

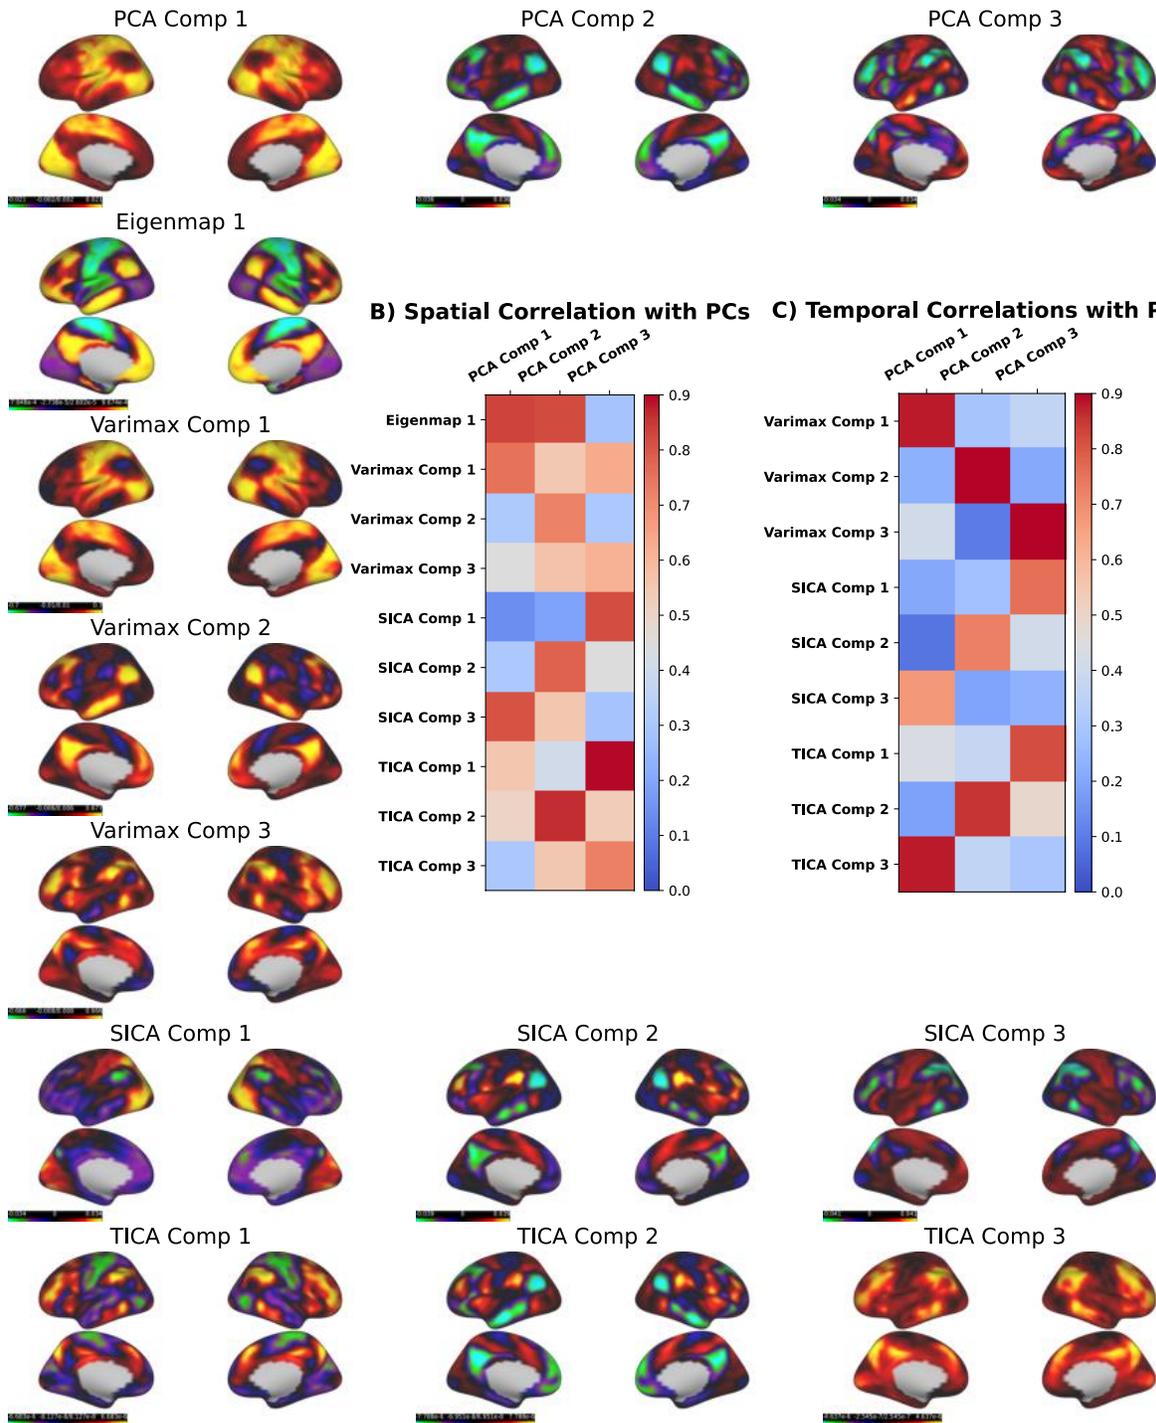
Supplementary Materials

Supplementary Results A - Latent Dimension Reduction Functional Connectivity Topographies

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

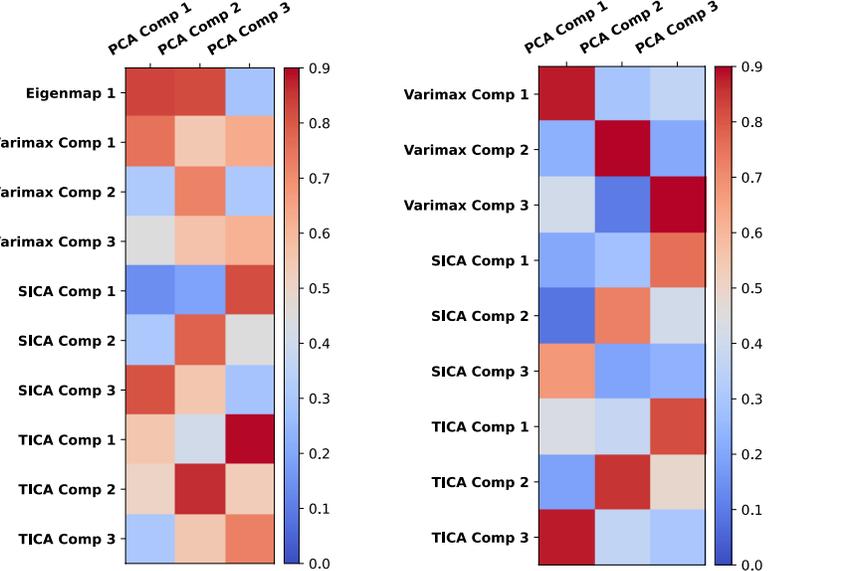
All of these rotation methods drop one or more of the restrictive constraints of the principal component solution: 1) orthogonality of the spatial eigenvectors, 2) uncorrelatedness of component time series, or 3) the criterion of maximum variance. The question is whether these rotation methods of the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

A) Low-Dimensional FC Topographies



B) Spatial Correlation with PCs

C) Temporal Correlations with PCs



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

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

