Slow Cortical Waves via Cyclicity

Ivan Abraham[1], Somayeh Shahsavarani[2], Benjamin Zimmerman[3], Fatima Husain[3,4] and Yuliy Baryshnikov[*,1,5]

Abstract

Fine-grained information about dynamic structure of cortical networks is crucial in unpacking brain function. Here, we introduced a novel analytical method to characterize the dynamic interaction between distant brain regions, based on cyclicity analysis, and applied it to data from the Human Connectome Project. Resting-state fMRI time series are aperiodic and, hence, lack a base frequency. Cyclicity analysis, which is time-reparametrization invariant, is effective in recovering dynamic temporal ordering of such time series along a circular trajectory without assuming any time scale. Our analysis detected the propagation of slow cortical waves across the brain with consistent shifts in lead-lag relationships between specific brain regions. We also observed short bursts of strong temporal ordering that dominated overall lead-lag relationships between pairs of regions in the brain, which were modulated by tasks. Our results suggest the possible role played by slow waves of ordered information between brain regions that underlie emergent cognitive function.

Keywords

cortical waves — functional MRI — dynamic networks — iterated integrals

[1] *Department of Electrical & Computer Engineering, University of Illinois at Urbana-Champaign*
[2] *Zuckerman Mind Brain Behavior Institute, Columbia University, New York*
[3] *Beckman Institute, University of Illinois at Urbana-Champaign*
[4] *Department of Speech & Hearing Science, University of Illinois at Urbana-Champaign*
[5] *Department of Mathematics, University of Illinois at Urbana-Champaign*
[*]**Corresponding author**: ymb@illinois.edu

Contents

1. Introduction

The brain spontaneously generates neural activity even in the absence of any sensory inputs, motor outputs, or attention-demanding cognitive tasks. Many studies have shown that this ongoing activity is not just random noise, but that it carries meaningful information, having a dynamic structure that can interact with perception and behavior [1, 2, 3, 4, 5, 6]. Conventionally, this structure has been characterized as synchronous patterns of neural activity in bilaterally symmetric, distant brain regions, forming brain-wide functional neural networks.

Over the past two decades, functional magnetic resonance imaging (fMRI) has emerged as an important brain-imaging tool to study such neural networks using a measure known as resting-state functional connectivity (FC). Studying the synchrony between spontaneous changes in fMRI blood oxygen level-dependent (BOLD) signals as a proxy for ongoing neural fluctuations provides researchers with a noninvasive tool to investigate functional organization of the brain at network level. It is well established that resting-state FC patterns are not stationary, but, rather, vary over the time course of an fMRI resting-state session [7, 8].

The temporal dynamics of spontaneous neural activity is not limited to synchronous activity. Neuroimaging studies in both humans and animal models suggest the propagation of cortical waves across the brain. For example, using EEG

recording in human subjects, Massimini and colleagues [9] showed that slow oscillations during sleep are cortical waves usually originating at anterior cortical regions. Using wide-field optical calcium imaging in mice, Matsui and colleagues [10] observed that transient neuronal coactivations were embedded within propagating waves of activity across the cortex, which was posited to carry important information underlying spatiotemporal neuronal dynamics of FC. Although cortical waves of neural activity at both mesoscopic and macroscopic scales are particularly well established at high frequencies [11], much remains to be explored to better understand to what extent these cortical waves would manifest at slower frequencies measured by the BOLD signal.

In the last decade, the latency structure of resting-state BOLD signals has been increasingly examined, and there is accumulating evidence for signal flow even at the low frequencies sampled by the BOLD signal. In one study, Majeed and colleagues [12] found signals that moved laterally to medially, primarily in sensorimotor cortex. In a follow up study including both rat and human data, the same group replicated their first findings [13]. More recently, Ma and Zhang [14] further explored the dynamic temporal characteristics of FC, and found that transitions between resting state FC patterns follow specific sequential ordering, which is conserved in rats. Additionally, Hindricks and colleagues [15] were able to use lag-based methods to observe cortical waves in the BOLD signal in early visual cortex. Recently, fMRI studies by Mitra and colleagues [16, 17, 18, 19] demonstrated the existence of slow resting-state inter- and intra-network propagation patterns, and the direction of such slow propagation was related to the level of arousal [17, 20]. Using lagged correlation analysis, they reported on patterns of propagated activity in resting-state BOLD signal and suggested FC synchrony as an emergent property of lag structure underlying such propagation [20]. An alternative perspective to the cortical wave was recently proposed by Huntenburg and colleagues [21] where they described how connectivity gradients predict hierarchical information flow through the cortex, and how these gradients could predict network connectivity. Altogether, although these studies have begun to expand our understanding of the brain networks in terms of inter-regional brain interactions beyond simultaneous brain activity, developing novel tools and techniques for characterizing the mechanisms underlying such interactions still remains imperative to advance this line of research.

The goal of this paper is to introduce a new tool to study such interactions, which provides fine-grained information about the temporal dynamics of BOLD signals for detecting and studying cortical waves propagating across cortical networks using resting-state fMRI data. Previously, we applied cyclicity analysis [22] on resting-state fMRI data and showed that spontaneous BOLD signals comprised temporal sequences, the temporal ordering of which could be recovered along a circular trajectory [23]. In the present study, using data from the Human Connectome Project [24], we expanded cyclicity analysis and introduced a new technique that reveals more complex dynamics of spontaneous BOLD signals. *Cyclicity analysis* (CA) is a novel technique that derives pairwise temporal relations between time series using iterated path integrals (for applications in fMRI studies, see [25]). While the findings resulting from both lagged correlations and cyclicity analyses overlap and provide evidence in favor of the propagation of slow brain activity, they have different underlying mathematical apparatuses and assumptions, and levels of granularity. The lagged correlation method infers lag threads by deriving singular vectors of the time-delay matrix, whereas the CA method recovers inherent ordering among BOLD time series through eigenvectors of a lead matrix (a representation of the strength of temporal ordering between pairs of regions, see Section 2). Moreover, lagged correlation relies on interpolation and windowing to capture the dynamics of FC, which is vulnerable to time delay estimation methods, autocorrelation [25], sampling variability [26], and parameters such as window length and window shift [27]. In contrast, cyclicity analysis offers a more robust approach with a higher level of granularity to study lag structure where there is no assumption regarding stationarity, latency estimation, state duration, and state transition.

To better understand the spatiotemporal dynamics of functional neural networks, it is essential to develop appropriate methods to fully characterize the interaction between distant brain regions. In the context of FC analysis, cyclicity analysis can provide complementary information by capturing the order of neuronal events as neural fluctuations propagate along cyclic trajectories. In the present study, our new method demonstrated a transient leader-follower relationship among brain areas. This provides further evidence in favor of the speculation that time-dependent information flow underlies resting-state FC besides neuronal synchrony and desynchrony. Our technique propels the FC field beyond correlation analysis, enabling researchers to gain more mechanistic understanding of brain organization.

