? Are the cofluctuation patterns fixed at the group level but expressed in different proportions from one individual to the next? Or are the proportions fixed while the cofluctuation patterns vary idiosyncratically? Here, we present a model to adjudicate between these and related hypotheses.

Motivated by previous studies showing that FC can be described using only high-amplitude frames (Tagliazucchi et al., 2012; Petridou et al., 2013; Allan et al., 2015; Cifre et al., 2020; Esfahlani et al., 2020), our model assumes that FC depends exclusively on the cofluctuation patterns expressed during events and that low and middle-amplitude frames make no contribution (Fig. 7a). Specifically, we model static FC as a linear combination of the centroids for communities 1, 2, and 3, weighted by how frequently those communities appear in the data. To fit this model, we must estimate centroids and frequencies.

These estimates can be carried out separately for each subject, yielding subject-specific centroids and frequencies. However, estimates can also be made at the group level, yielding a set of centroids and frequencies that are shared across individuals. In both cases, we define model fit as the correlation of edge weights in the observed and predicted FC matrices.

Here, we test five model variants that combine subject and group-level information in different configurations. The first two models make predictions of FC using subject-specific estimates of centroids and frequency estimates made at either the subject (model 1) or group level (model 2). Similarly, the next two models pair group-level estimates of centroids with either subject or group level estimates of frequencies (models 3 and 4). Finally, we test a fifth model that makes predictions of subject's FC using subject-specific centroids and frequencies estimated for the remaining seven subjects (different from the subject whose FC we are trying to predict), yielding seven independent predic-

tions. We define this model's fitness as the best prediction out of the seven.

Importantly, these models allow us to directly test the competing hypotheses that individualized FC is driven by subject-specific cofluctuation patterns or subject-specific frequencies (labeled H_1 and H_2 in Fig. 1). Here, we found that model performance under H_1 was significantly greater than performance for H_2 ($r_{H_1} = 0.85 \pm 0.06$, $r_{H_2} = 0.70 \pm 0.05$; paired sample t-test, $p < 10^{-15}$). Additionally, combining subject-specific centroids and frequencies yielded a small but statistically significant improvement in performance, largely by reducing the number of outlying points ($r = 0.85 \pm 0.05$; paired sample t-test, $p = 0.02$). Collectively, these findings suggest that the subject-specificity of high-amplitude cofluctuation patterns drive the organization of static FC. Finally, 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. 7d).

Additionally, we replicated these findings using MSC data processed using an alternative processing pipeline and parcellation (Esteban et al., 2019; fer et al., 2018) (Fig. S21), at different spatial scales (100-node parcellation) (Schaefer et al., 2018) (Fig. S22), and without including global signal regression in the processing pipeline (Fig. S23). One notable discrepancy, however, was that in the absence of global signal regression, high-amplitude co-fluctuation patterns did not exhibit anti-correlations. Rather, we find that putative “events” are driven by cohesive fluctuations of regional activity in the same direction, i.e. activity deflects positively or negatively but typically not both. This is in line with previous observations concerning the effect of the global signal on FC (Fox et al., 2009; Chai et al., 2012; Saad et al., 2012). In fact, we find that the groups of brain regions that had previously engaged in anticorrelated behavior now formed their own distinct community, yielding four communities instead of two (along with a much smaller fifth com-




Fig. 7. Modeling participant and scan-specific FC with confluctuation patterns. (a) We hypothesized that FC is driven by brief high-amplitude patterns of confluctuation. A scan session can be abstracted, then, as periods of time where confluctuation is close to zero punctuated by periods of time where confluctuation corresponds to one of the confluctuation communities. (b) We modeled FC as a linear combination of confluctuation patterns corresponding to communities 1, 2, and 3 as well as a blank state where confluctuation was treated as zeros (low-amplitude frames). We varied whether confluctuation patterns and the frequencies with which they appear are estimated at the subject or group level. We also tested a model that used subject-level estimates of centroids and frequencies from other subjects. (c) Performance of the five models. Overall, models that included participant-specific information about cluster centroids outperformed other models. (d) We performed a sensitivity analysis of model 2, which used subject-specific centroids and group-level frequencies to predict FC patterns. In this analysis, we selectively exclude each of the three communities from the model so that it does not contribute to the prediction. We find that removing community 1 yields the biggest decrease in model performance, suggesting that it drives the model performance more so than communities 2 or 3.

munity). As a consequence, we modeled FC as a linear combination of five states.

These results suggest that the confluctuation patterns 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 confluctuation 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.

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. 7b). It 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).

3. Discussion

Here, we extended our recent analyses of edge time series and putative high-amplitude confluctuation events (Esfahlani et al., 2020; Sporns et al., 2020). 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 confluctuation patterns expressed during these frames, discovering that across scans of the same individual, the confluctuation patterns expressed during frames with lower-than-expected amplitude were dissimilar. In contrast, we found that the confluctuation during high-amplitude frames was repeatable within an individual and dissimilar between individuals. We then clustered patterns of confluctuation expressed during high-amplitude frames, identifying a small number of confluctuation 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 confluctuation 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.

3.1. High-amplitude confluctuations 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 (Friston, 1994; Horwitz, 2003). It is often used to construct a graphical representation of the brain to be analyzed using tools from network science (Bullmore and Sporns, 2009; Rubinov and Sporns, 2010). Although inter-individual differences in FC have been linked to an individual's clinical (Lynall et al., 2010), cognitive (Cohen and D'Esposito, 2016), and developmental state (Chan et al., 2014), the dynamic origins of individualized FC remain unknown.

Recently, we presented a method for decomposing FC into its framewise (instantaneous) contributions (Esfahlani et al., 2020; Faskowitz et al., 2020). Our work, in agreement with other recent studies (Tagliazucchi et al., 2012; Petridou et al., 2013; Allan et al., 2015; Cifre et al., 2020), 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 confluctuations 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 confluctuations. Using a data-driven clustering method, we find evidence of two patterns of confluctuation 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 confluctuations 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

(Biswal et al., 1995) and coordinating flexible, goal-directed behavior (Dosenbach et al., 2007). 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 (Power et al., 2013) – a graph-theoretic measure that describes the extent to which a node’s connections are distributed across or concentrated within communities (Guimera and Amaral, 2005).

