

A Parsimonious Description of Global Functional Brain Organization in Three 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

²Department of Psychiatry and Biobehavioral Sciences, University of California Los Angeles,
Los Angeles, CA, USA

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

⁴Vanderbilt University, Nashville, TN, USA

⁵National University of Singapore, Singapore

Summary

Resting-state functional MRI has yielded many, seemingly disparate insights into large-scale organization of the human brain. Representations of the brain's large-scale organization can be divided into two broad categories - zero-lag representations of functional connectivity structure and time-lag representations of traveling wave or propagation structure. Here we sought to unify multiple observed phenomena in these two categories by applying concepts from the study of standing (zero-lag) and traveling (time-lag) waves. Using simulated data, we first described the behavior of zero-lag and time-lag analyses applied to spatiotemporal patterns with varying degrees of standing and traveling wave patterns. We then identified three low-frequency spatiotemporal patterns in empirical resting-state fMRI signals, composed of a mixture of standing and traveling wave dynamics, that describe both the zero-lag and time-lag structure of these signals. We showed that a wide range of previously observed empirical phenomena, including functional connectivity gradients, the task-positive/task-negative pattern, the global signal, time-lag propagation patterns, the quasiperiodic pattern, and the functional connectome network structure are manifestations of these three spatiotemporal patterns. These patterns account for much of the global spatial structure that underlies functional connectivity analyses, and therefore impact how we interpret everything derived from resting-state fMRI data from functional networks to graphs.

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

Introduction

Many of us are familiar with the Indian parable of the blind men and the elephant. A group of blind men encounter an elephant for the first time and try to acquire an overall description of the elephant by discussing the part of the elephant each blind man is touching (trunk, tusk, foot, etc.). Each blind man could access only a piece of the truth, and failed to grasp how their individual observations came together in a unified whole. This parable teaches the perils of missing the ‘bigger picture’ due to our own limited observations, and may be a suitable metaphor for the current state of our theoretical understanding of intrinsic functional brain organization as assessed using resting state fMRI. Since the discovery of spontaneous low-frequency blood-oxygenation-level dependent (BOLD) fluctuations in the 1990s¹, 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 approach², with little attempt to translate findings across different approaches.

We operationalize ‘intrinsic functional brain organization’ as the spatial correlation structure between cortical BOLD signals in the low-frequency range (~0.01 - 0.1 Hz). A notable feature of the brain’s spatial correlation structure is its organization into global, cortex-wide patterns that span across functional systems for cognition, perception and action³⁻⁶. This current study distinguished between two different characterizations of this global structure; zero-lag synchrony and time-lag synchrony between brain regions. Zero-lag synchrony is defined as instantaneous statistical dependence between two time courses - e.g. the correlation between two BOLD signals with no time-lag. Time-lag synchrony is defined as the statistical dependence between two time courses, where one time-course is lagged or delayed in time.

Zero-lag synchrony has often been labeled as ‘functional connectivity’ (FC). Starting from seed-based correlation methods^{1,4}, FC in the brain has been characterized by a wide variety of analytic methods^{3,6-10}, including dimension-reduction techniques^{3,11,12}, network or graph-theoretical techniques¹³, co-activation pattern analysis⁹ and latent state space models¹⁴. For convenience, we further distinguish between two types of zero-lag FC methods: topography-based methods and network-based methods. Topography-based methods represent patterns of zero-lag synchrony through low-dimensional latent spaces with weights assigned to brain regions (or voxels/vertices). Network-based methods analyze the zero-lag synchrony relationships (‘edges’) between brain regions, typically represented in the form of pairwise correlation matrices. We refer to these methods collectively as ‘zero-lag analyses.’ The zero-lag analysis approach has identified several global patterns spanning across functional networks that have generated sustained research interest: the global signal^{15,16}, the task-positive/task-negative pattern⁴, and more recently, the principal FC gradient³.

More recently, global patterns of time-lag synchrony in spontaneous BOLD fluctuations have been investigated. Two prominent global patterns with significant time-lag structure have been identified in previous studies. First, Mitra et al.^{17,18} have observed global propagation patterns at short-time scales in spontaneous BOLD fluctuations (~0 - 2s). Several types of short time-scale propagation patterns have been identified¹⁷, but of interest in this study is the brain's average 'lag projection'¹⁸, representing the average time-ordering of whole-brain spontaneous BOLD fluctuations¹⁸. The latency structure of the average lag projection extends across cortical and subcortical brain regions, and varies according to experimental manipulations of task demands and sensory inputs, suggesting that at least some of this structure is uncoupled from hemodynamic delays¹⁸. The time-ordering of brain regions is identified through the cross-correlation function between all pairs of brain region time courses. We refer to this analysis as the 'lag projection algorithm'. Second, Majeed et al.¹⁹ has described a pseudo-periodic spatiotemporal pattern at a longer time-scale (~20s) involving an alteration in BOLD amplitudes between the task-positive (TP) and default mode networks (DMN), known as the 'quasi-periodic pattern' (QPP). The shift in BOLD amplitudes between TP 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^{20,21}. The QPP is generally derived from spontaneous BOLD fluctuations using a recursive matching and averaging procedure, what we refer to simply as the 'QPP-algorithm.' We refer to the QPP-algorithm and lag projection algorithm collectively as 'time-lag analyses'.

We suggest that there may be an underlying unity to these representations that has heretofore remained overlooked. We hypothesized that the vast majority of widely-used zero-lag and time-lag representations of intrinsic functional brain organization are capturing different aspects of a small number of spatiotemporal patterns. As described below, these spatiotemporal patterns are large-scale, low-frequency fluctuations of spontaneous BOLD fluctuations that exhibit both zero-lag and time-lag structure. Our specific hypotheses were 1) global patterns of zero-lag and time-lag synchrony are describing different facets of the same spatiotemporal patterns, and 2) a small set of spatiotemporal patterns can explain a surprising number of previous findings in the literature describing spontaneous BOLD fluctuations.

Three lines of evidence provide support for the first hypothesis. First, recent studies have shown that time-lag representations of traveling waves in the cerebral cortex have a spatial distribution that precisely maps to the spatial weights of the principal FC gradient²²⁻²⁴. Second, the spatial topography of the cortical global signal is not entirely constituted by zero-lag spatial structure, but has significant time-lag structure^{24,25}. Third, removal of time-lag synchronous patterns, such as the QPP, from spontaneous BOLD fluctuations substantially alters patterns of zero-lag synchrony in FC network representations²⁶. These findings suggest that there may be a common pattern of global BOLD activity that unifies these zero-lag and time-lag representations.

To develop intuition on the proposed relationship between zero-lag and time-lag synchrony patterns, we translate conceptual insights from the physical study of complex waves in arbitrary mediums, and their application to geophysical phenomena^{27,28} (e.g., climate patterns). Specifically, we utilize the concepts of ‘standing’ and ‘traveling’ waves in geophysical fields and dedicated methods for their analysis^{29–31}. Standing waves refer to stationary oscillations (including non-sinusoidal) in a spatial field, exhibiting no time-lagged statistical dependencies across space. Due to the absence of time-lag statistical dependence, standing waves are entirely characterized by measures of zero-lag synchrony of the kind captured by FC analyses. Traveling waves refer to oscillations in a spatial field with non-zero time-lag statistical dependence across space. Traveling waves are characterized by measures of time-lag synchrony of the kind captured by the QPP and lag projection algorithm¹⁸.

Standing and traveling wave patterns often appear in the study of geophysical phenomena, particularly in the study of planetary ocean and atmospheric dynamics²⁸. For example, low-frequency (multi-decade) regional co-variations in climate anomalies (e.g. sea surface temperatures) across large distances, known as teleconnections, are composed of large-scale standing and traveling wave spatial structure^{32,33}. Several analytic approaches in the climate science community have been devised to decompose the spatial structure of these teleconnections into standing and traveling wave components^{34,35}. One flexible means to model the complex standing and traveling wave spatial structure of climate teleconnections is an extension of principal component analysis (PCA) into the complex-valued domain, known as complex PCA (CPCA)³⁰. Importantly, complex-valued dimension-reduction analyses, such as CPCA, allow the simultaneous modeling of both standing and traveling wave components in a single latent dimension.

Similar to complex climate phenomena, we suggest coherent, large-scale BOLD activity can be composed of a mixture of both standing and traveling wave spatial structure. The findings of similar spatial structure derived from time-lag and zero-lag analysis of spontaneous BOLD fluctuations supports this claim. Furthermore, CPCA offers a means to characterize the standing and traveling wave structure of these spatiotemporal patterns in spontaneous BOLD fluctuations. In addition, CPCA allows us to quantify the degree of traveling wave behavior in a spatiotemporal pattern. Through simulation and empirical study of spontaneous BOLD fluctuations, we demonstrated that at relatively low values of this metric, we should expect a convergence in results between the spatial patterns produced by zero-lag and time-lag analyses. In other words, to the extent that spontaneous BOLD fluctuations are characterized by standing waves, zero-lag FC analyses should provide a sufficient description of their dynamics. To the extent that spontaneous BOLD fluctuations are characterized by traveling waves, we should expect a divergence between zero-lag and time-lag representations of their dynamics.

In support of the second hypothesis, we begin with the observation that the resting-state fMRI literature reveals very similar patterns of global BOLD activity across analytic approaches, including FC gradients^{3,36}, co-activation patterns (CAPs)⁹, independent component analysis (ICA)¹⁰, latent brain state methods¹⁴, and seed-based correlation analyses⁴. We propose that these similar patterns across analysis methods are driven by the fact that they are descriptions of the same underlying spatiotemporal patterns proposed in our first hypothesis. To test the second hypothesis, we first conducted a systematic survey of zero-lag FC analyses and their outputs, and compared them to a set of spatiotemporal patterns derived from CPCA. We then compared the spatiotemporal patterns derived from CPCA to patterns derived from time-lag analyses - the average lag projection and the QPP. Finally, we compared the reconstruction of the brain's cortical functional connectome ('FC matrix') from the spatiotemporal patterns derived from CPCA to the observed functional connectome derived from raw, spontaneous BOLD fluctuations.

Our analyses revealed that three spatiotemporal patterns constitute the dominant spatial structure in spontaneous low-frequency BOLD fluctuations. With these three patterns, we can unify a range of previous findings in resting-state fMRI, including lag projections¹⁸, the QPP¹⁹, the topography of the global signal³⁷, the task-positive/task-negative pattern⁴, the principal FC gradient³, and FC network structure. We demonstrate that all of these previous observations are manifestations of three spatiotemporal patterns captured within a unifying framework capable of modeling standing and traveling oscillatory BOLD phenomena. This novel framework provides a parsimonious description of global functional brain organization that can inspire new hypotheses about the mechanisms underlying coordination of activity across the brain.

Results

1. Zero-lag and Time-lag Structure in Simulated Standing and Traveling Oscillations

A central claim of this study is that the dominant spatial structure of global spontaneous BOLD fluctuations is constituted by a small set of spatiotemporal patterns. Further, we suggest that the relative mixture of standing and traveling waves in these patterns explains the spatial similarity between outputs of zero-lag and time-lag analyses. In addition, we propose that CPCA provides the means to effectively capture the spatial and temporal properties of these spatiotemporal patterns. To test these claims, we conducted a simulation study of spatiotemporal patterns consisting of varying degrees of traveling and standing waves.

We generated data simulations of standing and traveling wave oscillations (see **Supplementary Modeling Note 1**). These consisted of a back-and-forth sinusoidal oscillation of

Gaussian curves on a two-dimensional grid (**Figure 1**). This synthetic data generation approach allowed us to systematically vary the degree of ‘traveling’ wave behavior in each oscillation by adjusting the distance between the peak locations of each Gaussian curve, from a distance of zero (‘pure standing’ motion), to a large distance (‘pure traveling’ motion) between peak locations. Zero-lag dimension-reduction techniques commonly used in the FC literature were applied to the time series of this two-dimensional grid to extract oscillation patterns between the two Gaussian curves.

As expected, zero-lag dimension-reduction techniques applied to pure standing motion of the two Gaussian curves effectively captured the oscillation in one latent factor (**Figure 1A**; zero peak distance). This was confirmed by inspection of the ordered eigenvalue plot (scree plot) of the correlation matrix from all time series in the grid. Using a heuristic known as Catell’s scree plot test^{38–41}, we identified the point in the scree plot at which the successive extraction of latent components exhibits a flattening (‘diminishing return’) in explained variance as the optimal number of components to extract. For pure standing motion, only the first eigenvalue was non-zero, confirming the presence of one latent factor. Further, correlations between the grid time courses and a seed time course extracted from the center of one Gaussian curve exhibited the same spatial pattern as the latent factors produced from dimension-reduction analyses.

At non-zero distance between the peak location of the Gaussian curves, zero-lag dimension reduction techniques separated the non-synchronous motion into two components (**Figure 1A**; peak distance: large and moderate). As reflected in the scree plots, the second eigenvalue of the correlation matrix increased with larger distance, indicating the presence of two latent components. The spatial patterns of the two components were largely consistent across methods. However, methods favoring sparsity in the spatial weights (varimax rotation) separated the traveling motion into isolated Gaussian curves. Methods favoring non-sparse, or spatially-overlapping patterns, such as PCA, temporal independent component analysis, and Laplacian eigenmaps, extracted separate phases or ‘snapshots’ of the overall oscillation.

To accurately characterize traveling wave behavior, we then extended dimension-reduction techniques into the complex-valued domain by augmenting the real time courses of the grid into a complex signal via the Hilbert transform. To demonstrate the utility of complex-valued dimension-reduction methods in extracting traveling wave oscillations, we applied complex-valued PCA (CPCA) to the same empirical data simulations of traveling motion. The application of CPCA to spatially-overlapping oscillations and simulated hemodynamic signals is detailed in **Supplementary Modeling Note 1**. Analogous to our diagnostics of the zero-lag correlation matrix, we also constructed a scree plot consisting of the ordered eigenvalues from the complex-valued correlation matrix constructed from the complex grid time courses.

For pure traveling motion, the scree plot from the complex-valued correlation matrix revealed that only the first eigenvalue was non-zero, indicating the existence of a single latent factor (**Figure 1A**; large peak distance). We extracted the amplitude and phase-lag of the first complex principal component's spatial weights, reflecting the component's amplitudes and phase-lag between grid time series, respectively. The grid amplitudes of the first component reflected the spatial distribution of the two Gaussian curves, indicating that the coherent fluctuations of the two Gaussian curves were captured by a single latent component. Importantly, the phase-lag values ($-\pi$ to π) of the first component precisely reflected the back-and-forth motion of the Gaussian curves.

In simulations with non-zero distance between peak locations, CPCA and zero-lag dimension reduction techniques were found to converge on a similar solution. Further, at smaller distances, the scree plots of the zero-lag and complex-valued correlation matrix agreed in terms of yielding a one-factor solution. In other words, the degree of agreement in the solutions of zero-lag and time-lag analyses depended on the degree of traveling waves in the simulation.

By construction, this simulation had access to the 'ground truth' mixture of traveling and standing wave components. No such ground truth is available in BOLD signals recorded from the brain. Thus, we invoked a 'traveling index' metric developed in Feeny⁴² that measures the presence of traveling waves in an extracted CPCA component, varying from 0 (pure standing waves) to 1 (pure traveling waves). In our data simulation, this metric was approximately linearly associated with the peak distance between the two Gaussian curves (**Figure 1B**). Such a metric provides a quantitative estimate of the mixture of traveling and standing waves in the absence of ground truth simulated data.

From the observation that zero-lag dimension-reduction methods tend to split traveling Gaussian waves into separate latent components, we used the percentage of explained variance of the first eigenvector as a quality metric. The greater the explained variance, the more effectively zero-lag dimension-reduction methods can capture traveling wave motion in a single latent component. We plotted the explained variance of the first eigenvector by the traveling index generated from separate runs of the Gaussian curve simulation at varying degrees of peak distance - from large distances (pure traveling) to zero distance (pure standing). We found that the explained variance exhibits a linear decrease moving from values beyond 0.2 to larger values of the traveling index (**Figure 1B**). For small values of the traveling index (less than 0.2), we observed that the explained variance remains largely stable ($> 95\%$). Overall, we found that for moderate values of the traveling index (< 0.5), the explained variance of the first eigenvector is greater than 80%. Thus, zero-lag FC methods can effectively capture a large majority of the variance of a spatiotemporal pattern with moderate traveling wave behavior. While the spatiotemporal dynamics of intrinsic BOLD signals can be expected to be more complex than our empirical simulation, the central insights from this simulation framework are broadly

applicable. For systems consisting predominantly of standing waves, or a moderate degree of traveling waves, we expect a convergence between the solutions of time-lag and zero-lag FC methods.

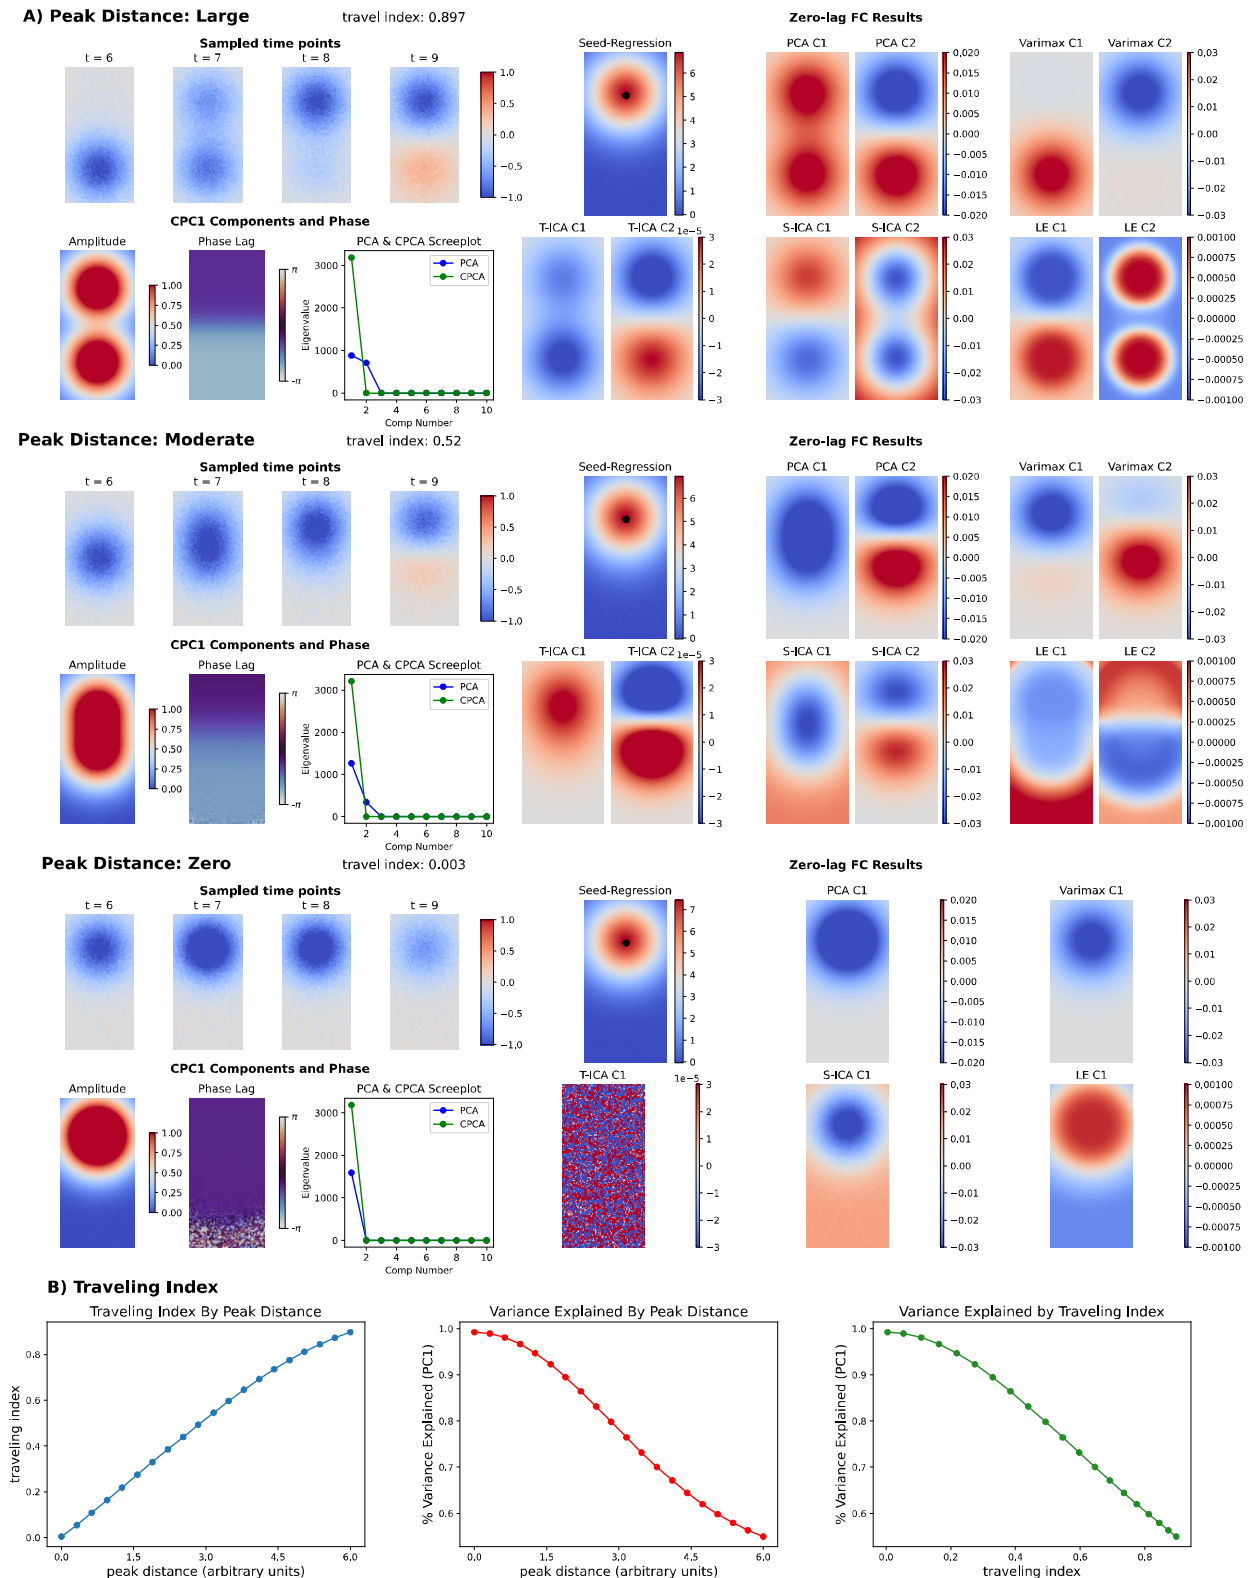


Figure 1. **Empirical Data Simulation to Analyze Standing and Traveling Wave Oscillations.** A) Visualization of simulation and analysis results of traveling two-dimensional Gaussian curves at three peak distances (large, moderate and zero distance). For each simulation scenario, the bottom Gaussian curve moves upward (bottom to top) towards the top Gaussian curve. In the top left panel of each simulation are four arbitrarily sampled time points.