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

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

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

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

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

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

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

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

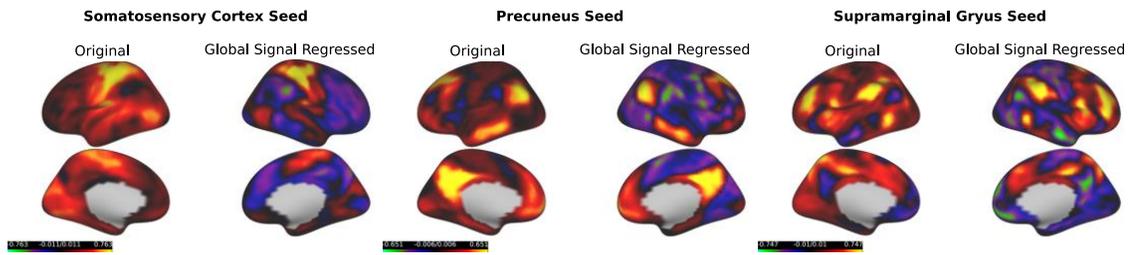
Jaccard similarity (Jaccard Index) between the pair of binary time courses. The Jaccard similarity between two CAP binary time courses varies from 0 to 1 and reflects the ratio of overlapping onset time points ($=1$) to the total number of time points ($N=60,000$). To account for potential time-lags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