Another interesting feature of these communities is their possible relationship to network states obtained by clustering sliding-window estimates of time-varying FC (Allen et al., 2014) or co-activation patterns (CAPs) (Liu and Duyn, 2013). Indeed, these approaches all yield estimates of repeating patterns of coactivity and connectivity suggesting a deep mathematical relationship, which was disclosed recently in a new paper (Novelli and Razi, 2021). Here, we analyze edge time series, a parameter-free method for estimating instantaneous confluctuations 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 (Hindriks et al., 2016; Shakil et al., 2016). 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 (Petridou et al., 2013; Liu et al., 2013; Tagliazucchi et al., 2012; Cifre et al., 2020; Allan et al., 2015). CAPs, on the other hand, which leverages a similar procedure as ours, identifies repeated and high-amplitude patterns of activity rather than confluctuation matrices (the product of instantaneous activity patterns). For this reason, CAPs traditionally does not offer the mathematical link between confluctuations and FC that edge time series does (Liu et al., 2013). Critically, then, 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 confluctuation 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 (Esfahlani et al., 2020) and open up opportunities for future studies to investigate interindividual differences in these patterns as well as the smaller and less frequent patterns.

3.2. High-amplitude confluctuation 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 confluctuation patterns are shared across individuals but expressed in different subject-specific proportions or b) confluctuation 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 (Mueller et al., 2013; Seitzman et al., 2019), has implications for brain-behavior studies (Rosenberg et al., 2016), and also informs our understanding of time-varying FC (Preti et al., 2017).

Here, we addressed these questions by aggregating and clustering high-amplitude confluctuation 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 confluctuations 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 (Tagliazucchi et al., 2012; Petridou et al., 2013; Allan et al., 2015; Cifre et al., 2020), this model assumed that FC is

driven exclusively by high-amplitude confluctuations 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 $\approx 50\%$ 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 $\approx 75\%$ 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 confluctuation patterns are individualized. This endorsement of hypothesis b comes with some caveats, however. Although FC is the product of individualized confluctuation 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 (Anderson et al., 2021; Kong et al., 2019; Kraus et al., 2020) to regional FC fingerprints (Finn et al., 2015). Notably, these findings also corroborate other “edge-centric” analyses of individualized patterns of brain connectivity. For instance, in (Jo et al., 2021), the authors demonstrated that the correlation structure of edge time series – a construct referred to as “edge functional connectivity” outperformed traditional FC in terms of identifiability. Moreover, reconstructions of eFC using principal components further improved its performance. Those results, paired with those of the present study, suggest that edge-based approaches may offer a useful framework for investigating individualization of connectivity and, eventually, linking it back to behavioral, cognitive, and clinical phenotypes.

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 confluctuation patterns. However, our sensitivity analysis (Fig. 7d) 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 (Allen et al., 2014; Shakil et al., 2016). 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.

3.3. 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 confluctuation patterns during events are shared versus individualized, the small number of participants precludes the possibility of investigating behavioral, cognitive, or disease correlates of events

(Marek et al., 2020). 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 co-fluctuations 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 (Cohen and D'Esposito, 2016). Do these changes reflect different patterns of co-fluctuations 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 (Greene et al., 2018)?

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 co-fluctuations make proportionally bigger contributions to time-averaged FC than low-amplitude co-fluctuations. 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 (Esfahlani et al., 2020; Faskowitz et al., 2020; Sporns et al., 2020; Jo et al., 2020a,b). 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.

Edge time series represent only the latest in a series of methods for tracking and modeling time-varying changes in networks that include time-frequency analysis (Chang and Glover, 2010), sliding windows (Hindriks et al., 2016), instantaneous phase-locking (Cabral et al., 2017), co-activation patterns (Liu and Duyn, 2013; Karahanoglu and Van De Ville, 2015; Petridou et al., 2013; Tagliazucchi et al., 2012), multiplication of temporal derivatives (Shine et al., 2015), quasi-periodic patterns (Majeed et al., 2009), and model-based frameworks (Lindquist et al., 2014). Broadly, these different approaches can be classified on the timescales over which they describe changes in brain activity or connectivity. Some, like CAPs, edge time series, and instantaneous phase-locking, characterize changes at a framewise timescale. Others require windowing or smoothing of data to track time-varying fluctuations. Of particular interest are quasi-period patterns, (QPPs), which refer to repeated sequences of brain activity (Majeed et al., 2011; Thompson et al., 2014; Yousefi et al., 2018). Recently, (Abbas et al., 2019) developed a method for assessing contributions to FC by repeating patterns. In essence, they “regress out” each QPP’s time course from the BOLD signal and calculate FC before and after doing so, allowing the authors to assess the contribution of the QPP to FC. More recently, (Bolt et al., 2021) identified three QPPs and demonstrated that, using only the time courses of those QPPs, could reconstruct an FC matrix that was correlated with the FC matrix estimated using the entire BOLD signal. Although these studies collectively suggest that QPPs contribute to the overall pattern of FC, the precise mathematical link remains inexact. Moreover, every QPP evolves over an extended period of time and includes multiple frames, further complicating its relationship with FC. In contrast, edge time series are an exact decomposition of FC into its framewise contributions. That is, the mean of a given edge time series is mathematically equivalent to that edge’s weight in the FC matrix. This makes it possible to assess the collective contribution of individual frames to the overall FC pattern (although the average of subsets

of frames, as we do here, generates a mean co-fluctuation value that does not correspond to a true correlation coefficient, once again making the link between co-fluctuations and FC inexact). In contrast, with QPPs and co-activation patterns (CAPs), the mathematical link to FC is not precise. Irrespective of whether their contributions can be determined precisely, QPPs, CAPs, iCAPs, “events” and related methods make clear that variation in activity and co-activity across time contribute to FC. Future work should focus not only on assessing the relative strengths and weakness of these methods but identifying underlying structural similarities between approaches (Novelli and Razi, 2021).

Here, we analyze two dense-sampling datasets in which a small number of individuals were scanned many times. Our work is naturally extended by examining inter-individual variation of high-amplitude co-fluctuation patterns and community frequency in large, cross-sectional datasets, e.g. the Human Connectome Project (Van Essen et al., 2013). In addition to imaging data, the HCP dataset includes rich cognitive, behavioral, and clinical assessments of participants. Future work should focus on linking inter-individual variation in those measures with high-amplitude events.