Note, the travel distance between the two Gaussian curves grows smaller at smaller peak distances (moving from the top to bottom panel). The amplitude and phase-lag maps of the first complex principal component from CPCA are shown in the bottom left panel of each simulation. For all simulations, the amplitude and phase maps of the first complex component accurately describe the spatial distribution and phase-lag between the two Gaussian curves. In addition, the scree plot of the zero-lag (blue) and complex-valued (green) correlation matrix is displayed. Results of various zero-lag FC analyses are displayed in the right panel. For the dimension-reduction techniques, two components were estimated at non-zero peak distances, motivated by the non-zero second eigenvalue of the correlation matrix. For zero-peak distance, only a single component was estimated. At non-zero peak distances, zero-lag dimension-reduction techniques tend to split the traveling wave oscillation into either 1) separate Gaussian curves (varimax), or 2) into separate phases of the oscillation. **B)** From left to right, plots of the traveling index by peak distance, the variance explained by the first eigenvalue (zero-lag correlation matrix) by peak distance, and the variance explained by the first eigenvalue by traveling index. PCA = principal component analysis; CPCA = complex PCA; S-ICA = spatial ICA; T-ICA = temporal ICA; LE = Laplacian eigenmaps

2. Zero-lag and Time-lag Structure in in Intrinsic BOLD Fluctuations Reveals Three Dominant Spatiotemporal Patterns

To understand the standing and traveling wave components that underlie empirical spontaneous BOLD fluctuations, we applied CPCA to a random subject sample ($n=50$) of resting-state fMRI scans (~15min each; TR = 0.72 sec; $n = 1200$ time points) collected as part of the Human Connectome Project Young Adult (HCP:YA S1200) database. As our interest was in global, cortex-wide patterns in spontaneous BOLD fluctuations, we downsampled the surface-based cortical time series to approximately 5,000 vertices. Subject time series were temporally concatenated and reshaped into a 2D matrix of time points (rows) by cortical vertices (columns) for group-wise analysis. All analyses in our investigation were successfully replicated in an independent sample ($n=50$) of unrelated HCP subjects (also unrelated to any subject in the primary sample) (**Supplementary Figure 1** and **Supplementary Figure 2**).

To inform the choice of the number of components to extract using CPCA, we utilized the same Catell's scree plot test applied in our simulations (**Figure 1A**). This widely-used criterion for component number selection indicated a clear flattening in explained variance beyond three principal components (**Supplementary Figure 3**). Thus, we committed to extracting three components using CPCA. The three leading complex principal components represent the top three dimensions of variability across complex-valued cortical BOLD time series. Associated with each complex principal component is a phase delay map, reflecting the time-delay (in seconds) between cortical vertices (see 'Methods and Materials' for construction of the phase delay maps). To examine the temporal progression of each complex principal component, we sampled the reconstructed BOLD time series from each complex principal component at multiple, equally-spaced phases of its cycle ($n_{timepoints} = 30$; see 'Methods and Materials'). Movies of the reconstructed BOLD time courses are displayed in **Movie 1**. Detailed descriptions of the spatiotemporal patterns of each complex principal component, and comparisons of the similarities and differences in their propagation patterns are provided in

Supplementary Figure 4. Phase delay maps of complex principal components beyond the first three are presented in **Supplementary Figure 5**.

The first complex principal component represents the leading axes of variance. The first component explains over 21% (21.4%) of the variance in intrinsic BOLD time series, greater than three times the variance explained by the second (6.8%) or third component (5.7 %). The traveling index of the first component was 0.25, indicating that its spatiotemporal pattern is largely stationary or standing, with some traveling wave behavior. Visual inspection of the first complex principal components suggested that its spatiotemporal pattern can be separated into two phases. In the first phase, strong negative BOLD amplitudes were observed in the sensorimotor cortex (SM), superior parietal lobule (SP), lateral visual cortex (LV), and superior temporal (ST) gyrus (**Figure 2A; Movie 1**). For brevity, we refer to this complex of brain regions as the somato-motor-visual (SMLV) complex, noting that this complex also includes some regions outside sensory-motor cortices (e.g. SP and ST). At peak negative amplitudes of the SMLV, the entire cerebral cortex also exhibits weak negative BOLD amplitudes in regions overlapping the default mode network (DMN) and primary visual cortex (V1). In other words, during the first phase of the spatiotemporal pattern, the cortex exhibited globally negative BOLD signals, with strong amplitudes in the SMLV regions, and weak amplitudes in the DMN. The second phase exhibited a propagation of strong negative BOLD amplitudes in the SMLV toward cortical regions overlapping primarily with the frontoparietal network (FPN), but also with the default mode network (DMN) and V1. Qualitative inspection suggested that the traveling wave behavior of the first component is largely restricted to the second phase of the spatiotemporal pattern. This entire spatiotemporal sequence of negative BOLD amplitudes was followed by a spatiotemporal sequence with positive BOLD amplitudes with the same dynamics. Because the explained variance of the first complex principal component is three times greater than the subsequent complex principal component, we also refer to this spatiotemporal pattern as the ‘dominant spatiotemporal pattern’ in intrinsic BOLD signals, or simply *pattern one*.

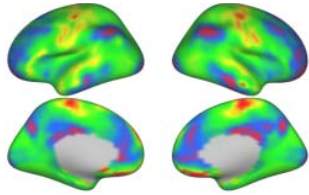
The second most explanatory complex component was the most stationary or standing component of the first three complex principal components, with a traveling index of 0.14. The overall spatiotemporal pattern can be succinctly described as an anti-correlated oscillation between SMLV regions and the DMN. Interestingly, this pattern of anti-correlation or bipolar contrast appeared to delineate the unipolar contrast in the first phase of the first complex principal component. The only discrepancy between the two patterns is in the SM region. In the spatiotemporal pattern of the first component, the SM region was in synchrony with other SMLV regions, while in the spatiotemporal pattern of the second component, it was in synchrony with the DMN. Visual inspection revealed a minor traveling wave pattern propagating out from the SM region to the premotor cortex (anterior direction) and SP (posterior direction) in-between peak amplitudes of the anti-correlated oscillation. For further reference, we henceforth refer to the spatiotemporal pattern of the second complex principal component as *pattern two*.

The spatiotemporal pattern of the third complex principal component had a traveling index of 0.27, similar in magnitude to the pattern of the first component. As with the first component, this implied that it is largely stationary or standing, with some traveling wave behavior. The spatiotemporal pattern of the third component can be described as an anti-correlated oscillation between regions consisting of the SM, ST, LV and DMN versus regions consisting of the inferior parietal lobule (IP), inferior temporal gyrus (IT), premotor cortex, dorsolateral prefrontal cortex (DLPFC) and V1. Visual inspection of the third component suggested that its spatiotemporal pattern can be split into two phases. In the first phase, strong negative amplitudes are observed in the SM, ST and LV, with weak negative amplitudes in the DMN; strong positive amplitudes are observed in the IP, IT, premotor cortex, DLPFC and V1. Over the course of the first phase, the strong negative amplitudes of the SM, ST and LV decreased in absolute value, approaching a similar magnitude to that of the DMN. The second phase was marked by propagation from the IP (anterior direction) and premotor (posterior direction) towards the SM, and the IT (posterior direction) towards the LV. For further reference, we henceforth refer to the spatiotemporal pattern of the third complex principal component as *pattern three*.

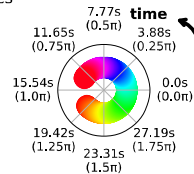
A) First Complex Principal Component - Pattern One

Travel index:0.25

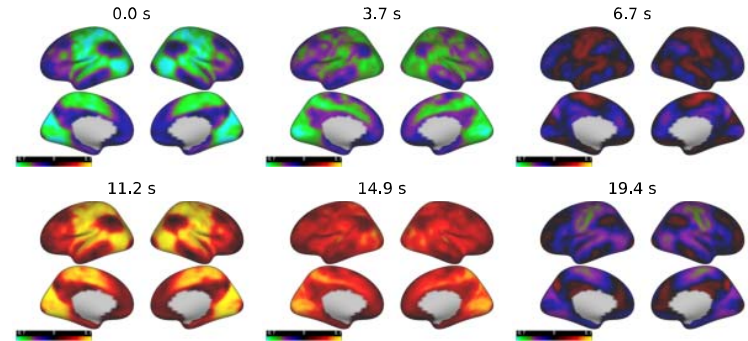
Phase Delay Map



Phase Delay Values



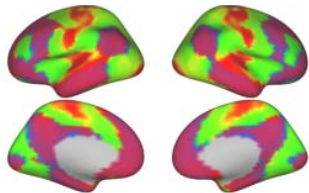
Reconstructed Time Points



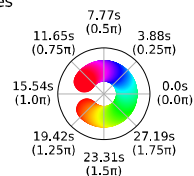
B) Second Complex Principal Component - Pattern Two

Travel index:0.14

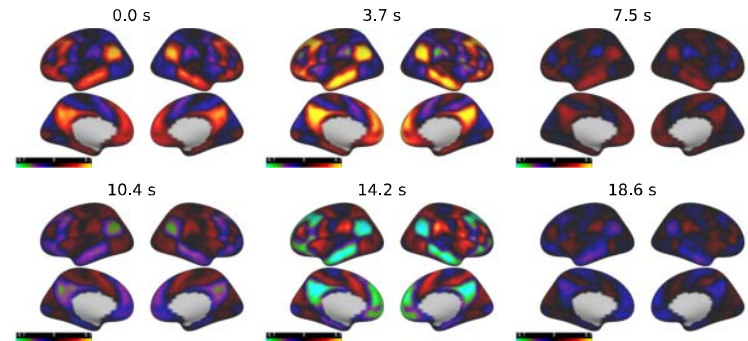
Phase Delay Map



Phase Delay Values



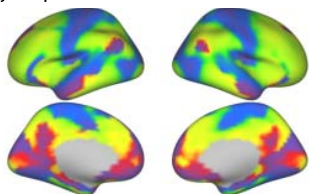
Reconstructed Time Points



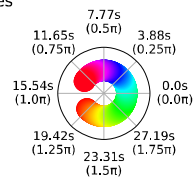
C) Third Complex Principal Component - Pattern Three

Travel index:0.27

Phase Delay Map



Phase Delay Values



Reconstructed Time Points

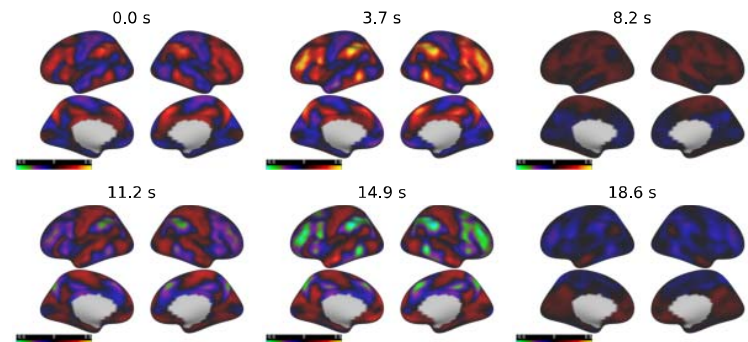


Figure 2. Form and Properties of Three Dominant Spatiotemporal Patterns. Phase delay maps and reconstructed time points of the first three complex principal components. Phase delay maps represent the temporal ordering (in seconds) of cortical vertex BOLD time series within the spatiotemporal pattern. Phase 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 phase 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. The phase delay map and reconstructed time points are shown for **A)** the first spatiotemporal pattern - 'pattern one', **B)** the second spatiotemporal pattern - 'pattern two', and **C)** the third spatiotemporal pattern - 'pattern three'.

Movie 1. Visualization of Spatiotemporal Patterns. Temporal reconstruction of all three spatiotemporal patterns displayed as movies in the following order - pattern one, pattern two, and pattern three. 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 (see **Supplementary Figure 4**) 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.

3. Zero-lag FC Topographies Effectively Capture the Three Spatiotemporal Patterns

The three spatiotemporal patterns derived from CPCA are predominantly composed of standing (stationary) waves. This finding, and the results of the simulation study (**Figure 2**), suggest that zero-lag FC methods should capture the majority of the variance in these spatiotemporal patterns. For example, traditional (non-complex) PCA applied to the zero-lag correlation matrix derived from the same time courses should yield similar spatial weights and time courses to that produced from CPCA. We selected a large number of commonly-used zero-lag FC methods and applied them to the original (non-complex) time courses from the same resting-state fMRI data. Our selection 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)^{43,44}, Laplacian Eigenmaps (LE)³⁶, the commonly used spatial independent component analysis (SICA)⁴⁵, and the more recent temporal independent component analysis (TICA)¹⁰. Hidden Markov models (HMM) are a commonly used latent state space model for estimating brain states¹⁴, and were also included in our study. Seed-based analysis methods included the traditional seed-based correlation/regression analysis⁴ and co-activation pattern (CAP) analysis⁹ 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 SMLV, default mode, and frontoparietal network (FPN) - the somatosensory cortex, precuneus, and dorsolateral prefrontal cortex⁶. Results were found to be identical for alternatively placed seed regions within these three networks (**Supplementary Figure 6**).

To determine a meaningful number of dimensions in our latent dimension-reduction analyses (PCA, varimax PCA, LE, SICA, TICA and HMM), we again used the scree plot

criterion. As with the scree plot of the complex-valued correlation matrix, the scree plot of the zero-lag correlation matrix indicated a clear flattening in explained variance after three principal components (**Figure 2C**).