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

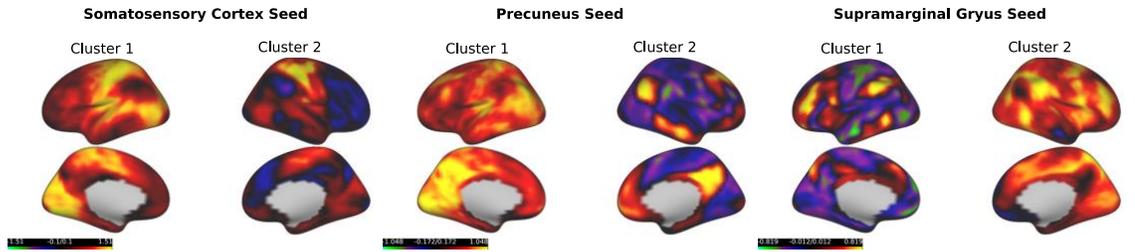
Normalizing BOLD values within suprathreshold time points removes scaling and mean differences between suprathreshold time points. In other words, information regarding global mean differences between suprathreshold time points is removed. In the standard k-means algorithm with Euclidean distances, normalization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

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

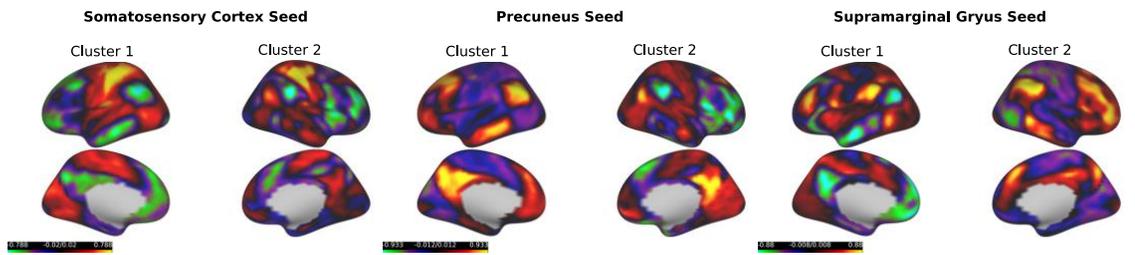
A) Seed-Based Regression Maps



B) Co-activation Pattern Clusters (N=2)



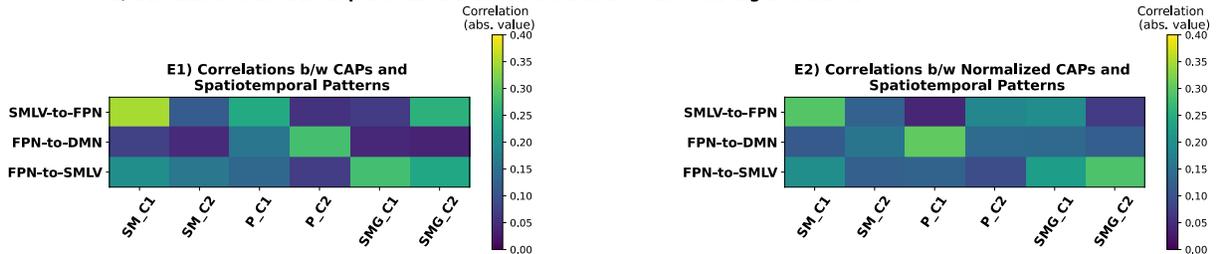
C) Time-point Normalized - Co-activation Pattern Clusters (N=2)



