

Large-Scale Intrinsic Functional Brain Organization Emerges from Three Canonical Spatiotemporal Patterns

Taylor Bolt¹, Jason S. Nomi², Danilo Bzdok³, Catie Chang⁴, B.T. Thomas Yeo⁵, Lucina Q. Uddin², Shella D. Keilholz¹

¹Emory University/Georgia Institute of Technology, Atlanta, GA, USA

²University of Miami, Coral Gables, FL, USA

³McGill University & Mila - Quebec Artificial Intelligence institute, Montreal, QC, Canada

⁴Vanderbilt University, Nashville, TN, USA

⁵National University of Singapore, Singapore

Summary

The characterization of intrinsic functional brain organization has been approached from a multitude of analytic techniques and methods. We are still at a loss of a unifying conceptual framework for capturing common insights across this patchwork of empirical findings. By analyzing resting-state fMRI data from the Human Connectome Project using a large number of popular analytic techniques, we find that all results can be seamlessly reconciled by three fundamental low-frequency spatiotemporal patterns that we have identified via a novel time-varying complex pattern analysis. Overall, these three spatiotemporal patterns account for a wide variety of previously observed phenomena in the resting-state fMRI literature including the task-positive/task-negative anticorrelation, the global signal, the primary functional connectivity gradient and the network community structure of the functional connectome. The shared spatial and temporal properties of these three canonical patterns suggest that they arise from a single hemodynamic mechanism.

Keywords: propagation, dynamics, time-lag, global signal, task-positive, task-negative, gradient, resting-state fMRI

Introduction

Many of us have heard the Indian parable of the blind men and the elephant. This anecdote teaches the perils of missing the ‘bigger picture’ due to our own limited observations. A group of blind men encounter an elephant for the first time and try to acquire an overall description of the elephant by discussing which part of the elephant each blind man is touching (trunk, tusk, foot, etc.). Each blind man knew a piece of the truth, but failed to grasp how their individual observations came together in a unified whole. This parable may be a suitable metaphor for the current state of our theoretical understanding of intrinsic functional brain organization.

Since the discovery of intrinsic low-frequency blood-oxygenation-level dependent (BOLD) fluctuations in the late 1990s, increasingly complex analytic techniques have been applied to understand the spatial and temporal structure of these signals. This has led to a fragmentation of the field into isolated conceptions of functional brain organization according to the researcher’s preferred analytic approach (Bijsterbosch et al., 2020), with little attempt to translate findings across different approaches. The field is lacking a unifying conceptual framework to translate findings across analytic approaches. Our present study aimed to develop such a framework to reconcile widely used, disparate analytic approaches. Surprisingly, we found that all the analytic approaches we surveyed are largely consistent in their results,

converging on a set of three, low-frequency spatiotemporal patterns of BOLD activity. Furthermore, with these three spatiotemporal patterns, we can account for a wide-variety of previously observed properties of intrinsic functional brain organization.

We operationalize ‘intrinsic functional brain organization’ as the spatial correlation structure between cortical BOLD signals in the low-frequency range ($\sim 0.01 - 0.1$ Hz). Spatial dependencies between neural activity fluctuations is by far the most studied property of low-frequency intrinsic BOLD signals. Importantly, the pattern of dependencies between neural activity across the brain are thought to reflect functional systems for cognition, perception and action (Biswal et al., 1995; Damoiseaux et al., 2006; Fox et al., 2005; Smith et al., 2009). For the purpose of the present study, we distinguish between two different descriptions of intrinsic functional brain organization - zero-lag synchrony between brain regions, and time-lag synchrony between brain regions. Zero-lag synchrony is defined as in-phase or anti-phase statistical dependence between intrinsic BOLD signals - e.g. the correlation between two intrinsic BOLD signals with no time-lag. Following the standard terminology of the functional magnetic resonance imaging (fMRI) literature, we refer to zero-lag synchrony among intrinsic BOLD fluctuations as ‘functional connectivity’ (FC) (Reid et al., 2019). FC between cortical brain regions organize into global, cortex-wide patterns, referred to as ‘FC topographies’. Examples of such global FC topographies include the task-positive/task-negative pattern (Fox et al., 2005), the primary FC gradient (Margulies et al., 2016) and global signal topography (Li et al., 2019).

More recently, relationships between BOLD time courses at temporal lags have been investigated. Studies of time-lag spatial structure in intrinsic BOLD signals have revealed 1) prominent wave-like, propagation patterns in cortical neural activity at short-time scales ($\sim 0 - 2$ s) (Mitra et al., 2014, 2015), and 2) recurring spatiotemporal patterns of BOLD activity changes at larger time scales (~ 20 s) (Majeed et al., 2011). At short time scales, Mitra and colleagues (2014) have observed replicable, time-lagged propagation of intrinsic BOLD activity that are uncoupled from hemodynamic delays. At longer time scales, Majeed et al. (2011) has described a pseudo-periodic spatiotemporal pattern involving an alteration in BOLD amplitudes between the frontoparietal (FPN) and default mode networks (DMN), known as the ‘quasi-periodic pattern’ (QPP). The shift in BOLD amplitudes between FPN and DMN regions is marked by a large-scale propagation of BOLD activity between the two networks. The QPP is thought to share a common mechanism with low-frequency local field potentials, according to multi-modal neuroimaging studies in rats and humans (Thompson et al., 2014). We refer to time-lag dynamics at both short- and long-time scales as ‘spatiotemporal patterns.’

The ever-increasing number and complexity of both zero- and time-lag analyses poses at least two challenges in creating a unified framework. The first challenge involves delineating the interrelationships between FC topographies derived from zero-lag analyses. For example, what

are the interrelationships between different approaches for constructing FC topographies, such as principal component analysis, non-linear manifold learning (Margulies et al., 2016), temporal independent component analysis (ICA) (Smith et al., 2012), spatial ICA (Calhoun et al., 2001), hidden markov models (Vidaurre et al., 2017), seed-based correlation analysis (Fox et al., 2005), and discrete co-activation patterns (Liu and Duyn, 2013b)?

The second challenge involves bridging the divide between zero-lag representations and time-lag representations of intrinsic functional brain organization. Specifically, is the relationship between FC topographies and spatiotemporal patterns one-to-many or one-to-one? If the relationship is one-to-many, FC topographies emerge from the shared dynamics of multiple spatiotemporal patterns. In this scenario, FC topographies describe distinct co-activation patterns that emerge from the additive, or potentially non-linear, combination of multiple spatiotemporal patterns. Consistent with this possibility, Mitra et al. (2015) have shown through simulation that the spatial structure of synchronous correlations in BOLD signals can emerge from the shared dynamics of multiple short time scale propagation patterns. If the relationship is one-to-one, distinct FC topographies correspond to complementary ‘static’ or ‘stationary’ descriptions of distinct spatiotemporal patterns. In other words, a single FC topography is a complementary description of a single spatiotemporal pattern. This possibility is also supported by recent research. Yousefi and Keilholz (2021) have shown that the time-lag relationships between BOLD time courses of the QPP correspond precisely to the ‘principal’ or ‘primary FC gradient’ discovered by (Margulies et al., 2016), suggesting that the primary FC gradient and QPP are a static and time-lagged description of the same spatiotemporal pattern, respectively.

We hypothesized that 1) most widely-used zero-lag analyses can be shown to produce a small set of fundamental FC topographies, and 2) that these fundamental FC topographies are complementary ‘static’ or ‘stationary’ descriptions of a small set of spatiotemporal patterns. Further, we hypothesize that the relationship between these fundamental FC topographies and spatiotemporal patterns is one-to-one, such that a single FC topography describes, at most, a single spatiotemporal pattern, and vice versa. To test these hypotheses, we conducted a systematic study of large-scale cortical FC topographies and spatiotemporal patterns in a cohort of participants with resting-state fMRI recordings from the Human Connectome Project (HCP, $n = 50$). We first carefully surveyed FC topographies produced from a variety of standard zero-lag FC analyses of intrinsic BOLD signals. We found that all zero-lag FC analyses surveyed consistently produced one of three spatially overlapping FC topographies. Using a novel application of a time-varying PCA algorithm, we find that these three FC topographies correspond to ‘static’ or time-invariant representations of three spatiotemporal patterns exhibiting dynamics divided into two main phases: a steady state period of relative stability in BOLD activity patterns between successive time points, and propagation events of rapid change in BOLD activity patterns between successive time points. Comparison of the dynamics between these spatiotemporal patterns revealed highly similar spatial patterns at certain phases of their

cycles. We then related these three spatiotemporal patterns to previously observed FC topographies and time-lag patterns, including lag projections (Mitra et al., 2014), the QPP (2011), the topography of the global signal (Li et al., 2019), the task-positive/task-negative pattern (Fox et al., 2005), and the primary FC gradient (Margulies et al., 2016). Finally, we show that the network structure of intrinsic BOLD signals can be explained by the shared dynamics of these three spatiotemporal patterns.

Results

1. Three Dominant Functional Connectivity Topographies in Intrinsic BOLD Fluctuations

We conducted a quantitative survey of widely-used zero-lag functional connectivity (FC) analyses with the primary aim of comparing the spatial overlap between the FC topographies produced by each analysis. These analyses were applied to a random sample ($n=50$) of human subject resting-state scans (~15min each; $n = 1200$ time points) from the Human Connectome Project (HCP). For input to all quantitative analyses, subject resting-state scans were temporally concatenated and reshaped into a 2D matrix of time points (rows) by cortical vertices (columns). All analyses in our investigation successfully replicated in an independent sample ($n=50$) of HCP subjects.

Our survey included several latent dimension-reduction methods, as well as seed-based correlation and co-activation methods. Latent dimension-reduction methods included principal component analysis (PCA), PCA with simple structure rotation (varimax) (Andersen et al., 1999; Thomas et al., 2002), Laplacian Eigenmaps (LE) (Vos de Wael et al., 2020), the commonly used spatial independent component analysis (SICA) (Calhoun et al., 2001), and the more recent temporal independent component analysis (TICA) (Smith et al., 2012). Hidden Markov models (HMM) are a commonly used latent state space model for estimating brain states (Vidaurre et al., 2017), and were also included in our study. Seed-based analysis methods included the traditional seed-based correlation/regression analysis (Fox et al., 2005) and co-activation pattern (CAP) analysis (Liu and Duyn, 2013b) with k-means clustering of suprathreshold time points into two clusters. Seed-based methods were run for three key seed locations corresponding to major hubs in the somatomotor network, default mode, and frontoparietal network - the somatosensory cortex, precuneus, and dorsolateral prefrontal cortex (Yeo et al., 2011). Results were found to be identical for alternatively placed seed regions within these three networks (**Supplementary Results B**).

To determine a useful number of dimensions in our latent dimension-reduction analyses (PCA, varimax PCA, LE, SICA, TICA and HMM) (Bzdok et al., 2016; Eickhoff et al., 2015), we examined the explained variance of the principal component solution at a range of dimension

numbers (see ‘Methods and Materials’). As we were interested in large-scale cortical FC topographies, our survey focuses on the lower end for the number of estimated latent dimensions. To estimate the number of dimensions, we examined the drop-off in explained variance (i.e. eigenvalues) associated with each successive principal component, a procedure known as Catell’s scree plot test (Carlson et al., 2011; Cattell, 1966; Ecker et al., 2007; Stetter et al., 2000). This widely used component number selection criterion indicated a clear drop-off in explained variance after three principal components (**Figure 1C**). Based on these assessments, we committed to a granularity of three latent neural activity dimensions for the dimension-reduction algorithms for comparability of subsequent analysis steps.