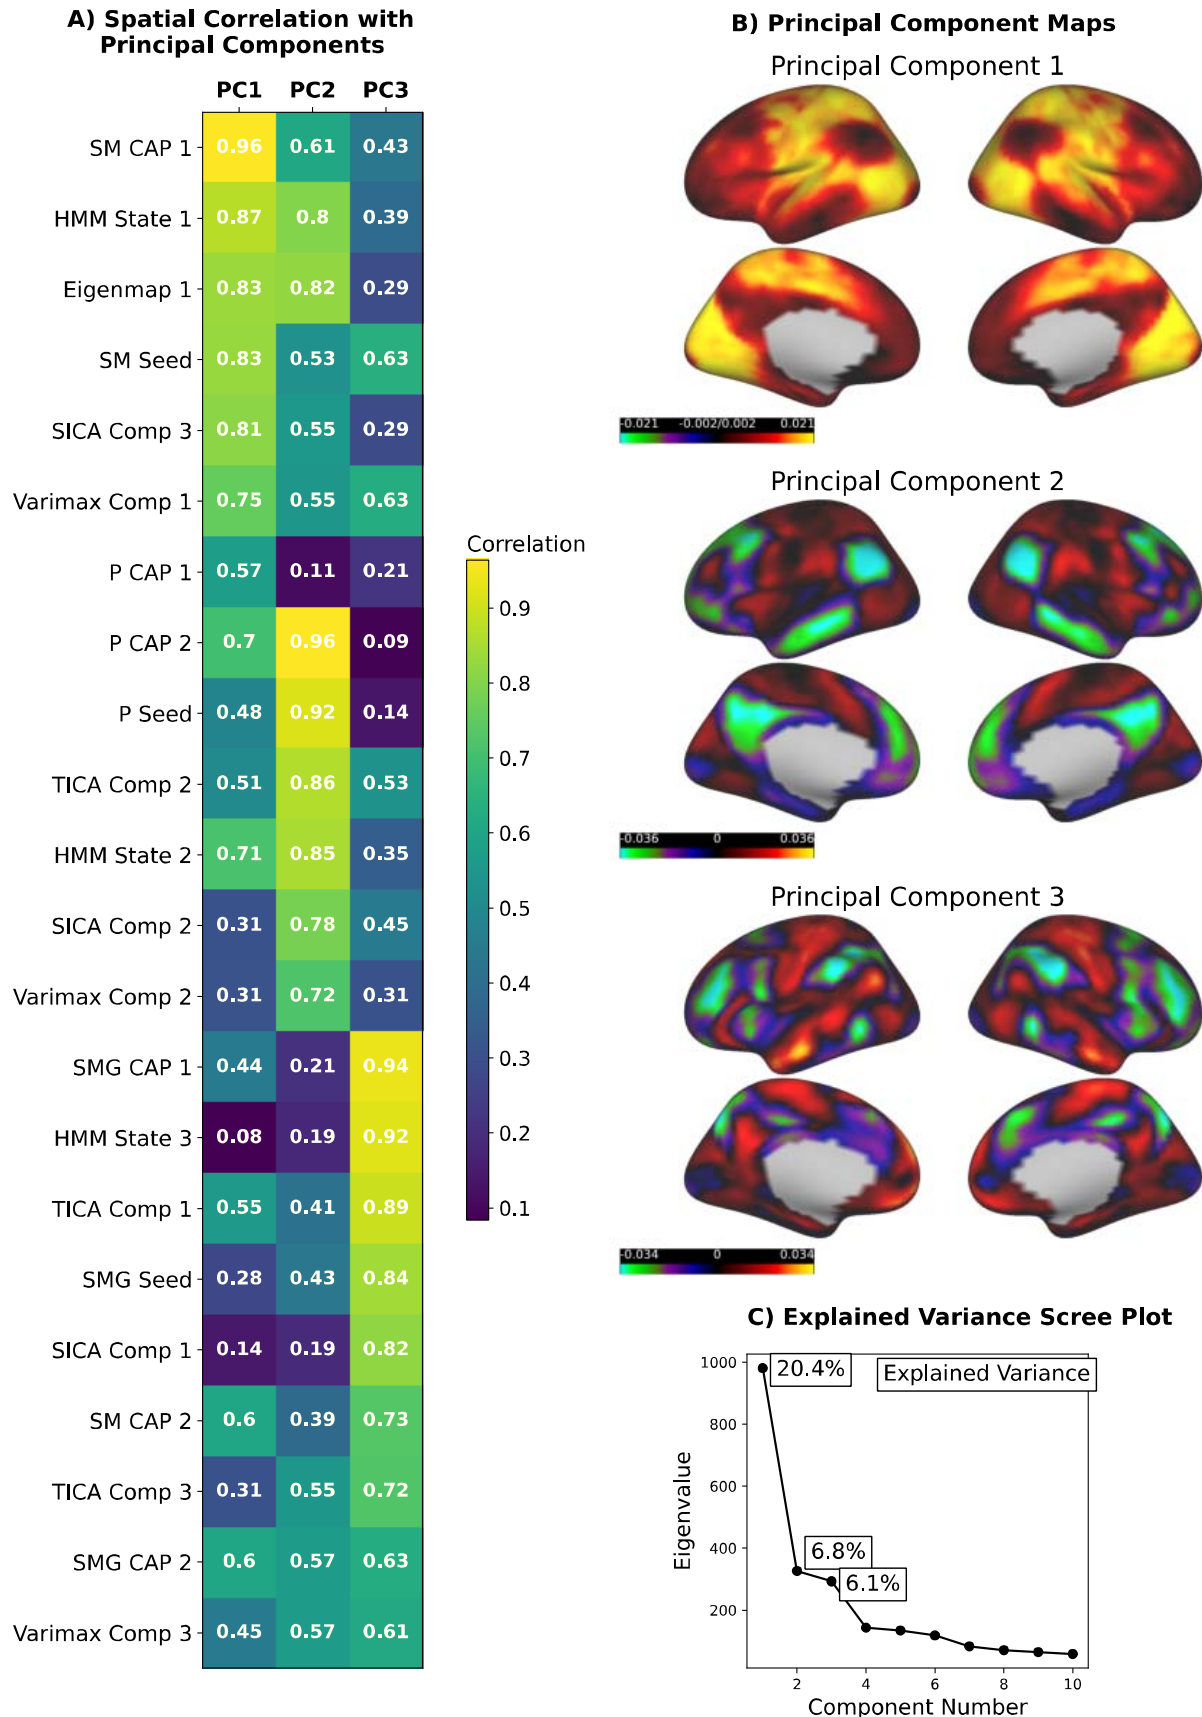


Figure 3. Form and Properties of Three Fundamental Functional Connectivity Topographies. **A)** The spatial absolute-valued correlation between the first three principal component maps and each FC topography displayed as a table. All correlations are rounded to the second decimal place. The color of 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. (HMM = hidden markov model; TICA = temporal ICA; SICA = spatial ICA; 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 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 standard (non-complex) PCA. Similar to the first complex principal component of CPCA, the first principal component explains 20.4% of the variance in BOLD time series. This is greater than three times the variance explained by the second (6.8%) or third principal component (6.1%). As can be observed from **Figure 3A**, each FC topography exhibits strong similarities with one or more of the first three principal components ($r > 0.6$). Statistical significance of the spatial correlation between each FC topography and its most strongly correlated principal component was computed using spin permutation tests⁴⁶ ($N_{\text{samples}} = 1000$). All correlation pairs were statistically significant ($p < 0.001$).

However, there was notable variability in the strength of correlations between each FC topography and one or more of the principal components. In one case, the first cluster centroid from CAP analysis of the precuneus seed ('P CAP 1'), the spatial similarity is strong ($r = 0.57$) but notably weaker than other FC topographies. Further, for many FC topographies, strong similarities are seen for more than one principal component. For topographies with strong similarity with both the first and second principal component, this can be explained by the spatial similarity between the first and second principal component ($r = 0.7$). Details of the output of seed-based regression/CAP and dimension-reduction analyses are provided in **Supplementary Figure 7** and **Supplementary Figure 8**, respectively. Overall, our survey revealed a considerable consistency in FC topographies across methods.

Visual comparison of the spatial weights of the three principal components (**Figure 3B**) with the spatiotemporal patterns (patterns one, two and three) of the three complex principal components (**Figure 2**) illustrates that both methods capture similar spatial patterns. The correspondence was found to be one-to-one, with the first principal component matching pattern one, the second principal component matching pattern two, and the third principal component matching pattern three.

To provide quantitative support to these observations, we computed the spatial correlation between the spatial weights of the three principal components and the reconstructed time courses of the three spatiotemporal patterns. We found that the static spatial weights of each PC appeared as a ‘snapshot’ within their corresponding spatiotemporal pattern: PC1-pattern one ($r = 0.998$, $t = 11.9s$), PC2-pattern two ($r = 0.986$, $t = 12.6s$), PC3-pattern three ($r = 0.972$, $t = 3.7s$). Further, the time courses of the three spatiotemporal patterns closely tracked the time courses of the first three principal components, respectively: PC1-pattern one ($r = 0.98$), PC2-pattern two ($r = 0.95$), and PC3-pattern three ($r = -0.83$, at a temporal lag of ~ 3 TRs). Given the similarity in results, we refer to the patterns produced by standard PCA and CPCA, as pattern one, two and three, interchangeably. Consistent with our simulation results, zero-lag FC methods capture the majority of the dynamics of these spatiotemporal patterns due to the predominance of standing waves in these patterns.

4. 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 projections¹⁸ and the QPP¹⁹ correspond to time-lagged phenomena at shorter ($\sim 2s$) and longer ($\sim 20s$) time scales, respectively. Lag projections represent the ‘average’ latency structure of spontaneous BOLD fluctuations. They are computed as the column-wise 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. An analogous column-wise averaging operation can be applied to the complex correlation matrix derived from the Hilbert transform of the BOLD time courses. Specifically, we computed the column-wise circular mean of the pairwise phase-delay values of the complex correlation matrix. The circular mean was computed due to the circular nature of the phase values of a complex correlation (i.e. $-\pi$ and π are identical phase differences). We found that the spatial correlation between the average latency structures were very similar ($r = 0.83$), and both exhibit the same direction of propagation (**Figure 4**). Interestingly, both average latency structures exhibited a strong spatial similarity with the phase delay map of pattern one (lag projection: $r = 0.81$; circular mean: $r = 0.98$) (**Figure 4**). This suggests the average latency structure of spontaneous BOLD fluctuations is largely driven by the first spatiotemporal pattern. Note, the lag projection we computed only partially resembles the average lag projection in Mitra et al.¹⁸ - the differences are due to preprocessing differences, which we discuss in **Supplementary Figure 9**.

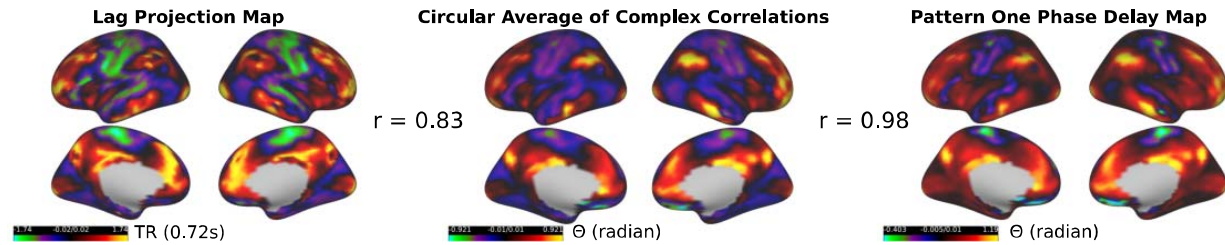


Figure 4. **Similar Propagation Patterns between Average Latency Structure and Pattern One.** Comparison between the average latency structure of spontaneous BOLD fluctuations and the phase delay map of pattern one. The lag projection map (left) represents the average time delay (in seconds) between each vertex of the cortex. The circular average map (middle) represents the average phase delay (in radians) between each vertex of the cortex. As in **Figure 2**, the pattern one phase delay map (right) represents the phase delay (in radians) between the vertices within the dynamics of pattern one. The two methods for computing average latency structure exhibiting strong agreement in their propagation patterns ($r = 0.83$). The strong similarity between the average latency structure (circular average) and pattern one phase delay map ($r = 0.98$) indicates that the average latency structure is primarily driven by the spatiotemporal dynamics of pattern one.

While lag projections describe short time scale propagation structures, the QPP is a much longer temporally-extended pattern ($>20s$). The QPP is conventionally derived from BOLD time courses after the application of global signal regression. As shown by Yousefi et al.⁴⁷, two types of QPP can be observed across individuals: one QPP pattern exhibits a global pattern of activity with positive correlation across brain regions ('global' QPP), and the other QPP exhibits a global pattern of anti-correlation between the TP and DMN ('anticorrelated' QPP). Global signal regression was found to consistently produce the anti-correlated QPP across all individuals.

We hypothesized that the global QPP corresponds to pattern one, and that the anti-correlated QPP corresponds to pattern two. To derive the global QPP, we applied the QPP algorithm without global signal regression. To derive the anti-correlated QPP, we applied the QPP algorithm after global signal regression was applied to all BOLD time courses. The QPP algorithm was used with a commonly used window size ($\sim 21s$; 30TRs)⁴⁷. Consistent with our hypotheses, the time course of the global QPP derived from non-global signal regressed time courses was strongly correlated with the time course of pattern one ($r = 0.71$) at a time-lag of 7TRs ($\sim 5s$). Further, the time course of the anti-correlated QPP was strongly correlated with pattern two ($r = 0.80$) at a time lag of 16 TRs ($\sim 12s$). The similarity between the spatiotemporal dynamics of the two QPPs, and patterns one and two can also be demonstrated visually. We visualized the spatiotemporal template of the global and anti-correlated QPP from the repeated template-averaging procedure, and compared it to the time point reconstruction (see above) of pattern one and pattern two, respectively. As can be seen in **Movie 2**, the temporal dynamics of the global QPP overlapped significantly with the dynamics of pattern one. Further, the temporal dynamics of the anti-correlated QPP overlapped significantly with the dynamics of pattern two (**Supplementary Figure 10**).

At peak amplitudes of pattern one, the distribution of weights is roughly unipolar, meaning it is either all positive or all negative (**Figure 2A**). This suggests that pattern one 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 pattern one 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 pattern one. 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 pattern one, we found that the temporal dynamics surrounding peaks of the global signal precisely match those of pattern one (**Movie 2**).

Movie 2. Dynamic Visualization of the Quasiperiodic Pattern, Pattern One, and Global Signal. The 30 time points (TR=0.72s) of the global QPP, pattern one, 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 pattern one are equally-spaced phase samples ($N=30$) of the time point reconstruction (see above). The time points of the global QPP are derived from the spatiotemporal template computed from the repeated-template-averaging procedure on non-global signal regressed data. 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 pattern two largely represents an anti-correlated pattern between the SMLV, excluding the SM region, and the DMN - i.e. when regions of the DMN exhibit negative BOLD activity, the regions of the SMLV exhibit positive BOLD activity (and vice versa). This resembles the “task-positive” vs. “task-negative” anti-correlation pattern originally observed by Fox et al.⁴ and Fransson⁴⁸. 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 SMLV and DMN (**Figure 4**). We compared the precuneus-seed correlation map to the time points of pattern two using spatial correlation. We found that the pattern of correlations in the precuneus-seed map precisely corresponds to the pattern of BOLD activity at peak amplitudes of pattern two (CPC1: $r = 0.92$, $t = 1.8s$). Thus, this would seem to suggest that the task-positive vs. task-negative pattern arises from the anti-correlated dynamics between the SMLV (task-positive) and DMN (task-negative) represented by pattern two.

A similar anti-correlation pattern to the task-positive/task-negative pattern has been observed in the FC gradient literature^{3,36}, known as the principal FC gradient (PG). Further, the PG has been referred to as the principal direction of variance in cortical functional connectivity³. In our zero-lag FC survey (**Figure 3**), we computed the PG as the first component derived from

the Laplacian Eigenmaps (LE) algorithm, consistent with Vos de Wael et al.³⁶. Visual comparison of the PG with pattern two revealed a superficial correspondence between the anti-correlated patterns, with the exception of somato-motor (SM) cortices. In pattern two, the SM region was in synchrony with the DMN, while in the PG the SM region is in synchrony with the SMLV, similar to that observed in pattern one. Further, our analysis suggested that the global unipolar pattern of BOLD activity exhibited by pattern one is the dominant axis of variance in FC, not the anti-correlated pattern observed in the PG.

With no clear correspondence between the PG and patterns one and two, we sought to identify the factors that might explain the uniqueness of the PG. We discovered that the spatial pattern of the PG is due to the confluence of two factors: 1) 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.¹⁵, 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 step⁴⁹ (**Supplementary Figure 11**). This is relevant because the global mean time course precisely tracks the time course of pattern one ($r = 0.97$). This would suggest that removal of the global mean time course through global signal regression and/or time-point centering effectively removes pattern one from BOLD time courses. What is left over is the second most dimension of variance in FC, pattern two. In fact, this would explain the strengthening of the task-positive vs. task-negative pattern after global signal regression in seed-based correlation analysis⁴.

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 pattern of BOLD activity at peak amplitude of pattern two (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 pattern two (CPC2: = 0.49 vs CPC1: = 0.10). Note, the correlation between the time-lag maps was computed using a circular correlation coefficient due to the circular properties of the phase map of the complex principal component (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 pattern one from BOLD time courses.

It is the dual effect of time-point centering and FC matrix thresholding that resolves the discrepancy between patterns one and two, 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^{3,36}. This preprocessing step is intended to remove noisy or artificial correlation values from the FC matrix. In our survey of zero-lag FC (**Figure 3**), 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 pattern two. In fact, if no thresholding of the FC matrix is performed, the first eigenmap produced from LE precisely resembles the pattern of BOLD activity at peak amplitudes of pattern two (PC2: $r = 0.83$; **Figure 5C**). As one raises the percentile threshold applied to the FC matrix, the spatial weights of vertices within the DMN and SMLV become more uniform (**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 pattern one (PC1: $r = 0.83$) and the anti-correlation contrast of pattern two (PC2: $r = 0.82$).

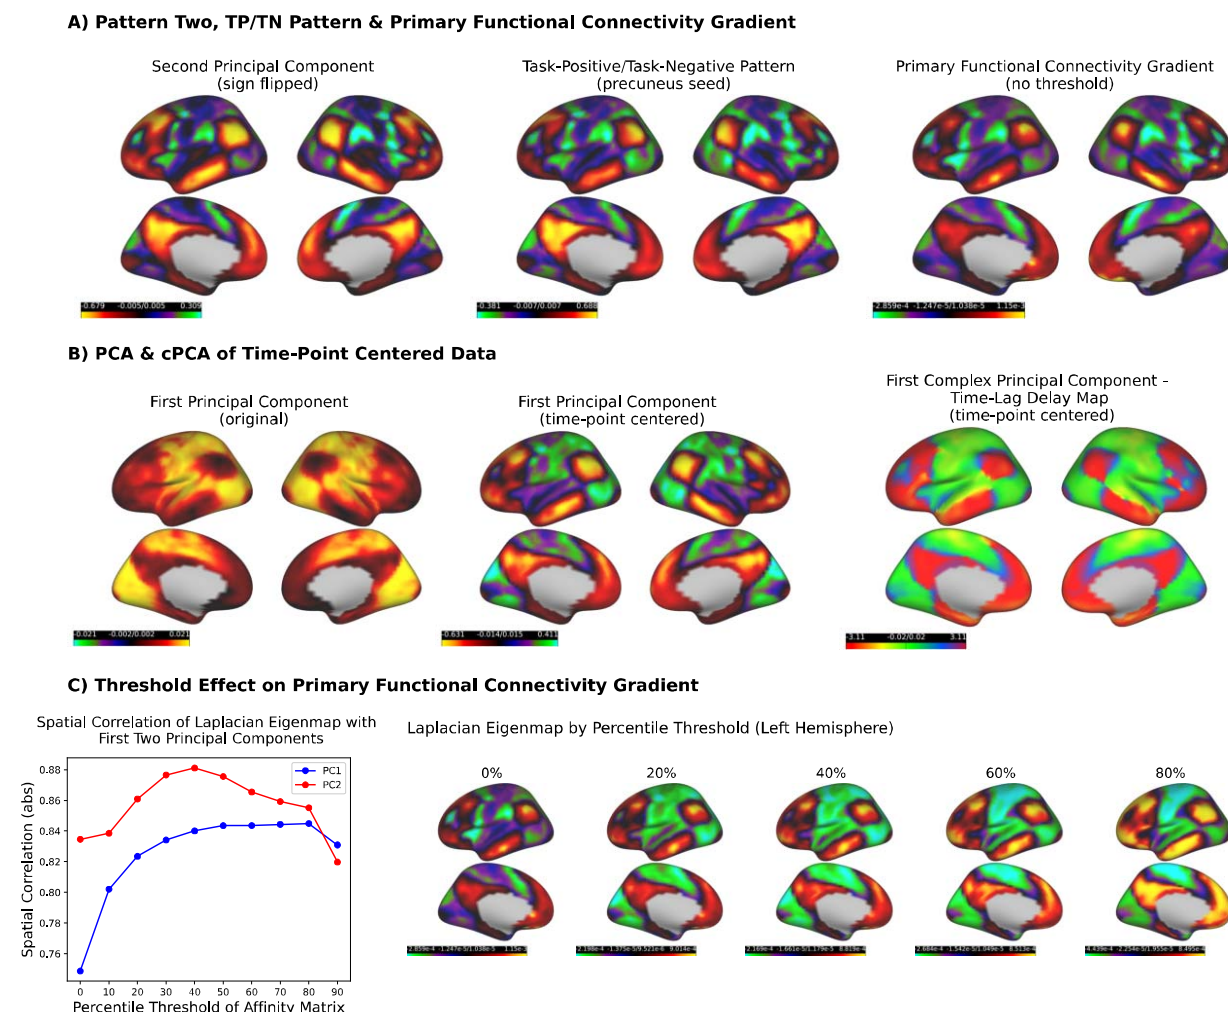


Figure 5. The Task-Positive/Task-Negative Pattern, Primary Functional Connectivity Gradient, and Pattern Two Describe the Same Spatiotemporal Pattern. A) From left to right, pattern two, task-positive/task-negative (TP/TN) pattern, and the PG represented by the spatial weights of the second principal 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. Similar spatial patterns are produced from all three analyses - pattern two: TP/TN ($r = 0.96$) and pattern two:PG ($r = 0.83$). **B**) From left to right, the first principal component from non-time-centered BOLD time courses (i.e. pattern one), the first principal component of time-centered BOLD time courses, and the first complex principal component phase 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; pattern one) with a bipolar (anti-correlated) principal component (middle) that resembles pattern two. 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 pattern two (**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 pattern two (PC2). As the threshold is raised, the spatial weights of vertices within the DMN and SMLV become more uniform. At higher thresholds this results in an Eigenmap that resembles pattern one.

5. 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 estimated network communities from the original and reconstructed FC matrix using the Louvain modularity-maximization algorithm with a resolution parameter value of 1 (with asymmetric treatment of negative weights). To compare the degree of similarity between the community structure of the original and reconstructed FC matrix, we computed the normalized mutual

information (NMI) between their community assignments from the Louvain algorithm. The NMI varies from 0 (completely independent communities) to 1 (completely identical communities). The NMI between the community assignments of the original and reconstructed FC matrix was 0.73, indicating a strong similarity. In other words, the community structure of the original FC matrix is very similar to the community structure of the FC matrix constructed from the shared dynamics of the three spatiotemporal patterns.