The results of our study suggest that co-fluctuation patterns are individualized and, in turn, contribute to the individualization of time-averaged, static FC. However, there remain several open questions about the origins of this individualization. Do co-fluctuation patterns appear individualized because of individualized differences in anatomy and subtle shifts in parcel boundaries (Bijsterbosch et al., 2018, 2019)? Here, we tried to address this issue by replicating our results using a coarser parcellation (100 nodes rather than 333). While the main results do, indeed, replicate, we also find that the size of the effect (difference in predicted FC using individualized versus group parcels) decreases. This observation suggests that, as expected, true differences in the underlying anatomy contribute to shaping the observed individualization of co-fluctuation patterns. However, a second contributing factor likely relates to the well-documented dependence of whole-brain FC (and derived measures) on choice of parcels is well-documented (Eickhoff et al., 2018; Gordon et al., 2016). Adopting coarser parcellations necessarily “washes out” idiosyncrasies of FC that may be evident with finer, more detailed parcellations. Adjudicating between these two possibilities is beyond the scope of this paper, but is a topic of intense ongoing research (Finn and Rosenberg, 2021; Bijsterbosch et al., 2020). Future studies should investigate the relative contributions of anatomy, parcellation choice, and other factors in shaping the apparent individualization of FC patterns.

Here, we focus on high-amplitude co-fluctuations. Our rationale is that, because the Pearson correlation (static FC) is a temporal average over framewise co-fluctuations, the frames with largest amplitude will necessarily have a greater impact on the mean value. That is, on a per-frame basis, frames with larger amplitudes mathematically will contribute more to the temporal mean. In focusing on high-amplitude frames, we necessarily neglect contributions of middle- and low-amplitude frames, which can also be used to approximate static FC (albeit not as accurately as high-amplitude; see Fig. S1). Nonetheless, an important direction for future research is to better characterize non-event frames. For instance, do they exhibit cluster structure? Are they simply “scaled down” analogs of the high-amplitude events? What are the physiological and psychological drivers of non-event frames?

Finally, a key overarching and open question concerns the origins of high-amplitude co-fluctuations. In our previous study, we demonstrated that movie-watching leads to synchronization of events across participants (Esfahlani et al., 2020), 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 (Newbold et al., 2020b), 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 (Newbold et al., 2020a). Another possibility is that burst activation of distinct systems, including default mode networks, co-occurs with the recall or previously-observed stimuli (Higgins et al., 2020). Yet another possibility, and one that is supported by recent results from electrophysiological studies (Kucyi et al., 2020), is that rapid and antagonistic (anti-correlated) relationships between specific subnetworks, including the default mode, are strongly associated with sustained attention, suggesting physiological origins. In all cases, these activation patterns are likely constrained by the underlying anatomical connectivity and reflect groups of mutually connected brain regions (Avena-Koenigsberger et al., 2018; Pope et al., 2021). A final possibility is that events are truly stochastic and are mere biproducts of modular correlated time series, where the activation of one element within a module implies the activation of the others (Novelli and Razi, 2021). 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.

3.4. 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 (Logothetis et al., 2001), intracranial EEG (Betzel et al., 2019), and optically recorded calcium imaging signals (Lake et al., 2020), 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 (Power et al., 2018; Lynch et al., 2020). 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 (Glasser et al., 2018, 2019; Power et al., 2019).

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 $\approx 25\%$ 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.

4. 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 co-fluctuation, broadly shared across individuals. However, these patterns undergo refinement at the level of individual participants, yielding reliable and individualized co-fluctuation patterns. Finally, we show that participants’ FC is more accurately predicted using participant-specific estimates of their high-amplitude co-fluctuation patterns compared to group-level estimates. Our study discloses high-amplitude, network-wide co-fluctuations as dynamical drivers of individualized FC and intro-

duces methodology for exploring their role in cognition, development, and disease in future studies.

5. Materials and methods

5.1. Datasets

5.1.1. Midnight scan club

The description of the Midnight Scan Club dataset acquisition, pre-processing, and network modeling is described in detail in (Gordon et al., 2017b). Here, we provide a high-level overview. Data were collected from ten healthy, righthanded, 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 (Gordon et al., 2017b; Gratton et al., 2018) 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 (Gordon et al., 2017b).

High-resolution structural MRI data were averaged together, and the average T1 images were used to generate hand-edited cortical surfaces using Freesurfer (Dale et al., 1999). The resulting surfaces were registered into fs LR 32k surface space as described in (Glasser et al., 2013). Separately, an average native T1-to-Talairach (Talairach, 1988) volumetric atlas transform was calculated. That transform was applied to the fs LR 32k surfaces to put them into Talairach volumetric space.

Volumetric fMRI pre-processing included including slice-timing correction, frame-to-frame alignment to correct 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 (Gordon et al., 2017b). Registration, atlas transformation, re-sampling to 3 mm isotropic resolution, and distortion correction were all combined and applied in a single transformation step (Smith et al., 2004). Subsequent steps were all completed on the atlas transformed and resampled data.

Several connectivity-specific steps were included (see (Power et al., 2014)): (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 (Gordon et al., 2017b)) 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 nonneuronal sources (Power et al., 2014; Ciric et al., 2017). 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 (Power et al., 2012)), 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 $28\% \pm 18\%$ (range: 6% – 67%) of the data; on average, participants retained 5929 $\pm$ 1508 volumes (range: 2733 – 7667). Note that in this paradigm, even the worst participant retained almost two hours of data. Nonetheless, we excluded two subjects from all analyses, both of whom

had fewer than 50% usable frames in at least five scan sessions (MSC08 in 7/10 and MSC9 in 5/10). See Fig. S24 for a summary of usable frames for each subject and scan.

Time courses were extracted from $N = 333$ cortical regions using a common (group) functional parcellation (Gordon et al., 2016). We also analyze time courses estimated from using individualized parcellations (see (Gordon et al., 2017a) for details). Both group and individualized time series were used for FC estimation and edge time series generation.

5.1.2. 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 (Laumann et al., 2015; Poldrack et al., 2015). 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.