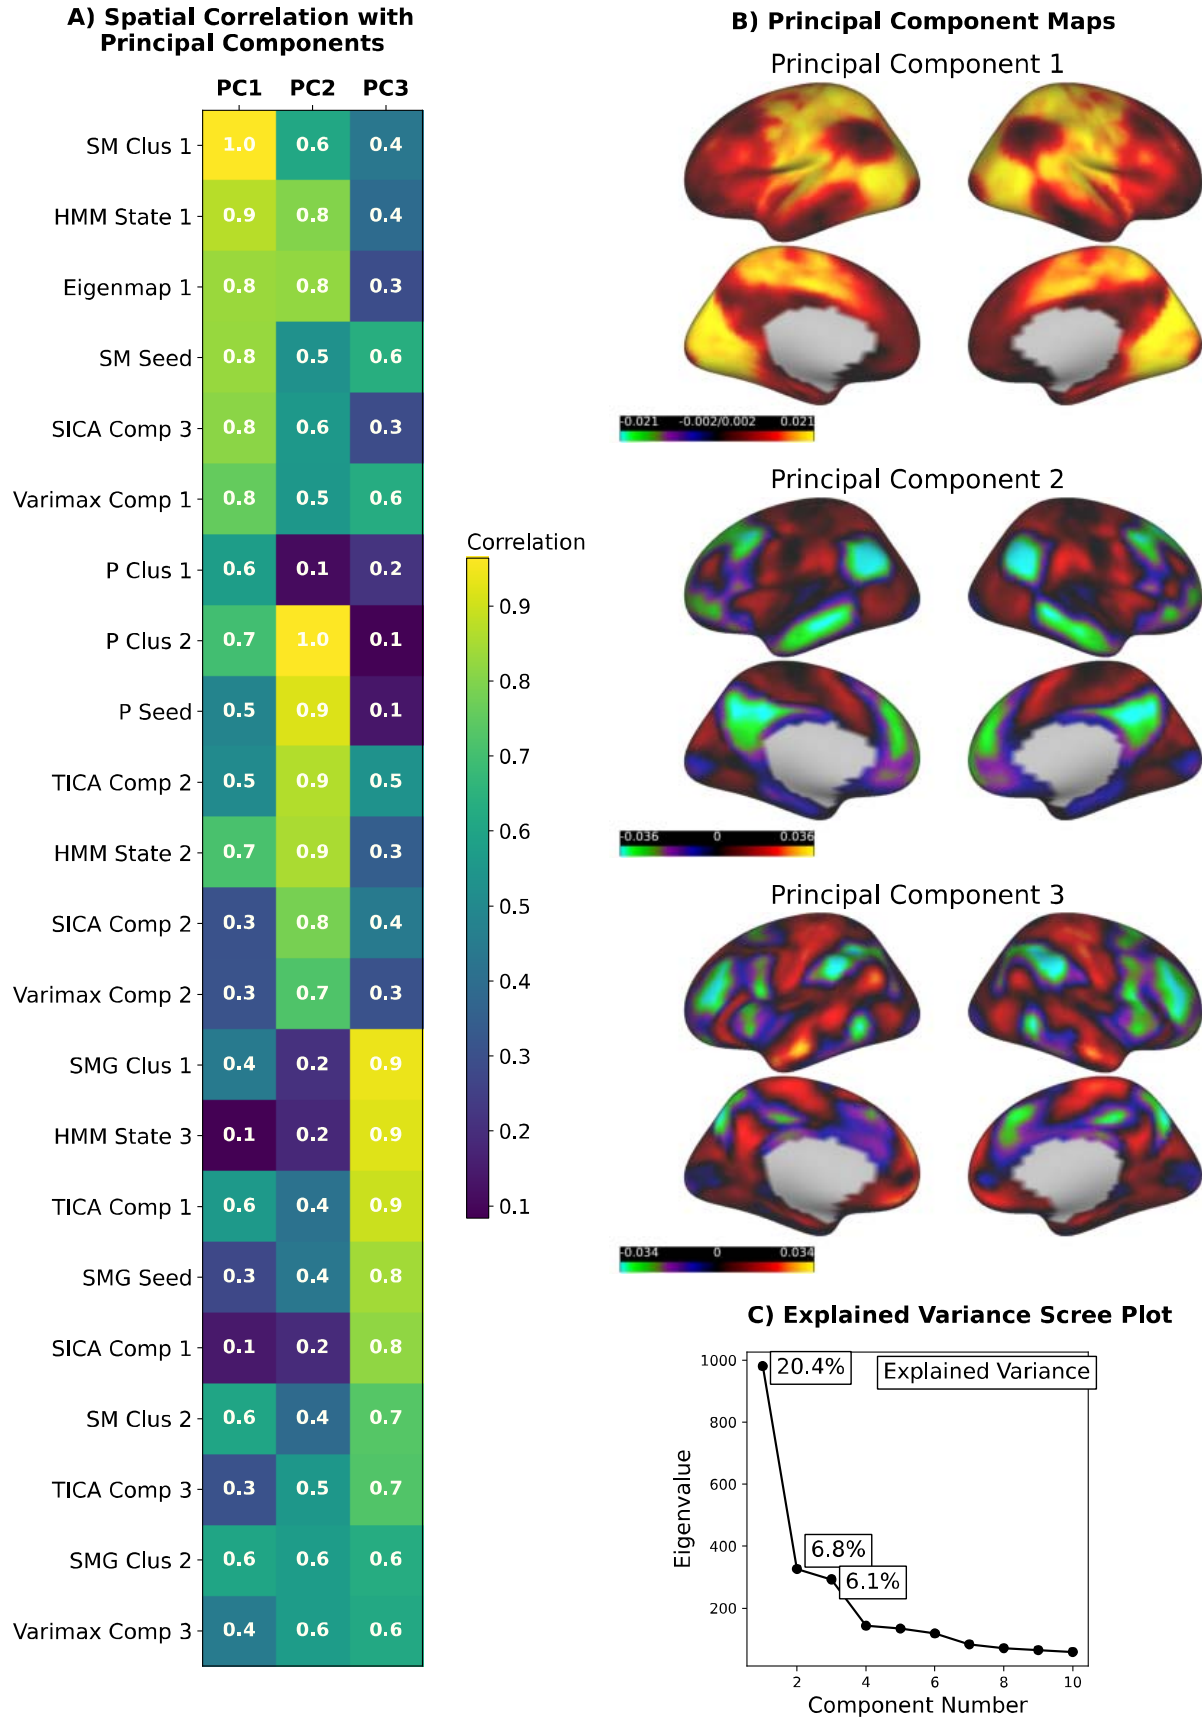


Figure 1. Form and Properties of Three Fundamental Functional Connectivity Topographies. **A)** The spatial correlation (abs. value) between the first three principal component maps and each FC topography displayed as a table. The color of each cell in the table is shaded from light yellow (strong correlation) to dark blue (weak correlation). All FC topographies in our survey exhibited strong spatial correlations (Pearson's correlation) with one (or two) of the first three principal components. **B)** The first three principal component spatial maps. **C)** The scree plot that displays the explained variance in cortical time series for each successive principal component. The scree plot indicates a clear elbow after the third principal component, indicating a 'diminishing return' in explained variance of extracting more components. (P=precuneus, SM= somatosensory; SMG=supramarginal gyrus; Clus=cluster; Comp=component; PC = Principal Component).

Each zero-lag analysis produced one or more FC topographies with weights for each cortical vertex, representing the degree to which that topography is expressed at that vertex. To compare the spatial similarity between two FC topographies, we used the spatial correlation (Pearson's correlation) between the cortical vertex weight values of each topography. To summarize the similarities among the FC topographies, we compared each topography to the first three principal component maps computed from PCA. The first three principal components represent the top three dimensions of variability across cortical BOLD time series. The first principal component represents the most dominant/leading axis of variance across cortical BOLD time series. The first principal component explains over 20% ($R^2 = 20.4\%$) of the variance in BOLD time series, greater than three times the variance explained by the second ($R^2 = 6.8\%$) or third principal component ($R^2 = 6.1\%$). As can be observed from **Figure 1A**, each FC topography exhibits strong similarities with one or more of the first three principal components ($r > 0.6$). In other words, all FC topographies in our survey resembled one of the top three dimensions of variability in cortical BOLD time series.

The three principal components can be differentiated most clearly with reference to three cortical brain networks: the default mode network (DMN), the frontoparietal or 'executive control' network (FPN) and the sensorimotor and medial/lateral visual cortices (SMLV) (**Figure 1C**). Note that the spatial extent of these three networks changes between principal components, and the reference to these networks is for descriptive simplicity. The first principal component is distinguished by its globally positive topography, with positive values (or negative, due to the sign ambiguity of PCA) in the SMLV and less positive values in the DMN. The second principal component is distinguished by negative values in the DMN and positive values in the FPN. The third principal component is distinguished by positive values in the SMLV and DMN, and negative values in the FPN. Detailed description of the output of dimension reduction analyses, as well as seed-based regression and CAP analyses are provided in **Supplementary Results A** and **Supplementary Results B**, respectively.

2. Three Dominant Spatiotemporal Patterns in Intrinsic BOLD Fluctuations

All FC topographies in our survey were found to resemble the first three principal components of cortical BOLD time series. FC topographies are produced from measures of zero-lag synchrony between BOLD time series (e.g. Pearson's correlation coefficient). Thus, FC topographies are unable to represent correlations between cortical BOLD time series at any time-lag. Time-lag relationships may reflect cortical propagation patterns or rapidly-changing sequences of cortical activity patterns. We refer to these time-lag structures as 'spatiotemporal patterns.' Examples of such spatiotemporal patterns include short time-scale lag projections (Mitra et al., 2014, 2015) and the quasi-periodic pattern (QPP) (Majeed et al., 2011). In the following section, we demonstrate that the three dominant FC topographies discovered from our survey of zero-lag FC analyses correspond to 'static' or zero-lag descriptions of three spatiotemporal patterns. We utilized a simple modification of PCA for detection of time-lag relationships between cortical BOLD time series. Specifically, we apply PCA to complex BOLD signals obtained by the Hilbert transform of the original BOLD signals (see 'Methods and Materials'). We refer to this analysis as complex PCA (cPCA).

We applied cPCA to the same resting state fMRI dataset. The scree plot criterion again motivated the choice of three complex principal components (**Supplementary Figure C**). The relative explained variance between the first three complex principal components from cPCA was similar to that observed between the first three principal components from PCA: component 1 ($R^2 = 21.4\%$), component 2 ($R^2 = 6.8\%$) and component 3 ($R^2 = 5.7\%$). Associated with each complex principal component is a time-lag delay map, reflecting the time-delay (in seconds) between cortical vertices (see 'Methods and Materials' for construction of the time-lag delay maps). To examine the temporal progression of each complex principal component, we sampled the reconstructed BOLD time courses from each complex principal component at multiple, equally-spaced phases of its cycle ($n=30$; see 'Methods and Materials'). Movies of the reconstructed BOLD time courses are displayed in **Movie 1**.


The time-lag delay maps and reconstructed time courses of each complex principal component are displayed in **Figure 2**. The first complex principal component describes a spatiotemporal pattern that begins with negative BOLD amplitudes in the SMLV complex (**Figure 2A**). This is followed by a propagation of negative BOLD amplitudes towards cortical regions overlapping primarily with the FPN, but also with the DMN and primary visual cortex. This is followed by a mirrored propagation of positive BOLD amplitudes with the same dynamics. Given this pattern of propagation, we refer to the first complex principal component as the 'SMLV-to-FPN' spatiotemporal pattern. Because the explained variance of the first complex principal component is three times greater than the subsequent complex principal components (**Supplementary Figure C**), we also refer to the SMLV-to-FPN as the 'dominant spatiotemporal pattern' in intrinsic BOLD fluctuations. The second principal component describes a spatiotemporal pattern that begins with positive BOLD amplitudes in the DMN and primary visual cortex, and negative BOLD amplitudes in the FPN. This is followed by the onset

of positive BOLD amplitudes in the SMLV that quickly propagates towards the FPN, with a simultaneous propagation of negative BOLD amplitudes from the FPN to the DMN. We refer to the second complex principal component as the ‘FPN-to-DMN’ spatiotemporal pattern. The third principal component describes a spatiotemporal pattern that begins with positive BOLD amplitudes in regions of the FPN and negative BOLD amplitudes in regions of the DMN and SMLV. This is followed by a fast propagation of positive BOLD amplitudes from the FPN to the SMLV, with a simultaneous propagation of negative amplitudes from the DMN to the FPN. Note, activation of the SMLV complex occurs slightly before the regions of the DMN. We refer to the third complex principal component as the ‘FPN-to-SMLV’ spatiotemporal pattern.


Examination of the reconstructed time courses reveals that the pattern of spatial weights in the three dominant FC topographies (i.e. principal component maps. **Figure 1B**) resembles the pattern of BOLD activity at individual time points of the three spatiotemporal patterns. The pattern of weights of the first principal component (PC1) occurs within the first spatiotemporal pattern ($r = 0.998$, $t = 11.9s$), PC2 occurs within the second spatiotemporal pattern ($r = 0.986$, $t = 12.6s$), and PC3 occurs within the third spatiotemporal pattern ($r = 0.972$, $t = 3.7s$). Further, the time courses of the three spatiotemporal patterns components closely tracks the time courses of the first three principal component time courses, respectively: PC1 ($r = 0.98$), PC2 ($r = 0.95$), and PC3 ($r = -0.83$, at a temporal lag of ~ 3 TRs). This finding suggests that the three dominant FC topographies are ‘static’ or ‘stationary’ representations of three temporally-extended, dynamic patterns of BOLD activity.

A) First Complex Principal Component - SMLV-to-FPN


Time-lag Delay Map



Time-lag Delay Values




Reconstructed Time Points




B) Second Complex Principal Component - DMN-to-FPN


Time-lag Delay Map



Time-lag Delay Values




Reconstructed Time Points




C) Third Complex Principal Component - FPN-to-SMLV

Time-lag Delay Map



Time-lag Delay Values



Reconstructed Time Points




Figure 2. Form and Properties of Three Fundamental Spatiotemporal Patterns. Time-lag delay maps and reconstructed time points of the first three complex principal components. Time-lag delay maps represent the temporal ordering (in seconds) of cortical vertex BOLD time series within the spatiotemporal pattern. Time-lag delay maps describe a repeating or cyclical pattern expressed in radians (0 to 2π) around a unit circle, where a phase value of 0 corresponds to the beginning of the spatiotemporal pattern, and 2π corresponds to the end of the

spatiotemporal pattern. For clarity, radians are converted to temporal units (seconds) (see ‘Methods and Materials’). The values in the time-lag delay map correspond to the temporal delay (in seconds) between two cortical vertices, such that smaller values occur *before* larger values. Values are mapped to a cyclical color map to emphasize the cyclical temporal progression of each spatiotemporal pattern. To illustrate the temporal progression of the spatiotemporal patterns, six reconstructed time points are displayed for each pattern. **A)** The time-lag delay map (top) and reconstructed time points (bottom) of the first spatiotemporal pattern - ‘SMLV-to-FPN’. **B)** The time-lag delay map and reconstructed time points of the second spatiotemporal pattern - ‘DMN-to-FPN’. **C)** The time-lag delay map and reconstructed time points of the third spatiotemporal pattern - ‘SMLV-to-FPN’.

3. Steady States and Propagation Events in Spatiotemporal Patterns

To visualize the temporal dynamics of the three spatiotemporal patterns, we projected the reconstructed time points (see above) into the 3-dimensional space formed by the first three principal components, corresponding to the three dominant FC topographies (**Figure 1B**). This projection allowed for a simple visualization of the temporal progression of BOLD activity within each spatiotemporal pattern. The structure of the resulting projection is interpreted as follows: reconstructed time points with a spatial pattern of BOLD activity resembling the spatial weights of one of the three principal components have higher scores on the axis of that principal component. The benefit of this representation of spatiotemporal patterns is two-fold: 1) examination of the movement of consecutive time points provides information regarding the ‘speed’ of change in BOLD activity between time points (e.g., steady states vs. rapid propagation or transition events), and 2) time points between two spatiotemporal patterns that are close together in this space indicate common patterns of BOLD activity between those spatiotemporal patterns (see below). The reconstructed time points of each spatiotemporal pattern is also displayed in **Movie 1**.

The temporal cycle of each spatiotemporal pattern forms an oval in the three-dimensional principal component space (**Figure 3B**), corresponding to a full cycle of the spatiotemporal pattern. For all three spatiotemporal patterns, most consecutive time points cluster closely together, indicating a ‘steady state’ of BOLD activity with relative stability of BOLD activity over that period. The steady states of the SMLV-to-FPN, FPN-to-DMN and FPN-to-SMLV vary strongest along the first, second and third principal component axes, respectively. This is apparent from the location of steady state time points (i.e. consecutive time points clustered closely together in space) in the 2-dimensional plots formed by two principal component axes in the three-dimensional space. For example, the steady state time points of the SMLV-to-FPN (time points displayed in blue) exhibit the highest positive and negative scores on the first principal component axis with smaller scores on the first and third principal component axes. This simply reflects the fact that the spatial pattern of BOLD activity during the steady states of the SMLV-to-FPN is strongly correlated with the pattern of spatial weights of the first principal component. These steady state periods are interrupted by large movement between consecutive time points that correspond to rapid propagation of BOLD activity towards another steady state.

All three spatiotemporal patterns spend most of their cycle in a period of steady synchronous activity that is interrupted by rapid propagation events.

In the case of the SMLV-to-FPN, the speed of propagation is relatively slower compared with the more abrupt propagation of the FPN-to-DMN and FPN-to-SMLV. This becomes apparent by tracing consecutive propagation time points (i.e. time points with long distances from their previous time point) between the steady states of each spatiotemporal pattern. The propagation events of the FPN-to-DMN and FPN-to-SMLV travel the same distance in approximately three time points as the SMLV-to-FPN in approximately five time points. For each spatiotemporal pattern, a full-cycle contains two mirrored steady-states and two mirrored propagation events. Mirrored steady-states and propagation events are the same spatial pattern of BOLD activity with a sign-flip - i.e. flipped positive and negative values. Another notable observation in this representation is that propagation time points of the SMLV-to-FPN and FPN-to-DMN vary most strongly along the third principal component axis. This indicates that the pattern of BOLD activity during propagation events of the SMLV-to-FPN and FPN-to-DMN resemble the pattern of spatial weights of the third principal component. The opposite is true of the FPN-to-SMLV. In this spatiotemporal pattern, steady-state time points vary strongest along the third principal component axis, and the propagation time points vary strongest along the second principal component axis.




Figure 3. Spatiotemporal Patterns Consist of Steady States and Propagation Events That Repeat Across Patterns. (PC = Principal Component) Illustration of the progression of BOLD activity over time in each spatiotemporal pattern. **A)** The spatial weights for the first three principal components from PCA (**Figure 1B**). For visualization of temporal dynamics, the reconstructed time points ($N=30$) from each spatiotemporal pattern were projected onto the 3-dimensional embedding space formed by the first three principal components. **B)** Two-dimensional slices of each spatiotemporal pattern in the 3-dimensional principal component space - PC1-PC2, PC1-PC3, and PC2-PC3 spaces. The time points of the SMLV-to-FPN, FPN-to-DMN, and FPN-to-SMLV are displayed

as blue, green and red points, respectively. Consecutive time points of each spatiotemporal pattern are linked by lines. The time points of each spatiotemporal pattern are colored from light to darker to visualize the progression of time ($N=30$). The score of each time point on a given principal component is proportional to the Pearson correlation coefficient between the BOLD activity at that time point with the spatial weights of the principal component. Examination of the movement of time points within the 3-dimensional space provides information regarding the temporal dynamics of the spatiotemporal pattern. **C)** The same two-dimensional slices of each spatiotemporal pattern in the 3-dimensional principal component space colored according to their cluster assignment by a k-means clustering algorithm. K-means clustering was used to identify recurring spatial patterns of BOLD activity across time points of the three spatiotemporal patterns. Six clusters were estimated. **D)** The cluster assignments (color) by time (*y-axis*) of each spatiotemporal pattern (*x-axis*). Note, that the same cluster assignment can occur across more than one spatiotemporal pattern. **E)** The cluster centroids from the k-means clustering algorithm, corresponding to the average spatial pattern of BOLD activity for the time points that belong to that cluster. Note, the cluster centroids of the first two clusters are mean-centered versions of the original unimodal (all-positive or all-negative) steady-state of the SMLV-to-FPN, as z-score normalization of the time-points across vertices was performed beforehand.