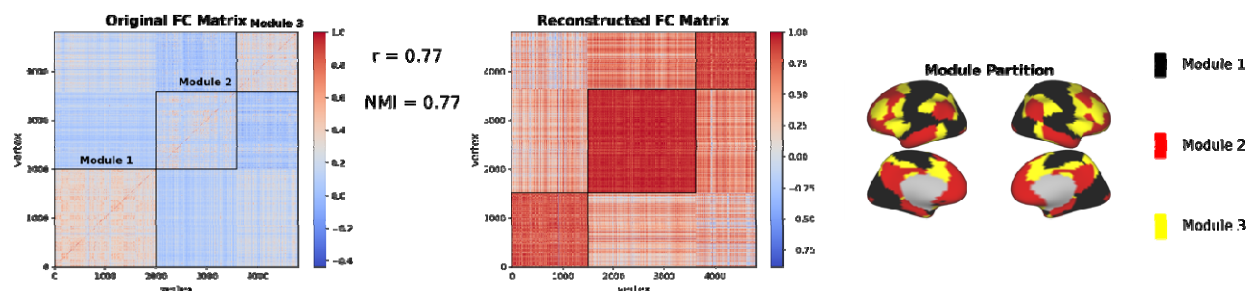


Figure 6. The Network Structure of Functional Connectivity Is Explained by the Three Spatiotemporal Patterns. Comparison of the correlation matrix of cortical BOLD time courses (left) with the correlation matrix of reconstructed cortical BOLD time courses (right) derived from the three spatiotemporal 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 estimated from the Louvain modularity algorithm. The algorithm identified three primary modules in the SMLV, DMN and FPN. Despite a higher mean value of correlations in the reconstructed correlation matrix, the pattern of correlations between the two correlation matrices is highly similar ($r = 0.77$). Further, the modular structure of the original correlation matrix exhibits a high degree of similarity with the modular structure of the reconstructed correlation matrix ($NMI = 0.73$).

Discussion

Approaches for analyzing resting-state functional MRI have yielded many, seemingly disparate insights into intrinsic organization of the human brain. Zero-lag representations of standing wave functional connectivity structure and time-lag representations of traveling wave structure have previously been used in isolation to characterize low-frequency BOLD fluctuations. Here we seek to unify multiple observed phenomena from these two categories under one analytic framework capable of simultaneously capturing standing and traveling oscillatory properties of the BOLD signal.

The goal of this study was to provide a parsimonious description of dominant global spatiotemporal patterns in spontaneous BOLD activity to provide insight into the functional architecture of the human brain. Utilizing methods and concepts from the study of geophysical phenomenon³¹, we provide a systematic description of the standing and traveling wave structure of these global spatial patterns. Utilizing a complex-valued extension of PCA, we identified three dominant spatiotemporal patterns in intrinsic BOLD signals consisting primarily of standing

waves with some traveling wave behavior. Consistent with our simulation study, the relative predominance of standing over traveling wave structure in these spatiotemporal patterns ensured that zero-lag FC methods effectively captured much of their temporal and spatial structure. This finding effectively explains previous observations that traveling waves of BOLD activity resemble patterns of large-scale FC topographies^{22–24}, such as that seen in FC gradients³. Rather than large-scale FC topographies emerging from the dynamics of traveling waves, our study suggests that the results from zero-lag FC and time-lag analyses are different representations of the same underlying spatiotemporal pattern. While the results from time-lag analysis, such as CPCA, accurately capture the time-lag structure produced by the traveling wave oscillations, the predominant standing wave oscillations of these spatiotemporal patterns are effectively captured by zero-lag FC analyses.

An important finding from this study was the ubiquitous presence of these three spatiotemporal patterns across a wide variety of zero-lag FC techniques. Our survey included several different analytic techniques, with markedly different mathematical assumptions, including latent dimension-reduction techniques, CAP analysis, seed-based correlation analysis, and latent brain state models. Two common features across these methods are worth noting. First, several zero-lag FC methods surveyed involved an initial PCA dimension-reduction step (varimax, temporal ICA, spatial ICA and HMM). Varimax and ICA dimension-reduction methods can be considered a rotation of the original PCA axes towards an orientation that maximizes a specific criteria - simple structure for varimax⁵⁰ and statistical independence in the time or spatial domain for ICA⁵¹. The similarity of the FC patterns derived from varimax and ICA suggest that the optimization of that criteria in our low-dimensional PCA solution (number of PCs=3) yielded solutions not ‘far’ from their original orientation. Second, both the Laplacian Eigenmap (LE) algorithm and PCA consist of an eigendecomposition of a similarity matrix (for PCA, the covariance/correlation matrix, for LE, the graph Laplacian). In fact, LE can be considered a variant of kernel PCA⁴⁹, an extension of PCA to arbitrary kernel/similarity matrices. In addition to these methods, HMM estimation of latent brain states was preceded by an initial PCA dimension reduction step for computational feasibility and to avoid overfitting. Thus, one might suggest that the common patterns resulting from application of zero-lag FC methods are due to a common reliance on PCA as a step in their algorithms. However, there are an ample number of applications in simulated and empirical data demonstrating that the similarity between the output of these methods is not guaranteed, but is dependent on the area of application^{31,52,53}. Further, zero-lag FC methods that do not rely on PCA - for example CAP analysis and seed-based correlation - yielded the same patterns. In addition, time-lag analyses with no relation to PCA (QPP-algorithm and lag projection algorithm) yielded the same patterns.

One property that distinguishes CAP analysis and the QPP-algorithm from other methods surveyed is their modeling of local (in time) changes in BOLD amplitude fluctuations; the rest of the methods model second- or higher-order statistics of BOLD fluctuations across the entire scan

(or a group of scans concatenated across time). CAP analysis detects prominent recurring patterns of BOLD activity through a clustering algorithm of suprathreshold time points identified from a seed time course. The QPP-algorithm detects a single prominent (temporally-extended) window of BOLD activity through a brute force matching and averaging procedure. Importantly, both these analyses identify local recurring changes in whole-brain amplitude fluctuations - single time-points for CAP analysis, and multiple consecutive time-points for the QPP algorithm. The identification of CAP analysis outputs at multiple seed locations (Figure 3; **Supplementary Figure 7**) with the three spatiotemporal patterns, and the output of the QPP algorithm with pattern one and two, suggests that the three spatiotemporal patterns are representations of local (in time) events in spontaneous BOLD fluctuations. This is consistent with previous observations using CAP analysis that the whole-brain correlation matrix computed across an entire scan can be reconstructed from a small number of high-amplitude (supra-threshold) BOLD activity events⁹.

Another important finding of this study is the identification of the brain's average latency structure (**Figure 4**) with pattern one. Mitra et al.¹⁸ first discovered this latency structure by a procedure of column-wise averaging pairwise time delays between brain regions. We found that an analogous procedure - column-wise (circular) averaging of pairwise phase delays between brain regions (derived from a complex correlation matrix) - yielded the same latency structure. Importantly, this latency structure was found to map directly to the phase delay pattern of pattern one. In other words, the average latency structure of spontaneous BOLD fluctuations simply reflects the traveling wave structure of the dominant spatiotemporal pattern (pattern one). Further, a repeated peak-averaging procedure demonstrated that this time-lag structure corresponds to local (in time) propagation patterns surrounding peaks of the global mean time course. This is consistent with previous findings of a global propagating wave event surrounding peaks of the global mean time course^{24,54}.

While time-varying or dynamic FC was not explicitly examined in this study, it is possible that the fluctuations between the spatiotemporal patterns identified in this study explain the variability in zero-lag FC structure over time. For example, a recently developed dynamic FC analysis, edge time series analysis⁷⁰, has found patterns of co-fluctuations in zero-lag FC that resemble the spatial structure of the three spatiotemporal patterns⁷¹. Future studies will be needed to determine the correspondence between variance in zero-lag FC over time and these spatiotemporal patterns.

As the primary aim of our study was descriptive, we have avoided any explanatory or causal explanation of the neuronal or non-neuronal origins of these three spatiotemporal patterns. However, the identification of pattern one with the global mean time course suggests potential candidate causal mechanisms. One promising candidate is a systemic circulatory origin⁵⁵. Work by Tong et al.^{25,56} has established that a significant portion of variance in low-

frequency intrinsic BOLD signals is correlated with systemic oxygenation fluctuations in the periphery (hands and toes) measured by concurrent near-infrared spectroscopy. This low-frequency peripheral oscillation tracks the global mean fMRI time course⁵⁴ and exhibits significant traveling wave behavior induced by differential blood transit time in the cerebral vasculature⁵⁷. Our observation that the time course of pattern one is virtually identical with the global mean time course suggests that this spatiotemporal pattern may originate in part from the systemic circulatory system. Potential sources of this systemic circulatory effect observed across the body include vasomotion, fluctuations in arterial CO₂ and/or Mayer waves⁵⁵. However, other studies support a neural origin for the global signal^{58,59}. We recently demonstrated that individual differences in global signal topography in the human brain are related to life outcomes and psychological function³⁷. This work follows studies using pharmacological intervention⁵⁹ and electrophysiological recordings in macaques⁵⁸ providing evidence for neural origins of the global signal. Momentary fluctuations in arousal and/or vigilance are known to be related to the global BOLD signal¹⁻³, and traveling wave signals^{10,13}. Some physiological processes strongly co-vary with cortical excitability^{60,61}, making it even more difficult to disentangle the neural vs. nonneural sources of the BOLD fluctuations. Taken together, these findings suggest that the global signal (and by extension pattern one), while influenced by vascular and other global phenomena, is at least partially driven by neuronal sources. The current study does not resolve these existing questions about the sources of the BOLD fluctuations, which continue to be actively investigated by the neuroimaging community. The current findings underscore the importance of understanding vascular components of the global signal in order to effectively denoise fMRI data while preserving neuronal signals⁶².

Such findings only provide candidate explanations for the origin of pattern one. Candidate explanations for the origin of pattern two and three are more difficult to identify from previous literature. Our study found that pattern two corresponds precisely to the task-negative/task-positive pattern originally discovered in resting-state by Fox et al.¹ and Fransson⁴⁸ (also see Greicius et al.⁶³). A similar anti-correlated pattern is consistently observed in response to task stimuli⁶⁴. Comparison of pattern two with pattern one revealed that they share similar patterns of BOLD activity. Further, the similar time-scales and dynamics of patterns two and three with pattern one suggests they may share a similar origin. Our analysis does not allow a characterization of the precise nature of this relationship. Some have argued that the task-positive/task-negative pattern arises from the application of global signal regression to BOLD time courses⁶⁵. It has been mathematically proven that global signal regression centers the distribution of pairwise correlations between brain regions and produces positive and negative correlation values⁶⁵. However, the appearance of all three patterns as distinct components across a wide variety of analytic methods (in data with no global signal regression) suggests that patterns two and three may be independent of pattern one. Pattern 2 strongly resembles the QPP obtained after global signal regression, which has been linked to infraslow electrical activity^{20,21,66}, and all patterns give rise to functional connectivity, which has been tied to infraslow and

higher frequency electrical activity⁶⁷. Patterns of propagating activity have also been observed in optical fluorescence imaging⁶⁸ and electrophysiological recordings⁶⁹, which demonstrates that time-lagged relationships can arise from neural as well as hemodynamic processes. 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.

Along similar lines, while these three spatiotemporal patterns explain a significant portion of variance in spontaneous BOLD fluctuations (~32%), there is a considerable portion of variance left unexplained. The exact value of this metric will depend on a variety of factors, including preprocessing/de-noising steps, parcellation size, scanner noise and intersubject variability. Despite this, it is clear there is much spatiotemporal dynamics in spontaneous BOLD fluctuations left unexplained by these three spatiotemporal patterns. Nevertheless, the three spatiotemporal patterns account for most ($r = 0.77$) of the functional connectivity observed. We interpret these reports as supporting our finding that functional connectivity is largely dictated by a few whole-brain spatiotemporal patterns of activity. The lesser variance explained in the individual time courses then suggests that there is an additional component that contributes substantially to the low frequency BOLD fluctuations, but is not reliably coordinated across areas, and therefore does not contribute as strongly to functional connectivity. The source of this signal component remains unclear. Some portion of it is likely to arise from non-neural noise, but there is an intriguing possibility that it also contains small changes in activity related to ongoing cognition. To speculate on the residual variance left unexplained by these three spatiotemporal patterns, we suggest that it may reflect more spatially focal variance in spontaneous BOLD fluctuations. This spatially local structure would be more accurately modeled by fine-grained, parcellation-based analysis approaches¹. In fact, this spatially local structure may be more related to neurovascular coupling than the global patterns observed in this study.

It is important to qualify the assumptions on which our comparisons were based. First, these spatiotemporal patterns are low-dimensional representations that capture global, or cortex-wide patterns of intrinsic BOLD activity. The primary metric that determined the number of latent dimensions to extract from the data was explained variance. Specifically, the number three was chosen based on the diminishing explained variance observed from the extraction of increasingly higher-dimensional solutions. This approach identified three latent dimensions, corresponding to three global, cortex-wide spatiotemporal patterns. Explained variance is not the only metric that may be used to judge the adequacy of a low-dimensional solution, nor is it desirable in all analysis contexts. However, at the level of a 3-dimensional solution, our survey has found remarkable consistency in the outputs of zero-lag FC analyses that yield solutions based on the optimization of criteria other than that of explained variance. Thus, while other approaches may identify different optimal numbers of latent dimensions, we found the explained variance approach highly desirable in terms of robustness. In a related manner, 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). Empirical examination of consistency in zero-lag FC analyses at solutions greater than three components confirms this expectation: higher dimensional solutions yield less consistency in FC topographies (**Supplementary Figure 12**). Despite this qualification, the consistency of analysis approaches at our low-dimensional level of analysis suggests that these three spatiotemporal patterns are robust phenomena in need of explanation.

Our findings underscore the benefits of capturing both zero-lag and time-lagged information for studies of the brain's intrinsic complex functional architecture. CPCA unites both analytical approaches within a single framework, and provides interpretable information about the relative contributions of standing and traveling waves of activity. With this approach we identified three spatiotemporal patterns that recapitulate the major findings from a wide range of analytical techniques while providing a unified and parsimonious description of zero-lag and time-lag structure of spontaneous BOLD fluctuations. Further, they account for much of the structure that underlies all FC based analysis, and therefore impact how we interpret everything from functional networks to graphs. As the study of the brain as a complex system advances, these patterns have potential applications in better constraining generative models of brain activity, predicting variability in response to external stimuli, informing targeted modulation of brain activity to achieve a particular state, and understanding interactions between the state of the brain as a whole and localized activity in neuronal circuits. The concise characterization of systems-level brain activity in three spatiotemporal patterns will facilitate the cross-scale research needed to link fundamental neuroscience studies and human behavior.

Supplementary Materials

Supplementary Modeling Note - CPCA on Simulated Data

Simulation Approach and Methodology

To understand the behavior of zero-lag FC analyses in the presence of traveling wave oscillations, we conducted a simulation of traveling Gaussian waves on a two-dimensional grid. The grid was constructed as a 100-by-50 two-dimensional array of pixels. To operationalize the notion of standing and traveling waves we use the mathematical formalization of complex sinusoidal wave motion developed in Feeny⁴². A complex wave motion can be represented by the following function:

$$\mathbf{x}(t) = 2e^{\gamma t} [\cos(\omega t)\mathbf{c} - \sin(\omega t)\mathbf{d}]$$

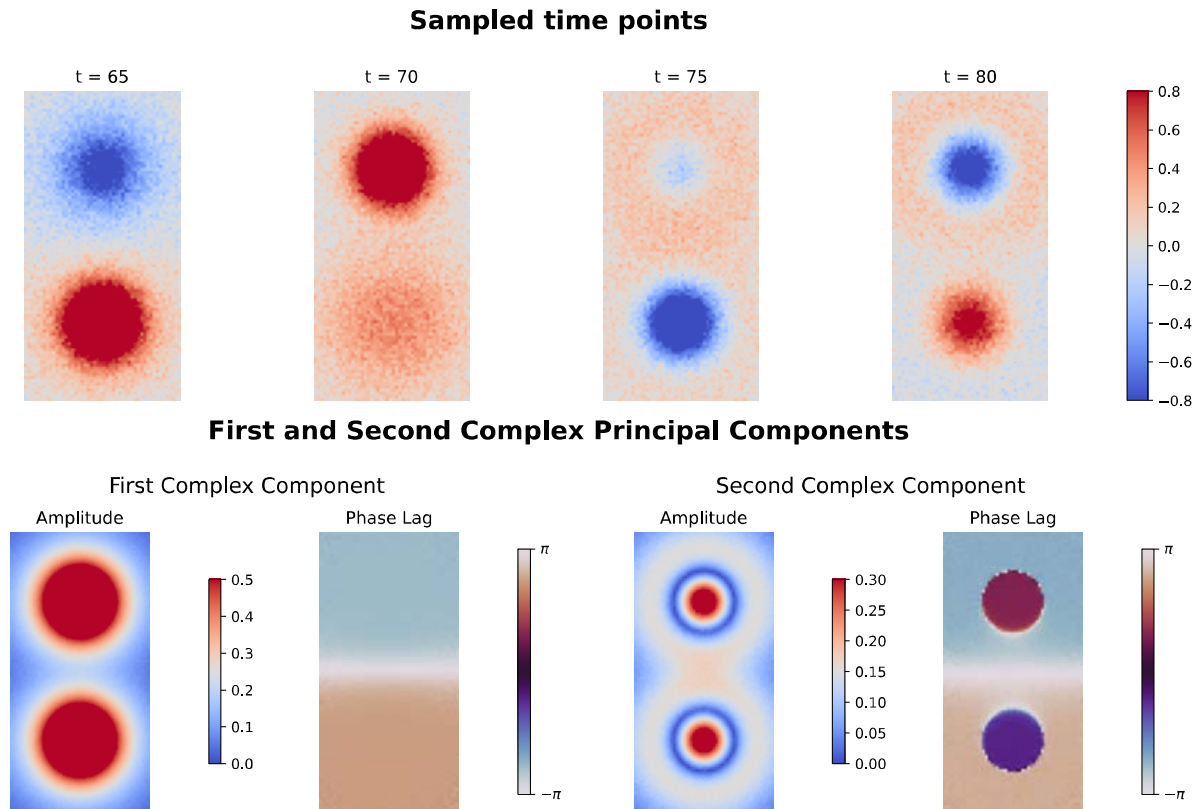
where t is the discrete-sampled time array, γ is an exponential decay term, ω is the angular frequency, and \mathbf{c} and \mathbf{d} are two vectors representing the two modes of the oscillation. The output of the function, \mathbf{x} , is a vector formed from the mixture of the two modes, \mathbf{c} and \mathbf{d} , weighted by a cosine and sine term (modulated by time, t). At the peak of the cosine oscillation, the sine oscillation is at zero, and the pattern encoded in vector \mathbf{c} is the dominant configuration. At the peak of the sine oscillation, the cosine oscillation is at zero, and the pattern encoded in vector \mathbf{d} is the dominant configuration. As a function of time, the output appears visually as cyclical ‘wave’ motion between the patterns defined by vectors \mathbf{c} and \mathbf{d} . These patterns can take any form or shape. For our simulation, the vectors represented flattened (i.e. converted to one-dimension) two-dimensional Gaussian curves with varying peak location. More specifically, \mathbf{d} is a two-dimensional Gaussian curve ($\sigma = 1$) with its peak at one position in the 100-by-50 grid, and \mathbf{c} is another two-dimensional Gaussian curve ($\sigma = 1$) with its peak lower down the 100-by-50 grid in the vertical direction. The simulation appears as a single two-dimensional Gaussian wave traveling from the Gaussian represented in vector \mathbf{c} upwards to the Gaussian represented in vector \mathbf{d} . The degree of traveling wave behavior of the complex motion is determined by the degree of statistical independence between the two Gaussian curve patterns. For example, for completely overlapping Gaussian curves (i.e. identical peak locations), the patterns are entirely dependent, which will appear as a pure standing oscillation. For completely non-overlapping Gaussian curves, the patterns are entirely independent, which will appear as a pure traveling wave oscillation.

To assess the behavior of zero-lag FC methods applied to oscillations with arbitrary mixtures of standing and traveling waves, we systematically vary the distance between the two Gaussian curves (vectors \mathbf{c} and \mathbf{d}) by moving the peak location of the starting Gaussian curve (vector \mathbf{c}) from the bottom of the 100-by-50 grid (no overlap) to the top of the grid (complete overlap with vector \mathbf{d}). In **Figure 1A**, we display the output of zero-lag FC analysis and CPCA at large (vertical peak distance of 50 pixels), moderate (peak distance of 25 pixels) and zero distances between the peaks of the two Gaussian curves. As noted above, large peak distance corresponds to approximately statistically independent Gaussian curve patterns (vectors \mathbf{c} and \mathbf{d}), producing a pure traveling Gaussian wave. Zero peak distance corresponds to complete statistically dependent Gaussian curve patterns, producing a pure standing Gaussian wave. A simulated oscillation is run for each peak distance with the following parameters ($\omega = 1$, $\gamma = -1$,

number of cycles = 40, number of samples = 800). For each simulation, a limited amount of Gaussian noise (variance = 0.2) is added to each pixel time course.