2. Cyclicity analysis

Mathematically, the idea of CA stems from the realization that all tools based on harmonic analysis (such as Fourier transform, autocorrelation functions, power spectra, etc.) would suffer if time is reparameterized in some fashion [22]. In physiology, one often encounters the processes where essentially the same time series is playing out at different speeds at different times (think, e.g. of the heart rhythm). One then turns, naturally, to data processing tools which return the same results if the time is reparameterized, or, more formally, will produce the same results on the time series related by a monotonic transformation $t \mapsto s(t)$, i.e. on

$$\mathbf{x}(t) \text{ and } \mathbf{x}'(t) := \mathbf{x}(s(t)). \tag{1}$$

Hence, a principle data analysis tool that withstands reparametrization of the timeline should output values that remain

invariant with respect to the change of parameters in (1). We will refer to such outputs, in the aggregate, as the reparamterization invariants (RI) of the trajectory $\mathbf{x}(t)$. The problem of finding RIs of trajectories

$$\mathbf{x}(t) = (x_1(t), x_2(t), \ldots, x_d(t))$$

in multidimensional space $\mathbb{R}^d$ was addressed by K.-T.Chen in [28], where he established that essentially (up to *detours*), the RIs are given by the *iterated integrals*. While iterated integrals form a countable family, having a rich underlying mathematical structure, we will focus on the first nontrivial instance, the *oriented areas*: below we will show that in a broad class of models, they allow recovery of subtle interactions between the components of the time series. We note, however, that the idea of using higher order iterated integrals was advanced by T. Lyons and his school, see e.g. [29].

2.1 Oriented areas

The analysis of this research builds on the iterated integrals of second order, better known as the *oriented areas*. It relies on an intuitive interpretation of the the oriented area encompassed by a curve in the plane, parametrically represented by a pair of trajectories, x and y. Consider a function of time $x(t)$ shaped as a pulse S (orange curve in the Figure 1(a)) and its time shifted copy, $y(t)$ (blue curve). We interpret the interrelationship between x and y as a *leader-follower* one: x leads; y follows. Now, to make this relationship an RI, we would like to render this relation using just the parametric plot (taking into account the orientation of the resulting curve in the $x-y$ plane). The observation that *when plotted against each other, these two trajectories enclose a region R_{xy} of the plane of positive area* (as shown in Figure 1) is one of the key intuitive premises of the cyclicity analysis.

Recall that a smooth curve partitions the plane into the open domains, such that the winding number of the curve around a point is constant within each of the domains (if the curve winds around a point *counterclockwise*, the point's winding number is positive, if *clockwise*, negative). The sum of the areas of those domains, weighted by their winding numbers is called the algebraic area (encircled by the curve). Alternatively, the algebraic area can be calculated using the Green's theorem, resulting in

$$\text{area}_{xy} = \frac{1}{2}\oint_C x\,dy - y\,dx, \tag{2}$$

where C is the curve, serving as the contour of integration (oriented by time). Note that this area is *signed*; in particular, reversing the orientation of the curve results in an integral with the opposite sign. The oriented area as expressed in (2) is an iterated integral of order 2.

Note that the quantity in (2) behaves differently from the more traditional measures of interactions between the time series, like the correlation coefficients. While the correlation

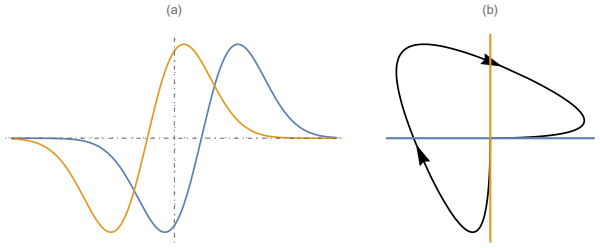

Figure 1. The *oriented* or *algebraic* area encompassed by a pair of time series. The left pane show a pair of time series x and y that are time-shifted copies of each other (abscissa is time). When plotted *against* each other, the right panel shows that they form a closed contour on the $x-y$ plane.

of a signal is largest with itself, the oriented areas are intrinsically antisymmetric, so that the autocorrelation necessarily vanishes. On the other hand, if two signals do not overlap in time oriented area measure is also zero. This naturally follows from the leader-follower relationship interpretation, because in either case, one is impossible to assign.

2.2 Lead matrices

In isolation, the oriented areas, while informative of the pairwise relationships between the time series, do not fully capture the collective phenomena the multidimensional trajectory represents. However, viewed together, those interactions reflect upon the collective, network structure among the components of the trajectory under some natural model assumptions.

For a trajectory in a multidimensional space, we arrange the oriented areas into the square $n \times n$ matrix (here n is the number of time series observed). This matrix, whose kl-th entry is the oriented area spanned by the pair x_k, x_l of the time series, is referred to as the *lead matrix* [22].We remark that this matrix is *skew symmetric*, and in particular, its eigenvalues form pairs of purely imaginary numbers, $\pm i\lambda, \lambda \in \mathbb{R}$, and the corresponding eigenvectors form complex-conjugated pairs with necessarily complex-valued components.

2.3 Chain of Offsets Model

Consider now the situation where the coordinates $x_j, j = 1, \ldots, n$ of the trajectory correspond to the same periodic function (which we will interpret here as the function on the internal clock space, a circle) $\phi : S^1 \to \mathbb{R}$, just offset by a different phase. In other words,

$$x_k(t) = a_k \phi(t - \alpha_k); \;\; \alpha_j \in S^1, k = 1, \ldots, n. \tag{3}$$

We will be referring to this model as the *Chain-Of-Offsets-Model* (COOM). The lead matrix can be readily computed in terms of the Fourier coefficients of ϕ: it is given by

$$A^{\phi}_{k,l} = 2\pi a_k a_l \sum_{m \geq 1} m |c_m|^2 \sin(m(\alpha_k - \alpha_l)) \tag{4}$$