4. Recurring Spatial Topographies in Spatiotemporal Patterns

There is overlap in cortex-wide BOLD activity at certain time points between the three spatiotemporal patterns. For example, the pattern of BOLD activity at 8.6 seconds into the SMLV-to-FPN corresponds closely ($r = 0.909$) to the pattern of activity at 12.2s into the FPN-to-SMLV. This suggests that the same spatial topography of BOLD activity may appear across more than one spatiotemporal pattern. To examine repeating spatial topographies across the three spatiotemporal patterns, we applied a clustering algorithm to the reconstructed time points from all three spatiotemporal patterns. We concatenated the reconstructed time points from each spatiotemporal pattern ($n=90$, 30 time points per spatiotemporal pattern) and clustered the time points from the concatenated matrix using a k-means clustering algorithm. To avoid scaling differences in the distance calculations between time points, the BOLD values within each time point were z-score normalized. We chose a six-cluster solution from the k-means algorithm, as this was found to capture the three mirrored pairs of steady states from each spatiotemporal pattern (**Figure 3E**). Examination of the cluster assignments of each time point across spatiotemporal patterns (**Figure 3C and D**) yields several important insights. First, the six clusters correspond to the three pairs of mirrored or sign-flipped steady-states of the three spatiotemporal patterns. The first two clusters correspond to the steady states of the SMLV-to-FPN. Clusters three and four correspond to the steady states of the FPN-to-DMN, and clusters five and six correspond to the steady states of the FPN-to-SMLV. Second, the pattern of BOLD activity in the steady-states of one spatiotemporal pattern occurs within propagation events of the other two spatiotemporal patterns. For example, the pattern of BOLD activity in the steady-states of the FPN-to-SMLV (cluster three and four) occur within propagation events of the SMLV-to-FPN and the FPN-to-DMN. Further, the pattern of BOLD activity in the steady-states of the FPN-to-DMN (cluster five and six) occur within propagation events of the FPN-to-SMLV. In other words, the same pattern of BOLD activity occurs as a steady state or a propagation event depending on the spatiotemporal pattern. Third, the FPN-to-DMN and FPN-to-SMLV

spatiotemporal patterns are mirror images of each other. The pattern of BOLD activity in the FPN-to-DMN steady-states occurs as propagation events in the SMLV-to-DMN, and vice versa. In fact, the FPN-to-DMN can be converted to the FPN-to-SMLV by swapping the steady-states (clusters five and six) and the propagation events (clusters three and four), and vice versa.

Movie 1. Visualization of Spatiotemporal Patterns. Temporal reconstruction of all three spatiotemporal patterns displayed as movies in the following order - SMLV-to-FPN, FPN-to-DMN, FPN-to-SMLV. The time points are equally-spaced samples ($N=30$) of the spatiotemporal patterns. The seconds since the beginning of the spatiotemporal pattern are displayed in the top left. In the bottom of the panel, the time points of the spatiotemporal pattern are displayed in three-dimensional principal component space (**Figure 3**). Two-dimensional slices of the three principal component space are displayed as the three 2-dimensional plots. The progression of time points in the principal component space is illustrated by a cyclical color map (light to dark to light). The movement of the spatiotemporal pattern through this space is illustrated by a moving red dot from time point-to-time point in synchronization with the temporal reconstruction in the movie.

5. Relationships With Previously Observed Phenomena in Intrinsic BOLD Fluctuations

A further aim of this study was to understand the relationship between these three spatiotemporal patterns and previously observed phenomena in intrinsic BOLD signals. First, we consider spatiotemporal patterns discovered by previous approaches. Lag projections (Mitra et al., 2014) and the QPP (Majeed et al., 2011) correspond to time-lagged phenomena at shorter (~ 2 s) and longer (~ 20 s) time scales, respectively. Lag projections are computed as the column average of the pairwise time-lag matrix. The time-lag between a pair of BOLD time courses is the difference in time at which the correlation between the BOLD time courses is maximal. The column average of the pairwise time-lag matrix, or lag projection, provides the average ‘ordering’ in time of cortical BOLD time courses. We hypothesized that the average time-delay, represented by the lag projection, would match the dominant spatiotemporal pattern, the SMLV-to-FPN. To test this hypothesis, we computed the lag projection of all cortical BOLD time courses from the same 50 subject sample of resting-state scans. We found that the spatial correlation between the time-lag delay map of the SMLV-to-FPN and the lag projection map is strong ($r = 0.81$), and both exhibit the same direction of propagation (**Figure 4**). Thus, the time-lag dynamics described by lag projections map closely onto that described by the SMLV-to-FPN. However, there is a discrepancy between the estimated duration between the two approaches - the estimated duration of the SMLV-to-FPN from the cPCA is ~ 22 s, and the full duration of the lag projection is ~ 2.5 s. This may suggest that the time-lag structure of the SMLV-to-FPN exists at shorter time scales. Note, the lag projection we computed only partially resembles the average lag projection in Mitra et al. (2014) - the differences are due to preprocessing differences, which we discuss in **Supplementary Results F**.




Figure 4. **Similar Time-lag Dynamics between SMLV-to-FPN and Lag Projection.** Comparison between the time-lag delay map of the SMLV-to-FPN (left) and the average lag projection (right). As in **Figure 2**, the SMLV-to-FPN time-lag delay map represents the phase delay (in radians) of cortical BOLD time series. The lag projection map represents the average time-lag delay (in seconds) between each vertex of the cortex. The spatial correlation between the SMLV-to-FPN time-lag delay map and lag projection is $r = 0.81$, indicating a strong similarity in time-lag dynamics.

While lag projections describe short-time scale propagation structures, the QPP is a much longer temporally-extended pattern ($>20s$). Visual comparison of the spatiotemporal pattern of the QPP (Majeed et al., 2011) with the SMLV-to-FPN revealed a superficial similarity. Thus, we hypothesized that both the QPP and the SMLV-to-FPN describe the same spatiotemporal dynamics. We derived the QPP from a repeated-template-averaging procedure with a commonly used window size ($\sim 21s$; $30TRs$; Yousefi et al., 2018) on the resting-state fMRI data. We then computed the correlation between the time course of the QPP and the time course of the SMLV-to-FPN. We found that the time courses of the SMLV-to-FPN and QPP were strongly correlated ($r = 0.72$) at a time-lag of $7TRs$ ($\sim 5s$). The similarity in spatiotemporal dynamics between the QPP and SMLV-to-FPN can also be illustrated visually. We visualized the spatiotemporal template of the QPP from the repeated template-averaging procedure, and compared it to the time point reconstruction of the SMLV-to-FPN described above (**Movie 2**). As can be seen from the visualization, the temporal dynamics of the SMLV-to-FPN overlap significantly with the dynamics of the QPP.

During the steady states of SMLV-to-FPN, the distribution of weights is roughly unipolar, meaning it is either all positive or all negative (**Figure 2A**). This suggests that the SMLV-to-FPN may track the global mean time course of cortical vertices, otherwise known as the 'global signal'. We found that this is indeed the case - the time course of the SMLV-to-FPN and the global mean time course are statistically indistinguishable ($r = 0.97$). This would also suggest that the temporal dynamics surrounding the time points before and after the peak of the global signal correspond to the dynamics of the SMLV-to-FPN. We constructed a dynamic visualization of the global signal through a peak-averaging procedure. Specifically, all BOLD time courses within a fixed window ($15TRs$ on each side) were averaged around randomly-sampled peaks ($N=200$, > 1 standard deviation above the mean) of the global mean time course.

Visually comparing the spatiotemporal visualization of the global signal to the SMLV-to-FPN, we found that the temporal dynamics surrounding peaks of the global signal precisely match those of the SMLV-to-FPN (**Movie 2**).

Movie 2. Dynamic Visualization of the Quasiperiodic Pattern, SMLV-to-FPN, and Global Signal. The 30 time points (TR=0.72s) of the QPP, SMLV-to-FPN, and peak-average global signal displayed as a movie (in that order). The time index of each sequence is displayed in the top left. The time points of the SMLV-to-FPN are equally-spaced phase samples (N=30) of the time point reconstruction (see above). The time points of the QPP are derived from the spatiotemporal template computed from the repeated-template-averaging procedure. The global signal visualization concatenates the left and right windows ($w=15$ TRs) of the global signal peak-average. The time points of the global signal visualization begin at TR=-15, corresponding to 15 TRs pre-peak, and proceed to TR=15, corresponding to 15TRs post-peak.

The temporal dynamics of the FPN-to-DMN largely represents an anti-correlated pattern between the FPN and DMN - i.e. when regions of the DMN exhibit negative BOLD activity, the regions of the FPN exhibit positive BOLD activity (and vice versa). This resembles the “task-positive” (i.e. FPN) vs. “task-negative” (i.e. DMN) anti-correlation pattern originally observed by Fox et al. (2005) and Fransson (2005). We reproduced these results by correlating each cortical vertices’ BOLD time course with a seed time course from the left and right precuneus, a key node of the DMN. Note that the same results were observed with a seed placed in the left and right inferior parietal cortex. As expected, an anti-correlated pattern was observed between the FPN and DMN (**Figure 4**). We compared the precuneus-seed correlation map to the time points of the FPN-to-DMN using spatial correlation. We found that the pattern of correlations in the precuneus-seed map precisely corresponds to the pattern of BOLD activity in the beginning phase of the FPN-to-DMN ($r = 0.92$, $t = 1.8s$). Thus, this would seem to suggest that the task-positive vs. task-negative pattern arises from the anti-correlated dynamics between the FPN (task-positive) and DMN (task-negative) represented by the FPN-to-DMN spatiotemporal pattern.

A similar anti-correlation pattern to the task-positive/task-negative pattern has been observed in the FC gradient literature (Margulies et al., 2016; Vos de Wael et al., 2020), known as the ‘principal’ or ‘primary’ FC gradient (PG). In our zero-lag FC topography survey (**Figure 1**), we computed the PG as the first component derived from the Laplacian Eigenmaps (LE) algorithm, consistent with Vos de Wael et al. (2020). As opposed to the task-positive/task-negative pattern, the PG exhibits an anti-correlated pattern of spatial weights between the SMLV complex and the DMN (**Figure 4**). Further, the PG has been referred to as the principal direction of variance in cortical functional connectivity (Margulies et al., 2016). However, the results from both PCA (**Figure 1**) and cPCA (**Figure 2**) identify the SMLV-to-FPN as the principal direction of variance in cortical functional connectivity, which does not exhibit the anti-correlated pattern between the SMLV complex and the DMN observed in the PG. In fact, none of the three


spatiotemporal patterns exhibit an anti-correlated dynamic between the SMLV complex and the DMN.

With no clear correspondence between the PG and the three spatiotemporal patterns, we sought to identify the factors that might explain the uniqueness of the PG. We discovered that the spatial topography of the PG is due to the confluence of two factors: 1) global signal regression and/or time point normalization (i.e., z-scoring or de-meaning without unit-variance scaling), and 2) thresholding of FC matrices. First, as has been previously observed by Liu et al. (2017), regression of the global mean time course, and de-meaning of cortex-wide BOLD values within a time point (i.e. time-point centering) have very similar effects on cortical time series. Implicit in the computation of LE for functional connectivity gradients, as well as other manifold-learning techniques, is a time-point centering step (Ham et al., 2004 see **Supplementary Discussion E**). This is relevant because the global mean time course precisely tracks the time course of the SMLV-to-FPN ($r = 0.96$). This would suggest that removal of the global mean time course through global signal regression and/or time-point centering effectively removes the SMLV-to-FPN from BOLD time courses. What is left over is the second most dimension of variance in FC, the FPN-to-DMN. In fact, this would explain the appearance of the task-positive vs. task-negative pattern after global signal regression in seed-based correlation analysis (Fox et al., 2005). We tested this possibility by comparing the output of PCA and cPCA with and without a time-point centering preprocessing step (**Figure 4B**). Consistent with our hypothesis, PCA of time-point centered BOLD time courses produces a pattern of spatial weights for the first principal component that resembles the second principal component from PCA of non-time-point centered BOLD time courses (PC2: $r = 0.94$). Further, the first complex principal component of cPCA on time-point centered BOLD time courses exhibits a time-delay map that resembles the second complex principal component time-delay map on non-time-point centered data (cPC2: $\theta(\square) = 0.49$ vs cPC1: $\theta(\square) = 0.10$). Note, the correlation between the time-lag maps was computed using a circular correlation coefficient due to the circular properties of the spatiotemporal patterns (e.g. 0 and 2π are identical angles). Thus, at least one effect of time-point centering and/or global signal regression of BOLD time courses is the removal of the first principal component and/or the SMLV-to-FPN from BOLD time courses.


It is the dual effect of time-point centering and FC matrix thresholding that resolves the discrepancy between the DMN-to-FPN and the PG observed in our study. The FC matrix represents Pearson's correlation of BOLD time courses between all pairs of cortical vertices (i.e. correlation matrix). It is standard practice in computation of the PG that a threshold is performed row or column-wise on the FC matrix (e.g. 90th percentile of correlation values within that row) before the creation of an affinity matrix to input to the manifold learning algorithm (Margulies et al., 2016; Vos de Wael et al., 2020). This preprocessing step is intended to remove noisy or artificial correlation values from the FC matrix. In our survey of zero-lag FC topographies (**Figure 1**), we applied a 90th percentile threshold across rows of the FC matrix prior to LE. We

found that this preprocessing step obscures the relationship between the PG and DMN-to-FPN. In fact, if no thresholding of the FC matrix is performed, the first eigenmap produced from LE precisely resembles the FPN-to-DMN contrast observed in the second principal component of non-time-point centered BOLD time courses ($r = 0.83$; **Figure 4A** and **4C**). As one raises the percentile threshold applied to the FC matrix, the spatial weights of vertices within the FPN, DMN and SMLV complex become more uniform, and the spatial weights of the vertices within the FPN fall to zero (**Figure 4C**). At the higher end of percentile thresholds (e.g. 90th percentile) a contrast between the SMLV and DMN begins to appear that is almost equally similar to the unipolar contrast of the first principal component ($r = 0.83$) and the anti-correlation contrast of the second principal component ($r = 0.82$).

A) FPN-to-DMN, TP/TN Pattern & Primary Functional Connectivity Gradient



B) PCA & cPCA of Time-Point Centered Data



C) Threshold Effect on Primary Functional Connectivity Gradient

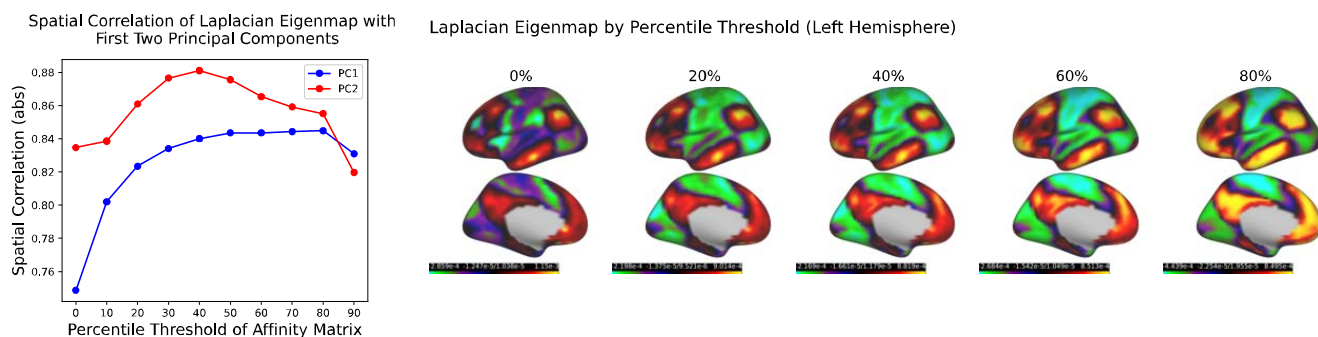


Figure 5. The Task-Positive/Task-Negative Pattern and the Primary Gradient Emerge from the FPN-to-DMN Spatiotemporal Pattern. **A)** From left to right, the FPN-to-DMN, task-positive/task-negative (TP/TN) pattern, and the PG represented by the spatial weights of the second principal component from PCA (sign flipped for consistency), seed-based correlation map (precuneus seed), and first Laplacian eigenmap with no thresholding of the affinity matrix, respectively. Note, the color mapping of each brain map is non-symmetric to emphasize the smaller

negative values within the FPN. As can be observed visually, similar spatial patterns are produced from all three analyses - FPN-to-DMN:TP/TN ($r=0.96$) and FPN-to-DMN:PG ($r=0.83$). **B**) From left to right, the first principal component from non-time-centered BOLD time courses, the first principal component of time-centered BOLD time courses, and the first complex principal component time-lag delay map from time-centered BOLD time courses. As can be observed visually, time-point centering of BOLD time courses replaces the original unipolar first principal component (left) with a bipolar (anti-correlated) principal component (middle). In the same manner, the first complex principal component of cPCA of time-centered BOLD time courses (right) exhibits a time-lag map resembling the time-lag map of the second complex principal component from cPCA of non-time-centered BOLD time courses (**Figure 2**). **C**) The effect of functional connectivity (FC) matrix percentile thresholding on the resulting spatial weights of the PG, computed as the first eigenmap of the Laplacian Eigenmap algorithm (only the left hemisphere presented for space). At zero to low-thresholding of the FC matrix, the first Laplacian Eigenmap resembles the second principal component of non-time-centered BOLD time courses. As the threshold is raised, the spatial weights of vertices within the FPN, DMN and SMLV complex become more uniform, and the spatial weights of the vertices within the FPN fall to zero. At higher thresholds this results in an Eigenmap that resembles the first principal component of non-time-centered BOLD time courses.