5.1.3. 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 $z_i = [z_i(1), \dots, z_i(T)]$, where T is the number of samples. The Pearson correlation is then calculated as:

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

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

5.1.4. Edge time series

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

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

where the value of $r_{ij}(t)$ indexes the instantaneous confluctuation 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 confluctuation 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 confluctuation between all node pairs as their root sum square, $RSS(t) = \sqrt{\sum_{i,j>i} r_{ij}(t)^2}$.

5.1.5. Multiresolution consensus clustering

We used a variation of modularity maximization (Newman and Girvan, 2004) to group confluctuation 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:

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

In this expression, S_{rs} is the similarity of confluctuation 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 multiresolution communities (Jeub et al., 2018). That is, a partition of confluctuation 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 on the distribution of S_{rs} values. At each value, we use a generalized version of the Louvain algorithm to optimize the corresponding $Q(\gamma)$ (Jutla et al., 2011) (<http://netwiki.amath.unc.edu/GenLouvain/GenLouvain>). 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:

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

In this expression, P_{rs}^c 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_r^c and z_s^c 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, a new co-assignment matrix is estimated and its modularity optimized. This procedure repeats until the algorithm converges. We use the same generalized version of the Louvain algorithm to optimize the consensus modularity function and detect consensus communities.

5.1.6. Principal component analysis of communities

The multiresolution consensus clustering procedure groups a series of vectorized $N \times N$ high-amplitude, confluctuation matrices into communities. Each clustered confluctuation matrix corresponds to the peak frame in a temporally contiguous series of frames whose RSS was significantly greater than that of a null model. While our primary aim was to understand how these matrices contribute to time-averaged FC, we were also interested in characterizing what types of activity, i.e. $N \times 1$ patterns of fMRI BOLD data, give rise to high-amplitude confluctuations.

Naively, one could address this question by identifying the frames corresponding to peak co-fluctuations and average over their corresponding activity patterns. This approach, however, can yield misleading results. This is because every co-fluctuation matrix can be generated by two different patterns of activity that are identical to one another except for their signs. To illustrate this, consider the toy case presented below. Suppose we have a four-node network with the following confluctuation matrix:

$$C = \begin{bmatrix} +1 & +1 & -1 & -1 \\ +1 & +1 & -1 & -1 \\ -1 & -1 & +1 & +1 \\ -1 & -1 & +1 & +1 \end{bmatrix}$$

The elements of this confluctuation matrix are given by $C_{ij} = z_i z_j$, where z_i is the z-scored activity of node i . Accordingly, this matrix could

have been generated by either of the following patterns of activity:

$$z = \begin{bmatrix} +1 \\ +1 \\ -1 \\ -1 \end{bmatrix}$$

or

$$z' = \begin{bmatrix} -1 \\ -1 \\ +1 \\ +1 \end{bmatrix}.$$

The simple average of z and z' is a vector of zeros, which is unrelated to the co-fluctuation matrix. Averaging peak activity patterns can yield a similarly misleading result. Fortunately, both patterns are co-linear and satisfy the relationship $z = -1 \cdot z'$. That is, these patterns can be described by a single “mode” of activity, which we can detect by applying principal components analysis (PCA) to the activity patterns directly.

With this in mind, our strategy for characterizing the patterns of activity underpinning high-amplitude co-fluctuations was as follows. For each of the co-fluctuation matrices assigned to community 1, 2, and 3, we identified the scan and frame number in which they originated. Separately for each community, we aggregated the corresponding patterns of activity into matrices with dimensions of $N \times N_{\text{peaks}}$. Here, N_{peaks} is the number of high-amplitude frames assigned to a given community. Then, we applied PCA to these matrices, yielding a series of orthogonal components of dimension $N \times 1$ and the variance explained by each component. Note that these components are used mostly for visualization and to better understand the link between co-fluctuation matrices and brain activity. In the main text, we also describe a second approach for uncovering the dominant mode of activity underpinning the centroids of communities. Each centroid represents the average over many co-fluctuation matrices. To discover the optimal mode of activity, we aimed to determine the elements of $z = [z_1, \dots, z_N]$ that minimized the following cost function:

$$P = \sum_{ij} (C_{ij} - z_i z_j)^2.$$

To optimize P , we used a greedy algorithm which we repeated 100 times. Briefly, we initialized the algorithm with a $z \in \mathbb{R}^{N \times 1}$ vector whose elements were drawn independently from $\mathcal{N}(0, 1)$ and calculated the corresponding cost of P . Then we randomly selected a node, i , replaced its current value with another value randomly drawn from the same distribution. We denote the resulting vector as z' and its cost as P' . If $P' < P$ then we retained z' . We repeated this procedure 25000 times, gradually reducing P . In practice, we found that the algorithm converged to highly similar solutions (mean similarity across 100 runs of $r = 0.993$).

5.1.7. Predictive model of FC

In the main text we described a procedure for modeling FC in terms of co-fluctuation community centroids. In this section, we provide more details of how the model works. In our previous work (Esfahlani et al., 2020), 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 co-fluctuation patterns are expressed at high-amplitude frames. Here, we take this claim one step further and state that the co-fluctuation 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:

$$FC_{\text{subject,scan}} = f_1 c_1 + f_2 c_2 + f_3 c_3 + f_0 0$$

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 c_1 , c_2 , and c_3 represent the average pattern of co-fluctuation for the three communities. The final parameter, 0 is a node-by-node matrix where all elements are zero.

The first two models make predictions of FC using subject-specific estimates of centroids and frequency estimates made at the subject and group-level (models 1 and 2). Similarly, the next two models pair group-level estimates of centroids with subject and group-level estimates of frequencies (models 3 and 4). Finally, we test a fifth model that makes predictions of subject's FC using subject-specific centroids and frequencies estimated for the remaining seven subjects. For this model, we retain the best fit of the seven.

Note that when predicting FC for a given scan, both subject and group-level centroids and frequencies are estimated while excluding observations from that scan. Note also that frequency is estimated as the total number of frames associated with a given community. To obtain this number, we first map the community assignments of event co-fluctuation patterns back to the segment they originated (a temporally contiguous set of frames whose amplitude is significantly greater than that of the null model). We then assign the same community label to all frames that make up that segment. Finally, we calculate the frequency of each community as the total number of frames assigned to that community divided by the total possible number of frames.