1) Simulation of Two-Mode Traveling Wave Oscillations

The simulation illustrated in **Figure 1** consisted of a single traveling wave oscillation of two gaussian curves. To assess the ability of CPCA to separate multiple, overlapping traveling wave oscillations, we slightly modify the simulation presented in **Figure 1** (see ‘Methods and Materials’). In addition to the existing traveling wave oscillation between the two gaussian curves on the two-dimensional grid, we embed a higher-frequency ($\omega = 2$; see ‘Methods and Materials’) traveling wave oscillation of two smaller gaussian curves at identical peak locations. For this simulation, the two gaussian curves of each oscillation are completely separated, or statistically independent, thereby creating a pure traveling wave between the gaussian curves of each simulation. Identical parameters to simulation in the main text were used, excluding the angular frequency (ω ; see ‘Methods and Materials’). This two mode simulation appears as overlapping traveling wave oscillations between two lower frequency, large gaussian curves, and two higher frequency, small gaussian curves (**Supplementary Model Figure 1**). We applied CPCA to the pixel time courses to assess its ability to separate these oscillations into two latent factors.



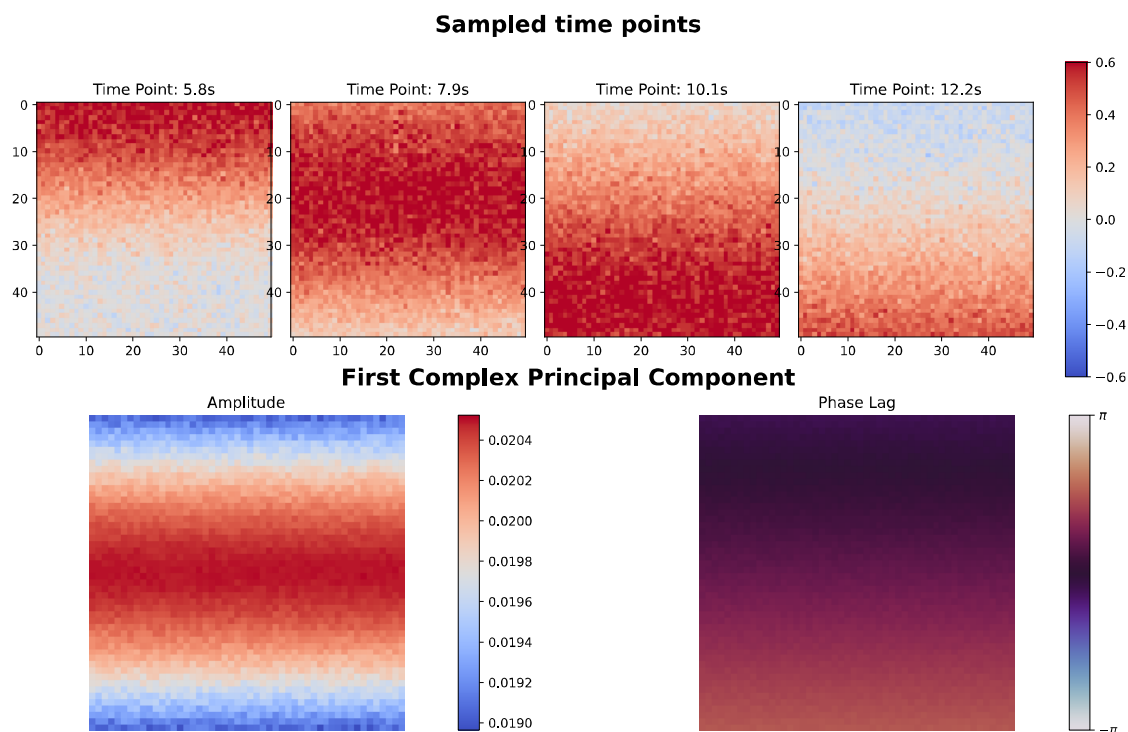
Supplementary Model Figure 1. **Simulation and Analysis of Two-Mode Traveling Wave Oscillations.** Simulation of two spatially overlapping, traveling wave oscillations. The simulation consists of a high frequency oscillation between two smaller gaussian curves, and a low frequency oscillation between larger gaussian curves. Four sampled time points are displayed in the top panel. The amplitude and phase delay maps of the first two complex principal components estimated by CPCA are displayed in the bottom panel. As can be observed from the amplitude and phase delay map, CPCA accurately separated the two traveling wave oscillations into separate components.

We extracted two complex principal components from the complex-valued pixel time courses. We plotted the amplitude and phase lag of each component, representing the strength of ‘participation’ and time-lag of the pixels in the complex component, respectively. As indicated by the amplitude maps, the spatial outline of the larger and smaller gaussian curves is accurately separated into the first and second complex principal components, respectively. Further, the phase maps of the two complex principal components accurately distinguish the higher frequency oscillations of the smaller gaussian curves from the lower frequency oscillations of the larger gaussian curves.

II) Simulation of Traveling Wave Hemodynamic Signals

To demonstrate the ability of complex principal component analysis (cPCA) to extract spatiotemporal patterns from more complex signals, we applied cPCA to propagating fields of

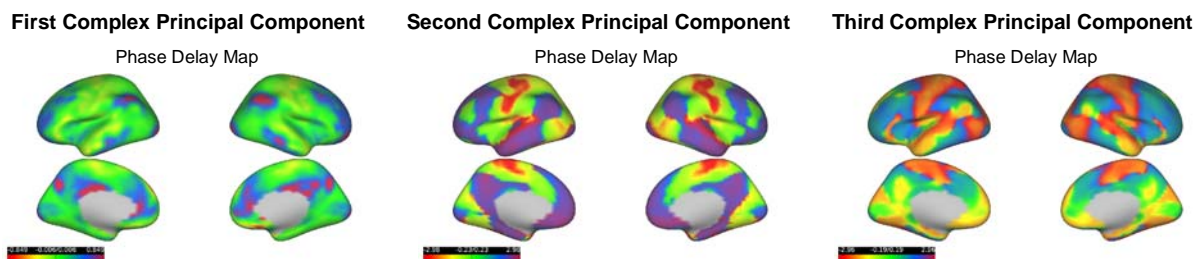
signals convolved with a hemodynamic response function. Impulse time series (1 for activation, 0 otherwise) convolved with the canonical hemodynamic response are spatially arranged along a square grid. The time series are arranged such that the time series in the top part of the square grid peaks early, 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. 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. 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 Model Figure 2**)



Supplementary Model Figure 2. **Global BOLD Propagation Simulation and Analysis.** **Top panel:** Four sampled time points 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. **Bottom Panel:** The amplitude and phase delay weights of the first complex principal component.

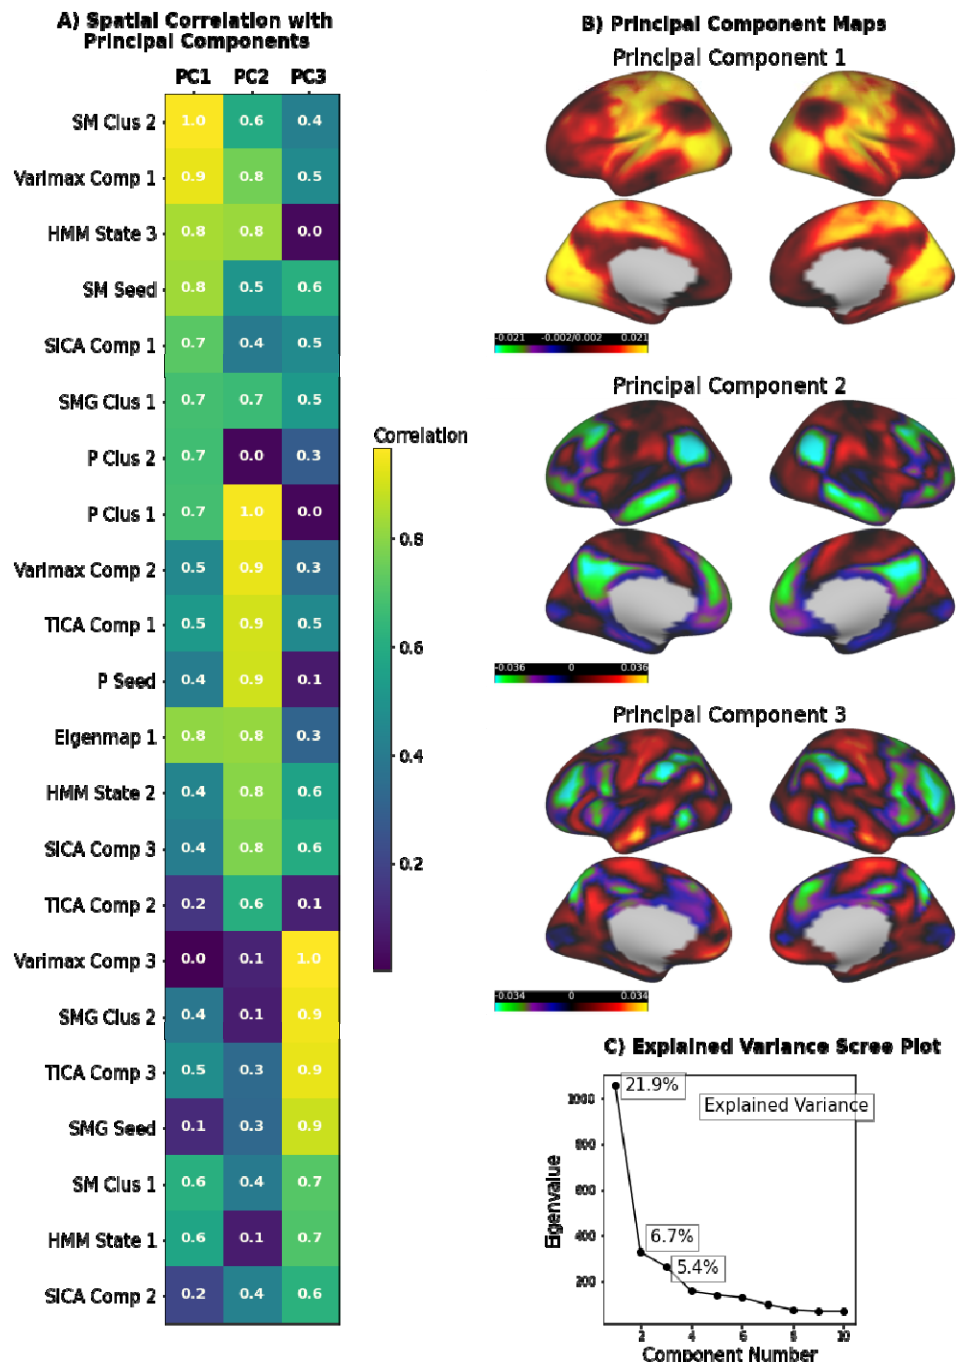
The sampled time points of the simulation (**Supplementary Model Figure 2**) illustrate the simulated global propagation event: peak BOLD amplitudes are first observed in the top of the square grid, followed by a subsequent propagation of peak BOLD amplitudes down the grid. Following the globally positive BOLD propagation event, there is a mirrored negative BOLD propagation event due to the post-response undershoot of the canonical hemodynamic response function. We extracted the first complex principal component from the complex time courses of the simulation. As illustrated in **Supplementary Model Figure 2**, the first complex principal component from CPCA accurately recovers the spatiotemporal pattern of the global propagation simulation. The phase delay map accurately describes the spatiotemporal pattern as BOLD activity that travels at a steady rate down the grid.

Supplementary Figure 1 - Replication of CPCA Phase Delay Maps in Independent Sample



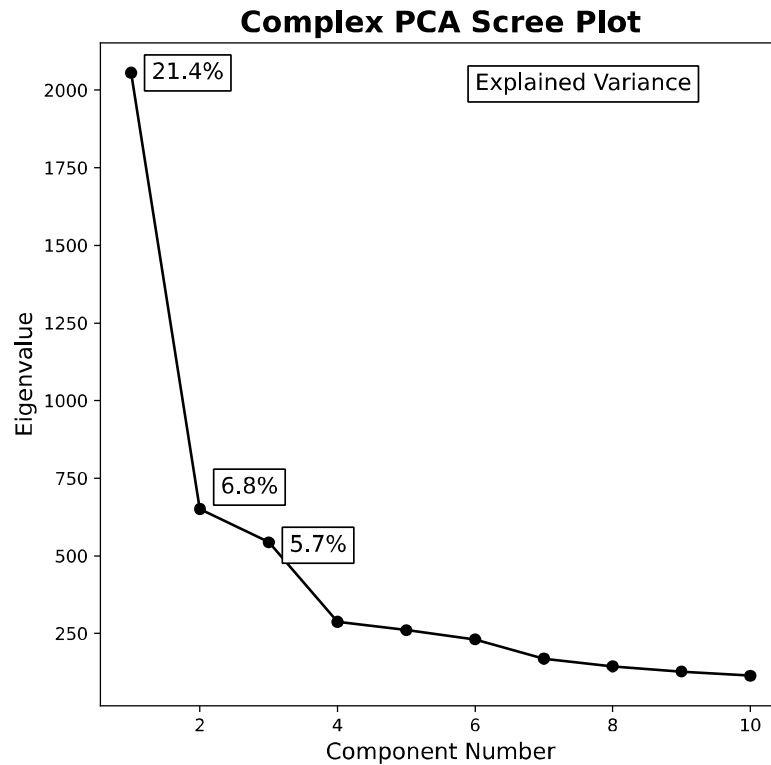
Supplementary Figure 1. Phase Delay Maps of Complex Principal Components in an Independent Sample. To give a sense of the robustness of our findings in an independent sample, we display the CPCA phase delay maps derived from the independent sample. Displayed are the phase delay maps of the first three complex principal components in the independent replication sample of 50 subjects. Phase delay values are displayed in radians.

Supplementary Figure 2 - Replication of Zero-lag FC Survey in Independent Sample



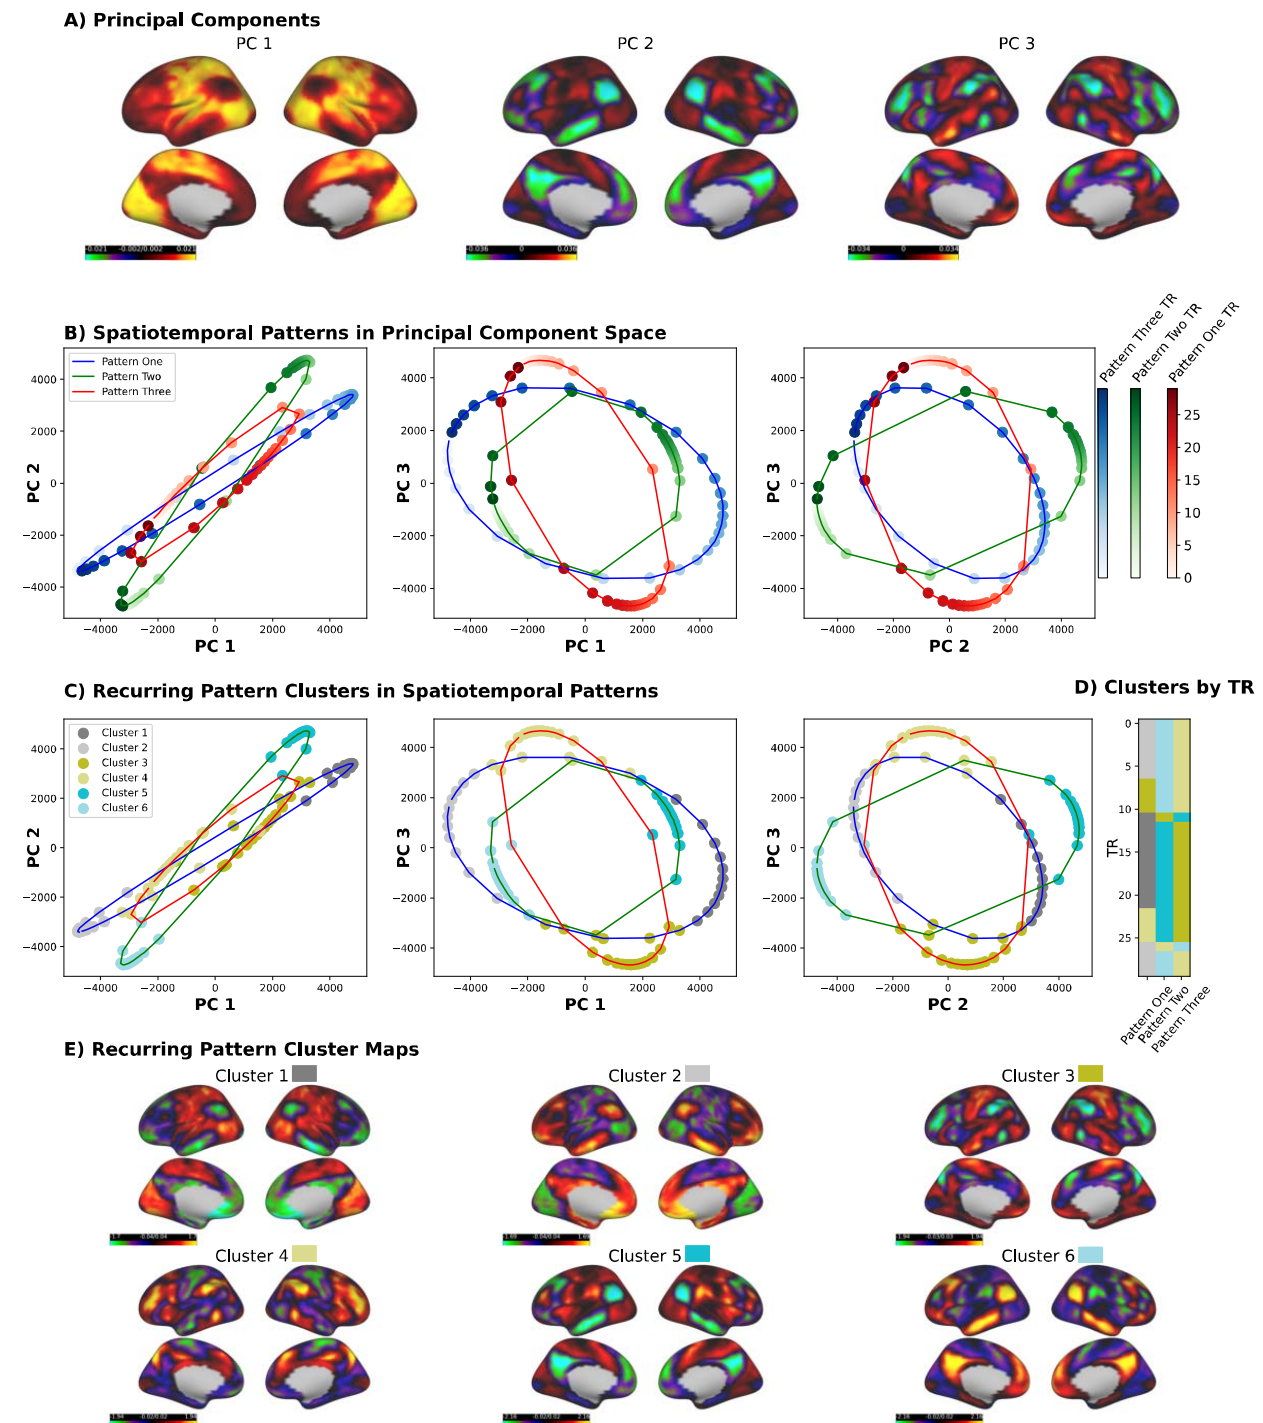
Supplementary Figure 2. **Zero-lag FC Survey in Independent Sample.** To give a sense of the robustness of our findings in an independent sample, we display zero-lag FC survey results (**Figure 3**) for the independent sample of 50 Subjects.

Supplementary Figure 3 - Complex Principal Component Scree Plot



Supplementary Figure 3. **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. 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).