In particular, the lead matrix given by (4) decomposes into the sum of *rank two* skew symmetric matrices with entries

$$A^{(m)}_{k,l} = m|c_m|^2 a_k a_l \sin\left(m(\alpha_k - \alpha_k)\right)$$

If one of the coefficients in the Fourier series for ϕ dominates, the skew symmetric matrix A^{ϕ} is well approximated (in Frobenius norm) by the rank 2 matrix $A^{(m)}$ with coefficients

$$\left(|c_m|^2 a_k a_l \sin(\alpha_k - \alpha_l)\right)_{kl} = |c_m|^2 \left(v_k u_l - v_l u_k\right)_{kl}, \qquad (5)$$

where $u_k = a_k \cos(\alpha_k), v_k = a_k \sin(\alpha_k)$. In general, if a skew-symmetric operator of rank 2 is represented as $Q = u \otimes v' - v \otimes u'$ (where u, v are linearly independent, and v' denotes the conjugate vector to v), then one can see that $w = -e^{-i\theta} u/|u| + v/|v|$ is an eigenvector of Q with the eigenvalue $i \sin\theta |u||v|$ (here θ is the angle between u, v). This implies that the real and imaginary components of the eigenvector $w_k = p_k + iq_k$ are obtained from the real and imaginary components u_k and v_k defining the matrix $A^{(m)}$ via the linear transformation of the plane

$$\begin{pmatrix} p_k \\ q_k \end{pmatrix} = \begin{pmatrix} Lu_k + Mv_k \\ Nu_k \end{pmatrix}$$

where $L = -\cos\theta/|u|, M = 1/|v|$ and $N = \sin\theta/|u|$. Therefore, if

$$p_k + iq_k$$

are the components of the eigenvector corresponding to the rank 2 matrix are stemming from a purely harmonic COOM with the offsets $\alpha_k, k = 1, \ldots, n$, they are just obtained from the complex numbers $a_k \exp i\alpha_k, k = 1, \ldots, n$ by a linear transformation of the complex plane, and therefore *the cyclic order defined by the arguments of the of these components* is the same as *the cyclic order of the collection of points on the unit circle* $\{\cos(\alpha_k) + i\sin(\alpha_k)\}_{k=1,\ldots,n}$.

Thus the spectral decomposition of the lead matrix can lead to the recovery of the order in which the signals are represented by the components of the time series $\mathbf{x}(t)$.

3. Model example: Signal propagation in networks

To illustrate how the cyclicity analysis works, we introduce a model example, where the signals are propagating through a network by gossiping.

3.1 Networks and Signals

Consider a weighted undirected graph G with n vertices, and interpret it as a model of a network, where internode communication delays are proportional to the edge lengths. If a node broadcasts a signal, it passes it to its neighbors, which receive it after the corresponding delay, and immediately rebroadcast it. If the signal was already received by a node, it ignores it (so that each signal reaches any given node along the shortest path connecting it to the source, and then replicates it perfectly).

Such peer-to-peer propagation models are often referred to as *gossip, epidemic* or *first passage percolation* algorithms [30, 31], and are relevant in studies of social networks, peer-to-peer broadcast algorithms, distributed resource location, etc. [32, 33, 34] Assuming such a model, one can easily recover the structure of the underlying network, if the shortest path lengths between the pairs of the nodes are known. Suppose a signal $s_i(t) = f(t)$ propagates from a *source* node i. If the underlying graph G is a connected one, then a node $j \neq i$ of the network will eventually receive the signal from i as the delayed shapeform, so that

$$s_k = f(t - d(i,k)) \qquad (6)$$

where $d(i,j)$ is the shortest path distance between nodes i and j. For connected graphs, this distance is well-defined and finite even if there is no direct connection between nodes i and j. Quite often one is only interested only in the ordinal information about the edgelengths (e.g. the minimal spanning trees depend only on the ordering of the edgelengths). Can this ordinal information be reconstructed using only the observations of the (perhaps, noisy) signals at the nodes?

Conceptually, this model, of the waves propagating in a network, from an unknown source to an undefined destination, is clearly a caricature of the cortex waves. Before presenting the results on the (tentative) recovery of those waves from the fMRI readings, it is enlightening to see what the situation is in the model example, where the ground truth is known.

3.2 Cyclicity analysis

We will apply the cyclicity computational pipe to recover the network structure from the observations of the signals $s_k, k = 1, \ldots, n$, as it manifestly fits the COOM. For each of the $\binom{n}{2}$ pairs of signals (6) observed at the nodes of the network, we form the oriented area of the corresponding 2-dimensional projection as in (2). As previously mentioned, the orientation is indicated by attaching a sign to the computed area value, with a positive sign corresponding to counter-clockwise integration; interpreted as $j(t)$ *following* $i(t)$ in time.

The spectral decomposition of square matrix A whose entries are given by (2), i.e. the *lead matrix*, allows us to determine the relative distances of the nodes of the network to the seed, where the signal originates. Indeed, as we argued above, the cyclic order of the arguments of the components of the (complex) eigenvector corresponding to the leading (in absolute value) eigenvalue of the lead matrix will reflect the cyclic order of the propagating wave, if the rank two skew-symmetric lead matrix corresponding to the purely harmonic signal propagating through the network is a good enough approximation of the sampled lead matrix A. In this case, spectral analysis of the lead matrix recovers the lag-structure between the source node and every other node.

One measure of how well the assumptions and results hold is to compute the eigenvalue ratio $|\lambda_1(A)|/|\lambda_3(A)|$ of the matrix, with larger ratios leading to better results. Repeating this analysis for various source nodes would then enable one to reconstruct the structure of the network.

3.2.1 Model example

In the Figure 2 we introduce a network on $n = 12$ nodes. The graph of Figure 2(a) is constructed by starting with the C_3 graph and adding nodes with $k = 3$ edges at a time. The edges are attached to vertices at random following a distribution proportional to the vertex degree [35]. The shown matrices represent the internode distances, both as given by the edge lengths in Figure 2(b), and by the shortest path distances between pairs of nodes in Figure 2(c). (Note that the shortest distance using relays can be shorter than the direct link.)

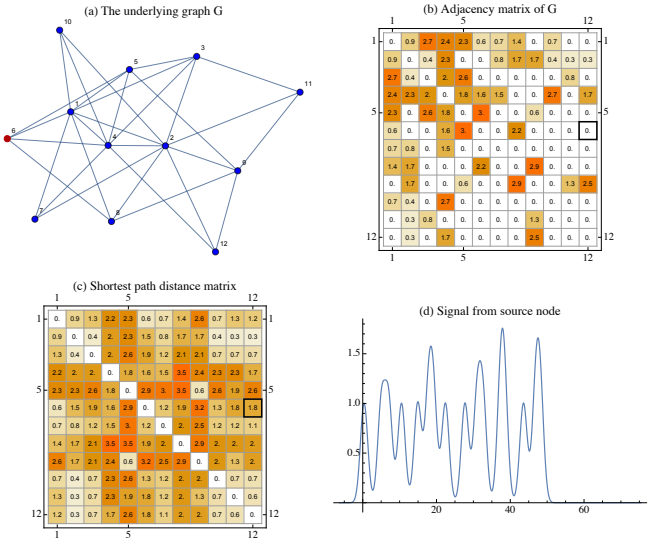

Figure 2. Details of the model example. **(a)** The graph G with $N = 12$ and source node $i = 6$, in red. **(b)** Random edge lengths. Zeros indicate the absence of the direct link between the nodes. **(c)** The shortest-path length matrix. Note that the connectivity of the graph results in all non-diagonal elements being nonzero. **(d)** The randomly generated the signal emanating from the source node.