6. Network-Based Representations of Functional Connectivity

FC topographies are low-dimensional representations of zero-lag synchronous relationships among BOLD time courses. In the network or graph-based approach to FC analysis, the unit of analysis is pairwise relationships between BOLD time courses. Rather than reducing pairwise relationships to lower-dimensional representations, the network-based approach analyzes the structure of these relationships. We sought to identify the degree to which the structure of pairwise zero-lag synchronous relationships between BOLD time courses arises from the shared dynamics of the three distinct spatiotemporal patterns. A network representation of FC was constructed by computing the correlations between all pairs of cortical BOLD time courses to create a correlation or FC matrix (**Figure 5**). We compared this FC matrix to a FC matrix that was reconstructed from the three spatiotemporal patterns. Reconstructed cortical BOLD time courses were created from the spatiotemporal patterns by projecting the time courses of each pattern back into the cortical vertex space. A ‘reconstructed’ FC matrix was computed from these reconstructed time courses in the same manner as the original BOLD time courses. We estimated the similarity between the two FC matrices by computing the correlation coefficient between the lower triangles of each matrix. Despite a larger mean correlation value in the reconstructed FC matrix, we found that the patterns of correlations between the FC matrices were highly similar ($r=0.77$).

We also sought to determine whether the community structure of cortical BOLD time courses can arise from the shared dynamics of the three spatiotemporal patterns. We first thresholded the FC matrices by setting those correlation values below the top 80% of correlation values to zero. We estimated network communities from the original FC matrix using the Louvain modularity-maximization algorithm with a resolution parameter value of 1. To assess the degree of community structure in the original FC matrix, we computed the modularity value

of the partition of vertices into communities from the Louvain algorithm. The modularity value varies from -1 to 1 and represents the ratio of the summed intra-community correlation coefficients to that expected at random, such that higher values indicate a ‘higher quality’ partition. The modularity value of the original FC matrix was $Q = 0.34$. We then examined whether the same community structure of the original FC matrix was present in the FC matrix reconstructed from the spatiotemporal patterns. We assigned the vertices of the reconstructed FC matrix to the community assignments derived from the original FC matrix and re-calculated the modularity value. We found that the modularity value of the community assignments applied to the reconstructed FC matrix was almost as strong as the original FC matrix ($Q = 0.29$). In other words, the community structure of the original FC matrix is present in the FC matrix constructed from the shared dynamics of the three spatiotemporal patterns.

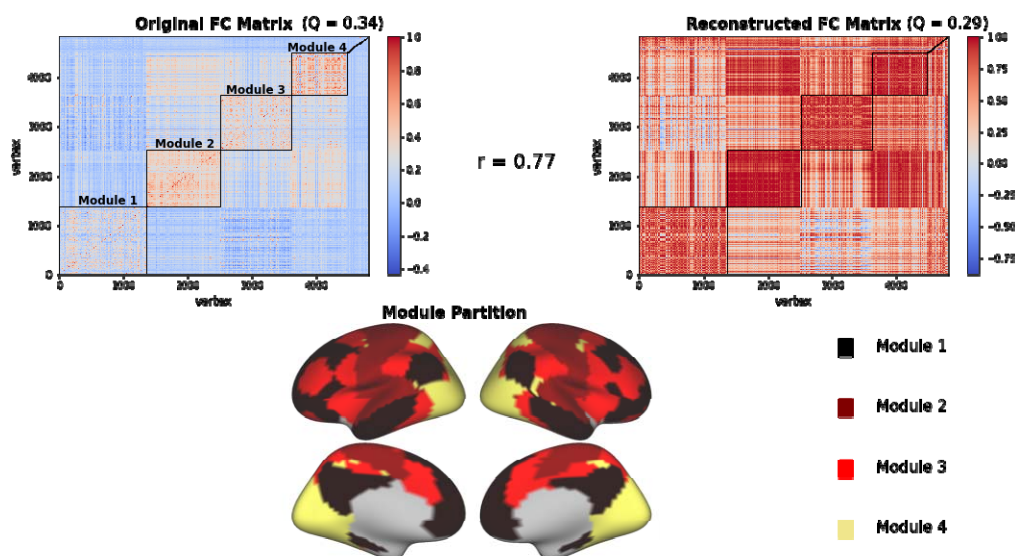


Figure 6. The Network Structure of Functional Connectivity Is Explained by the Three Fundamental Spatiotemporal Patterns. Comparison of the correlation matrix of cortical BOLD time courses (left) with the correlation matrix of reconstructed cortical BOLD time courses (right) derived from the three spatiotemporal patterns, and the module assignments of each vertex (bottom). The rows and columns of the original and reconstructed correlation matrix are sorted and outlined (in black) according to the modular structure of the original correlation matrix estimated from the Louvain modularity algorithm. The algorithm identified four primary modules in the DMN, sensory-motor cortices, FPN, and the posterior parietal/visual cortex. Modules with < 100 vertices were considered ‘junk’ modules and are placed in the upper right hand corner of each sorted correlation matrix. ‘Junk’ modules were not represented in the module assignments (bottom). Beside the title in parentheses of each correlation matrix is the modularity value. The modularity value represents the degree of community structure for the modular partition of that correlation matrix (i.e. degree of intra vs. inter module correlation weights). The modular structure estimated from the original correlation matrix ($Q = 0.34$) retains a high degree of modularity when used to partition the reconstructed correlation matrix ($Q = 0.29$), indicating the three spatiotemporal patterns explain a significant amount of the modular structure in the original correlation matrix. Further, despite a higher mean value of correlations in the reconstructed correlation matrix, the pattern of correlations between the two correlation matrices is highly similar ($r = 0.77$).

Discussion

Over the past decade, intrinsic functional brain organization has been characterized by a myriad of analytic methods (Bijsterbosch et al., 2020). This study attempted to synthesize this complex landscape into a set of fundamental patterns and organizational principles. We found three canonical spatiotemporal patterns of BOLD activity that are consistently observed across analytic methods, what we refer to as the ‘SMLV-to-FPN’, ‘FPN-to-DMN’ and ‘FPN-to-SMLV’ spatiotemporal patterns. These spatiotemporal patterns are best represented by time-lag analyses capable of characterizing temporally-extended patterns of BOLD activity. However, we also found that many zero-lag FC analyses are capable of representing parts of this spatiotemporal pattern, albeit in the form of ‘static’ or ‘stationary’ snapshots. Further, we found that a significant proportion of cortex-wide FC network structure is explained by the shared dynamics of these three spatiotemporal patterns.

The three spatiotemporal patterns account for a wide-variety of previous findings in resting-state fMRI. Previous research has demonstrated that travelling waves or propagatory patterns are ubiquitous features of intrinsic cortical BOLD activity (Gu et al., 2020; Majeed et al., 2011; Mitra et al., 2015). We found that the ‘SMLV-to-FPN’, representing the dominant axis of variance in cortical BOLD time courses, captures the same pattern of propagatory activity in lag projections originally discovered by Mitra et al. (2014). Further, the ‘SMLV-to-FPN’ precisely matches the spatiotemporal dynamics of the QPP originally discovered by (Majeed et al., 2011) (though recent findings suggest there may be more than one QPP, Yousefi and Keilholz, 2021). Consistent with findings that the time course of the QPP is closely correlated with the global BOLD signal (Yousefi et al., 2018), we find that the time course of the SMLV-to-FPN is statistically indistinguishable from that of the global BOLD signal. Further, the spatiotemporal pattern of BOLD activity around peaks of the global signal precisely matches the spatiotemporal pattern of the SMLV-to-FPN.

Another ubiquitous feature of intrinsic cortical BOLD activity is anti-correlated FC between the DMN and FPN, sometimes referred to as the ‘task-positive’ and ‘task-negative’ pattern (Fox et al., 2005; Fransson, 2005). Some have argued that the anti-correlated pattern of FC between the DMN and FPN is artificially introduced by regression of the global BOLD time course (Murphy et al., 2009). Others have argued that features of the anti-correlated pattern are independent of global signal regression (Fox et al., 2009). Our findings suggest that the global signal and DMN-FPN correspond to descriptions of two distinct spatiotemporal patterns, the ‘SMLV-to-FPN’ and the ‘DMN-to-FPN’, respectively. We found that global signal regression effectively eliminates the dominant axis of variance in cortical BOLD signals, corresponding to the ‘SMLV-to-FPN’. This leaves the second most dominant axis of variance in cortical BOLD signals, the ‘DMN-to-FPN’ as the primary contributor to functional connectivity. Thus, global signal regression does not produce the anti-correlated pattern between the DMN and FPN, but merely eliminates the variance in BOLD signals associated with the SMLV-to-FPN. The anti-correlated pattern of FC between the DMN and FPN that emerges after this procedure is not

artificial but corresponds to the anti-correlated dynamics of the DMN-to-FPN. The independence of this anti-correlated pattern of FC from the global signal (or SMLV-to-FPN) is further confirmed by the observation that both patterns are represented as distinct FC topographies across a wide variety of analytic methods.

Despite their independence, the FPN-to-DMN shares at least two features in common with the SMLV-to-FPN. First, the all-positive contrast between SMLV (high amplitude) and DMN (low amplitude) in the beginning phase of the SMLV-to-FPN has a strong degree of spatial similarity with the anti-correlated contrast in the later phase of the FPN-to-DMN ($r = 0.7$). The discrepancy between the two patterns is the lack of high amplitude values in the SMLV in the beginning phase of the FPN-to-DMN. However, the SMLV does appear at a later phase of the FPN-to-DMN, during the propagation of the SMLV to the FPN. In the SMLV-to-FPN, SMLV activation is simultaneous with activation of the FPN. In the FPN-to-DMN, SMLV activation occurs slightly after that of the FPN. This suggests that the primary difference between the two spatiotemporal patterns is a differential onset time of BOLD activity within the SMLV. Second, the direction of propagation is consistent between the two spatiotemporal patterns. Both patterns exhibit propagation from the SMLV towards the FPN. However, as noted above, the propagation of BOLD activity from the SMLV to FPN occurs later in the FPN-to-DMN than the SMLV-to-FPN. The similarity between the two spatiotemporal patterns suggests that they may share common physiological or neuronal mechanisms.

The temporal dynamics of the three spatiotemporal patterns can be divided into steady-states and propagation events. Steady-states are periods of relative stability of BOLD activity across consecutive time points of the spatiotemporal pattern. These steady-states are interrupted by propagation events exhibiting rapid changes in BOLD activity between consecutive time points. The three spatiotemporal patterns can be distinguished by their ratio of time points classified into propagation events versus steady-states. The FPN-to-DMN and FPN-to-SMLV spend most of their time in steady-states with very few time points exhibiting propagation. The SMLV-to-FPN exhibits a smooth, spatially continuous propagation of BOLD activity and contains a higher ratio of propagation time points versus steady-state time points than the FPN-to-DMN and FPN-to-SMLV.

Analysis of the temporal evolution of the three spatiotemporal patterns revealed a surprising pattern: the same spatial topography of BOLD activity can appear across more than one spatiotemporal pattern. For example, the spatial topography of BOLD activity that occurs during the propagation events of the SMLV-to-FPN and FPN-to-DMN appears in the steady-state time points of the FPN-to-SMLV. Comparison of the temporal evolution between the FPN-to-DMN and SMLV-to-FPN revealed another striking observation - the dynamics of the FPN-to-DMN and SMLV-to-FPN appear to be mirror images of one another. Specifically, the steady-state time points of the FPN-to-DMN resemble the propagation time points of the FPN-to-SMLV,

and vice versa. However, these spatiotemporal patterns exhibit opposite directions of propagation: propagation of BOLD activity from the FPN to DMN and SMLV to FPN in the FPN-to-DMN, and propagation of BOLD activity from the DMN to FPN and FPN to SMLV in the FPN-to-SMLV.

Our findings suggest that intrinsic functional brain representations are markedly consistent across analytic methods. This fact is even more surprising considering the wide array of mathematical and statistical assumptions of the analytic methods surveyed in this study. This observation does not imply that all analytic methods we surveyed produce the same *insights*. Each method affords a unique perspective on the spatial and temporal properties of the three spatiotemporal patterns discovered in this survey. Further, the level of functional brain organization explored in this study is an important qualification of our findings. This study has shown consistency at the level of widely-distributed cortical BOLD activity patterns among analytic approaches. Most of the analysis approaches we surveyed can reveal finer-grained spatial insights at higher component or cluster numbers. Thus, we do not expect the same consistency in analytic approaches at finer-grained levels of analysis (e.g. an ICA solution of 50 components vs. a PCA solution of 50 components). Despite this qualification, the consistency of widely-distributed representations of cortical functional brain organization, in the form of three common spatiotemporal patterns, attests to the dominance of these structures at higher spatial levels of analysis. Another limitation of this study was the exclusion of subcortical areas from our analyses. This study was primarily interested in cortical functional brain organization. However, previous research suggests that the large-scale spatiotemporal patterns observed in this study have significant subcortical contributions as well (Yousefi and Keilholz, 2021; Yousefi et al., 2018).

The appearance of these three large-scale spatiotemporal patterns across a wide-variety of analytic methods leads to the question of what these patterns may represent in terms of neuronal or hemodynamic processes. The similar time-scales and temporal dynamics between these three spatiotemporal patterns suggest they may emerge from similar mechanisms. Momentary fluctuations in arousal and/or vigilance are known to be related to the global BOLD signal (Liu et al., 2015, 2018; Schölvinck et al., 2010), BOLD propagation dynamics (Gu et al., 2020), and the anti-correlated dynamic between the FPN and DMN (Kucyi et al., 2020). All three of these features of intrinsic BOLD activity were found to be related to one or more of three spatiotemporal patterns. This may suggest that the three spatiotemporal patterns emerge from neurometabolic processes associated with momentary fluctuations in arousal. Future research may be directed towards a more complete understanding of the common or distinct neuronal or physiological mechanisms that give rise to these spatiotemporal patterns.