D) Overlap in Suprathreshold Time Points

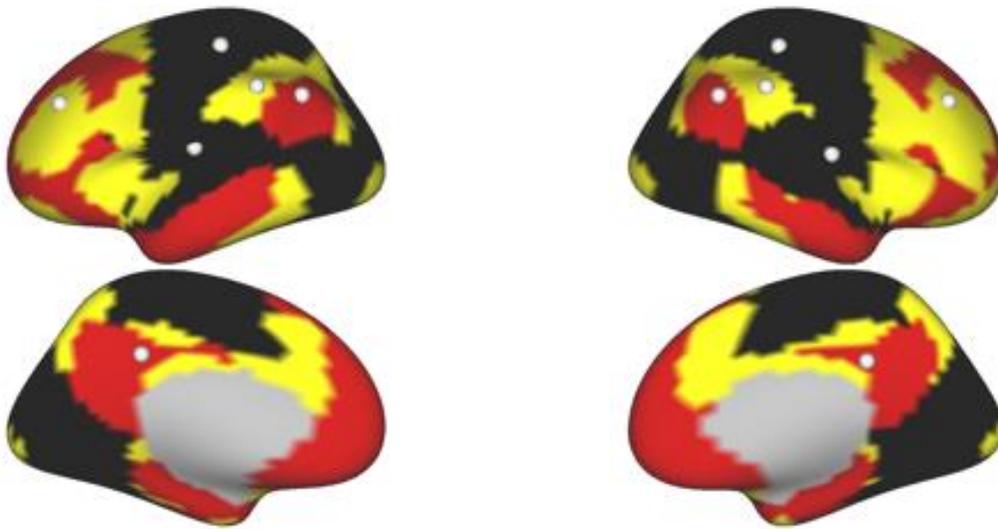


E) Correlation between Suprathreshold Time Points and Three Time-lag Structures



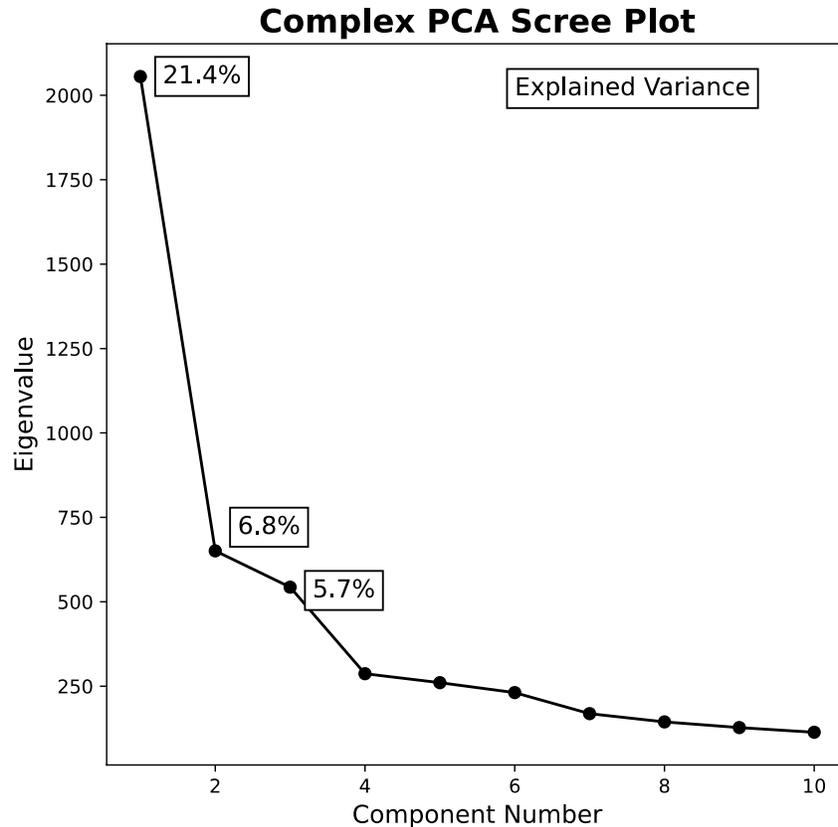
Supplementary Figure B1. **Seed-Based Regression and CAP Analysis.** (SM = somatosensory cortex; P=Precuneus; SMG=Supramarginal Gyrus). FC topographies of seed-based regression maps and CAP centroids from somatosensory (SM), precuneus and supramarginal gyrus seeds. **A)** Seed-based regression maps with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