To generate signal of Figure 2(d) emitted by the seed note (node 6 in this case), we use a convex combination of randomly modulated Gaussian functions, randomly displaced,

$$f(t) = \sum_{k=1}^{s} r_k g(t,k), \quad g(t,k) = \exp\left(-(t - k\pi + c\mathbf{R}_k)^2\right).$$

Here the displacements $\mathbf{R}_k$ and amplitudes r_k were drawn uniformly from the appropriate intervals (implementation of the simulation is available on the Github repository). The signal propagates through the network according to the equation (6), to which we add a small noise, realized as the scaled Brownian motion, independent at each node.

3.2.2 Recovering the network structure

The results of the cyclicity processing of the time series generated in the example on Figure 2 are presented in Figure 3. Figure 3(a) shows the sampled lead matrix. The interpretation of the entries is quite intuitive; thus one can see, for example, the vanishing oriented area between 5th and 6th nodes; this is consistent with the fact that the network distance from the seed to either of them is approximately the same.

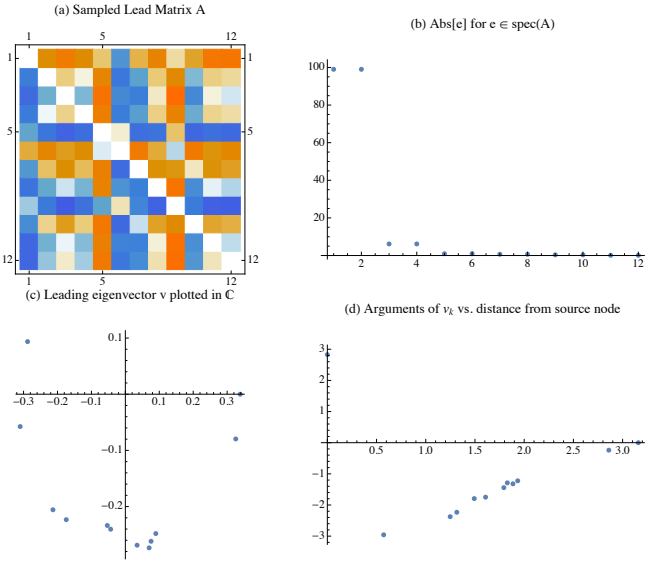

Figure 3. Results of the cyclicity analysis of the model example. **(a)** The skew-symmetric lead matrix. **(b)** (Absolute values of) the eigenvalues of the lead matrix. The rank 2 matrix corresponding to the first two (complex-conjugated) eigenvectors approximates the lead matrix well, as the first pair of eigenvalues dominates: $|\lambda_1|/|\lambda_3| \approx 20$. **(c)**The "constellation" of the components of the leading eigenvector winds around the origin; the circular order of their arguments (phases) indicates the propagation of the gossip wave through the network. **(d)** More precisely, the scatter plot of these arguments against the *known* distances from the source node shows a near perfect linear dependence, implying that lag-structure is encoded with eigenvector. Typically the distances between the nodes and the source will be unknown, however repeated analysis with multiple sources and careful consideration of obtained complex arguments from the eigenvectors enables inference of the underlying pairwise relative distances.

The first pair of the conjugate eigenvalues dominates in the spectrum of the lead matrix, indicative of the high resolution of the analysis as shown in Figure 4(b). The eigenvector corresponding to (one of) the leading eigenvectors has components shown in Figure 3(c). The arguments of these components, that is the angles formed by the rays pointing towards them and the ray of positive real numbers are shown on the display Figure 3(d) as the scatter plot, against the shortest distance to the seed node.

The strong, essentially linear dependence between these phases and the distances to the seed shows that cyclicity can be used to detect the latter from the former.

3.2.3 Areas dynamics

Besides recovering the (cyclic) order of the signals represented by the time series, one can extract additional information from the oriented areas, as they are accumulated in time. Namely, one can consider the integrals (2) with the variable upper end, resulting in the functions $A_{ij}(t), i, j \leq n$. These functions give a richer characterization of the collective behavior of the network. As a visualization tool, we will use the area gain plots, as shown on the Figure 4. Since $A_{ij} = -A_{ji}$ these plots are visualized in an upper triangular matrix, with a few of the plots magnified to emphasize some details.

Some of the plots are flat-lines, which mean there was no inferred *leader-follower* relationship between that particular pair of signals. Some plots show a sustained steady climb - *implying a strongly inferred leader-follower relationship* - whereas many others climb up and then climb down - *implying a reversal of the leader-follower mid period*. The maximum gain (or drop) of the plots characterize the strength of the inferred directed relationship. The insets in the image show the time evolution of the pair of signals that generated them.

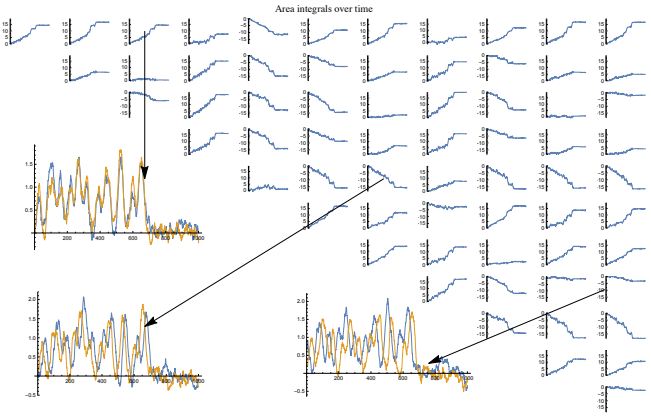

Figure 4. A panel showing the dynamics of the pairwise area-integrals computed as entries of the lead matrix. Since oriented areas are skew symmetric, the lower half shows insets with pairs of BOLD signals from ROIs that together generate the area dynamics. Shown are three typical instances: (top) strong directed relation as exhibited by increase, (bottom) strong directed relation as indicated by decrease and (right) a weak directed relationship as indicated my relatively flat line.

4. Analysis of the low frequency waves in cortex

4.1 Dataset

From the Human Connectome Project 1200 Release[36, 24], we considered 889 de-noised minimally preprocessed participants who completed all the structural, resting state, and task fMRI sessions using a customized Siemens Skyra 3T scanner. Of those, 27 were excluded from the analysis for the segmentation issues noted in the HCP Quality Control process or functional preprocessing errors reported in the HCP Data Released Updates. After exclusion, 862 participants remained for analysis who ranged in age from 22-45 years and included 464 females.