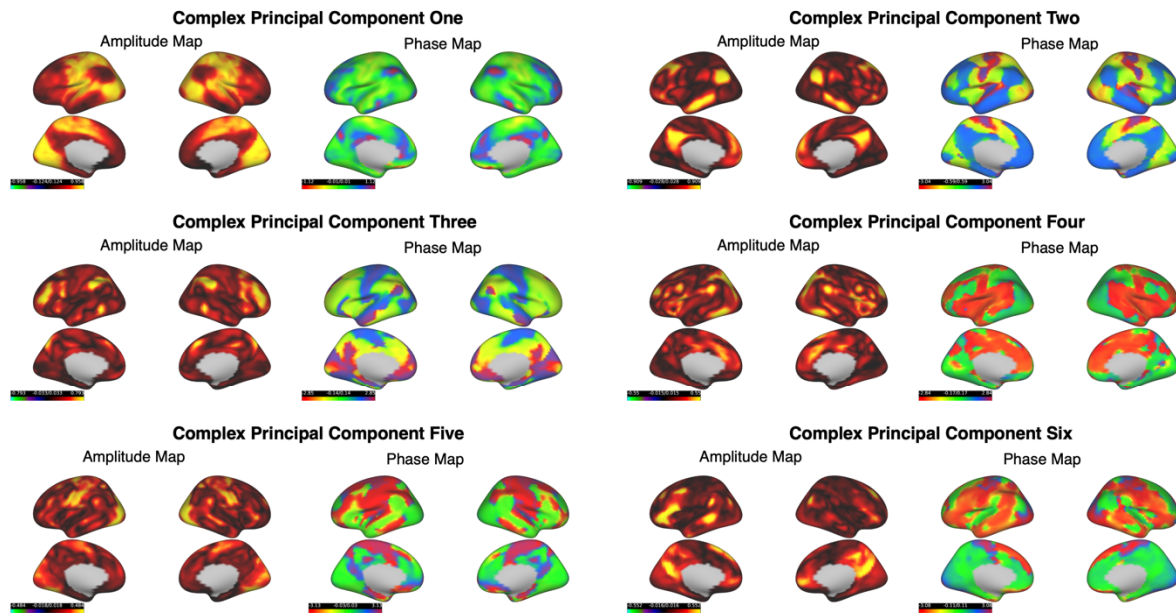
Supplementary Figure 4 - Steady States and Propagation Events in Spatiotemporal Patterns



Supplementary Figure 4. **Spatiotemporal Patterns Consist of Steady States and Propagation Events That Repeat Across Patterns.** (PC = Principal Component). To visualize the temporal dynamics of the three spatiotemporal patterns (patterns one, two and three), we projected their reconstructed time points (see ‘Methods and Materials’) into the 3-dimensional space formed by the first three principal components (**Panel A**) of the zero-lag FC matrix (**Panel B**). The time points are displayed on two-dimensional slices of each spatiotemporal pattern in the

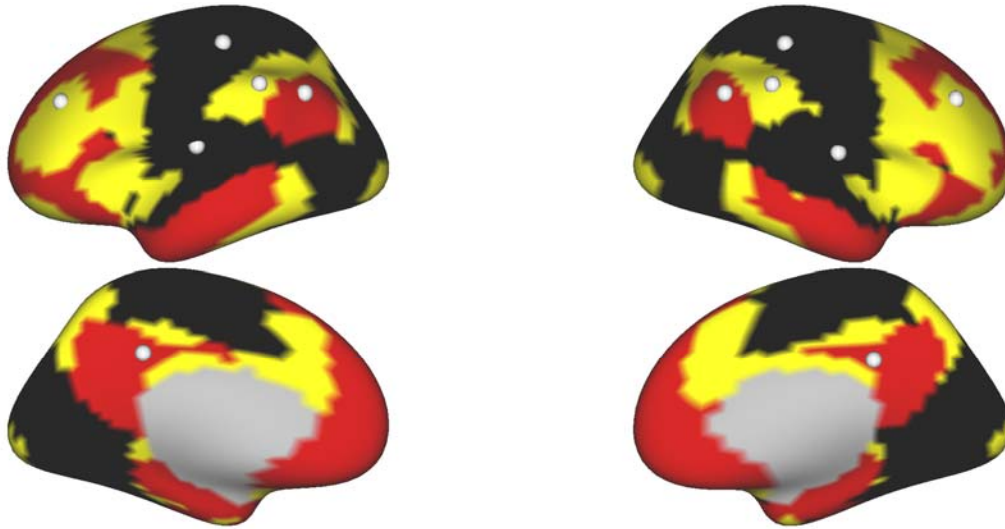
3-dimensional principal component space - PC1-PC2, PC1-PC3, and PC2-PC3 spaces. The time points of patterns one, two and three are displayed as blue, green and red points, respectively. Consecutive time points of each spatiotemporal pattern are linked by lines. The time points of each spatiotemporal pattern are colored from light to darker to visualize the progression of time (N=30). The temporal cycle of each spatiotemporal pattern forms an oval in the three-dimensional principal component space (**Panel B**), 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 pattern one, pattern two and pattern three vary strongest along the first, second and third principal component axes, respectively. 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. **C)** To examine repeating spatial patterns across the three spatiotemporal patterns, we applied a clustering algorithm (K-Means) to the reconstructed time points from all three spatiotemporal patterns. To avoid scaling differences in the distance calculations between time points, the BOLD values within each time point were z-score normalized. In **Panel C**, the same two-dimensional slices of each spatiotemporal pattern in the 3-dimensional principal component space are 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 pattern one, as z-score normalization of the time-points across vertices was performed beforehand. As revealed by the cluster solution, 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 pattern one. Clusters three and four correspond to the steady states of pattern three, and clusters five and six correspond to the steady states of pattern two.

Supplementary Figure 5 - Complex Principal Components Beyond Three



Supplementary Figure 5. **Amplitude and Phase Delay Maps of First Six Complex Principal Components.** While the application of the scree plot criterion identified three components as an optimal solution, there is still a large portion of variance left unexplained in BOLD time courses (~68%). To further explore this remaining variance, we present the phase delay and amplitude maps of the first six complex principal components. Phase delay values are displayed in radians.

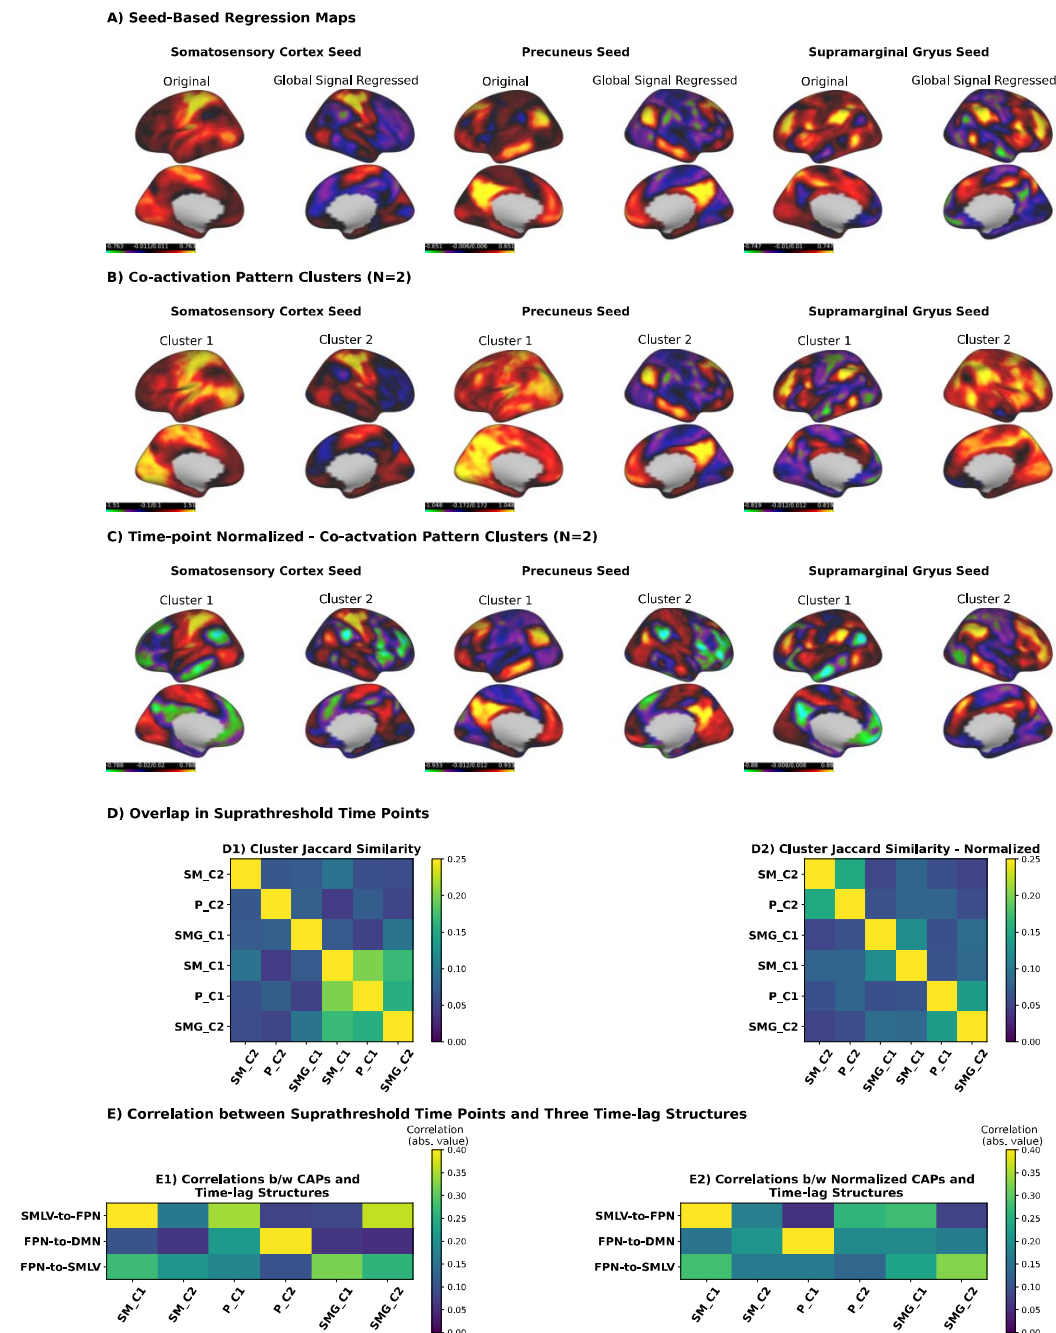
Supplementary Figure 6 - Location of Seed Regions in Seed-Based Regression and Co-Activation Pattern Analysis



Supplementary Figure 6. Location of Seed Regions in Seed-Based Regression and Co-Activation Pattern

Analysis. 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 - somatomotor (SM) and lateral visual (LV) networks, frontoparietal network (FPN) and DMN. 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. Six bi-lateral seed locations are displayed for the SMLV (black), FPN (yellow) and DMN (red). 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 Figure 7 - Seed-Based Analysis and Their Relationship to the Three Spatiotemporal Patterns.

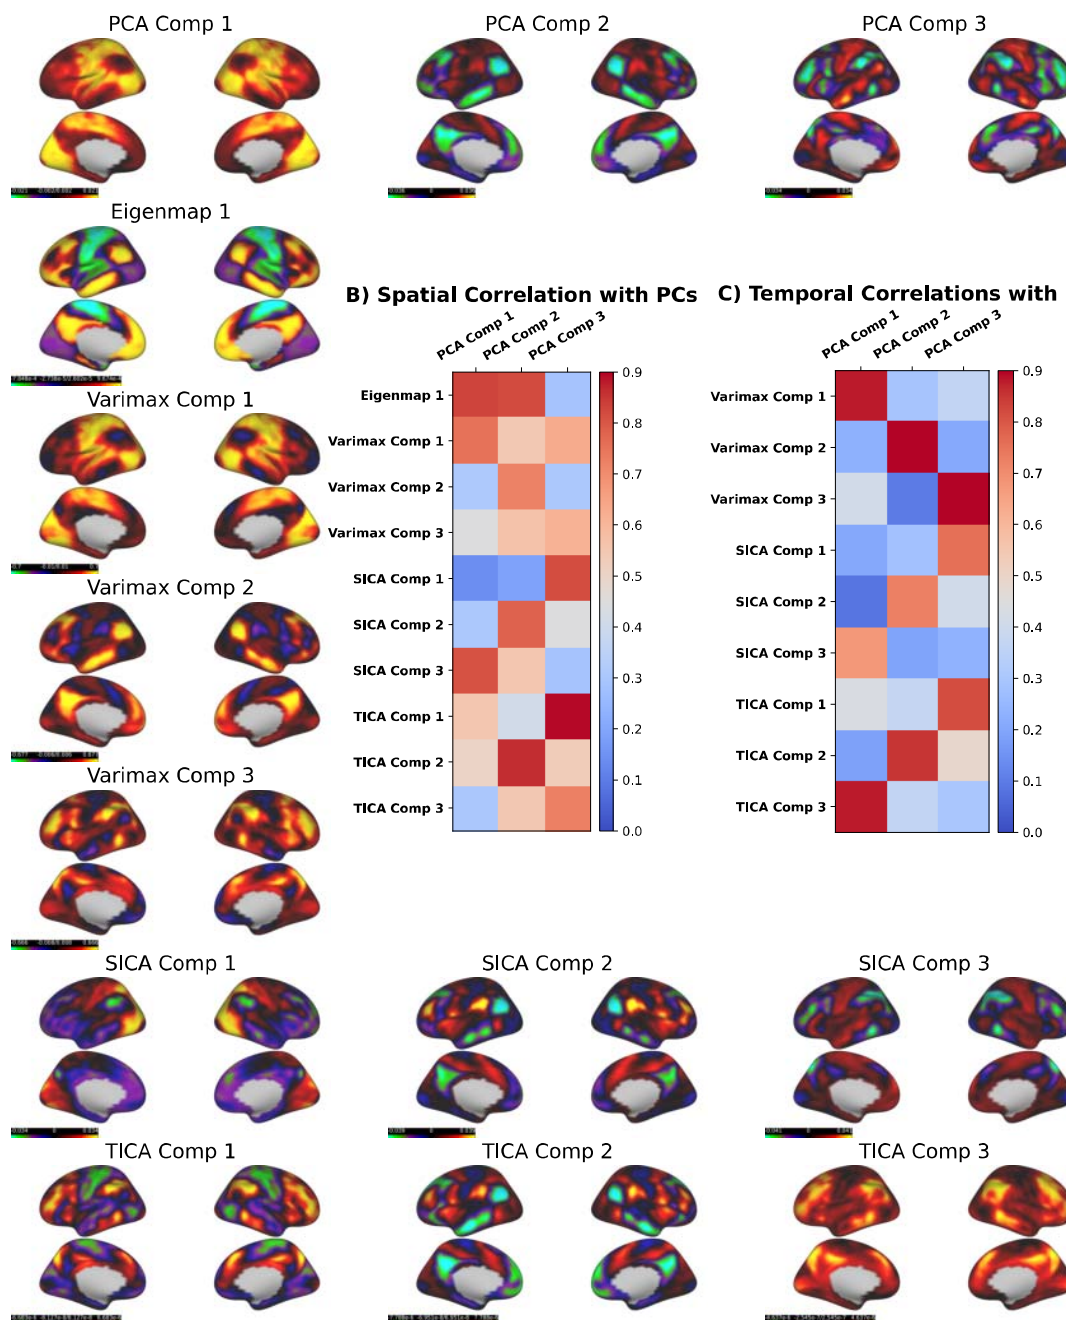


Supplementary Figure 7. **Seed-Based Regression and CAP Analysis.** (SM = somatosensory cortex; P=Precuneus; SMG=Supramarginal Gyrus). FC topographies of seed-based regression maps and CAP centroids from somatosensory (SM), precuneus and supramarginal gyrus seeds. **A)** Seed-based regression maps with (left hemisphere) and without global signal regression (right hemisphere) for SM, precuneus and supramarginal gyrus

seeds. 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 ($r = 0.92$). For CAP analysis, we chose a threshold equal to the 85th percentile of the seed time course BOLD values, consistent with previous applications⁹. **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. Normalization (**Panel C**) results in two anti-correlated CAPs per seed, as opposed to a globally-positive vs. anti-correlated CAP per seed (**Panel B**) in the non-normalized solution. **D1)** Temporal overlap between binary time courses 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). 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: = 0.173 > anti-correlated CAPs: = 0.073). This is consistent with the global synchronization effect associated with the global mean time course. **D2)** Temporal overlap between CAP binary time courses from the normalized solutions of each seed analysis. The temporal overlap observed between globally-positive CAP time points in the non-normalized solution disappears in the normalized solution. **E)** Temporal correlation between the time course of the three spatiotemporal patterns (patterns one, two and three) and the CAP binary time courses for the non-normalized (**E1**) and normalized (**E2**) solutions. We found that CAPs with globally-positive BOLD activation patterns (**E1**) somatosensory C1, precuneus C1 and supramarginal C2) were most strongly correlated with the time course of pattern one: somatosensory cortex (SM; $r = 0.49$), precuneus ($r = 0.35$), and supramarginal gyrus ($r = 0.37$). The temporal correlation between the pattern one and normalized CAPs across the three seeds is reduced (**E2**), excluding the somatosensory CAP (C1).

Supplementary Figure 8 - Dimension Reduction Functional Connectivity Topographies

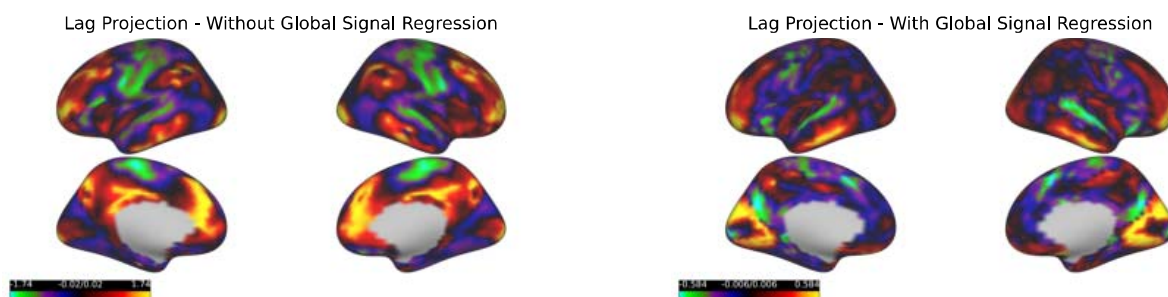
A) Low-Dimensional FC Topographies



Supplementary Figure 8. **Functional Connectivity Topographies of 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. 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^{31,72}. In other words, ICA rotates the original PCA solution to maximize a different criterion: statistical independence in the temporal or spatial domain. 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). We chose a popular simple structure rotation, Varimax rotation - an orthogonal rotation of the principal component loadings that maximizes simple structure. It is important to emphasize that the total cumulative variance explained by the principal component axes remains the same before and after rotation. 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. 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 Figure 9 - Global Signal Regression and Lag Projections

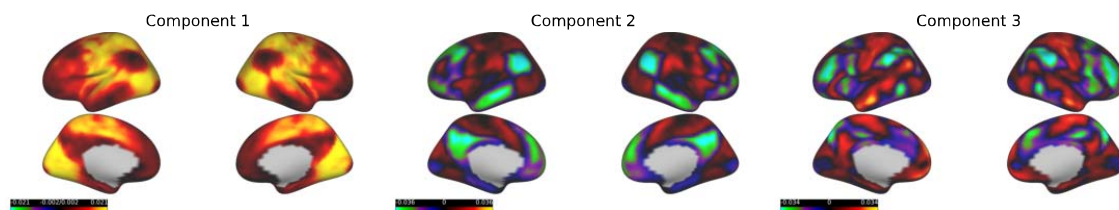


Supplementary Figure 9. Comparison of Lag Projections With and Without Global Signal Regression. The lag projection from our study partially resembles the group average lag projection observed in Mitra et al.¹⁸. However, our data differs in one important respect: Mitra et al.¹⁸ performed global signal regression as a preprocessing step. We display 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.¹⁸: 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. This is consistent with the observation by Mitra et al. that global signal regression reduces the range of observed latencies between BOLD time courses.