In general, the model can generate matrices that are not positive semidefinite (all positive eigenvalues) and therefore not possible correlation matrices. Accordingly, we transform each matrix (translation/rotation/scaling) to match the nearest admissible matrix by minimizing the Frobenius norm using the MATLAB function `nearcorr.m`. 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 and code 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> Mew7h-yjtAC4ObzG VaKa?dl=0. Processed and parcellated MyConnectome data is available here: <http://myconnectome.org/wp/data-sharing/>. Code for calculating edge time series, detecting events, clustering events, and predicting FC from event clusters is available here: <https://github.com/brain-networks/event> detection.

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.

RFB designed experiments, wrote code, carried out analyses, and wrote and revised manuscript.

SC, SG, JF, and OS wrote and revised manuscript.

JF processed data.

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).

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2022.118993.

References

- Abbas, A., Belloy, M., Kashyap, A., Billings, J., Nezafati, M., Schumacher, E.H., Keilholz, S., 2019. Quasi-periodic patterns contribute to functional connectivity in the brain. *Neuroimage* 191, 193–204.
- Alexander-Bloch, A.F., Shou, H., Liu, S., Satterthwaite, T.D., Glahn, D.C., Shinohara, R.T., Vandekar, S.N., Raznahan, A., 2018. On testing for spatial correspondence between maps of human brain structure and function. *Neuroimage* 178, 540–551.
- Allan, T.W., Francis, S.T., Caballero-Gaudes, C., Morris, P.G., Liddle, E.B., Liddle, P.F., Brookes, M.J., Gowland, P.A., 2015. Functional connectivity in MRI is driven by spontaneous bold events. *PLoS One* 10, e0124577.
- Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking whole-brain connectivity dynamics in the resting state. *Cereb. Cortex* 24, 663–676.
- Amico, E., Goñi, J., 2018. Mapping hybrid functional-structural connectivity traits in the human connectome. *Netw. Neurosci.* 2, 306–322.
- Anderson, K.M., Ge, T., Kong, R., Patrick, L.M., Spreng, R.N., Sabuncu, M.R., Yeo, B.T., Holmes, A.J., 2021. Heritability of individualized cortical network topography. *Proc. Natl. Acad. Sci.* 118.
- Avena-Koenigsberger, A., Misis, B., Sporns, O., 2018. Communication dynamics in complex brain networks. *Nat. Rev. Neurosci.* 19, 17.
- Betzel, R.F., Medaglia, J.D., Kahn, A.E., Soffer, J., Schonhaut, D.R., Bassett, D.S., 2019. Structural, geometric and genetic factors predict interregional brain connectivity patterns probed by electrocorticography. *Nat. Biomed. Eng.* 3, 902–916.
- Bijsterbosch, J., Harrison, S.J., Jbabdi, S., Woolrich, M., Beckmann, C., Smith, S., Duff, E.P., 2020. Challenges and future directions for representations of functional brain organization. *Nat. Neurosci.* 23, 1484–1495.
- Bijsterbosch, J.D., Beckmann, C.F., Woolrich, M.W., Smith, S.M., Harrison, S.J., 2019. The relationship between spatial configuration and functional connectivity of brain regions revisited. *Elife* 8, e44890.
- Bijsterbosch, J.D., Woolrich, M.W., Glasser, M.F., Robinson, E.C., Beckmann, C.F., Van Essen, D.C., Harrison, S.J., Smith, S.M., 2018. The relationship between spatial configuration and functional connectivity of brain regions. *Elife* 7, e32992.
- Biswal, B., Zerrin Yetkin, F., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magn. Reson. Med.* 34, 537–541.
- Bolt, T., Nomi, J.S., Bzdok, D., Chang, C., Yeo, B.T., Uddin, L.Q., Keilholz, S.D., 2021. Large-scale intrinsic functional brain organization emerges from three canonical spatiotemporal patterns. *bioRxiv*.
- Bolton, T.A., Kebets, V., Glerean, E., Z'oller, D., Li, J., Yeo, B.T., Caballero-Gaudes, C., Van De Ville, D., 2020. Agito ergo sum: correlates of spatiotemporal motion characteristics during fmri. *Neuroimage* 209, 116433.
- Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.* 10, 186–198.
- Cabral, J., Vidaurre, D., Marques, P., Magalhães, R., Moreira, P.S., Soares, J.M., Deco, G., Sousa, N., Kringelbach, M.L., 2017. Cognitive performance in healthy older adults relates to spontaneous switching between states of functional connectivity during rest. *Sci. Rep.* 7, 1–13.
- Chai, X.J., Castan˘on, A.N., O˘ngur, D., Whitfield-Gabrieli, S., 2012. Anticorrelations in resting state networks without global signal regression. *Neuroimage* 59, 1420–1428.
- Chan, M.Y., Park, D.C., Savalia, N.K., Petersen, S.E., Wig, G.S., 2014. Decreased segregation of brain systems across the healthy adult lifespan. *Proc. Natl. Acad. Sci.* 111, E4997–E5006.
- Chang, C., Glover, G.H., 2010. Time–frequency dynamics of resting-state brain connectivity measured with fmri. *Neuroimage* 50, 81–98.
- Chen, J.E., Chang, C., Greicius, M.D., Glover, G.H., 2015. Introducing coactivation pattern metrics to quantify spontaneous brain network dynamics. *Neuroimage* 111, 476–488.
- Cifre, I., Zarepour, M., Horovitz, S., Cannas, S., Chialvo, D., 2020. Further results on why a point process is effective for estimating correlation between brain regions. *Pap. Phys.* 12, 120003–120003.
- Ciric, R., Wolf, D.H., Power, J.D., Roalf, D.R., Baum, G.L., Ruparel, K., Shinohara, R.T., Elliott, M.A., Eickhoff, S.B., Davatzikos, C., et al., 2017. Benchmarking of participant-level confound regression strategies for the control of motion artifact in studies of functional connectivity. *Neuroimage* 154, 174–187.
- Cohen, J.R., D'Esposito, M., 2016. The segregation and integration of distinct brain networks and their relationship to cognition. *J. Neurosci.* 36, 12083–12094.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis: I. segmentation and surface reconstruction. *Neuroimage* 9, 179–194.
- Demeter, D.V., Engelhardt, L.E., Mallett, R., Gordon, E.M., Nugiel, T., Harden, K.P., Tucker-Drob, E.M., Lewis-Peacock, J.A., Church, J.A., 2020. Functional connectivity fingerprints at rest are similar across youths and adults and vary with genetic similarity. *Science* 23, 100801.
- Dosenbach, N.U., Fair, D.A., Miezin, F.M., Cohen, A.L., Wenger, K.K., Dosenbach, R.A., Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E., et al., 2007. Distinct brain networks for adaptive and stable task control in humans. *Proc. Natl. Acad. Sci.* 104, 11073–11078.
- Dubois, J., Adolphs, R., 2016. Building a science of individual differences from fmri. *Trends Cogn. Sci.* 20, 425–443.
- Eickhoff, S.B., Yeo, B.T., Genon, S., 2018. Imaging-based parcellations of the human brain. *Nat. Rev. Neurosci.* 19, 672–686.