ROI #	ROI full name	ROI #	ROI full name
1	Banks of superior temporal S.	18	Pars orbitalis C.
2	Caudal anterior cingulate C.	19	Pars triangularis C.
3	Caudal middle frontal C.	20	Pericalcarine C.
4	Cuneus	21	Postcentral C.
5	Entorhinal C.	22	Posterior cingulate C.
6	Fusiform C.	23	Precentral C.
7	Inferior parietal C.	24	Precuneus
8	Inferior temporal C.	25	Rostral anterior cingulate C.
9	Isthmus cingulate C.	26	Rostral middle frontal C.
10	Lateral occipital C.	27	Superior frontal C.
11	Lateral orbitofrontal C.	28	Superior parietal C.
12	Lingual C.	29	Superior temporal C.
13	Medial orbitofrontal C.	30	Supramarginal C.
14	Middle temporal C.	31	Frontal P.
15	Parahippocampal C.	32	Temporal P.
16	Paracentral C.	33	Traverse temporal C.
17	Pars opercularis C.	34	Insula

Table 1. Table of 68 ROIs involved in the analysis, showing numbers for left-side regions; right-side regions run their indices as 35 through 68. Here **C** - cortex, **S** - sulcus and **P** - pole.

Analysis of the data started by downloading resting state fMRI scans from the Human Connectome Dataset's Young Adult Dataset. Towards this we made use of the Open Access to the data made available by the HCP Consortium via the Amazon Web Services hosting servers. Connectome Workbench software was used to extract regions of interest (ROI) time series information from the fMRI parcellations available for download.

The ROIs extracted were 34 in number (bilateral, therefore total $N = 68$) and are listed in the table below. The parcellated output files were further processed in Python/R to create time-courses for the regions listed in Table 1 . At a $TR = 720$ ms, this resulted in a set of arrays D, with elements of dimension 68×1200 that were fed into the cyclicity analysis. Further details of the scan protocol are available here. The left panel in Figure 5 shows representative BOLD signals in D.

5. Results

5.1 Cyclicity analysis on dataset

For cyclicity analysis, lead matrices were generated from the time courses and their eigenstructure analyzed. These matrices have dimension 68×68 with each (i, j) entry denoting the average leader follower relationship between ROI i and ROI j. See the right panel in Figure 5 for a representative example of a generic lead matrix. The $|\lambda_1|/|\lambda_3|$ ratio , where λ_k are eigenvalues, and the ratio is a measure of the quality of analysis,- was computed for all lead matrices. Note that $\lambda_1 = \overline{\lambda_2}$ since the lead matrix is real and skew-symmetric by construction. A subset, $S \subset D$, of the data was identified by restricting the lead matrices to have $|\lambda_1|/|\lambda_3|$ ratios one standard deviation above the group average $\mu_D(|\lambda_1|/|\lambda_3|)$. For a visualization of the distribution of the four leading eigenvalues in D see the left

image in Figure 6. The right image shows the distribution of $|\lambda_1|/|\lambda_3|$ ratio across D. We can see that S roughly amounts to 1/6th of all scans analyzed. The leading eigenvectors for this subset of the data was further examined. Recall that each entry v^i of the eigenvector v of a lead matrix corresponds to one of the 68 ROIs whose time courses it was created from. Since skew-symmetric matrices only admit zero or purely imaginary eigenvalues λ_k and corresponding eigenvectors v_k, the elements of an eigenvector $v_k$can be visualized as a *constellation,* $v_k^i = \left(x_k^i + j \cdot y_k^i\right)_{i=1}^{68}$, in the complex plane with each of the 68 points in it standing in for an ROI.

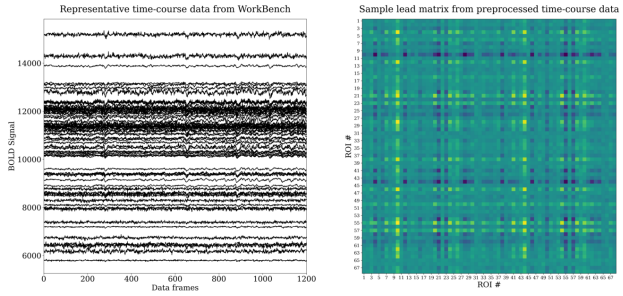

Figure 5. Representative time course data from the HCP dataset after processing using Connectome Workbench (left) and corresponding lead matrix (right). The lead matrix is generated after an appropriate normalization of the BOLD signal (see [23, 25]).

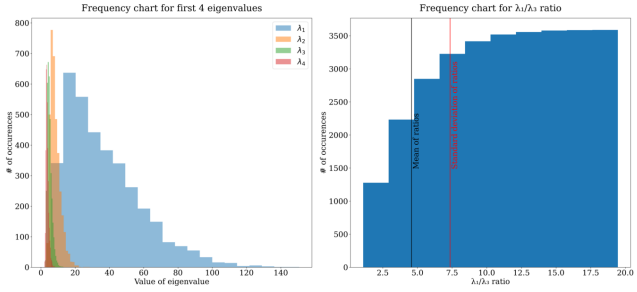

Figure 6. Frequency of the absolute values of four largest non-conjugate eigenvalues of the lead matrices in the data (left) along with the cumulative histogram of observed λ_1/λ_3 values (right). Higher λ_1/λ_3 values indicate more reliable outcomes from the cyclicity analyiss pipeline.

5.2 Significant ROIs

In such a constellation, the points farthest from the origin correspond to ROI time courses that dominate the multi-dimensional time series. To identify such dominant ROIs in a principled fashion, consider the single harmonic case in [22] presented in the previous section where perfect recovery of phases is possible. In this case, the constellations are ellipses in the complex plane. Adding noise and disturbances to the signal distorts the ellipse. Since cyclicity assumes periodicity in the internal clock of the underlying generative process, for a general multi-dimensional signal, one can assume that ellipses underlie the observed configuration of the constellations. Therefore, one can fit ellipses (for example, via least square regression) to the components of the eigenvector, to distinguish points in the constellation corresponding to dominant (outside the ellipse) vs. relatively immaterial (inside the ellipse) ROIs - see Figure 7 for a representative example. Such analysis resulted in the identification of a subset R, of 14 consistently dominant ROIs for/across the time courses in S.

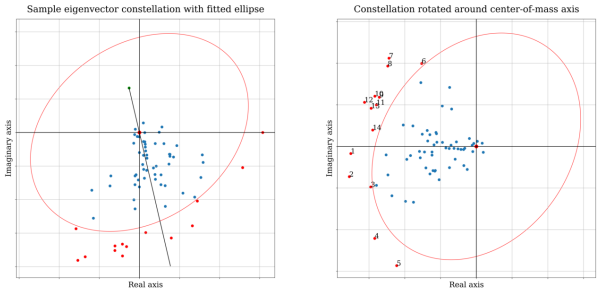

Figure 7. This figure shows one step in the determination of dominant ROIs in cyclicity analysis. The left figure shows components of the leading eigenvector visualized on the complex plane. Each point corresponds to an ROI and its BOLD signal; with greater absolute values indicating greater dominance in the multidimensional time-series. The right panel shows the constellation of points rotated to provide a consistent ordering across samples.