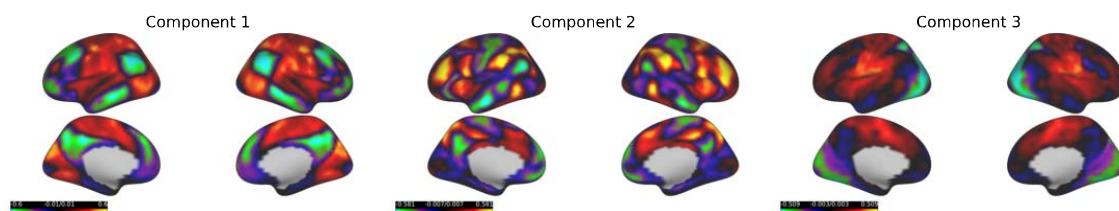
Supplementary Figure 10 - Anti-Correlated QPP and Pattern Two Visualization

Supplementary Movie 10. **Dynamic Visualization of the Anti-Correlated Quasiperiodic Pattern and Pattern Two.** The 30 time points (TR=0.72s) of the anti-correlated QPP and pattern two. The time index of each sequence is displayed in the top left. The time points of pattern two are equally-spaced phase samples (N=30) of the time point reconstruction (see 'Methods and Materials'). The time points of the anti-correlated QPP are derived from the spatiotemporal template computed from the repeated-template-averaging procedure on global signal regressed data.

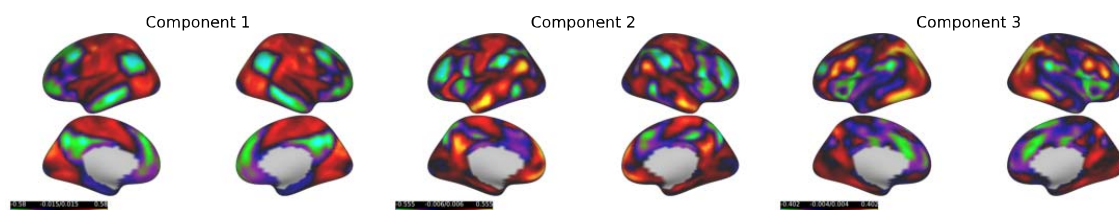
Principal Component Analysis



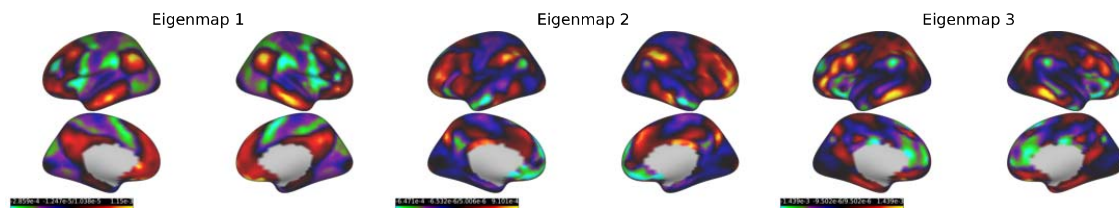
Principal Component Analysis - Global Signal Removed



Principal Component Analysis - Time-Point Centered



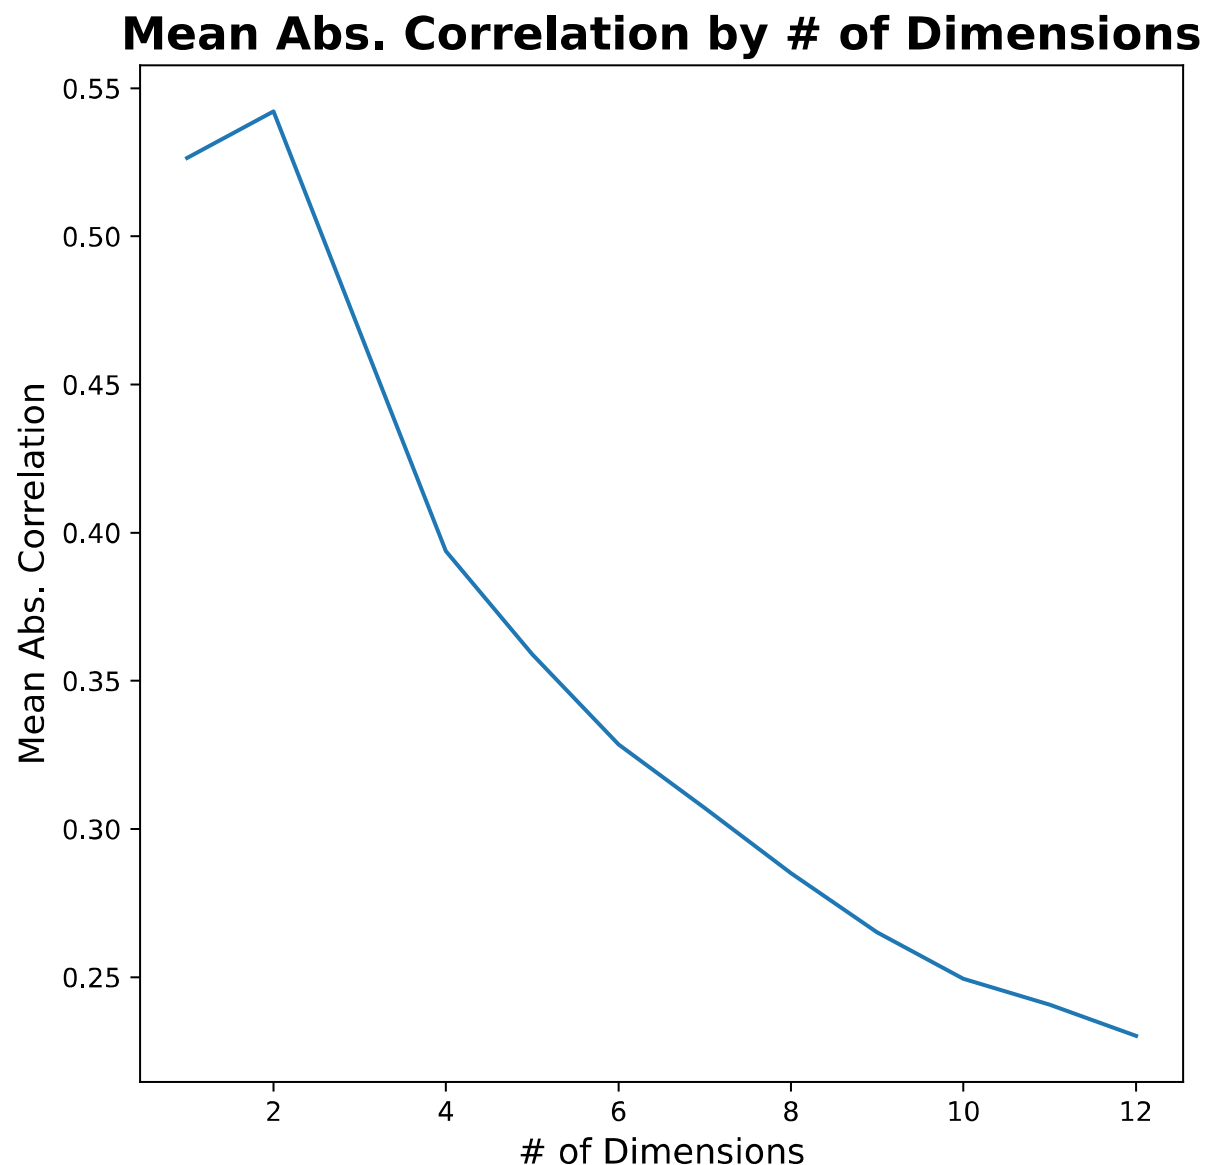
Laplacian Eigenmaps - Manifold Learning



to the principal functional connectivity gradient (PG)³. However, we note that the exact spatial pattern of the PG

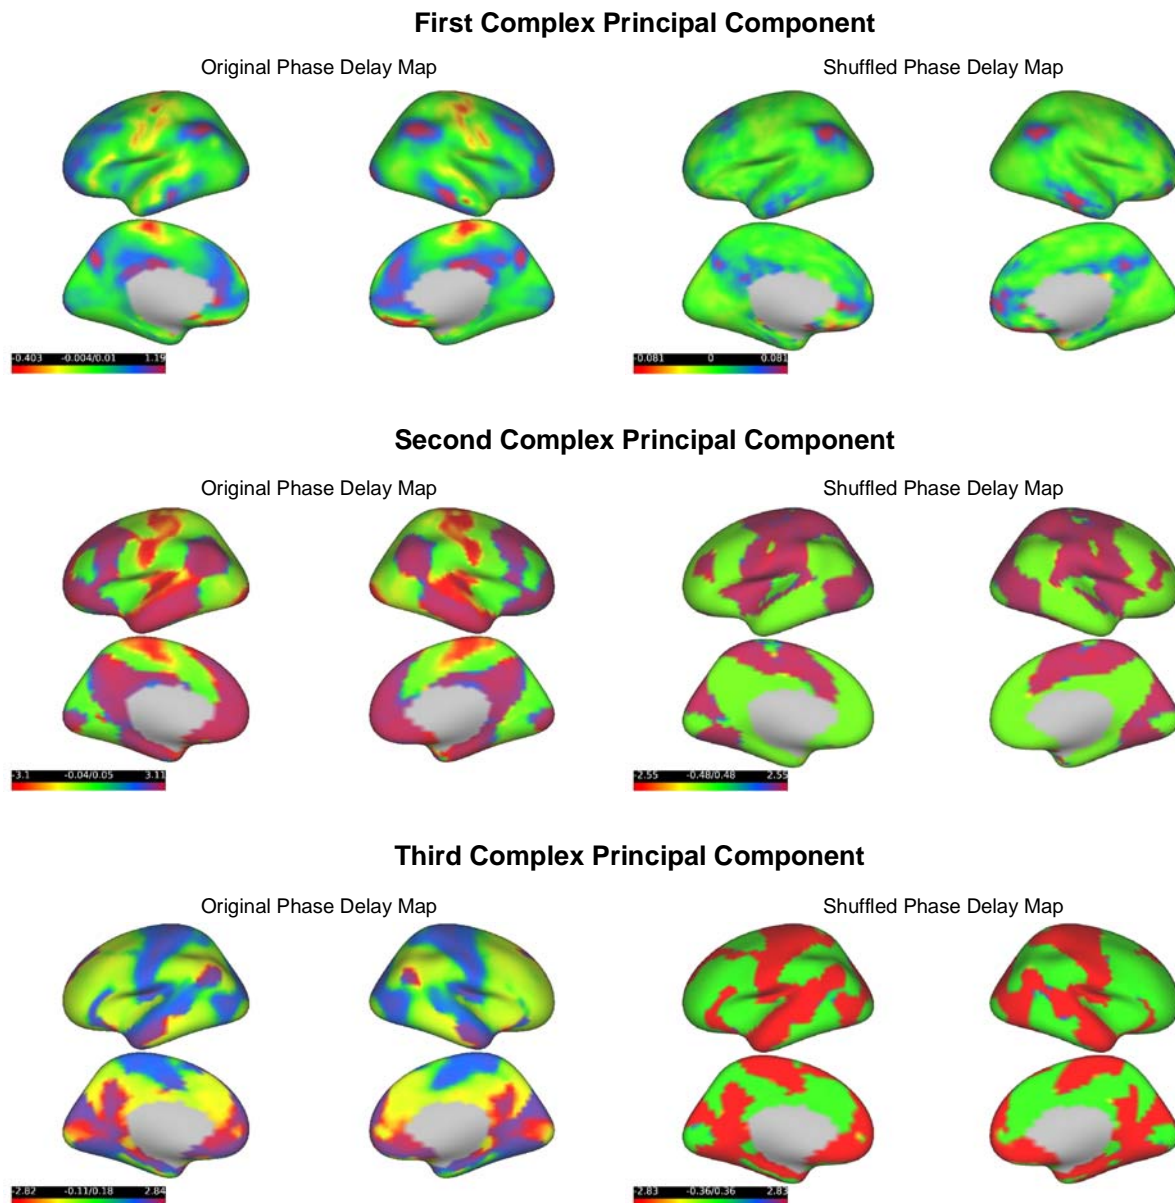
depends on the level of thresholding applied to the FC matrix (**Figure 5**). In this LE solution, no thresholding was applied. The first and second Laplacian Eigenmaps match the second and third principal components from the PCA solution, respectively (with an arbitrary sign-flip). The difference between the PCA and LE solution is that the first principal component seems to be missing from the LE solution. However, the first three components from the PCA on global-signal regressed data and time-point centered do match the three Laplacian Eigenmaps (with an arbitrary sign flip). These similarities between the spatial maps produced by PCA and Laplacian Eigenmaps have been previously observed by Vos de Waal et al.³⁶. 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, pattern one is precisely tracked by the global signal. Consistent with previous findings¹⁵, time-point centering and global signal regression have similar effects - reducing or eliminating the variance of the global signal time course. In the Laplacian Eigenmap solution, a time-point centering is not as explicit. Manifold learning of a time vertex-by-vertex kernel matrix operates on a mean centering of the feature space, i.e. each time point is centered. As illustrated, 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 patterns (some with an arbitrary sign difference).

Supplementary Figure 12 - Consistency in Zero-lag FC Topographies at Finer-Grained Solutions



Supplementary Figure 12. **Consistency in FC Topographies by Granularity of Solutions.** One of the findings of our study is that there is a surprising consistency in FC topographies estimated from ‘three-dimensional’ zero-lag analyses. However, it is worth exploring the degree of consistency in FC topographies at higher component numbers. We estimated higher number component (cluster) solutions for several latent dimension-reduction and clustering methods - including PCA, varimax-rotated PCA, temporal ICA, spatial ICA, HMM, and CAP analysis (all three seeds from main text). We examined the consistency in FC topographies at component numbers varying from one to 12. Note, the same number of components (clusters) was estimated across methods at each value. While it is difficult to quantify the degree of consistency in outputs from multiple methods with a single number, we used the mean absolute-valued correlation between all pairs of FC topographies produced by all methods as a rough estimate. We display the mean absolute-valued correlation by component number solution. As illustrated in the plot, the consistency of FC topographies drops significantly at higher component (cluster) numbers.

Supplementary Figure 13 - Complex PCA on Shuffled BOLD Time Courses

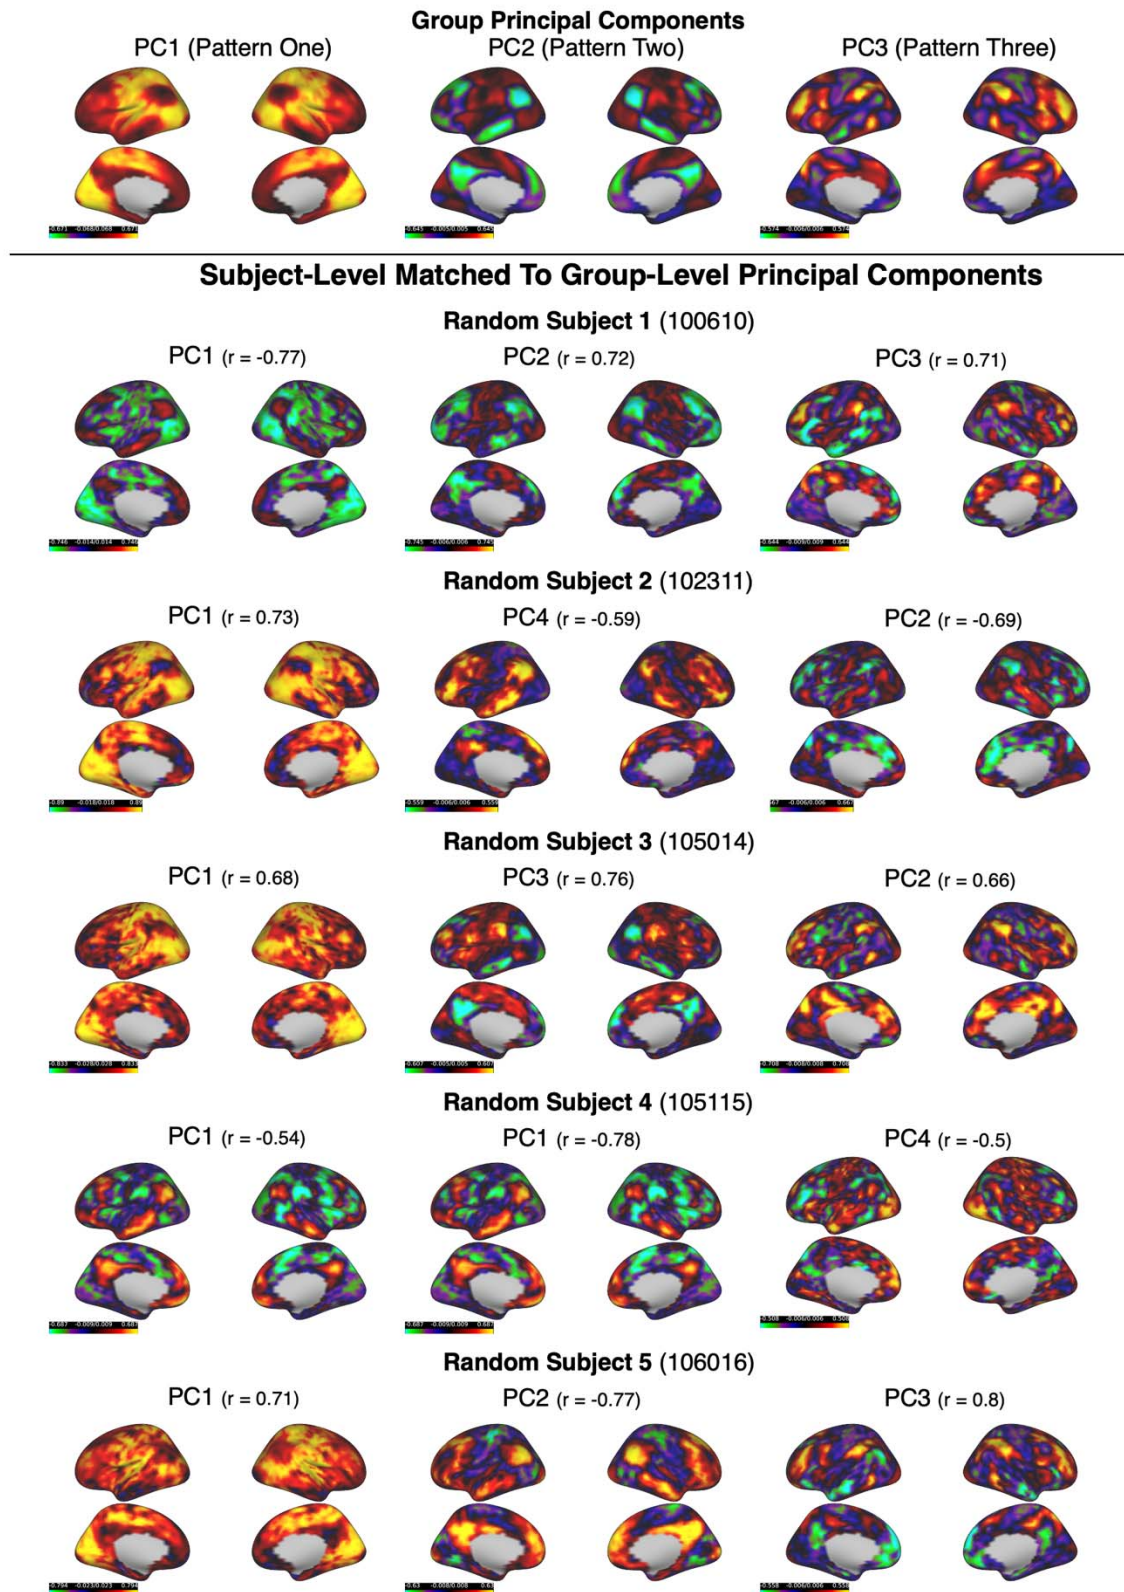


Supplementary Figure 13. **Complex PCA On Shuffled Time Courses.** One means to illustrate the time-lag information captured by CPCA, as opposed to zero-lag analyses, is through the use of a null model. Specifically, we can preserve the zero-lag correlation structure of spontaneous BOLD fluctuations and remove time-lag (autocorrelation) structure by randomly shuffling their time courses. We randomly shuffled the group-concatenated BOLD time courses and estimated three complex principal components from CPCA. We display the phase delay maps of the first three complex principal components from the shuffled time courses and original time courses (displayed in radians). Random shuffling of BOLD time courses effectively eliminates the time-lag structure present in the original phase delay maps of the first three complex principal components. What is preserved is the zero-lag structure, i.e. in-phase and anti-phase statistical dependence. Visually, the phase delay maps of the shuffled BOLD

time courses display a strictly bi-modal phase structure, corresponding to the anti-correlated fluctuations of the standing wave components of the complex principal components. The intermediate phase values present in the original complex principal components, corresponding to the traveling wave components, are eliminated.