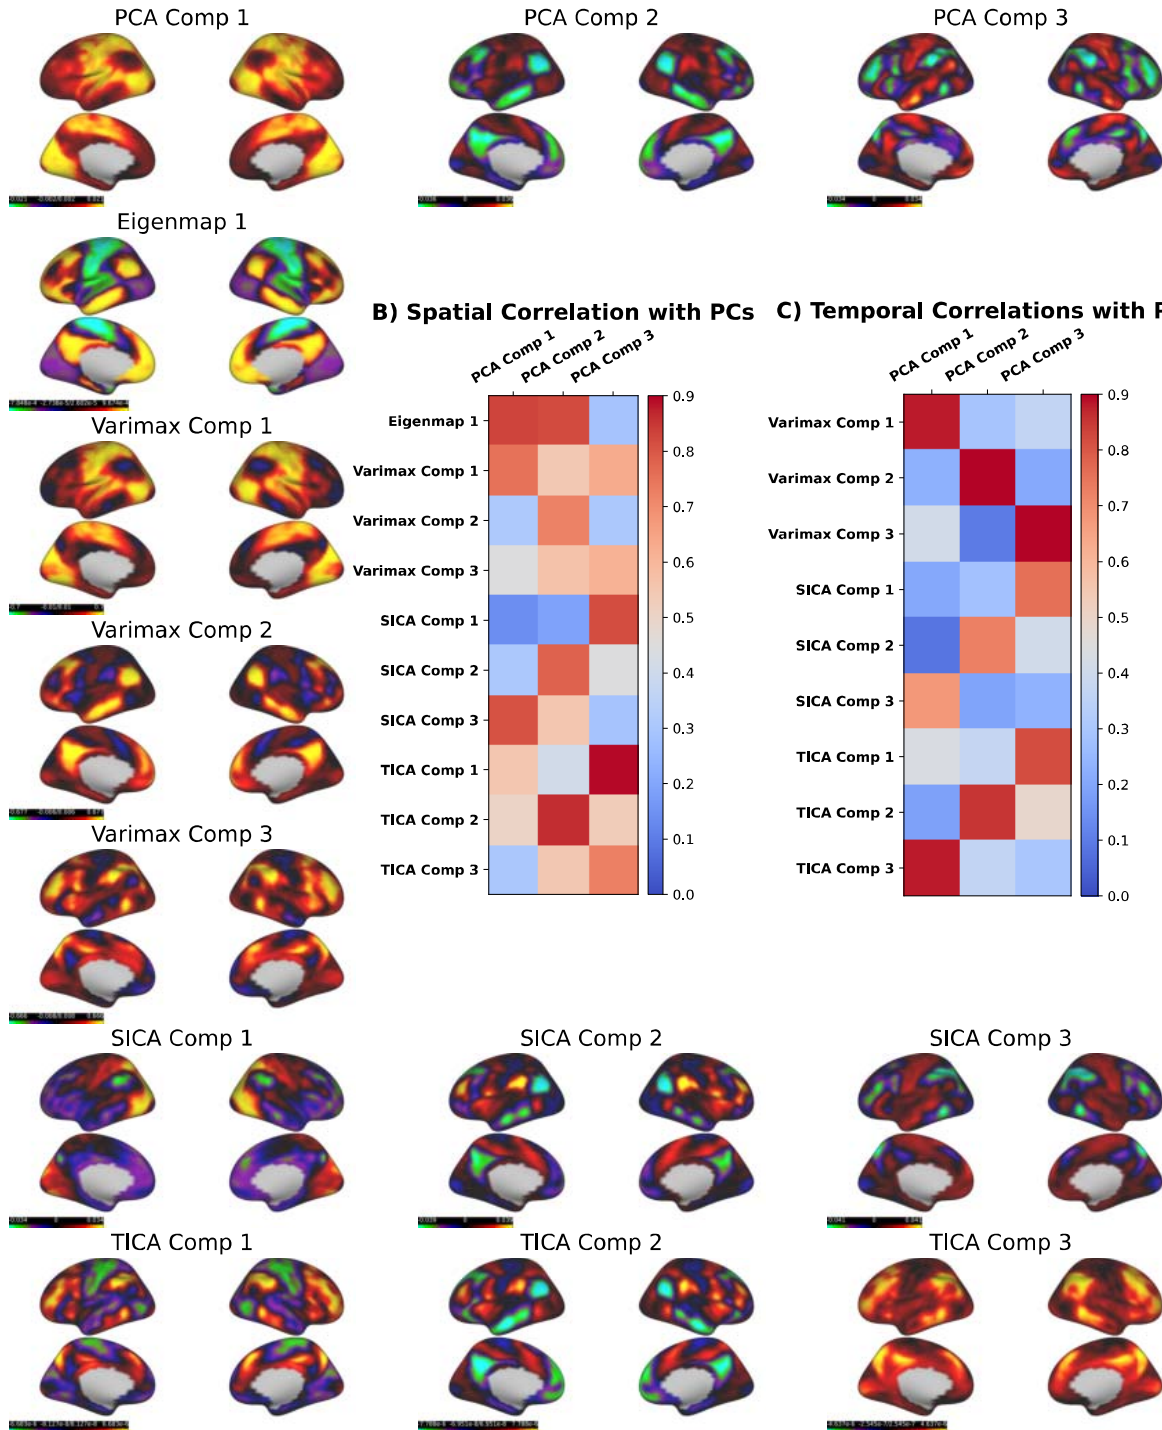
Supplementary Materials

Supplementary Results A - Latent Dimension Reduction Functional Connectivity Topographies

PCA represents only one of the many popular FC analyses applied to resting-state BOLD time courses. Spatial and temporal independent component analysis (ICA) are also popular analysis approaches in the resting-state fMRI literature. In relation to principal components derived from PCA, spatial and temporal ICA amounts to a rotation (i.e. unmixing matrix) of the whitened temporal or spatial principal component axes such that the statistical independence between the axes is maximized, respectively (Aires et al., 2002; Hannachi et al., 2009). In other words, ICA rotates the original PCA solution to maximize a different criterion: statistical independence in the temporal or spatial domain. Statistical independence is a stronger criterion than the orthogonal or uncorrelatedness assumption of PCA. Independence implies uncorrelatedness, but the converse is not always true. Thus, for the original three component PCA solution, we could estimate a three component temporal ICA solution by rotating the principal component axes (**Figure 3**) in the temporal domain towards statistical independence. If we wanted to estimate a three component spatial ICA solution, we could rotate the first three whitened principal component axes derived from PCA of the time point-by-time point covariance/correlation matrix. In this conception, ICA is one of a larger family of principal component rotation methods that also includes rotations towards so-called 'simple structure'. Simple structure rotations rotate the principal component loadings so that the parsimony of the loadings are maximized (each vertex loads strongly on only one component). One popular simple structure rotation is the so-called Varimax rotation - an orthogonal rotation of the principal component loadings that maximizes simple structure.

All of these rotation methods drop one or more of the restrictive constraints of the principal component solution: 1) orthogonality of the spatial eigenvectors, 2) uncorrelatedness of component time series, or 3) the criterion of maximum variance. The question is whether these rotation methods of the principal component axes or loadings matrix return spatial/temporal components that capture distinct spatiotemporal dynamics. To provide a more quantitative assessment of similarities, we display the spatial weights of the three-component rotated solutions, along with their spatial and temporal correlations with the first three principal components. In addition, we display the first three principal components and the first component of the Laplacian Eigenmaps algorithm for comparison.

A) Low-Dimensional FC Topographies



Supplementary Figure A. **FC Topographies of Latent Dimension Reduction Analyses.** (SICA=Spatial ICA; TICA = Temporal ICA). The spatial weights of components from PCA (N=3), Laplacian Eigenmaps (N=1), varimax rotation of principal components (N=3), spatial ICA (N=3) and temporal ICA (N=3). The temporal and spatial correlations (absolute value) between the components of dimension-reduction analyses and the first three principal components are shown in the middle of the plot. Note, due to the nature of the Laplacian Eigenmap algorithm as a non-linear manifold learning algorithm, time courses cannot be extracted for their components. As illustrated in the

spatial and temporal correlations table, the dimension-reduction analyses are largely consistent in their spatial topographies and temporal dynamics with the first three principal components.

As shown in **Supplementary Figure A**, low-dimensional solutions of each dimension-reduction algorithm are generally similar to the first three principal components. In all cases, the rotation methods (SICA, TICA and Varimax) return components more or less similar in spatial and temporal dynamics to the first three principal components. In other words, despite the differing mathematical assumptions and objective criteria of these dimension-reduction methods, the results produced from each method for low-dimensional solutions are roughly consistent.

Supplementary Results B - Seed-Based Regression and Co-Activation Pattern Analysis

Latent dimension-reduction methods decompose the pair-wise synchronous relationships between all cortical BOLD time courses. In contrast, seed-based approaches examine the synchronous relationships between cortical BOLD time courses and the time course of a single brain region - i.e. a seed time-course. Thus, seed-based approaches require an a priori choice of seed region. There are various methods for determining the location of seed regions. In our analysis, we chose seed regions within the three central networks of the three dominant spatiotemporal patterns - SMLV, FPN and DMN. We chose seeds in the somatosensory cortex (SMLV), precuneus (DMN), and supramarginal gyrus (FPN) (**Supplementary Figure B2**). The spatial outline of the SMLV, DMN and FPN for guiding the selection of seed regions were determined through a k-means clustering analysis of cortical vertices based on the similarity in their BOLD time courses. We found that a three-cluster k-means clustering solution precisely delineated the spatial outline of the three networks. This spatial outline was used to ensure the seeds were placed within their appropriate location of each network. In addition, we also tested the robustness of our results for different seed locations in the three networks - medial insula (SMLV), inferior parietal cortex (DMN) and dorsolateral prefrontal cortex (FPN) - and found that the results were identical. In **Supplementary Figure B1**, we present the results of seed-based regression analysis with and without global signal regression, and CAP analysis with and without time-point normalization, for the somatosensory, precuneus and supramarginal seeds.

Seed-based regression maps represent the correlation between the BOLD time courses of all cortical vertices and the time course of the seed region. Consistent with previous findings (Fox et al., 2009), the seed-based regression maps without global signal regression exhibit mostly positive correlations across the cortex. However, the vertices of the SMLV, FPN and DMN stand out with stronger positive correlations relative to vertices outside these networks.

It is common practice in the resting-state fMRI literature to perform global signal regression before seed-based regression analysis (Fox et al., 2009). To ensure the correspondence of seed-

based regression maps with the three dominant spatial topographies was robust to global signal regression we examined the spatial correspondence between seed-based regression analysis with and without global signal regression. Consistent with previous findings (Murphy et al., 2009), global signal regression of the cortical BOLD time courses before seed-based regression analysis returns anti-correlated patterns of correlations. These anti-correlated patterns exhibit positive values in the SMLV, DMN and FPN, and negative values for vertices outside these networks. We found that the effect of global signal regression is primarily a centering operation of correlation values, with the pattern of correlations largely consistent between the original and global-signal regressed correlation patterns - somatosensory ($r = 0.96$), precuneus ($r = 0.89$), supramarginal gyrus ($r = 0.92$). Overall, regardless of global signal regression, seed-based regression analysis returns patterns consistent with the three dominant FC topographies we identified in our survey of zero-lag FC topographies (**Figure 1**).

The seed-based regression maps characterize zero-lag, synchronous correlations between the seed time-course and cortical BOLD time courses. However, this approach ignores the possibility of time-varying, zero-lag relationships at shorter time-scales. CAP analysis allows the identification of recurring whole-brain BOLD activity patterns at the level of a single time point. CAP analysis begins with the identification of a subset of time points where the seed time course exceeds a pre-specified threshold. We chose a threshold equal to the 85th percentile of the seed time course BOLD values, consistent with previous applications (Liu and Duyn, 2013b). After the selection of suprathreshold time points, cortex-wide BOLD activity at those time points are clustered into recurring activity patterns using a clustering algorithm. We estimated two clusters (activity patterns) per seed using a k-means clustering algorithm. A central preprocessing step is the choice to normalize (i.e. center and rescale) the BOLD activation values within each time point before application of the k-means clustering algorithm. This step is particularly relevant for BOLD time courses that undergo periods of global BOLD synchronization associated with the global mean time course. In **Supplementary Figure B1** (Panel B and C), we present the two cluster centroids for the non-normalized and normalized analyses of each seed, representing two recurring activation patterns during suprathreshold time points for that seed.

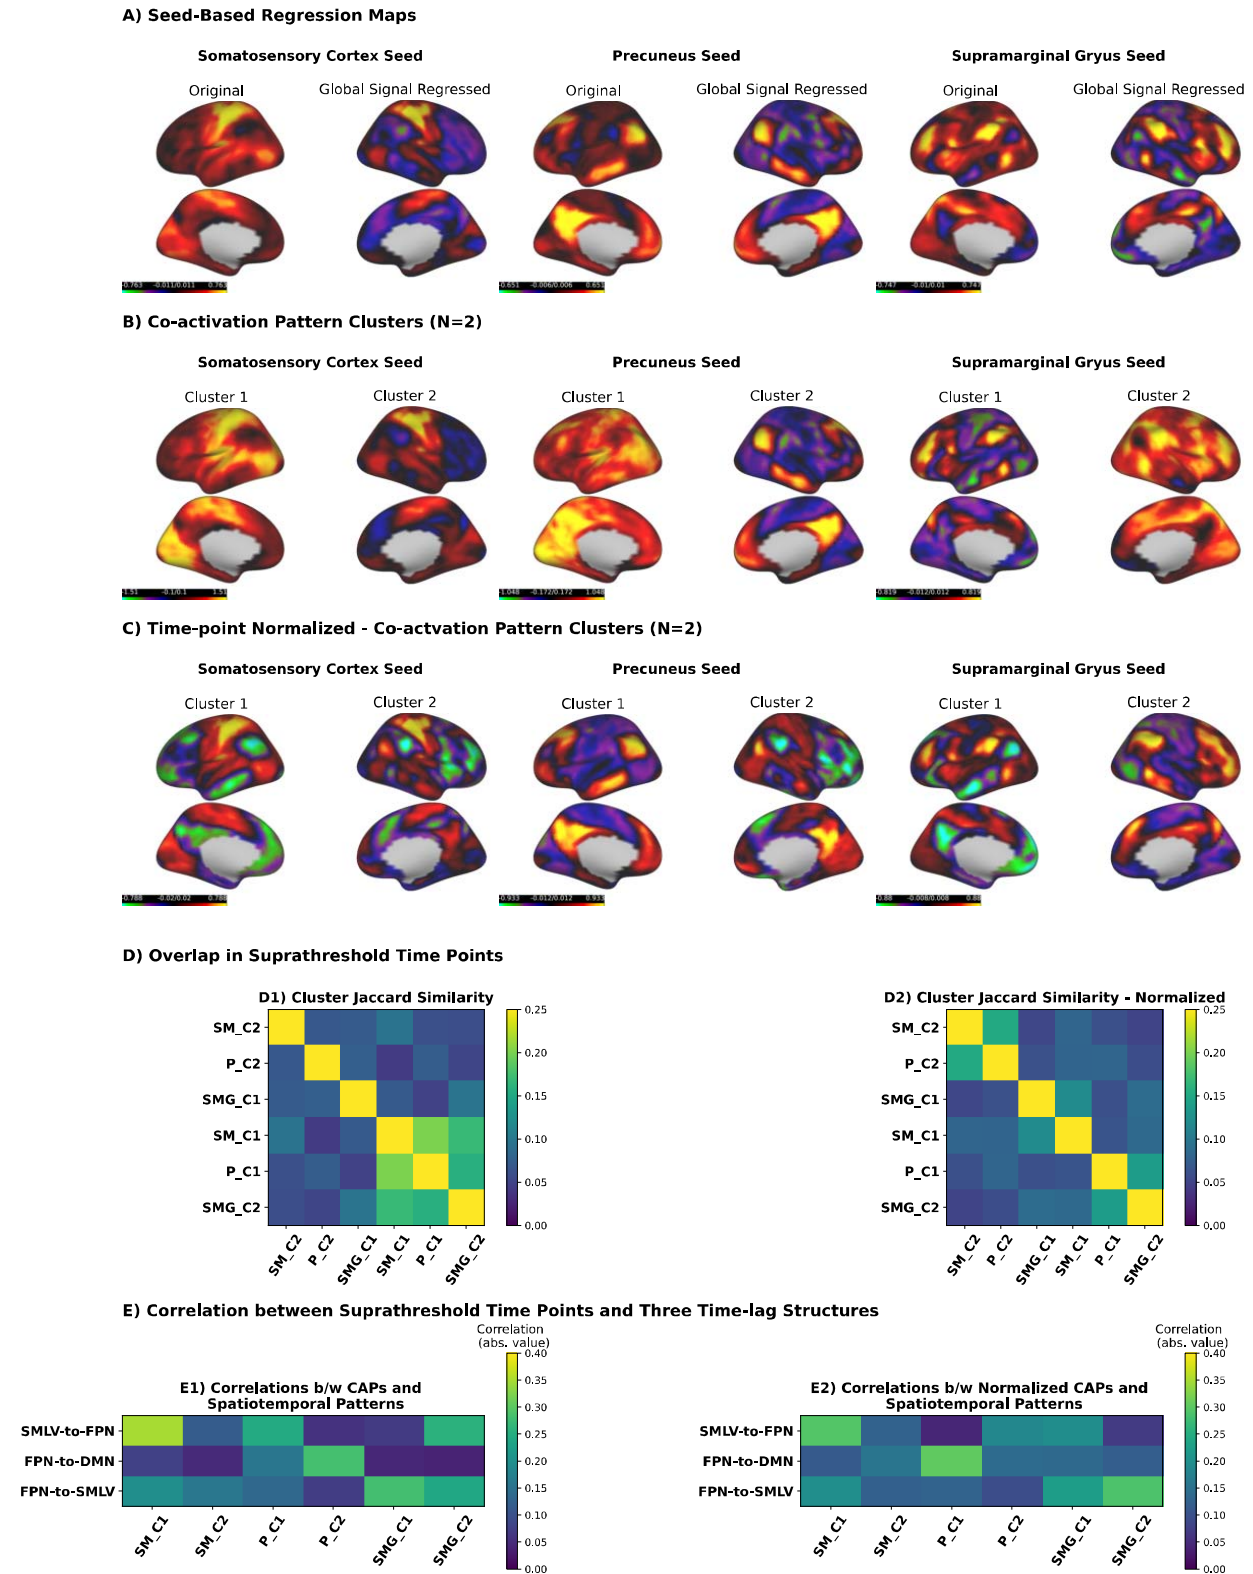
We first examined the CAP patterns of the non-normalized suprathreshold time points. Interestingly, the two cluster solution approximately separates a globally-positive pattern of BOLD activation from an anti-correlated pattern for each seed. This is consistent with previous findings that CAP analysis can separate global topographies associated with global synchronization effects from more focal, anti-correlated patterns (Liu and Duyn, 2013a). To examine the temporal overlap between the CAP patterns of different seeds, we computed binary time courses for each CAP by assigning a binary indicator (i.e., 1) to time points of a length $N=60,000$ time course, where a suprathreshold time point was classified as that CAP by the k-means clustering algorithm. The binary time course of each CAP reflects the onsets of that CAP pattern at various time points of the full duration of the group-concatenated time series. We then assessed the temporal overlap between pairs of CAP binary time courses by computing the