5.3 Robustness of dominant ROIs

Since the 14 ROI subset R was generated by considering the restriction of the data set to S it pertinent to ask whether this dominance was a feature common to the entire dataset D. To this end, random subsets $\overline{S}$ of the data of size $\overline{N} = 299$ were run through the pipeline instead of S and the reported dominant ROIs were tracked. Figure 8 below is a visualization of this as a matrix - each row corresponds to a random subset $\overline{S}$ and the values along the rows represent the dominant ROIs reported from examination of $\overline{S}$ - with the color scheme set to aide visual identification of change in values between rows. The average Levenshtein distance between each pair of rows in such a matrix (across multiple trials) was found to be less than 2 illustrating that the set R is robust to the choice of $\overline{S}$.

5.4 Average leader-follower relationship

As noted in [22, 23], cyclicity analysis is able to extract an approximate *cyclic ordering* by pairwise comparison of the leader-follower relationships between pairs of time series. It is *cyclic* because the order of events has no *fixed* directionality or *starting point,* rather, only their relative positions are of import - consider successively visiting points along a circle when it is inconsequential whether the traversal is performed clockwise or counter-clockwise. Therefore, to compare between obtained orderings, one needs a consistent way of presenting them. Towards this end, one can construct an axis joining the origin to the center of mass for each constellation. Choosing

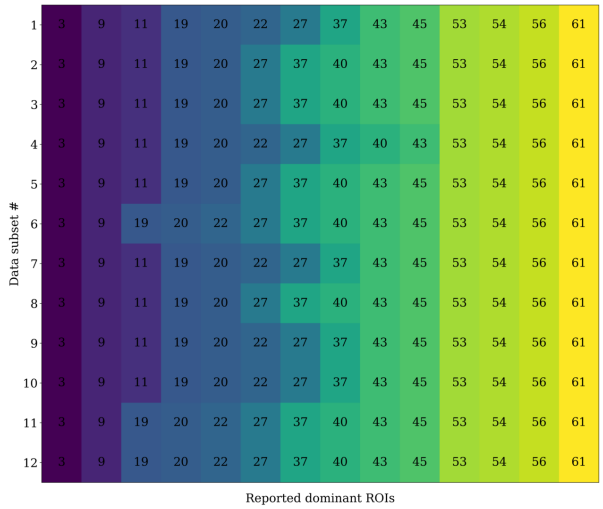

Figure 8. Figure showing reportedly dominant ROIs in 12 random (disjoint) partitions of the dataset. Each row of this matrix represents the most dominant ROIs reported by restricting the analysis to a random subset of the data (size $N = 299$). One can see that regardless of the subset, the same ROIs are reported to be dominant.

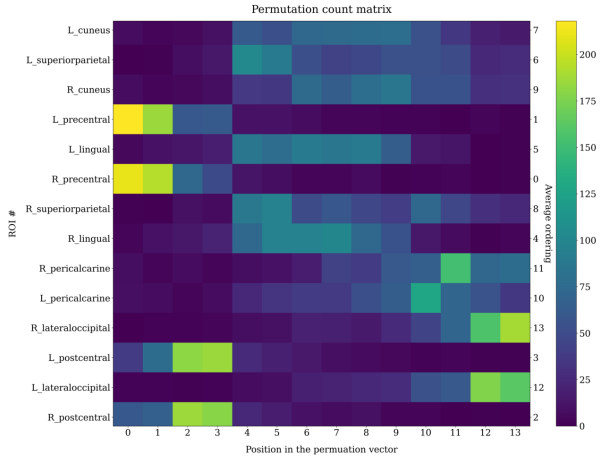

Figure 9. The *permutation count matrix* obtained from S by restricting ROIs considered to R. Each row of this matrix represents observed positions of the ROI in the cyclic ordering. Highlights indicate that ROI corresponding to that row *consistently* showed up in a particular position (along the columns) in obtained cyclic orderings.

this as the abscissa, the cycle is then assumed to be proceeding counter clockwise from this axis. This was shown in Figure 7 where the left figure shows a sample constellation of the eigenvector components along with the above-mentioned axis. The right figure shows the same constellation after the standard rotation. Applying this procedure for the regions in R to S resulted in estimated cyclic orderings for these dominant ROIs.

These obtained cyclic orderings can be visualized in a square *permutation count matrix* shown on Fig. 9. In this matrix the rows represent fixed ROIs and the columns represent positions in the cyclic ordering. Therefore the (i, j) entry of this matrix represents a value showing how many times ROI i showed up in position j in all the cyclic orderings obtained. Using this matrix, it was possible to estimate the *average cyclic ordering* for the ROIS in R as shown in Table 2. To do

ROI full name	Ordering in cycle	
bilateral order	Right	Left
Precentral cortex	1	2
Postcentral cortex	3	4
Lingual cortex	5	6
Superior parietal cortex	9	7
Cuneus	10	8
Pericalcarine	12	11
Lateral occipital cortex	14	13

Table 2. The fourteen dominant ROIs (including bilateral pairs) obtained and their average cyclic ordering in the dataset. The ordering shows that bilateral pairs frequently occur together in the cycle.

so, one first obtains a list of average positions by considering the values in each row of the matrix as samplings from a distribution and estimating its mean. Then, the permutation vector that sorts this list of mean values will correspond to the average cyclic order across the subset S.

5.5 Pairwise time-series analysis

Having established that the 14 dominating ROIs are a feature prevalent in the dataset, the time-series (i.e. the BOLD signal) itself of the 14 ROIs were examined in depth. Recall that cyclicity methods depends on pairwise analysis of time-series signals via iterated integrals defined between them [22]. For each of the 91 pairs of BOLD time series that can be formed from R, one can examine the area integral between the pairs as in (2). This is visualized in the Figure 10 below. The top-left panel shows a pair of time-series corresponding to a pair of ROIs from R. The top right panel shows the value of the above mentioned area integral computed between the pair of time-series signal. A definite increasing (or decreasing) trend over the period of observation shows that there is an average leader (or follower) relationship between the pair of time series - i.e. activity in one region precedes (or lags) activity in the other region. While the overall trend is indicative of of this relationship, it was noted that in a large number of instances, *the increase in the value of the area integral happened in short bursts as opposed to a continuous increase*.