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



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

Supplementary Results C - Complex Principal Component Scree Plot

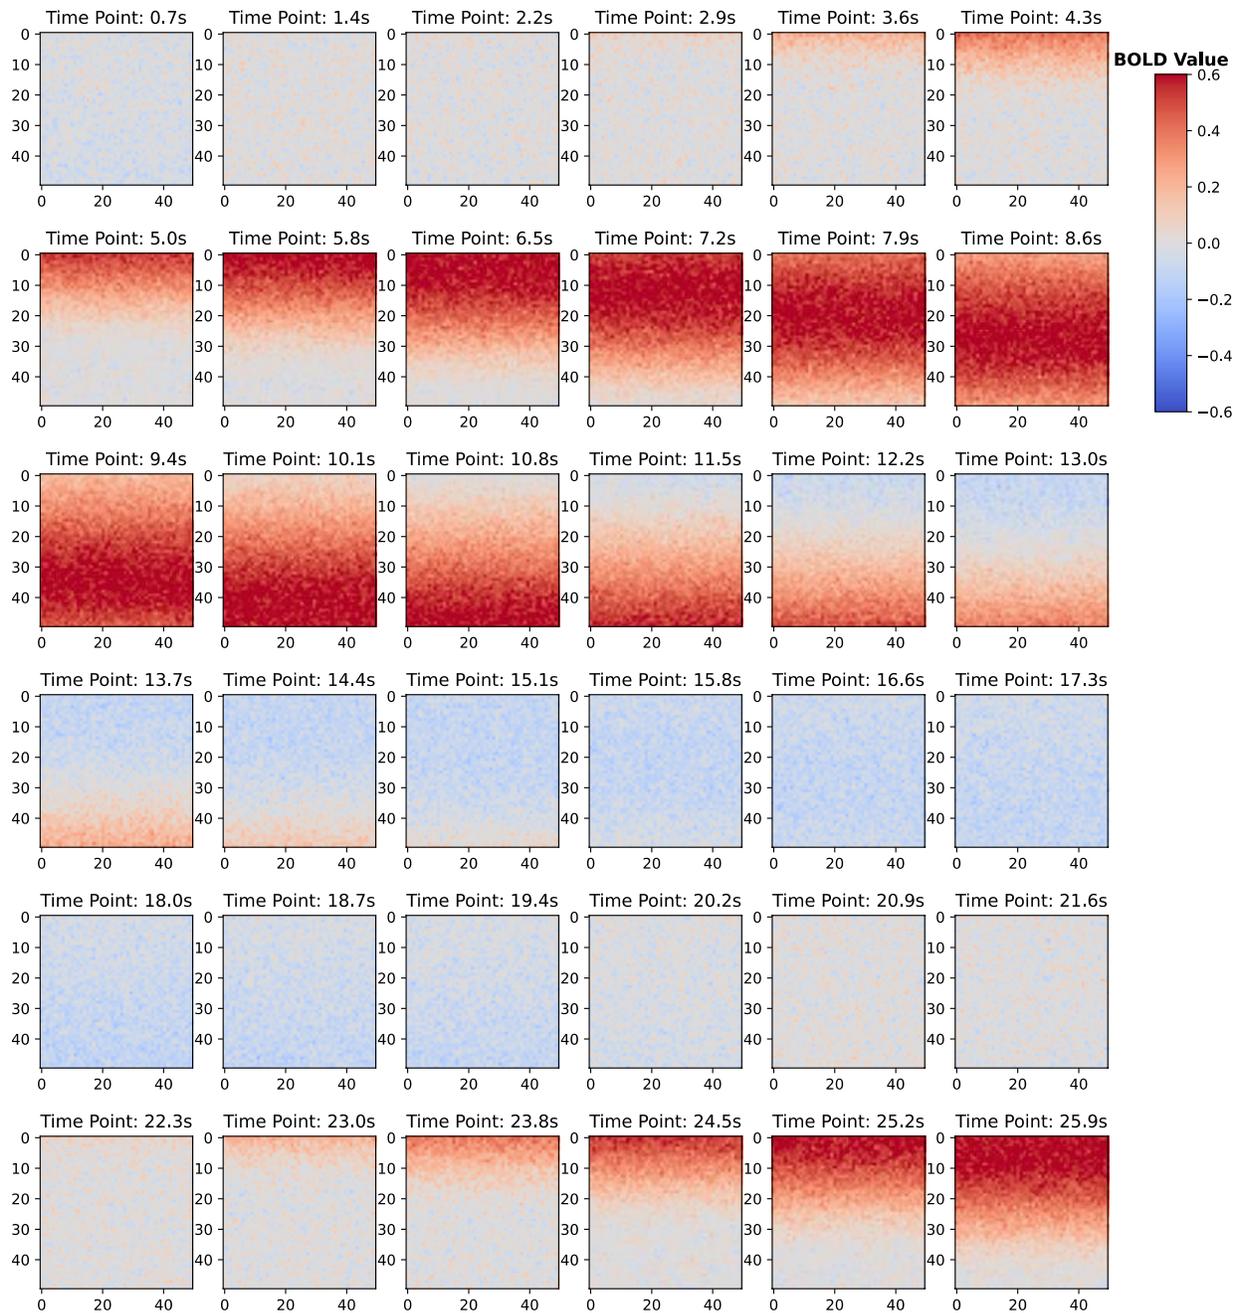


Supplementary Figure C. **Scree Plot from Complex Principal Component Analysis.** The eigenvalue by component number plot (i.e. scree plot) used to determine the number of components to extract. As with the original PCA solution (**Figure 1C**), there are clear elbows in the plot after one and three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

Supplementary Results D - Complex Principal Component Analysis Simulation

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

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

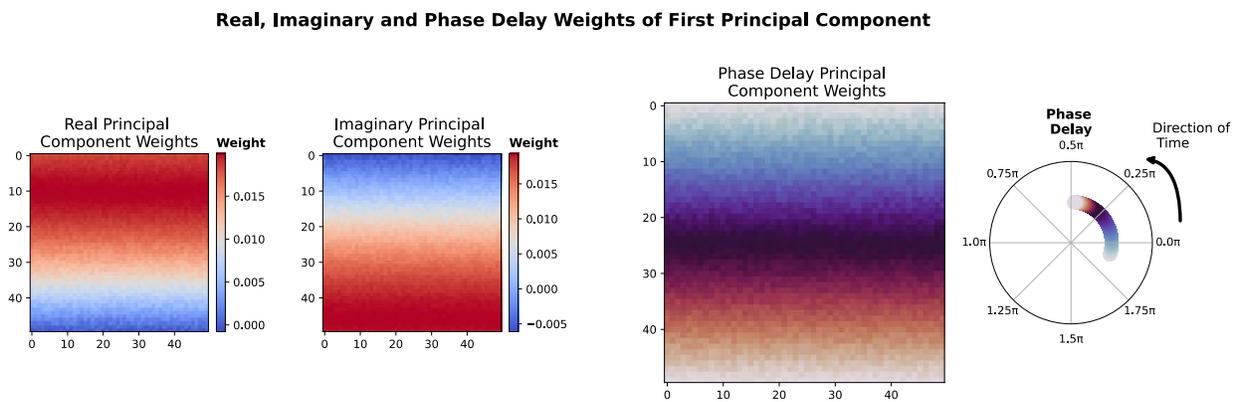