- Esfahlani, F.Z., Byrge, L., Tanner, J., Sporns, O., Kennedy, D., Betzel, R., 2021. Edge-centric analysis of time-varying functional brain networks with applications in autism spectrum disorder. *bioRxiv*.
- Esfahlani, F.Z., Jo, Y., Faskowitz, J., Byrge, L., Kennedy, D., Sporns, O., Betzel, R., 2020. High-amplitude co-fluctuations in cortical activity drive functional connectivity. In: *Proceedings of the National Academy of Sciences*.
- Esteban, O., Markiewicz, C.J., Blair, R.W., Moodie, C.A., Isik, A.I., Erramuzpe, A., Kent, J.D., Goncalves, M., DuPre, E., Snyder, M., et al., 2019. fmriprep: a robust pre-processing pipeline for functional MRI. *Nat. Methods* 16, 111–116.
- Faskowitz, J., Esfahlani, F.Z., Jo, Y., Sporns, O., Betzel, R.F., 2020. Edgcentric functional network representations of human cerebral cortex reveal overlapping system-level architecture. *Nat. Neurosci.* 23, 1644–1654.
- Finn, E.S., Rosenberg, M.D., 2021. Beyond fingerprinting: Choosing predictive connectomes over reliable connectomes. *Neuroimage*, 118254.
- Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nat. Neurosci.* 18, 1664–1671.
- Fong, A.H.C., Yoo, K., Rosenberg, M.D., Zhang, S., Li, C.S.R., Scheinost, D., Constable, R.T., Chun, M.M., 2019. Dynamic functional connectivity during task performance and rest predicts individual differences in attention across studies. *Neuroimage* 188, 14–25.
- Fox, M.D., Zhang, D., Snyder, A.Z., Raichle, M.E., 2009. The global signal and observed anticorrelated resting state brain networks. *J. Neurophysiol.* 101, 3270–3283.
- Friston, K.J., 1994. Functional and effective connectivity in neuroimaging: a synthesis. *Hum. Brain Mapp.* 2, 56–78.
- Glasser, M.F., Coalson, T.S., Bijsterbosch, J.D., Harrison, S.J., Harms, M.P., Anticevic, A., Van Essen, D.C., Smith, S.M., 2018. Using temporal ica to selectively remove global noise while preserving global signal in functional mri data. *Neuroimage* 181, 692–717.
- Glasser, M.F., Coalson, T.S., Bijsterbosch, J.D., Harrison, S.J., Harms, M.P., Anticevic, A., Van Essen, D.C., Smith, S.M., 2019. Classification of temporal ica components for separating global noise from fMRI data: Reply to power. *Neuroimage* 197, 435–438.
- Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., et al., 2013. The minimal preprocessing pipelines for the human connectome project. *Neuroimage* 80, 105–124.
- Gordon, E.M., Laumann, T.O., Adeyemo, B., Huckins, J.F., Kelley, W.M., Petersen, S.E., 2016. Generation and evaluation of a cortical area parcellation from resting-state correlations. *Cereb. Cortex* 26, 288–303.
- Gordon, E.M., Laumann, T.O., Adeyemo, B., Petersen, S.E., 2017a. Individual variability of the system-level organization of the human brain. *Cereb. Cortex* 27, 386–399.
- Gordon, E.M., Laumann, T.O., Gilmore, A.W., Newbold, D.J., Greene, D.J., Berg, J.J., Ortega, M., Hoyt-Drazen, C., Gratton, C., Sun, H., et al., 2017b. Precision functional mapping of individual human brains. *Neuron* 95, 791–807.
- Gratton, C., Kraus, B.T., Greene, D.J., Gordon, E.M., Laumann, T.O., Nelson, S.M., Dosenbach, N.U., Petersen, S.E., 2019. Defining individual-specific functional neuroanatomy for precision psychiatry. *Biol. Psychiatry*.
- Gratton, C., Laumann, T.O., Nielsen, A.N., Greene, D.J., Gordon, E.M., Gilmore, A.W., Nelson, S.M., Coalson, R.S., Snyder, A.Z., Schlaggar, B.L., et al., 2018. Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. *Neuron* 98, 439–452.
- Greene, A.S., Gao, S., Scheinost, D., Constable, R.T., 2018. Task-induced brain state manipulation improves prediction of individual traits. *Nat. Commun.* 9, 1–13.
- Greenwell, S., Faskowitz, J., Pritschet, L., Santander, T., Jacobs, E.G., Betzel, R.F., 2021. High-amplitude network co-fluctuations linked to variation in hormone concentrations over menstrual cycle. *bioRxiv*.
- Guimera, R., Amaral, L.A.N., 2005. Functional cartography of complex metabolic networks. *Nature* 433, 895–900.
- Higgins, C., Liu, Y., Vidaurre, D., Kurth-Nelson, Z., Dolan, R., Behrens, T., Woolrich, M., 2020. Replay bursts in humans coincide with activation of the default mode and parietal alpha networks. *Neuron*.
- Hindriks, R., Adhikari, M.H., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N.K., Deco, G., 2016. Can sliding-window correlations reveal dynamic functional connectivity in resting-state fmri? *Neuroimage* 127, 242–256.
- Horwitz, B., 2003. The elusive concept of brain connectivity. *Neuroimage* 19, 466–470.
- Jeub, L.G., Sporns, O., Fortunato, S., 2018. Multiresolution consensus clustering in networks. *Sci. Rep.* 8, 1–16.
- Jo, Y., Esfahlani, F.Z., Faskowitz, J., Chumin, E.J., Sporns, O., Betzel, R.F., 2020a. The diversity and multiplexity of edge communities within and between brain systems. *Cell reports* 37 (7), 110032.
- Jo, Y., Faskowitz, J., Esfahlani, F.Z., Sporns, O., Betzel, R.F., 2021. Subject identification using edge-centric functional connectivity. *Neuroimage*, 118204.
- Jutla, I.S., Jeub, L.G., Mucha, P.J., 2011. A generalized louvain method for community detection implemented in MATLAB. URL <http://netwiki.amath.unc.edu/GenLouvain>.
- Karahanoğlu, F.I., Van De Ville, D., 2015. Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. *Nat. Commun.* 6, 1–10.
- Kong, R., Li, J., Orban, C., Sabuncu, M.R., Liu, H., Schaefer, A., Sun, N., Zuo, X.N., Holmes, A.J., Eickhoff, S.B., et al., 2019. Spatial topography of individual-specific cortical networks predicts human cognition, personality, and emotion. *Cereb. Cortex* 29, 2533–2551.
- Kraus, B.T., Perez, D., Ladwig, Z., Seitzman, B.A., Dworetzky, A., Petersen, S.E., Gratton, C., 2020. Network variants are similar between task and rest states. *Neuroimage*, 117743.
- Kucyi, A., Daitch, A., Raccach, O., Zhao, B., Zhang, C., Esterman, M., Zeineh, M., Halpern, C.H., Zhang, K., Zhang, J., et al., 2020. Electrophysiological dynamics of autonomic brain networks reflect attentional fluctuations. *Nat. Commun.* 11, 1–14.
- Lake, E.M., Ge, X., Shen, X., Herman, P., Hyder, F., Cardin, J.A., Higley, M.J., Scheinost, D., Papademetris, X., Crain, M.C., et al., 2020. Simultaneous cortex-wide fluorescence ca 2+ imaging and whole-brain fMRI. *Nat. Methods* 17, 1262–1271.
- Laumann, T.O., Gordon, E.M., Adeyemo, B., Snyder, A.Z., Joo, S.J., Chen, M.Y., Gilmore, A.W., McDermott, K.B., Nelson, S.M., Dosenbach, N.U., et al., 2015. Functional system and areal organization of a highly sampled individual human brain. *Neuron* 87, 657–670.