These short periods of time where there is a significant contribution to the increase in the area integral can be termed "events" or significant periods of "directed activity" between the two pairs of brain signals. To methodically extract such instances in time, for each frame f_k of the time series, successive intervals of lengths 3-30 (i.e. intervals:

$$[f_k, f_{k+3}], [f_k, f_{k+4}], \ldots, [f_k, f_{k+30}]$$

were obtained and their slope computed). These slopes ob-

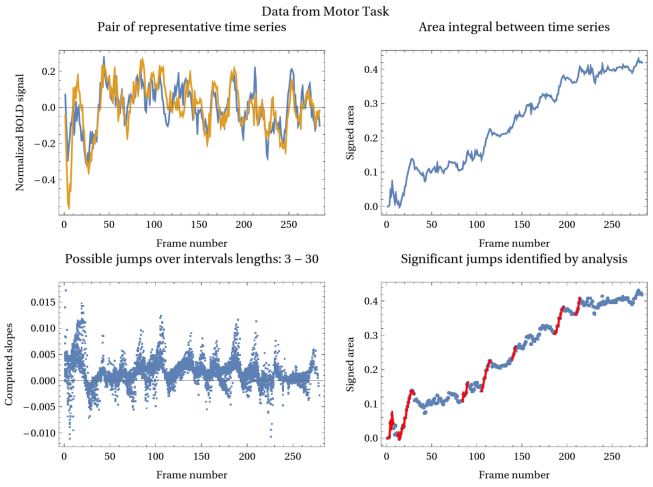

Figure 10. Procedure for identifying significant jumps in area integral corresponding to leader-follower activity. For each pair of timeseries as shown in the top-left panel, the area integral of Eq. (3) is computed as in the top-right figure. The slopes for possible jumps at each frame are identified and a threshold applied to designate significant jumps that contribute to the greatest overt increase in the value of the area integral.

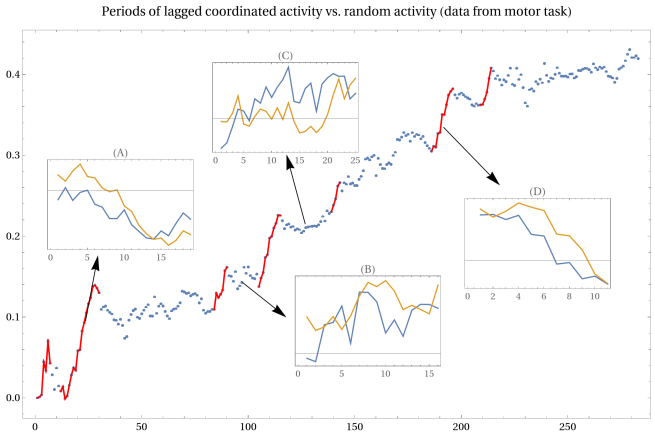

Figure 11. Panel showing instances of significant contribution to area integral highlighted in bottom-right image of Figure 10. Insets are titled (A) - (D) and referred to as Figure 11(A) - Figure 11(D) in the text. The inset abscissa label shows frame length of sub-interval considered from the major axis. Each pair of time-series here is the normalized BOLD signal over the titled frame. Instances when a region follows the BOLD signal in another region with a time lag correspond to greatest increases in value of the signed area integral.

tained at each frame, are visualized in the bottom left image of Figure 10. Half of the maximum slope observed among all such intervals over all frames was taken to be a nominal threshold and *significant* jumps or contributions were deemed to be those instances when the slope of the area integral was higher than said threshold. Periods of jumps identified in the manner are shown marked in red in the bottom right image of Figure 10 and the behavior of the corresponding time series during the marked intervals show in the panel of Figure 11.

The first inset Figure 11(A) corresponds to a period where there is a strong directed *leader-follower* relationship between the pair since the blue curve leads the activity of the orange one. A similar observation holds for Figure 11(D). On the other hand the relatively plateaued dynamics of the area integral shown in Figure 11(B) and 11(C) correspond to time periods between when the pair either act in concert simultaneously or a leader follower relationship is difficult to assign.

One can visualize all the area integrals coming from pairs constituted from the set R at once by representing them in a matrix. Figure 12 below shows one such image where only the upper half of the matrix is shown for clarity (lower half is the mirror image). This allows us to eyeball collective trends across pairs of time-series from R. For example in Figure 12, it is clear that the last two ROIs (along the columns) generally exhibit the same qualitative relationship with all other ROIs (area integral is always increasing). The insets in Figure 12 show the corresponding BOLD signals that generated the area dynamics and are specifically chose to show *increasing, decreasing* and *plateaued* dynamics. While Figure 12 was generated using BOLD signals recorded during a motor-task activity, Figure 13 was generated using BOLD signals recorded during the progression of a social/cognitive task. The two task protocols are as follows - for the **motor task** [37]:

> Participants are presented with visual cues that ask them to tap their left or right fingers, squeeze their left or right toes, or move their tongue to map motor areas. Each block of a movement type lasts 12 s (10 movements), and is preceded by a 3 s cue. In each of the two runs, there are 13 blocks, with 2 of tongue movements, 4 of hand movements (2 right and 2 left), 4 of foot movements (2 right and 2 left) and three 15 s fixation blocks per run.

whereas for the **social task**[37]:

> Participants are presented with short video clips (20 s) of objects (squares, circles, triangles) either interacting in some way, or moving randomly. These videos were developed by either Castelli and colleagues (Castelli et al., 2000) or Martin and colleagues (Wheatley et al., 2007). After each video clip, participants chose between 3 possibilities: whether the objects had a social interaction (an interaction that appears as if the shapes are taking into account each other's feelings and thoughts), Not Sure, or No interaction (i.e., there is no obvious interaction between the shapes and the movement appears random). Each of the two task runs has 5 video blocks (2 Mental and 3 Random in one run, 3 Mental and 2 Random in the other run) and 5 fixation blocks (15 s each).

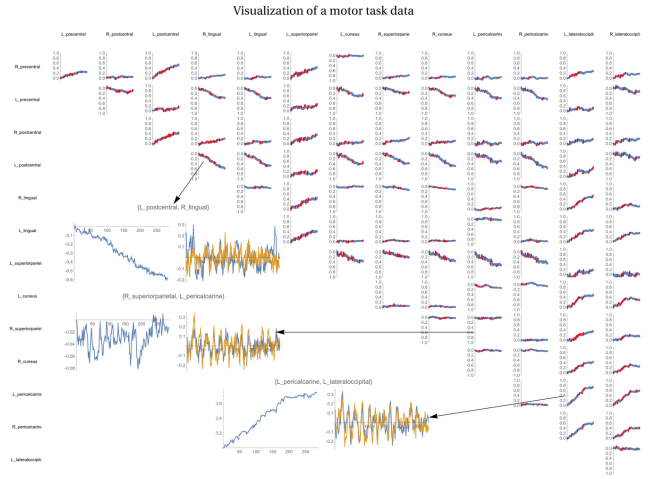

Figure 12. A way to visualize directed leader-follower activity between pairs of dominant ROIs obtained in the data set. A definitely increasing or decreasing trend indicates a strongly constrained directed-leader follower activity between two pairs while a variable trend indicates activity that is less constrained. In the top inset image, the blue signal leads the orange one whereas, in the bottom inset, the relationship is reversed.