Supplementary Figure D1. **Global BOLD Propagation Simulation.** The first ~25 sec of an artificially constructed BOLD propagation simulation. Artificial ‘vertices’ are arranged along a 50-by-50 spatial grid. Vertex time series are created by convolving simple impulse time series with the canonical hemodynamic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

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

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

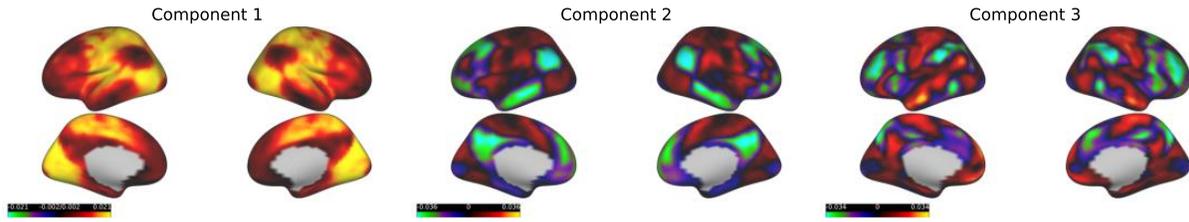
As illustrated in **Supplementary Figure D2**, the first complex principal component from cPCA accurately recovers the spatiotemporal pattern of the global propagation simulation. A sketch of this spatiotemporal pattern can be derived by comparing the activity patterns of the real and imaginary parts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