Jaccard similarity (Jaccard Index) between the pair of binary time courses. The Jaccard similarity between two CAP binary time courses varies from 0 to 1 and reflects the ratio of overlapping onset time points ($=1$) to the total number of time points ($N=60,000$). To account for potential time-lags between CAP binary time courses, we took the maximum Jaccard similarity between the CAP binary time courses at a max temporal lag of 30 time points. Examination of the temporal overlap between CAP binary time courses revealed that the onsets of globally-positive CAP patterns (somatosensory C1, precuneus C1 and supramarginal C2) tended to co-occur at much greater rate than the anti-correlated CAPs (global CAPs: $\chi^2_{(1)} = 0.173 >$ anti-correlated CAPs: $\chi^2_{(1)} = 0.073$) (**Supplementary Figure B1** - Panel D1). This is consistent with the global synchronization effect associated with the global mean time course.

As illustrated in our survey of FC topographies (**Figure 1**) the spatial topographies of the two CAP centroids map variously onto the three dominant FC topographies identified by other techniques. We also sought to determine whether the binary CAP time courses are correlated with the time courses of the three spatiotemporal patterns (**Figure 2**). To relate the CAP pattern from each seed to a spatiotemporal pattern, we computed the temporal correlation between the time course of the spatiotemporal pattern and the binary time course of each CAP. We took the absolute value of the correlation to ignore potential sign-differences between the CAP and the spatiotemporal patterns. To account for potential time-lags, we took the maximum absolute value correlation between the CAP binary time course and spatiotemporal pattern time course at a max temporal lag of 30 time points. We found that CAPs with globally-positive BOLD activation patterns (somatosensory C1, precuneus C1 and supramarginal C2) were most strongly correlated with the time course of the SMLV-to-FPN: somatosensory cortex (SM; $r = 0.49$), precuneus ($r = 0.35$), and supramarginal gyrus ($r = 0.37$) (**Supplementary Figure B1** - Panel E1). This is consistent with the observation that the global mean time course closely tracks the time course of the SLMV-to-FPN ($r = 0.96$). The temporal correlations between the anti-correlated CAPs and three spatiotemporal patterns are primarily driven by seed location. The anti-correlated precuneus CAP (C2) correlates most strongly with the FPN-to-DMN ($r = 0.41$), consistent with the anti-correlated pattern of activity between the FPN and DMN in this spatiotemporal pattern. The anti-correlated supramarginal gyrus CAP (C1) correlates most strongly with the FPN-to-SMLV ($r = 0.41$), consistent with the anti-correlated pattern between the FPN and SMLV in this spatiotemporal pattern. Interestingly, the correlation between the anti-correlated somatosensory CAP and the SMLV-to-FPN ($r = 0.17$) and FPN-to-SMLV ($r = 0.23$) is roughly similar. This is consistent with the observation in the main text (**Figure 4**) that the propagation phase of the SMLV-to-FPN briefly exhibits the same pattern of BOLD activity as observed in the steady-state of the FPN-to-SMLV.


Normalizing BOLD values within suprathreshold time points removes scaling and mean differences between suprathreshold time points. In other words, information regarding global mean differences between suprathreshold time points is removed. In the standard k-means algorithm with Euclidean distances, normalization as a preprocessing step means that the

difference between a suprathreshold time point and a cluster centroid is measured by the dissimilarity in their *relative pattern* of BOLD values. Note, the k-means algorithm with Pearson correlation distance (Liu and Duyn, 2013b) does this scaling implicitly in the computation of dissimilarity (the Pearson correlation is proportional to the dot product of the mean-centered and standardized, or z-scored, BOLD values of two time points). This preprocessing step is non-trivial for BOLD activation patterns, as globally-positive and/or globally-negative values across the majority of cortical vertices occur regularly across BOLD time courses. As shown in the **Supplementary Figure B1** (Panel C), normalization results in two anti-correlated CAPs per seed, as opposed to a globally-positive vs. anti-correlated CAP per seed in the non-normalized solution (Panel B). Normalization also affects the binary time courses of each CAP pattern (see description above). While the suprathreshold time points input to the K-means algorithm are identical between the non-normalized and normalized CAP solutions (i.e. both use the same time point indices that exceed the seed time course threshold value - 85th percentile), the assignment of a suprathreshold time point to one cluster over another may change between non-normalized and normalized solutions. Thus, the onset times between the two CAPs of a normalized and non-normalized solution will not generally be identical. Thus, as we might expect, the temporal overlap observed between globally-positive CAP time points in the non-normalized solution disappears in the normalized solution (**Supplementary Figure B1 - Panel E2**). In addition, the temporal correlation between the SMLV-to-FPN and CAPs across the three seeds is reduced, excluding the somatosensory CAP (C1), consistent with the two facts that the SMLV-to-FPN tracks the global signal, and normalization removes global differences.




Supplementary Figure B1. **Seed-Based Regression and CAP Analysis.** (SM = somatosensory cortex; P=Precuneus; SMG=Supramarginal Gyrus). FC topographies of seed-based regression maps and CAP centroids from somatosensory (SM), precuneus and supramarginal gyrus seeds. **A)** Seed-based regression maps with (left

hemisphere) and without global signal regression (right hemisphere) for SM, precuneus and supramarginal gyrus seeds. **B**) CAP cluster centroids (N=2) from k-means clustering of non-normalized (i.e. not z-scored) suprathreshold time points from SM, precuneus and supramarginal seeds. **C**) CAP cluster centroids (N=2) of the same suprathreshold time points with normalization (i.e. z-scored) before input to the k-means clustering algorithm. **D1**) Temporal overlap between binary time courses (see main text) of the two CAPs from each seed using the Jaccard similarity (Jaccard index). The Jaccard similarity between two CAP binary time courses varies from 0 to 1, and reflects the ratio of overlapping onset time points (=1) to the total number of time points (N=60,000). **D2**) Temporal overlap between CAP binary time courses from the normalized solutions of each seed analysis. **E**) Temporal correlation between the beginning phase time course of the three spatiotemporal patterns (SMLV-to-FPN, FPN-to-DMN and FPN-to-SMLV) and the CAP binary time courses for the non-normalized (**E1**) and normalized (**E**) solutions.



Supplementary Figure B2. **Seed Locations in SMLV, FPN and DMN.** The two seed locations for the SMLV, FPN and DMN. The spatial outline of the SMLV, FPN and DMN were identified by clustering vertices based on the similarity in their BOLD time courses using a k-means clustering algorithm (N=3). Seed locations were bilateral, with one seed in the left hemisphere and one seed in the right hemisphere. For the analyses in this study, we presented the results from seeds placed in the somatosensory cortex, precuneus and supramarginal gyrus. To test the robustness of our analyses to seed location, we also ran seed-based regression and CAP analyses with seeds placed medial insula (SMLV), inferior parietal cortex (DMN) and DLPFC (FPN). Because the results were found to be identical with the somatosensory cortex, precuneus and DLPFC, respectively, we do not present results for these seeds.

Supplementary Results C - Complex Principal Component Scree Plot




Supplementary Figure C. **Scree Plot from Complex Principal Component Analysis.** The eigenvalue by component number plot (i.e. scree plot) used to determine the number of components to extract. As with the original PCA solution (**Figure 1C**), there are clear elbows in the plot after one and three components, indicating a preferred solution of one or three principal components (three were chosen).

Supplementary Results D - Complex Principal Component Analysis Simulation

To demonstrate the ability of complex principal component analysis (cPCA) to extract spatiotemporal patterns, we applied cPCA to propagating fields in a simple simulation framework. Impulse time series (1 for activation, 0 otherwise) convolved with the canonical hemodynamic response are spatially arranged along a square grid. The time series are arranged such that the time series in the top part of the square grid peaks early, and peaks later and later (time steps of 0.1 secs per row) as one moves down the grid. This arrangement provides a simple illustration of a global propagation event, where activity in one location travels to all other locations in a spatially continuous fashion. Gaussian noise is added to every time point of each time series, and slight phase and amplitude jitter are applied to each time series within a row of the square grid drawn from a uniform distribution. The global propagation event in the simulation is re-run 1000 times and temporally concatenated. Time points are sampled at the sampling frequency of the HCP scans (0.72 TRs). It is important to note that this simulation is

not intended to be a biologically realistic simulation of the mechanisms that produce observed BOLD propagation patterns. The intent of this simulation is to illustrate the properties of cPCA applied to fields containing globally propagating phenomena. While this simulation is designed to have superficial similarities to spatial and temporal properties of BOLD propagation, it does not intend to simulate its underlying *data generating process*. Below we display the first 20 time points of the global BOLD propagation simulation (**Supplementary Figure D1**)




Supplementary Figure D1. **Global BOLD Propagation Simulation.** The first ~25 sec of an artificially constructed BOLD propagation simulation. Artificial ‘vertices’ are arranged along a 50-by-50 spatial grid. Vertex time series are created by convolving simple impulse time series with the canonical hemodynamic response function. The time-series of each vertex in the grid are time-lagged such that vertex time series in the upper part of the grid peak first, while those in the bottom part peak last. As observed from the first ~25 sec snapshots of the spatial grid, a positive BOLD propagation sweeps from the top of the grid towards the bottom in the first 10 seconds. Due to the undershoot following the peak in the canonical hemodynamic response function, this positive propagation event is followed by a mirrored slightly negative propagation event.

The first 20 time points of the simulation illustrate the simulated global propagation event: peak BOLD amplitudes are first observed in the top of the square grid, followed by a subsequent propagation of peak BOLD amplitudes down the grid. Following the globally positive BOLD propagation event, there is a mirrored negative BOLD propagation event due to the post-response undershoot of the canonical hemodynamic response function. As emphasized above, this is not intended as a biologically realistic simulation of the mechanism by which the QPP arises, but as a simple simulation of a globally propagating BOLD field.

To demonstrate the ability of cPCA to extract time-lagged information from a propagating field, we applied cPCA to the simulated global propagation patterns. As explained in the main text, cPCA is simply PCA applied to complex-valued time courses (i.e. analytic signals) derived from the Hilbert transform. The original time courses and their Hilbert transforms are complex vectors with real and imaginary components, corresponding to the non-zero-lagged time course ($t=0$) and the time course phase shifted by $t=\frac{\pi}{2}$ (i.e. 90 degree), respectively. The complex time course allows us to define the instantaneous amplitude envelope and phase of the signal at any point in time (via Euler’s identity). In the same manner that the complex time courses admit of a phase and amplitude representation, the complex correlation between two time courses is complex-valued with phase and amplitude components, corresponding to the magnitude of the correlation and phase shift between the time courses. Thus, in the same manner that PCA decomposes a real-numbered correlation matrix into orthogonal components, cPCA decomposes a complex-valued correlation matrix, reflecting correlation magnitudes and phase shifts, into complex orthogonal components. Just as with complex time courses and correlations, the complex principal components can be expressed in real and imaginary components, corresponding to the spatial weights of the principal component at zero-lag ($t=0$) and the spatial weights of the principal component at a $\pi/2$ phase shift, respectively. Examination of the spatial weights between the real and imaginary component gives a partial representation of the temporal progression of the time-delay structure captured by the principal component. The full temporal evolution of a spatiotemporal pattern can be best captured by the phase of the complex principal component weights, representing the phase-offsets or -delay values between time courses within that component.


As illustrated in **Supplementary Figure D2**, the first complex principal component from cPCA accurately recovers the spatiotemporal pattern of the global propagation simulation. A sketch of this spatiotemporal pattern can be derived by comparing the activity patterns of the real and imaginary parts of the cPC, corresponding to snapshots of the grid values at phases $t=0$ and $t=0+\frac{\pi}{2}$, respectively. Comparing the activity patterns between the two phases indicates the cPC captures the progression of activity down the grid over time. Consistent with the interpretation of the real and imaginary part of the cPC, the phase delay map accurately describes the spatiotemporal pattern as BOLD activity that travels at a steady rate down the grid. Overall, this simulation illustrates the interpretive advantages of cPCA applied to propagation or traveling-wave phenomena in BOLD time courses, and spatiotemporal dynamics in BOLD signals more generally.




Supplementary Figure D2. Real, Imaginary and Phase Delay Weights of First Complex Principal Component. The real part, imaginary part and phase delay weights of the first complex principal component (cPC) spatial weights derived from the propagation simulation. The real and imaginary parts (left) of the cPC correspond to snapshots of the grid values at phases $t=0$ and $t=0+\frac{\pi}{2}$. A sketch of the temporal dynamics of the propagating simulation can be derived from comparing the activity patterns of the real and imaginary parts of the cPC. Comparing the activity patterns between the two phases indicates the cPC captures the progression of activity down the grid over time. The phase delay representation of the cPC (right) represents the phase-delay between different points of the grid of the cPC. The phase delay representation is colored according to the phase delay of each point in the grid, and the progression of time can be read by moving ‘up’ the colormap, or in terms of a color wheel - moving counterclockwise. This provides a more complete description of the spatiotemporal pattern of the propagation simulation.

PCA is a well-known eigendecomposition approach in fMRI research. More importantly, it is related to a wide variety of analytic approaches popular in the fMRI literature. In fact, PCA is closely related to the recent resurgence of eigendecomposition methods in resting-state fMRI analysis that has gone under the name of 'functional connectivity gradients'. The seminal study of functional connectivity gradients by Margulies et al. (2016) introduced the concept of connectivity gradients to the study of resting-state fMRI. To derive functional connectivity gradients, manifold learning algorithms, such as diffusion map embedding and Laplacian Eigenmaps (LE), are applied to a vertex-by-vertex (or voxel-by-voxel) similarity/kernel matrix. A fixed number of dimensions or manifolds are returned from the manifold learning algorithm, representing orthogonal subspaces that preserve the local distances (in a Euclidean sense) between vertices/voxels. PCA has a very close relationship with these approaches. Both diffusion maps and LE involve an eigendecomposition of a modified similarity matrix, the graph Laplacian, analogous to the covariance/correlation matrix used by PCA. Further, these techniques can be viewed as special cases of kernel PCA - PCA extended to an arbitrary kernel/similarity matrix (Ham et al., 2004). We computed the first three eigenmaps from the vertex-by-vertex cosine similarity matrix (with no thresholding of the FC matrix, see main text), and compared them to the first three principal components from PCA, PCA of global-signal regressed data and PCA of time-point centered data (**Supplementary Figure E**). Note, LE analysis with a radial basis function kernel (non-linear kernel) was also tried and the results were very similar.