- Lindquist, M.A., Xu, Y., Nebel, M.B., Caffo, B.S., 2014. Evaluating dynamic bivariate correlations in resting-state fmri: a comparison study and a new approach. *Neuroimage* 101, 531–546.
- Liu, J., Liao, X., Xia, M., He, Y., 2018. Chronnectome fingerprinting: identifying individuals and predicting higher cognitive functions using dynamic brain connectivity patterns. *Hum. Brain Mapp.* 39, 902–915.
- Liu, X., Chang, C., Duyn, J.H., 2013. Decomposition of spontaneous brain activity into distinct fmri co-activation patterns. *Front. Syst. Neurosci.* 7, 101.
- Liu, X., Duyn, J.H., 2013. Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proc. Natl. Acad. Sci.* 110, 4392–4397.
- Liu, Z.Q., Vazquez-Rodriguez, B., Spreng, R.N., Bernhardt, B., Betzel, R.F., Misis, B., 2021. Time-resolved structure-function coupling in brain networks. *bioRxiv*.
- Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412, 150–157.
- Lynall, M.E., Bassett, D.S., Kerwin, R., McKenna, P.J., Kitzbichler, M., Muller, U., Bullmore, E., 2010. Functional connectivity and brain networks in schizophrenia. *J. Neurosci.* 30, 9477–9487.
- Lynch, C.J., Silver, B.M., Dubin, M.J., Martin, A., Voss, H.U., Jones, R.M., Power, J.D., 2020. Prevalent and sex-biased breathing patterns modify functional connectivity MRI in young adults. *Nat. Commun.* 11, 1–14.
- Majeed, W., Magnuson, M., Hasenkamp, W., Schwab, H., Schumacher, E.H., Barsalou, L., Keilholz, S.D., 2011. Spatiotemporal dynamics of low frequency bold fluctuations in rats and humans. *Neuroimage* 54, 1140–1150.
- Majeed, W., Magnuson, M., Keilholz, S.D., 2009. Spatiotemporal dynamics of low frequency fluctuations in bold fMRI of the rat. *J. Magn. Reson. Imaging: Off. J. Int. Soc. Magn. Reson. Med.* 30, 384–393.
- Marek, S., Tervo-Clemmens, B., Calabro, F.J., Montez, D.F., Kay, B.P., Hatoum, A.S., Donohue, M.R., Foran, W., Miller, R.L., Feczko, E., et al., 2020. Towards reproducible brain-wide association studies. *BioRxiv*.
- Markello, R.D., Misis, B., 2021. Comparing spatial null models for brain maps. *Neuroimage*, 118052.
- Mueller, S., Wang, D., Fox, M.D., Yeo, B.T., Sepulcre, J., Sabuncu, M.R., Shafee, R., Lu, J., Liu, H., 2013. Individual variability in functional connectivity architecture of the human brain. *Neuron* 77, 586–595.
- Newbold, D.J., Gordon, E.M., Laumann, T.O., Montez, D.F., Seider, N.A., Gross, S.J., Zheng, A., Nielsen, A.N., Hoyt, C.R., Hampton, J.M., et al., 2020a. Cingulo-opercular control network supports disused motor circuits in standby mode. *bioRxiv*.
- Newbold, D.J., Laumann, T.O., Hoyt, C.R., Hampton, J.M., Montez, D.F., Raut, R.V., Ortega, M., Mitra, A., Nielsen, A.N., Miller, D.B., et al., 2020b. Plasticity and spontaneous activity pulses in disused human brain circuits. *Neuron* 107, 580–589.
- Newman, M.E., Girvan, M., 2004. Finding and evaluating community structure in networks. *Phys. Rev. E* 69, 026113.
- Novelli, L., Razi, A., 2021. A mathematical perspective on edge-centric functional connectivity. *arXiv preprint arXiv:2106.10631*.
- Petridou, N., Gaudes, C.C., Dryden, I.L., Francis, S.T., Gowland, P.A., 2013. Periods of rest in fMRI contain individual spontaneous events which are related to slowly fluctuating spontaneous activity. *Hum. Brain Mapp.* 34, 1319–1329.
- Poldrack, R.A., Laumann, T.O., Koyejo, O., Gregory, B., Hover, A., Chen, M.Y., Gorgolewski, K.J., Luci, J., Joo, S.J., Boyd, R.L., et al., 2015. Longterm neural and physiological phenotyping of a single human. *Nat. Commun.* 6, 1–15.
- Pope, M., Fukushima, M., Betzel, R., Sporns, O., 2021. Modular origins of high-amplitude co-fluctuations in fine-scale functional connectivity dynamics. *bioRxiv*.
- Power, J.D., 2019. Temporal ica has not properly separated global fMRI signals: a comment on Glasser et al (2018). *Neuroimage* 197, 650–651.
- Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity mri networks arise from subject motion. *Neuroimage* 59, 2142–2154.
- Power, J.D., Mitra, A., Laumann, T.O., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2014. Methods to detect, characterize, and remove motion artifact in resting state fmri. *Neuroimage* 84, 320–341.
- Power, J.D., Plitt, M., Gotts, S.J., Kundu, P., Voon, V., Bandettini, P.A., Martin, A., 2018. Ridding fmri data of motion-related influences: Removal of signals with distinct spatial and physical bases in multi-echo data. *Proc. Natl. Acad. Sci.* 115, E2105–E2114.