Supplementary Figure 14 - Inter-Subject Variability in Principal Components

Principal Components at the Subject Level



Supplementary Figure 14. **Intersubject Variability of Spatiotemporal Patterns.** To assess the degree of intersubject variability in the three spatiotemporal patterns we randomly selected five subjects from our group analysis for further study. As illustrated in **Figure 3**, PCA effectively captures the zero-lag correlation structure of the three spatiotemporal patterns (patterns one, two and three). We conducted PCA on the BOLD time courses of the five subject scans separately, and compared the spatial structure of the first five principal components to the first three group-level principal components. As the ordering of the subject-level principal components by explained variance may not correspond precisely to the ordering of the group-level principal components, we conducted a matching procedure to map the subject-level principal components to the first three group-level principal components. This matching procedure was guided both by the spatial correlation between the subject-level and group-level principal components and by visual comparison. Beside each subject's header is their corresponding HCP ID. For each subject, their matching principal component was displayed in the same column as the group-level principal component (PC1: left, PC2: middle, PC3: right). Beside the header of each subject's principal component is the spatial correlation between that principal component and its matched group-level principal component (in the same column). Overall, patterns one, two and three were present in all five subjects, with notable intersubject variability. Several types of variability were observed: 1) the ordering by explained variance of the subject-level principal components was sometimes different than that observed at the group-level, 2) the same spatiotemporal patterns exhibited global and anti-correlated types across subjects, and 3) two subject-level principal component matched most closely to a single spatiotemporal pattern (for one subject - subject 4).

Methods and Materials

All code for preprocessing and analysis is provided at https://github.com/tsb46/BOLD_WAVES.

Resting-State fMRI Data

Our study utilized resting-state fMRI scans from the Human Connectome Project (HCP) S1200 release⁷⁴. Participants were unrelated, 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 (MSMall registered) that had been previously preprocessed with the HCP's ICA-based artifact removal process⁷⁵ 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 0.7mm isotropic. The resting-state fMRI data were at 2mm isotropic spatial resolution and with TR=0.72s temporal resolution. Further details of the data collection and preprocessing pipelines of the HCP can be found elsewhere^{74,75}. 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 Freesurfer⁷⁶. 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.

Complex Principal Component Analysis

To extract traveling wave patterns, 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 literature³⁰, or complex orthogonal decomposition in the engineering/physics literature⁴².

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 $t=\pi/2$ radians (i.e. 90 degrees), 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. CPCA applied to the complex-valued correlation matrix 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 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 delay 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. 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/fbpc>).

Estimating The Time-Scale of Complex Principal Components

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 (2 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 (2pi radians) would be approximately 12.6 TRs. Using this procedure, we found that the average duration of the first three complex principal components are ~28s (38.7 TRs), ~27s (37.4 TRs) and ~28s (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 (2pi radians) phase difference between two cortical BOLD time series would correspond to a ~28s time-lag between the time series. A smaller phase difference between two cortical BOLD time series, such as 1 radian, would correspond to ~14s time-lag between the time series, and so forth.

Temporal Reconstruction of Complex Principal Components

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 2pi 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.

Traveling Index of Complex Principal Components

The real and imaginary parts of a complex principal component correspond to the spatial weights of the component at zero and pi/2 (90 degree) phase shift of the original time courses. In a sense, they encode the temporal evolution of the complex principal component from one configuration (the real part) to a subsequent configuration (the imaginary part). By definition, a pure standing wave would exhibit the same spatial configuration from zero to pi/2 (90 degree) phase shifts of its cycle. A pure traveling wave would exhibit a different spatial configuration from zero to pi/2 (90 degree) phase shifts of its cycle. This observation suggests a means to quantify the degree of ‘traveling’ wave behavior of a complex principal component using the statistical dependence between its real and imaginary parts. A coherent traveling wave (i.e. propagation) of BOLD amplitudes across the cortex would exhibit one spatial configuration at one point in time and a different spatial configuration at another point in time. Thus, a complex

principal component that encodes this traveling wave behavior would exhibit differing spatial configurations in its real and imaginary spatial weights. Utilizing a metric developed by Feeny⁴², we define the ‘traveling’ index of a complex principal component as the reciprocal of the condition number of the matrix whose two columns are the real and imaginary parts of the complex principal component. This metric simply encodes the statistical dependence between the real and imaginary parts of the complex principal component. Pure traveling waves would exhibit completely orthogonal real and imaginary parts, and a traveling index of one. Pure standing waves would exhibit completely dependent real and imaginary parts, and a traveling index of zero.

Zero-Lag Functional Connectivity Analyses

Description of Zero-lag FC Analyses

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)⁷⁷. 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 literature³⁶. The input to the Laplacian eigenmaps algorithm was the vertex-by-vertex cosine similarity matrix³, representing the similarity in the 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 6**). 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 studies⁹, 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^{78,79}. 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 flattening (or diminishing return) in explained variance (i.e. eigenvalues) associated with neighboring principal components, a procedure known as Catell's scree plot test³⁹. 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 3C**). 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.

Quasiperiodic Pattern and Lag Projections

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)^{17,18} for the detection of lag/latency projections (~0-2s) and 2) a repeated-template-averaging algorithm of similar spatiotemporal segments¹⁹ for detection of the QPP (~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.¹⁹ 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^{19,47}, 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.

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).

Data Availability

Data from the Human Connectome Project is publicly available at <http://www.humanconnectomeproject.org/data/>

References

1. Biswal, B., Yetkin, F. Z., Haughton, V. M. & Hyde, J. S. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* **34**, 537–541 (1995).
2. Bijsterbosch, J. *et al.* Challenges and future directions for representations of functional brain organization. *Nature Neuroscience* 1–12 (2020) doi:10.1038/s41593-020-00726-z.
3. Margulies, D. S. *et al.* Situating the default-mode network along a principal gradient of macroscale cortical organization. *PNAS* **113**, 12574–12579 (2016).
4. Fox, M. D. *et al.* The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *PNAS* **102**, 9673–9678 (2005).
5. Smith, S. M. *et al.* Correspondence of the brain’s functional architecture during activation

- and rest. *PNAS* **106**, 13040–13045 (2009).
6. Yeo, B. T. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol* **106**, 1125–1165 (2011).
 7. Bassett, D. S. *et al.* Dynamic reconfiguration of human brain networks during learning. *PNAS* **108**, 7641–7646 (2011).
 8. Hutchison, R. M. *et al.* Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage* **80**, 360–378 (2013).
 9. Liu, X. & Duyn, J. H. Time-varying functional network information extracted from brief instances of spontaneous brain activity. *PNAS* **110**, 4392–4397 (2013).
 10. Smith, S. M. *et al.* Temporally-independent functional modes of spontaneous brain activity. *Proc Natl Acad Sci U S A* **109**, 3131–3136 (2012).
 11. Beckmann, C. F. & Smith, S. M. Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans Med Imaging* **23**, 137–152 (2004).
 12. Billings, J. C. W. *et al.* Instantaneous brain dynamics mapped to a continuous state space. *Neuroimage* **162**, 344–352 (2017).
 13. van den Heuvel, M. P. & Hulshoff Pol, H. E. Exploring the brain network: a review on resting-state fMRI functional connectivity. *Eur Neuropsychopharmacol* **20**, 519–534 (2010).
 14. Vidaurre, D., Smith, S. M. & Woolrich, M. W. Brain network dynamics are hierarchically organized in time. *Proc Natl Acad Sci USA* **114**, 12827–12832 (2017).
 15. Liu, T. T., Nalci, A. & Falahpour, M. The Global Signal in fMRI: Nuisance or Information? *Neuroimage* **150**, 213–229 (2017).
 16. Fox, M. D., Zhang, D., Snyder, A. Z. & Raichle, M. E. The Global Signal and Observed Anticorrelated Resting State Brain Networks. *J Neurophysiol* **101**, 3270–3283 (2009).

17. Mitra, A., Snyder, A. Z., Blazey, T. & Raichle, M. E. Lag threads organize the brain's intrinsic activity. *Proc Natl Acad Sci U S A* **112**, E2235–2244 (2015).
18. Mitra, A., Snyder, A. Z., Hacker, C. D. & Raichle, M. E. Lag structure in resting-state fMRI. *J Neurophysiol* **111**, 2374–2391 (2014).
19. Majeed, W. *et al.* Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans. *NeuroImage* **54**, 1140–1150 (2011).
20. Grooms, J. K. *et al.* Infralow Electroencephalographic and Dynamic Resting State Network Activity. *Brain Connect* **7**, 265–280 (2017).
21. Thompson, G. J., Pan, W.-J., Magnuson, M. E., Jaeger, D. & Keilholz, S. D. Quasi-periodic patterns (QPP): Large-scale dynamics in resting state fMRI that correlate with local infralow electrical activity. *NeuroImage* **84**, 1018–1031 (2014).
22. Raut, R. V. *et al.* Global waves synchronize the brain's functional systems with fluctuating arousal. *Sci Adv* **7**, eabf2709 (2021).
23. Yousefi, B. & Keilholz, S. Propagating patterns of intrinsic activity along macroscale gradients coordinate functional connections across the whole brain. *NeuroImage* **231**, 117827 (2021).
24. Gu, Y. *et al.* Brain Activity Fluctuations Propagate as Waves Traversing the Cortical Hierarchy. *Cereb Cortex* **31**, 3986–4005 (2021).
25. Tong, Y., Hocke, L. M., Licata, S. C. & deB. Frederick, B. Low-frequency oscillations measured in the periphery with near-infrared spectroscopy are strongly correlated with blood oxygen level-dependent functional magnetic resonance imaging signals. *J Biomed Opt* **17**, 106004 (2012).
26. Abbas, A. *et al.* Quasi-periodic patterns contribute to functional connectivity in the brain.

- NeuroImage* **191**, 193–204 (2019).
27. Roundy, P. E. & Frank, W. M. A Climatology of Waves in the Equatorial Region. *Journal of the Atmospheric Sciences* **61**, 2105–2132 (2004).
 28. Ghil, M. & Mo, K. Intraseasonal Oscillations in the Global Atmosphere. Part I: Northern Hemisphere and Tropics. *Journal of the Atmospheric Sciences* **48**, 752–779 (1991).
 29. Hayashi, Y. A Generalized Method of Resolving Transient Disturbances into Standing and Traveling Waves by Space-Time Spectral Analysis. *Journal of the Atmospheric Sciences* **36**, 1017–1029 (1979).
 30. Horel, J. D. Complex Principal Component Analysis: Theory and Examples. *J. Climate Appl. Meteor.* **23**, 1660–1673 (1984).
 31. Hannachi, A., Jolliffe, I. T. & Stephenson, D. B. Empirical orthogonal functions and related techniques in atmospheric science: A review. *International Journal of Climatology* **27**, 1119–1152 (2007).
 32. Madden, R. A. Observations of large-scale traveling Rossby waves. *Reviews of Geophysics* **17**, 1935–1949 (1979).
 33. Visbeck, M. H., Hurrell, J. W., Polvani, L. & Cullen, H. M. The North Atlantic Oscillation: Past, present, and future. *PNAS* **98**, 12876–12877 (2001).
 34. Hayashi, Y. On the Coherence between Progressive and Retrogressive Waves and a Partition of Space-Time Power Spectra into Standing and Traveling Parts. *Journal of Applied Meteorology and Climatology* **16**, 368–373 (1977).
 35. Watt-Meyer, O. & Kushner, P. J. Decomposition of Atmospheric Disturbances into Standing and Traveling Components, with Application to Northern Hemisphere Planetary Waves and Stratosphere–Troposphere Coupling. *Journal of the Atmospheric Sciences* **72**, 787–802

- (2015).
36. Vos de Wael, R. *et al.* BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. *Communications Biology* **3**, 1–10 (2020).
 37. Li, J. *et al.* Topography and behavioral relevance of the global signal in the human brain. *Scientific Reports* **9**, 14286 (2019).
 38. Carlson, J. M., Foti, D., Mujica-Parodi, L. R., Harmon-Jones, E. & Hajcak, G. Ventral striatal and medial prefrontal BOLD activation is correlated with reward-related electrocortical activity: A combined ERP and fMRI study. *NeuroImage* **57**, 1608–1616 (2011).
 39. Cattell, R. B. The description of personality: basic traits resolved into clusters. *The Journal of Abnormal and Social Psychology* **38**, 476–506 (1943).
 40. Ecker, C., Reynaud, E., Williams, S. C. & Brammer, M. J. 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 (2007).
 41. Stetter, M. *et al.* Principal Component Analysis and Blind Separation of Sources for Optical Imaging of Intrinsic Signals. *NeuroImage* **11**, 482–490 (2000).
 42. Feeny, B. F. A complex orthogonal decomposition for wave motion analysis. *Journal of Sound and Vibration* **310**, 77–90 (2008).
 43. Andersen, A. H., Gash, D. M. & Avison, M. J. Principal component analysis of the dynamic response measured by fMRI: a generalized linear systems framework. *Magn Reson Imaging* **17**, 795–815 (1999).
 44. Thomas, C. G., Harshman, R. A. & Menon, R. S. Noise reduction in BOLD-based fMRI using component analysis. *Neuroimage* **17**, 1521–1537 (2002).

45. Calhoun, V. D., Adali, T., Pearlson, G. D. & Pekar, J. J. A method for making group inferences from functional MRI data using independent component analysis. *Human Brain Mapping* **14**, 140–151 (2001).
46. Alexander-Bloch, A. F. *et al.* On testing for spatial correspondence between maps of human brain structure and function. *NeuroImage* **178**, 540–551 (2018).
47. Yousefi, B., Shin, J., Schumacher, E. H. & Keilholz, S. D. Quasi-periodic patterns of intrinsic brain activity in individuals and their relationship to global signal. *Neuroimage* **167**, 297–308 (2018).
48. Fransson, P. 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 (2005).
49. Ham, J., Lee, D. D., Mika, S. & Schölkopf, B. A kernel view of the dimensionality reduction of manifolds. in *Proceedings of the twenty-first international conference on Machine learning* 47 (Association for Computing Machinery, 2004). doi:10.1145/1015330.1015417.
50. Thurstone, L. L. An analytical method for simple structure. *Psychometrika* **19**, 173–182 (1954).
51. Hannachi, A., Unkel, S., Trendafilov, N. T. & Jolliffe, I. T. Independent Component Analysis of Climate Data: A New Look at EOF Rotation. *J. Climate* **22**, 2797–2812 (2009).
52. Espadoto, M., Martins, R. M., Kerren, A., Hirata, N. S. T. & Telea, A. C. Toward a Quantitative Survey of Dimension Reduction Techniques. *IEEE Transactions on Visualization and Computer Graphics* **27**, 2153–2173 (2021).
53. Bartenhagen, C., Klein, H.-U., Ruckert, C., Jiang, X. & Dugas, M. Comparative study of unsupervised dimension reduction techniques for the visualization of microarray gene

- expression data. *BMC Bioinformatics* **11**, 567 (2010).
54. Erdoğan, S. B., Tong, Y., Hocke, L. M., Lindsey, K. P. & deB Frederick, B. Correcting for Blood Arrival Time in Global Mean Regression Enhances Functional Connectivity Analysis of Resting State fMRI-BOLD Signals. *Frontiers in Human Neuroscience* **10**, 311 (2016).
55. Tong, Y., Hocke, L. M. & Frederick, B. B. Low Frequency Systemic Hemodynamic “Noise” in Resting State BOLD fMRI: Characteristics, Causes, Implications, Mitigation Strategies, and Applications. *Frontiers in Neuroscience* **13**, 787 (2019).
56. Tong, Y. & Frederick, B. deB. Time lag dependent multimodal processing of concurrent fMRI and near-infrared spectroscopy (NIRS) data suggests a global circulatory origin for low-frequency oscillation signals in human brain. *Neuroimage* **53**, 553–564 (2010).
57. Tong, Y., Yao, J. (Fiona), Chen, J. J. & Frederick, B. deB. The resting-state fMRI arterial signal predicts differential blood transit time through the brain. *J Cereb Blood Flow Metab* **39**, 1148–1160 (2019).
58. Schölvinck, M. L., Maier, A., Ye, F. Q., Duyn, J. H. & Leopold, D. A. Neural basis of global resting-state fMRI activity. *PNAS* **107**, 10238–10243 (2010).
59. Turchi, J. *et al.* The Basal Forebrain Regulates Global Resting-State fMRI Fluctuations. *Neuron* **97**, 940-952.e4 (2018).
60. Yuan, H., Zotev, V., Phillips, R. & Bodurka, J. Correlated slow fluctuations in respiration, EEG, and BOLD fMRI. *NeuroImage* **79**, 81–93 (2013).
61. Özbay, P. S. *et al.* Sympathetic activity contributes to the fMRI signal. *Commun Biol* **2**, 1–9 (2019).
62. Colenbier, N. *et al.* Disambiguating the role of blood flow and global signal with partial information decomposition. *NeuroImage* **213**, 116699 (2020).

63. Greicius, M. D., Krasnow, B., Reiss, A. L. & Menon, V. Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *PNAS* **100**, 253–258 (2003).
64. Raichle, M. E. *et al.* A default mode of brain function. *PNAS* **98**, 676–682 (2001).
65. Murphy, K., Birn, R. M., Handwerker, D. A., Jones, T. B. & Bandettini, P. A. The impact of global signal regression on resting state correlations: Are anti-correlated networks introduced? *Neuroimage* **44**, 893–905 (2009).
66. Pan, W.-J., Thompson, G. J., Magnuson, M. E., Jaeger, D. & Keilholz, S. Infraslow LFP correlates to resting-state fMRI BOLD signals. *NeuroImage* **74**, 288–297 (2013).
67. Murta, T., Leite, M., Carmichael, D. W., Figueiredo, P. & Lemieux, L. Electrophysiological correlates of the BOLD signal for EEG-informed fMRI. *Hum Brain Mapp* **36**, 391–414 (2014).
68. Liang, Y. *et al.* Cortex-Wide Dynamics of Intrinsic Electrical Activities: Propagating Waves and Their Interactions. *J. Neurosci.* **41**, 3665–3678 (2021).
69. Takeda, Y., Hiroe, N. & Yamashita, O. Whole-brain propagating patterns in human resting-state brain activities. *NeuroImage* **245**, 118711 (2021).
70. Esfahlani, F. Z. *et al.* High-amplitude cofluctuations in cortical activity drive functional connectivity. *PNAS* **117**, 28393–28401 (2020).
71. Betzel, R. F., Cutts, S. A., Greenwell, S. & Sporns, O. *Individualized event structure drives individual differences in whole-brain functional connectivity*. 2021.03.12.435168 <https://www.biorxiv.org/content/10.1101/2021.03.12.435168v2> (2021) doi:10.1101/2021.03.12.435168.
72. Aires, F., Rossow, W. B. & Chédin, A. 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 (2002).
73. Liu, X. & Duyn, J. Resting-State fMRI Signal Anti-Correlation Exists in Absence of Global Signal Regression. in *Proceedings of the 21st ISMRM Annual Meeting* 1 (2013).
74. Van Essen, D. C. *et al.* The WU-Minn Human Connectome Project: an overview. *Neuroimage* **80**, 62–79 (2013).
75. Smith, S. M. *et al.* Resting-state fMRI in the Human Connectome Project. *Neuroimage* **80**, 144–168 (2013).
76. Dale, A. M., Fischl, B. & Sereno, M. I. Cortical surface-based analysis. I. Segmentation and surface reconstruction. *Neuroimage* **9**, 179–194 (1999).
77. Reid, A. T. *et al.* Advancing functional connectivity research from association to causation. *Nature Neuroscience* **22**, 1751–1760 (2019).
78. Bzdok, D. *et al.* Formal Models of the Network Co-occurrence Underlying Mental Operations. *PLOS Computational Biology* **12**, e1004994 (2016).
79. Schaefer, A. *et al.* Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex* **28**, 3095–3114 (2018).