Principal Component Analysis




Principal Component Analysis - Global Signal Removed



Principal Component Analysis - Time-Point Centered



Laplacian Eigenmaps - Manifold Learning



Supplementary Figure E. **Principal Component and Laplacian Eigenmap Topographies.** Displayed are the FC topography spatial weights from PCA, PCA on global-signal regressed data, PCA on time-point centered data, and Laplacian Eigenmaps. Note, we observed that the eigenmaps were highly positively skewed. To make the negative values of the eigenmaps more visible the colormap is made non-symmetric. The first and second eigenmaps match the second and third principal component from PCA. The first principal component is missing from the LE, global-signal regressed, and time-point centered PCA solutions.


The first eigenmap corresponds to the primary functional connectivity gradient (PG) described by (Margulies et al., 2016). However, we note that the exact spatial pattern of the PG depends on the level of thresholding applied to the FC matrix (**Figure 5**). In this LE solution, no thresholding was applied. The first and second Laplacian Eigenmaps match the second and third

principal components from the PCA solution, respectively (with an arbitrary sign-flip). The difference between the PCA and Laplacian Eigenmap solution is that the first principal component seems to be missing from the Laplacian Eigenmap solution. However, the first three components from the PCA on global-signal regressed data and time-point centered do match the three Laplacian Eigenmaps (w/ an arbitrary sign flip). These similarities between the spatial maps produced by PCA and Laplacian Eigenmaps has been previously observed by Vos de Waal et al. (2020).

As we note in the main text, the first principal component corresponds to the beginning phase of the SMLV-to-FPN spatiotemporal pattern. An important question is why the first principal component, the variance upon which BOLD time courses vary the greatest, is not returned by LE and PCA applied to global-signal regressed and time-point centered data? In all three cases, the difference is due to the same mechanism: mean-centering along the time domain (i.e. mean centering vertex BOLD values within a time point). As discussed in the main text, the onset of the SMLV-to-FPN is precisely tracked by the first principal component and global signal. Consistent with previous findings (Liu et al., 2017), time-point centering and global signal regression have similar effects - reducing or eliminating the variance of the global signal time course. In global signal regression, this is done in a voxel-specific manner, according to the optimal linear fit between each vertex time course and the global signal. In the Laplacian Eigenmap solution, a time-point centering is not as explicit. Consider a matrix $\square_{\square\square}$, where i indexes observations and j indexes variables. Manifold learning algorithms are conventionally described as projection of the *observations* (rows) into a low-dimensional space that preserves the local and/or global distances between the observations. Implicit in the steps of many of these algorithms is a mean-centering in the feature space, or mean-centering of the *variables* (Ham et al., 2004). Note, in most manifold learning algorithms this mean centering is done directly on the kernel/similarity matrix rather than the original matrix X . This is analogous to the double centering operation of a distance matrix to produce the centered Gram matrix in classical multidimensional scaling. In other words, manifold learning of a time point-by-time point kernel matrix operates on a mean centering of the feature space, i.e. each vertex/voxel is mean centered. In the context of functional connectivity gradients, non-linear manifold learning algorithms are applied directly to the *vertex-by-vertex* affinity/kernel matrix. Thus, this implies an implicit mean centering of the *time points*, i.e. the vertex BOLD values within a time point are mean-centered. As illustrated in **Supplementary Figure E**, this has the same practical effect as global signal regression and time-point centering. Thus, LE, PCA of global-signal regressed data and PCA of time-point centered data return very similar spatial topographies (some with an arbitrary sign difference).

Supplementary Results F - Global Signal Regression and Lag Projections

The lag projection from our study partially resembles the group average lag projection observed in (2014), and our lag projection is nearly twice as long (~2.5 sec compared to ~1 sec). However, our data differs in one important respect: Mitra et al. (2014) performed global signal regression as a preprocessing step. In fact, when we regress out the global signal of our data, we see the following lag projection compared to the original:



Supplementary Figure F **Comparison of Lag Projections With and Without Global Signal Regression.** Lag projections with and without global signal regression as a preprocessing step. Values on each cortical map represent the average time-delay between each cortical vertex and all others. Time-delay values are colored from light green/blue (earlier in time) to bright yellow/green (later in time). The range between the earliest and latest time-delay values are significantly shorter for lag projections on global-signal regressed data.

The lag projection of the global signal regressed data resembles the spatial distribution time-lags observed in Mitra et al. (2014): BOLD activity beginning in superior medial prefrontal cortex, inferior precuneus, motor cortex, anterior cingulate cortex, and temporal gyrus and ending in the DMN and visual cortex. In addition, the length of the lag projection is now 1 sec (cut in time by half from non-globally regressed data), closely matching the duration found by Mitra et al. (2014). This is consistent with the observation by Mitra et al. (2014) that global signal regression reduces the range of observed latencies between BOLD time courses.

Methods and Materials

Resting-State fMRI Data

Our study utilized resting-state fMRI scans from the Human Connectome Project (HCP) S1200 release (Van Essen et al., 2013). Participants were non-related, healthy young adults (ages 22–37). Resting-state fMRI data was collected over two consecutive days for each subject and two sessions, each consisting of two 15 minute runs, amounting to four resting-state scans per subject. Within a session, the two runs were acquired with opposite phase encoding directions: L/R encoding and R/L encoding. We selected a single 15 min scan from a random sample of participants ($n=50$; 21 males) on the first day of scanning. We balanced the number of L/R and R/L phase encoding scans across our participants ($n=25$ for each encoding direction) to ensure results were not biased by acquisition from any given phase encoding direction. We chose a

single 15 min scan per participant to ensure that the phase encoding/decoding parameter and the imaging session (two resting-state scans per imaging session) did not differ within the same participant. A second independent random sample of participants ($n=50$, 22 males) was used as a validation sample. We selected surface-based CIFTI resting-state fMRI scans that had been previously preprocessed with the HCP's ICA-based artefact removal process (Smith et al., 2013) to minimize effects of spatially structured noise in our analysis. All brain-imaging data were acquired on a customized Siemens 3T Skyra at Washington University in St. Louis using a multi-band sequence. The structural images were 0.7 mm isotropic. The resting-state fMRI data were at 2 mm isotropic spatial resolution and with $TR=0.72$ s temporal resolution. Further details of the data collection and preprocessing pipelines of the HCP can be found elsewhere (Smith et al., 2013; Van Essen et al., 2013). Informed consent was obtained from all subjects. All methods were carried out in accordance with relevant guidelines and the University of Miami Institutional Review Board approved the study.

Resting-State fMRI Preprocessing

Resting-state fMRI scans were spatially smoothed with a 5mm FWHM kernel using the surface-based smoothing algorithm in Connectome Workbench Version 1.4.2. Resting-state fMRI signals from each vertex were then temporally filtered to the conventional low-frequency range of resting-state fMRI studies using a Butterworth bandpass zero-phase filter (0.01-0.1Hz). Due to 1) the computational complexity of our analytic pipeline, owing to the large number of analyses studied, and 2) our interest in global, spatially distributed patterns, resting-state fMRI scans were then resampled to the *fs4* average space from Freesurfer (Dale et al., 1999). This step down-sampled the total number of vertices in the left and right cortex to 4800 vertices. In group analyses, we z-scored (to zero mean and unit variance) the BOLD time series from all vertices before temporal concatenation of individual scans. All analyses were applied to group-level data formed by temporal concatenation of subject resting-state scans.

Survey of Functional Connectivity Topographies

Description of Zero-lag FC Analyses

Our study distinguished between two different descriptions of intrinsic functional brain organization - zero-lag synchrony between brain regions, and time-lag synchrony between brain regions. Zero-lag synchrony is defined as in-phase or anti-phase statistical dependence between intrinsic BOLD signals - e.g. the correlation between two intrinsic BOLD signals with no time-lag. Following the standard terminology of the functional magnetic resonance imaging (fMRI) literature, we refer to zero-lag synchrony among intrinsic BOLD fluctuations as 'functional connectivity' (FC) (Reid et al., 2019). FC between cortical brain regions organize into global, cortex-wide patterns, referred to as 'FC topographies'. All analyses were conducted so as to be consistent as possible with previous studies. For some of these analyses, results were compared with and without global signal regression. Global signal regression was performed by regression

of the global mean time series (averaged across all cortical vertices) on all cortical time series. Residual time series from each regression were then used for subsequent analysis. All analyses were conducted using custom Python scripts, and are publicly-available at https://github.com/tsb46/BOLD_WAVES. The following zero-lag FC analyses were conducted:

- Principal component analysis (PCA): consists of eigendecomposition of the empirical covariance matrix of the vertices' time series, or alternatively, singular value decomposition of the mean-centered group data matrix (time series along rows, vertices as columns). The first T principal components represent the top T dimensions of variance among cortical BOLD time courses. By construction, the first principal component is the latent direction of variation with the largest explained variance across all input variables, followed by the second most explanatory component, and so forth. The principal component spatial weights on each vertex were used to interpret the spatial patterns of each principal component. Principal component scores were obtained from the projection of the temporally-concatenated group time series onto the principal component space, and represent the time course of each principal component.
- Varimax rotation of principal components: consists of an orthogonal rotation of the principal component spatial weights, such that the simple structure of the spatial weights are maximized. Simple structure is defined such that each vertex loads most strongly one component, and weakly on all others. We used the implementation of varimax rotation in the *FactorAnalyzer* Python package (https://github.com/EducationalTestingService/factor_analyzer).
- Laplacian Eigenmaps (spectral embedding): is a nonlinear manifold learning algorithm popular in the FC gradient literature (Vos de Wael et al., 2020). The input to the Laplacian eigenmaps algorithm was the vertex-by-vertex cosine similarity matrix (Margulies et al., 2016), representing the similarity in BOLD time series between all cortical vertices. Of note, cosine similarity is equivalent to Pearson correlation in mean-centered and unit normalized time series (i.e. z-score normalization), as was the case with our data. Laplacian Eigenmaps performs an eigendecomposition of the transformed similarity matrix, known as the normalized Laplacian matrix. We also computed Laplacian Eigenmaps with a Gaussian radial basis function ($\gamma=1$), and the results were virtually identical to the cosine similarity metric. We used the spectral embedding algorithm implemented in the *scikit-learn* (V0.23) Python package, and details can be found at (<https://scikit-learn.org/stable/modules/generated/sklearn.manifold.SpectralEmbedding.html>).
- Spatial and temporal independent component analysis (ICA): estimates linearly mixed, statistically independent sources from a set of input variables. In the case of spatial ICA, principal component axes derived from PCA of the *time point-by-time point covariance matrix* are rotated to enforce statistical independence in the spatial domain. In the case of temporal ICA, principal component axes derived from PCA of the *vertex-by-vertex covariance* are rotated to enforce statistical independence in the temporal domain. As

with varimax rotation, we input a three principal component solution for both temporal and spatial ICA. We used the FastICA algorithm implemented in the *scikit-learn* (V0.23) Python package. Details can be found at (<https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.FastICA.html>).

- Seed-based correlation analysis: consists of correlations between a seed brain region time course and time courses of all cortical vertices. Seed-based correlation analysis was performed for three seed locations. There are various methods for determining the location of seed regions. In our analysis, we chose seed regions within the three most prominent networks in the three dominant spatiotemporal patterns - SMLV, FPN and DMN. We chose seeds in the somatosensory cortex (SMLV), precuneus (DMN), and supramarginal gyrus (FPN) (**Supplementary Figure B2**). The spatial outline of the SMLV, DMN and FPN for guiding the selection of seed regions were determined through a k-means clustering analysis of the temporally-concatenated group time series with cortical vertices as observations and BOLD values at each time points as input variables (i.e. features). We found that a three-cluster k-means clustering solution precisely delineated the spatial outline of the three networks. This spatial outline was used to ensure the seeds were placed within their appropriate location of each network. In addition, we also tested the robustness of our results for different seed locations in the three networks - medial insula (SMLV), inferior parietal cortex (DMN) and dorsolateral prefrontal cortex (FPN) - and found that the results were identical.
- Co-activation pattern (CAP) analysis: Three CAP analyses were performed for the same three seed regions used in the seed-based regression analysis. CAP analyses first identify time points with the highest activation for a seed time course. Consistent with previous studies (Liu and Duyn, 2013b), we chose the top 15% of time points from the seed time course. The BOLD values for all cortical vertices in the top 15% time points are then input to a k-means clustering algorithm to identify recurring CAPs of BOLD activity. We chose a two cluster solution for all CAP analyses. For each seed, the two cluster centroids from the k-means clustering analysis represent two CAPs associated with the seed time course.
- Hidden Markov modeling (HMM): is a probabilistic generative model used to infer the sequence and form of discrete hidden states, as well as their transition probabilities from an unobserved sequence of latent states. HMM construes the data-generating process based on multivariate Gaussian distributions conditioned on unknown latent ‘brain states’ that are assumed to generate the observed cortical BOLD time series. Each brain state represents a recurring pattern of BOLD co-activations/deactivations, somewhat similar to CAPs. To avoid overfitting and to reduce noise in the high-dimensional input data, we conducted a PCA of the cortical BOLD time series. The first 100 principal component projections of the time series served as input to the HMM algorithm. Associated with each brain state is a mean amplitude vector with a value for each principal component ($N = 100$), and a covariance matrix between the 100 principal component time courses. The

mean amplitude vector represents the pattern of BOLD activity amplitudes associated with that brain state. For interpretation, the mean amplitude vector is projected back into cortical vertex space for interpretation. A variety of potential ‘observation models’ are frequently used in HMM models. As cortical time series are measured on a continuous scale (as opposed to discrete measurements), the probability of a time point conditional on a hidden brain state (i.e. emission probabilities) is modeled as a mixture of Gaussian distributions. We used the HMM algorithm with Gaussian mixture emission probabilities implemented in the Python package *hmmlearn* (V0.2.5) (<https://github.com/hmmlearn/hmmlearn>).

Model Selection: Choice of Number of Dimensions in Dimension-Reduction Algorithms

The dimension-reduction algorithms used in this study, including PCA, PCA with varimax rotation, spatial and temporal ICA, and Laplacian Eigenmaps, as well as HMMs, require a choice of the number of latent dimensions/hidden states to estimate. For PCA with varimax rotation, spatial and temporal ICA, and HMM, this controls the degree of richness and/or fine-grained distinctions of the data description - i.e. how many separate unobserved hidden phenomena are assumed and quantitatively modeled to underlie each given data point or observation. We did not assume or try to derive a single ‘best’ number of latent dimensions to represent intrinsic functional brain organization (Bzdok et al., 2016; Schaefer et al., 2018). As we were interested in large-scale cortical patterns of FC, our survey focuses on low-dimensional latent solutions. As an initial estimate of the number of latent dimensions for all choices of dimension reduction algorithms, we examined the first T dominant axes of variation (i.e. principal components) of the correlation matrix formed between all pairs of cortical BOLD time series. Specifically, we examined the drop-off in explained variance (i.e. eigenvalues) associated with neighboring principal components, a procedure known as Cattell’s scree plot test (Cattell, 1966). According to this test, the number of components to extract is indicated by an ‘elbow’ in the plot, representing a ‘diminishing return’ in extracting more components. Clear elbows in the scree plot were observed after a principal component solution of one and three (**Figure 1A**). We chose the higher-dimensional solution of three components. Note, the elbow in explained variance after three components was independent of the functional resolution (i.e. vertex size) of the cortex - we found the same elbow after three components in a scree plot constructed from high-resolution functional scans (~60,000 vertices without downsampling to 4,800 vertices as described above in our preprocessing pipeline). Thus, three latent dimensions were estimated for all dimension-reduction algorithms, and three hidden states were estimated for the HMM.

Time-lagged Analyses of Resting-State fMRI Data

Quasiperiodic Pattern and Lag Projections

Time-lag analyses capture relationships between two or more time series at past or future lags of the time series. We refer to representations of time-lag relationships between cortical time

series as ‘spatiotemporal patterns’. There are two widely-used algorithms for the study of spatiotemporal patterns in BOLD signals: 1) interpolated cross-covariance functions (Mitra et al., 2014, 2015) for the detection of lag/latency projections ($\sim 0-2s$) and 2) a repeated-template-averaging algorithm of similar spatiotemporal segments (Majeed et al., 2011) for detection of the QPP ($\sim 20s$).