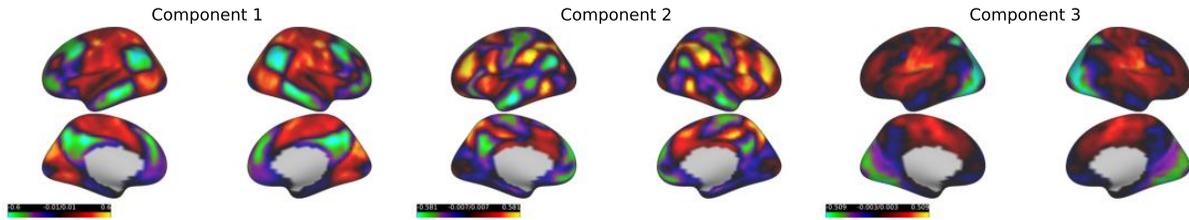
Supplementary Figure D2. Real, Imaginary and Phase Delay Weights of First Complex Principal Component. The real part, imaginary part and phase delay weights of the first complex principal component (cPC) spatial weights derived from the propagation simulation. The real and imaginary parts (left) of the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

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

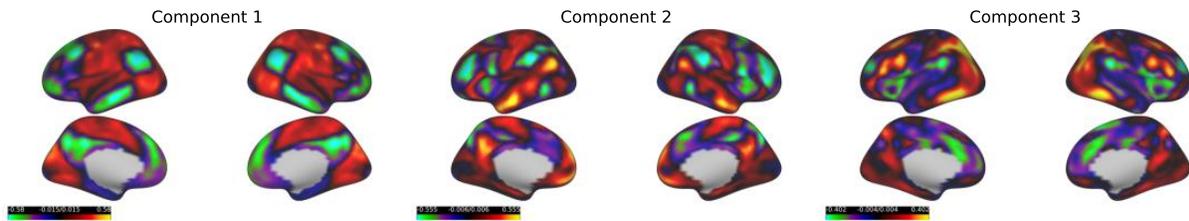
Principal Component Analysis



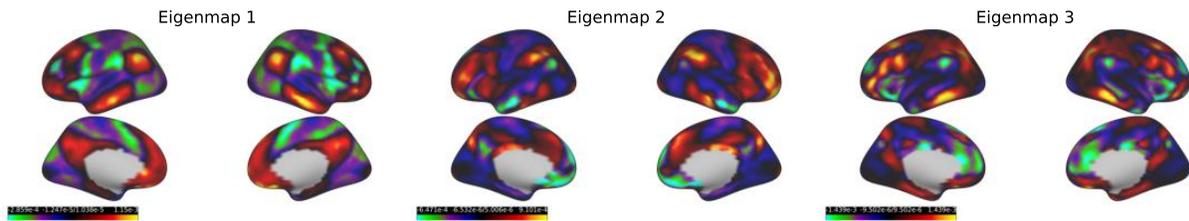
Principal Component Analysis - Global Signal Removed



Principal Component Analysis - Time-Point Centered



Laplacian Eigenmaps - Manifold Learning



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

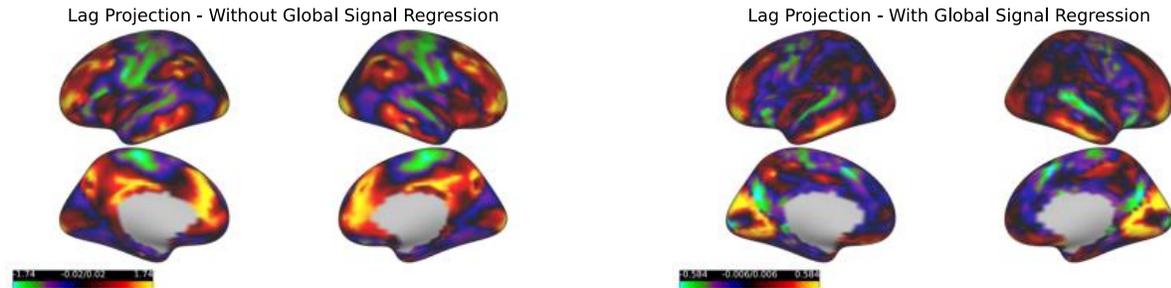
The first eigenmap corresponds to the primary functional connectivity gradient (PG) described by (Margulies et al., 2016). However, we note that the exact spatial pattern of the PG depends on the level of thresholding applied to the FC matrix (**Figure 5**). In this LE solution, no thresholding was applied.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

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

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

Supplementary Results F - Global Signal Regression and Lag Projections

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



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

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