- Power, J.D., Schlaggar, B.L., Lessov-Schlaggar, C.N., Petersen, S.E., 2013. Evidence for hubs in human functional brain networks. *Neuron* 79, 798–813.
- Preti, M.G., Bolton, T.A., Van De Ville, D., 2017. The dynamic functional connectome: State-of-the-art and perspectives. *Neuroimage* 160, 41–54.
- Rogers, B.P., Morgan, V.L., Newton, A.T., Gore, J.C., 2007. Assessing functional connectivity in the human brain by fmri. *Magn. Reson. Imaging* 25, 1347–1357.
- Rosenberg, M.D., Finn, E.S., Scheinost, D., Papademetris, X., Shen, X., Constable, R.T., Chun, M.M., 2016. A neuromarker of sustained attention from whole-brain functional connectivity. *Nature neuroscience* 19 (1), 165–171.
- Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* 52, 1059–1069.
- Saad, Z.S., Gotts, S.J., Murphy, K., Chen, G., Jo, H.J., Martin, A., Cox, R.W., 2012. Trouble at rest: how correlation patterns and group differences become distorted after global signal regression. *Brain Connectivity* 2, 25–32.
- Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T., 2018. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cereb. Cortex* 28, 3095–3114.
- Seitzman, B.A., Gratton, C., Laumann, T.O., Gordon, E.M., Adeyemo, B., Dworetzky, A., Kraus, B.T., Gilmore, A.W., Berg, J.J., Ortega, M., et al., 2019. Trait-like variants in human functional brain networks. *Proc. Natl. Acad. Sci.* 116, 22851–22861.
- Shakil, S., Lee, C.H., Keilholz, S.D., 2016. Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states. *Neuroimage* 133, 111–128.
- Shine, J.M., Koyejo, O., Bell, P.T., Gorgolewski, K.J., Gilat, M., Poldrack, R.A., 2015. Estimation of dynamic functional connectivity using multiplication of temporal derivatives. *Neuroimage* 122, 399–407.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., et al., 2004. Advances in functional and structural MRI image analysis and implementation as fsl. *Neuroimage* 23, S208–S219.
- Sporns, O., Faskowitz, J., Teixeira, S., Betzel, R., 2020. Dynamic expression of brain functional systems disclosed by fine-scale analysis of edge time series. *bioRxiv*.
- Tagliazucchi, E., Balenzuela, P., Fraiman, D., Chialvo, D.R., 2012. Criticality in large-scale brain fmri dynamics unveiled by a novel point process analysis. *Front. Physiol.* 3, 15.
- Tagliazucchi, E., Siniatchkin, M., Laufs, H., Chialvo, D.R., 2016. The voxelwise functional connectome can be efficiently derived from co-activations in a sparse spatio-temporal point-process. *Front. Neurosci.* 10, 381.
- Talairach, J., 1988. Co-planar stereotaxic atlas of the human brain-3dimensional proportional system. An approach to cerebral imaging.
- Thompson, G.J., Pan, W.J., Magnuson, M.E., Jaeger, D., Keilholz, S.D., 2014. Quasi-periodic patterns (qpp): large-scale dynamics in resting state fMRI that correlate with local infraslow electrical activity. *Neuroimage* 84, 1018–1031.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.M.H., et al., 2013. The Wu-Minn human connectome project: an overview. *Neuroimage* 80, 62–79.
- Váša, F., Seidlitz, J., Romero-Garcia, R., Whitaker, K.J., Rosenthal, G., Vértes, P.E., Shinn, M., Alexander-Bloch, A., Fonagy, P., Dolan, R.J., et al., 2018. Adolescent tuning of association cortex in human structural brain networks. *Cereb. Cortex* 28, 281–294.
- Wig, G.S., Laumann, T.O., Cohen, A.L., Power, J.D., Nelson, S.M., Glasser, M.F., Miezin, F.M., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2014. Parcellating an individual subject's cortical and subcortical brain structures using snowball sampling of resting-state correlations. *Cereb. Cortex* 24, 2036–2054.
- Yousefi, B., Shin, J., Schumacher, E.H., Keilholz, S.D., 2018. Quasi-periodic patterns of intrinsic brain activity in individuals and their relationship to global signal. *Neuroimage* 167, 297–308.
- Zeng, L.L., Wang, D., Fox, M.D., Sabuncu, M., Hu, D., Ge, M., Buckner, R.L., Liu, H., 2014. Neurobiological basis of head motion in brain imaging. *Proc. Natl. Acad. Sci.* 111, 6058–6062.

Further reading

- Behzadi, Y., Restom, K., Liao, J., Liu, T.T., 2007. A component based noise correction method (compcor) for bold and perfusion based fmri. *Neuroimage* 37, 90–101.