Lag projections represent the average time-lag between a brain region’s time course and all other brain regions. It provides an estimate of the average temporal ‘ordering’ of brain region time courses, such that a brain region with a greater average time-lag occurs after a brain region with a smaller average time-lag. For our study, we applied the lag projection algorithm to all cortical vertex time courses. The time-lag between a pair of cortical vertex time courses is defined as the peak of their lagged cross-covariance function. Lag projections are derived as the column average of the pairwise time-lag matrix between all cortical vertex time courses.

To estimate the QPP, the template-autoregressive matching algorithm of Majeed et al. (2011) was used. The algorithm operates in the following manner: start with a random window of BOLD TRs, compute a sliding window correlation of the window across the temporally concatenated group data at each time point, and then average this segment with similar segments of BOLD TRs (defined using a correlation threshold). This process is repeated iteratively until a level of convergence is reached. The result is a spatiotemporal averaged template of BOLD dynamics (that could be displayed in a movie, for example), along with the final sliding window correlation time series. The final sliding window time series is the same length as the original subject or group concatenated time series and provides a time index of the appearance of the QPP in BOLD data. Python code for this analysis was modified from the C-PAC toolbox (<https://fcp-indi.github.io/>). Consistent with previous studies (Majeed et al., 2011; Yousefi et al., 2018), the following parameters were chosen for the template matching algorithm: the window length was 30 TRs, the maximum correlation threshold for identifying similar segments was $r > 0.2$, and the algorithm was repeated 10 times. The template with the highest average sliding window correlation time series across the 10 runs was chosen as the final result.

Complex Principal Component Analysis

To tie the three dominant FC topographies to spatiotemporal processes in the cortex, we used a simple modification of PCA for detection of spatiotemporal patterns. Specifically, we apply PCA to complex BOLD signals obtained by the Hilbert transform of the original BOLD signals. We refer to this analysis as complex PCA (cPCA). This technique has been referred to as complex Hilbert empirical orthogonal functions in the Atmospheric and Climate sciences literature (Horel, 1984).

cPCA allows the representation of time-lag relationships between BOLD signals through the introduction of complex correlations between the Hilbert transformed BOLD signals. The original time courses and their Hilbert transforms are complex vectors with real and imaginary

components, corresponding to the non-zero-lagged time course ($t=0$) and the time course phase shifted by $t=\frac{\pi}{2}$ (i.e. 90 degree), respectively. The correlation between two complex signals is itself a complex number (composed of a real and imaginary part), and allows one to derive the phase offset (and magnitude) between the original time courses - i.e. the time-lag at which the correlation is maximum. In the same manner, cPCA produces complex spatial weights for each principal component that can give information regarding the time-lags between BOLD time courses. In the same manner that a complex signal is composed of real and imaginary signal components, the complex principal component's spatial weights are composed of real and imaginary components. and the phase delay of the spatial weights. The real part of the complex principal component corresponds to the 'beginning phase' of the complex principal component ($\phi=0$, where θ represents phase in radians), corresponding to the spatial map of BOLD activity at the beginning of the spatiotemporal pattern. The imaginary component corresponds to a 'later phase' of the complex principal component ($\phi=0+\frac{\pi}{2}$), corresponding to the spatial map of BOLD activity several time points (a quarter cycle) after the beginning of the complex principal component.

In the same manner that a complex signal can be represented in terms of amplitude and phase components (via Euler's transform), the real and imaginary components of the complex principal component can be represented in terms of amplitude and phase spatial weights. Of interest in this study is the phase spatial map that represents the time-lag between pairs of BOLD time courses - i.e. those cortical vertices with a low phase value activate earlier than cortical vertices with a high phase value. Importantly, the principal components from the cPCA retain the same interpretive relevance as the original PCA - the first N principal components represent the top N dimensions of variance in the Hilbert transformed BOLD signals. An illustration of cPCA applied to simulated globally propagated fields is provided in **Supplementary Results D**. cPCA was implemented with singular value decomposition of the groupwise temporally-concatenated complex-valued time series using the fast randomized SVD algorithm developed by *Facebook* (<https://github.com/facebookarchive/fbpca>).

For simplicity, the phase spatial maps of each complex principal component are displayed in seconds (**Figure 2**), as opposed to radians. However, the conversion of phase values (in radians) to time-units (seconds) requires an estimation of the time-scale of each complex principal component. The phase spatial maps of the complex principal components have no characteristic time scale other than that imposed by our band-pass filtering operation (0.01 - 0.1 Hz, i.e. 100 to 10s) in the preprocessing stage. To approximate a unique time scale within this frequency range for each component, we calculated the average duration for a full oscillation of each complex principal component using the temporal phase of the complex component time series. This was calculated by fitting a linear curve to the unwrapped temporal phases of the complex principal component time series. The slope of the curve was then used as an estimate of the average duration in radians of a TR (0.72s) or time-points. To estimate the average duration

in TRs of a full oscillation, we divided a full oscillation (2π radians) by the duration in radians of a TR. For example, for a TR duration of 0.5 radians, the duration of a full oscillation (2π radians) would be approximately 12.6 TRs. Using this procedure, we found that the average duration of the first three complex principal components are ~ 28 s (38.7 TRs), ~ 27 s (37.4 TRs) and ~ 28 s (39.1 TRs), respectively. Using this duration as an estimate of the characteristic time scale of each complex principal component, allows us to provide an estimate of the time-delay in seconds of the spatial phase map. For example, for the first complex principal component, a 360 degree (2π radians) phase difference between two cortical BOLD time series would correspond to a ~ 28 s time-lag between the time series. A smaller phase difference between two cortical BOLD time series, such as 1π radians, would correspond to ~ 14 s time-lag between the time series, and so forth.

To examine the temporal progression of each complex principal component, we sampled the reconstructed BOLD time courses from each complex principal component at multiple, equally-spaced phases of its cycle ($N=30$; **Figure 2**). For each complex principal component, the reconstruction procedure was as follows: 1) the complex principal component time series was projected back into the original vertex-by-time space to produce time courses of the complex principal component at each vertex, 2) the temporal phase of the complex principal component time course was segmented into equal-width phase bins ($N=30$) spanning a full oscillation of the spatiotemporal pattern (0 to 2π radians), and 3) the vertex values within each bin were averaged to produce a ‘snapshot’ of BOLD activity at each phase bin ($N=30$) of the spatiotemporal pattern. The end result is a spatiotemporal representation of each complex principal component in terms of time-varying BOLD activity at equally spaced phases of its cycle.

Acknowledgements

This work was supported by grants from the Canadian Institute for Advanced Research, a Gabelli Senior Scholar Award from the University of Miami, and R01MH107549 from the National Institute of Mental Health (NIMH) (to LQU), and an NIMH award (R03MH121668), and a NARSAD Young Investigator Award to (to JSN).

References

Aires, F., Rossow, W.B., and Chédin, A. (2002). Rotation of EOFs by the Independent Component Analysis: Toward a Solution of the Mixing Problem in the Decomposition of Geophysical Time Series. *J. Atmos. Sci.* 59, 111–123.

Andersen, A.H., Gash, D.M., and Avison, M.J. (1999). Principal component analysis of the dynamic response measured by fMRI: a generalized linear systems framework. *Magn Reson Imaging* 17, 795–815.

Beckmann, C.F., and Smith, S.M. (2004). Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans Med Imaging* 23, 137–152.

Bijsterbosch, J., Harrison, S.J., Jbabdi, S., Woolrich, M., Beckmann, C., Smith, S., and Duff, E.P. (2020). Challenges and future directions for representations of functional brain organization. *Nature Neuroscience* 1–12.

Biswal, B., Yetkin, F.Z., Haughton, V.M., and Hyde, J.S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* 34, 537–541.

Bzdok, D. (2017). *Classical Statistics and Statistical Learning in Imaging Neuroscience*. Front Neurosci 11.

Bzdok, D., Varoquaux, G., Grisel, O., Eickenberg, M., Poupon, C., and Thirion, B. (2016). Formal Models of the Network Co-occurrence Underlying Mental Operations. *PLOS Computational Biology* 12, e1004994.

Calhoun, V.D., Adali, T., Pearlson, G.D., and Pekar, J.J. (2001). A method for making group inferences from functional MRI data using independent component analysis. *Human Brain Mapping* 14, 140–151.

Carlson, J.M., Foti, D., Mujica-Parodi, L.R., Harmon-Jones, E., and Hajcak, G. (2011). Ventral striatal and medial prefrontal BOLD activation is correlated with reward-related electrocortical activity: A combined ERP and fMRI study. *NeuroImage* 57, 1608–1616.

Cattell, R.B. (1966). The Scree Test For The Number Of Factors. *Multivariate Behavioral Research* 1, 245–276.

Dale, A.M., Fischl, B., and Sereno, M.I. (1999). Cortical surface-based analysis. I. Segmentation and surface reconstruction. *Neuroimage* 9, 179–194.

Damoiseaux, J.S., Rombouts, S. a. R.B., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., and Beckmann, C.F. (2006). Consistent resting-state networks across healthy subjects. *PNAS* 103, 13848–13853.

Ecker, C., Reynaud, E., Williams, S.C., and Brammer, M.J. (2007). Detecting functional nodes in large-scale cortical networks with functional magnetic resonance imaging: a principal component analysis of the human visual system. *Hum Brain Mapp* 28, 817–834.

Eickhoff, S.B., Thirion, B., Varoquaux, G., and Bzdok, D. (2015). Connectivity-based parcellation: Critique and implications. *Hum Brain Mapp* 36, 4771–4792.

Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Essen, D.C.V., and Raichle, M.E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *PNAS* 102, 9673–9678.

Fox, M.D., Zhang, D., Snyder, A.Z., and Raichle, M.E. (2009). The Global Signal and Observed Anticorrelated Resting State Brain Networks. *J Neurophysiol* 101, 3270–3283.

Fransson, P. (2005). Spontaneous low-frequency BOLD signal fluctuations: an fMRI

investigation of the resting-state default mode of brain function hypothesis. *Hum Brain Mapp* 26, 15–29.

Gu, Y., Sainburg, L.E., Kuang, S., Han, F., Williams, J.W., Liu, Y., Zhang, N., Zhang, X., Leopold, D.A., and Liu, X. (2020). Brain activity fluctuations propagate as waves traversing the cortical hierarchy. *BioRxiv* 2020.08.18.256610.

Ham, J., Lee, D.D., Mika, S., and Schölkopf, B. (2003). A Kernel View Of The Dimensionality Reduction Of Manifolds.

Hannachi, A., Unkel, S., Trendafilov, N.T., and Jolliffe, I.T. (2009). Independent Component Analysis of Climate Data: A New Look at EOF Rotation. *J. Climate* 22, 2797–2812.

Horel, J.D. (1984). Complex Principal Component Analysis: Theory and Examples. *J. Climate Appl. Meteor.* 23, 1660–1673.

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 antagonistic brain networks reflect attentional fluctuations. *Nature Communications* 11, 325.

Li, J., Bolt, T., Bzdok, D., Nomi, J.S., Yeo, B.T.T., Spreng, R.N., and Uddin, L.Q. (2019). Topography and behavioral relevance of the global signal in the human brain. *Scientific Reports* 9, 14286.

Liu, X., and Duyn, J. (2013a). Resting-State fMRI Signal Anti-Correlation Exists in Absence of Global Signal Regression. In *Proceedings of the 21st ISMRM Annual Meeting*, (Salt Lake City, Utah), p. 1.

Liu, X., and Duyn, J.H. (2013b). Time-varying functional network information extracted from brief instances of spontaneous brain activity. *PNAS* 110, 4392–4397.

Liu, T.T., Nalci, A., and Falahpour, M. (2017). The Global Signal in fMRI: Nuisance or Information? *Neuroimage* 150, 213–229.

Liu, X., Yanagawa, T., Leopold, D.A., Chang, C., Ishida, H., Fujii, N., and Duyn, J.H. (2015). Arousal Transitions in Sleep, Wakefulness, and Anesthesia are Characterized by an Orderly Sequence of Cortical Events. *Neuroimage* 116, 222–231.

Liu, X., de Zwart, J.A., Schölvink, M.L., Chang, C., Ye, F.Q., Leopold, D.A., and Duyn, J.H. (2018). Subcortical evidence for a contribution of arousal to fMRI studies of brain activity. *Nature Communications* 9, 395.

Majeed, W., Magnuson, M., Hasenkamp, W., Schwarb, H., Schumacher, E.H., Barsalou, L., and Keilholz, S.D. (2011). Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans. *NeuroImage* 54, 1140–1150.

Margulies, D.S., Ghosh, S.S., Goulas, A., Falkiewicz, M., Huntenburg, J.M., Langs, G., Bezgin, G., Eickhoff, S.B., Castellanos, F.X., Petrides, M., et al. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. *PNAS* 113, 12574–12579.

Mitra, A., Snyder, A.Z., Hacker, C.D., and Raichle, M.E. (2014). Lag structure in resting-state fMRI. *J Neurophysiol* *111*, 2374–2391.

Mitra, A., Snyder, A.Z., Blazey, T., and Raichle, M.E. (2015). Lag threads organize the brain's intrinsic activity. *Proc Natl Acad Sci U S A* *112*, E2235–2244.

Murphy, K., Birn, R.M., Handwerker, D.A., Jones, T.B., and Bandettini, P.A. (2009). The impact of global signal regression on resting state correlations: Are anti-correlated networks introduced? *Neuroimage* *44*, 893–905.

Reid, A.T., Headley, D.B., Mill, R.D., Sanchez-Romero, R., Uddin, L.Q., Marinazzo, D., Lurie, D.J., Valdés-Sosa, P.A., Hanson, S.J., Biswal, B.B., et al. (2019). Advancing functional connectivity research from association to causation. *Nature Neuroscience* *22*, 1751–1760.

Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.-N., Holmes, A.J., Eickhoff, S.B., and Yeo, B.T.T. (2018). Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex* *28*, 3095–3114.

Schölvinck, M.L., Maier, A., Ye, F.Q., Duyn, J.H., and Leopold, D.A. (2010). Neural basis of global resting-state fMRI activity. *PNAS* *107*, 10238–10243.

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., et al. (2009). Correspondence of the brain's functional architecture during activation and rest. *PNAS* *106*, 13040–13045.

Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J., Glasser, M.F., et al. (2012). Temporally-independent functional modes of spontaneous brain activity. *PNAS* *109*, 3131–3136.

Smith, S.M., Andersson, J., Auerbach, E.J., Beckmann, C.F., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D.A., Griffanti, L., Harms, M.P., et al. (2013). Resting-state fMRI in the Human Connectome Project. *Neuroimage* *80*, 144–168.

Stetter, M., Schießl, I., Otto, T., Sengpiel, F., Hübener, M., Bonhoeffer, T., and Obermayer, K. (2000). Principal Component Analysis and Blind Separation of Sources for Optical Imaging of Intrinsic Signals. *NeuroImage* *11*, 482–490.

Thomas, C.G., Harshman, R.A., and Menon, R.S. (2002). Noise reduction in BOLD-based fMRI using component analysis. *Neuroimage* *17*, 1521–1537.

Thomas Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., et al. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol* *106*, 1125–1165.

Thompson, G.J., Pan, W.-J., Magnuson, M.E., Jaeger, D., and 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.J., Yacoub, E., Ugurbil, K., and WU-

Minn HCP Consortium (2013). The WU-Minn Human Connectome Project: an overview. *Neuroimage* 80, 62–79.

Vidaurre, D., Smith, S.M., and Woolrich, M.W. (2017). Brain network dynamics are hierarchically organized in time. *Proc Natl Acad Sci USA* 114, 12827–12832.

Vos de Wael, R., Benkarim, O., Paquola, C., Lariviere, S., Royer, J., Tavakol, S., Xu, T., Hong, S.-J., Langs, G., Valk, S., et al. (2020). BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. *Communications Biology* 3, 1–10.

Yousefi, B., and Keilholz, S. (2021). Propagating patterns of intrinsic activity along macroscale gradients coordinate functional connections across the whole brain. *NeuroImage* 231, 117827.

Yousefi, B., Shin, J., Schumacher, E.H., and Keilholz, S.D. (2018). Quasi-periodic patterns of intrinsic brain activity in individuals and their relationship to global signal. *Neuroimage* 167, 297–308.