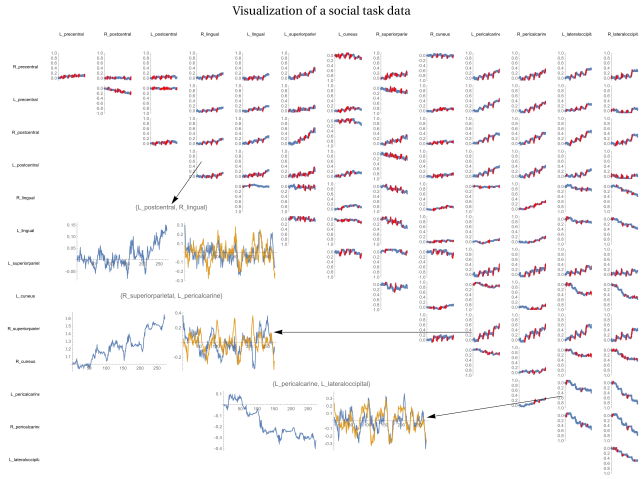

Figure 13. Visualization of directed activity between ROIs observed in the social cognition task. Compared to resting state or motor task scans, intermittent bursts of directed activity are a more prevalent feature in this analysis.

In Table 2 we see that precentral cortices lead the average cycle whereas lateral occitpial cortices end it. To examine the area-integral dynamics of the corresponding lead-lag pairs across different kinds of fMRI paradigms in the Connectome dataset, Figure 14 shows data computed from (a) motor task, (b) a social cognition task and (c) resting state.

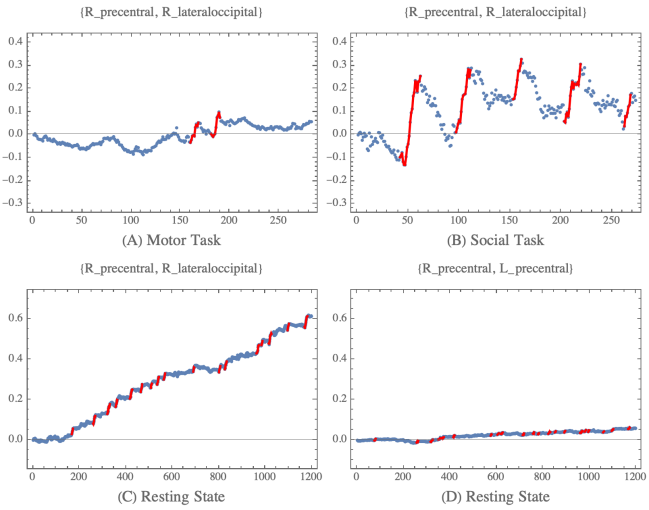

Figure 14. Some examples from analysis of data - **(A)** Shows the jump in the area integral for a pair of time series obtained from a motor task (`183034_r1`) and similarly, **(B)** shows jumps in a social task for the same pair in the same participant. Notice that the social task exhibits switching, i.e. there are periods of increase followed by a corresponding decrease. **(C)** shows resting state analysis (`970764_s2_r2`) while, **(D)** shows the behavior in the same motor task for the bilateral pair in the same scan.

6. Discussion

Cyclicity analysis holds the potential to provide new insights into dynamics of FC, building an improved understanding of brain networks and the interplay among them. Based on this technique, we introduced an effective method in revealing the directionality of the propagation of spontaneous BOLD signals across cortex. Our results provided supporting evidence of cortical waves propagating along the cortex between primary visual and transmodal areas of the brain. In addition, we found that the lead-lag relationships between pairs of regions reflected in the lead matrix were driven mostly by short periods of constrained temporal ordering. Figure 9 shows that certain brain regions appear at earlier or later times in a cycle, suggesting a time dependent flow of information underlying the dynamic aspect of resting states hemodynamic measures.

Although cortical waves have been noted at much higher frequencies, observing them in functional MRI data at low frequencies is still unexpected. At higher frequencies, cortical waves of oscillatory activity are thought to contribute to a spatiotemporal framework for neuronal communication by coordinating a range of synchronizing patterns [38, 39, 40, 41, 42, 11]. However, the temporal duration of these more typical cortical waves tends to be on the order of tens to hundreds of milliseconds rather than the super-second durations to which the BOLD signal should be sensitive. Nevertheless, over the past decade, a growing body of studies has observed BOLD signal flow across cortex [13, 14, 15, 16]. Taken together, these studies have yielded evidence that there is likely a propagation of the BOLD signal throughout the brain that dynamically carries or reflects longer periods of stable information flow between brain regions. The cyclicity analysis method that we have outlined here may be a useful tool for

investigating these dynamics due to its time invariance to variable lag-times between pairs of regions of the brain, which may limit correlational methods.

One of the more striking patterns our analysis uncovered was that the overall lead-lag relationships between pairs of regions in the brain were facilitated by short bursts of strong temporal ordering, rather than long consistent stretches of moderate ordering. In a recent paper by Esfahlani and colleagues [43], the authors investigated the contributions of moment-to-moment BOLD activity to the overall pattern of functional connectivity. Similar to our observation, they saw that only a small fraction of frames in the time-series explained a significant amount of the variance in the network connectivity. These fluctuations corresponded with activity patterns corresponding to fluctuations in the level of default mode and attentional control network activity, which are often viewed as in opposition to each other [44].

This opposition between the default mode and attention networks has a strong overlap with the idea of a principle gradient of macroscopic cortical organization in the brain [45]. According to this framework, a topography of connectivity patterns is reflected in the spatial distances between so-called "higher" areas of the brain, where more multi-modal, abstract, predictive information is encoded, and "lower" areas, such as the primary sensory/motor regions. This primary connectivity gradient predicts the positions of canonical resting-state networks, which are viewed in this framework as reflecting representational hierarchies rather than distinct modules. In other words, resting state networks are reflective of the temporal phase of propagating patterns, rather than as independent networks. The functions associated with various cortical networks are correlated to the level of the hierarchy of sensory or motor processing.

Note that, given the hierarchical framework that we are presenting, there are some unexpected results in the ordering that we observe, which should be resolved or investigated in future studies. The first is that the precentral and postcentral gyri, which are usually considered functionally as the primary motor and primary somatosensory cortices, sit at one end of the temporal ordering in the dominant cycle in the brain. We would expect that a cycle that reflected hierarchical processing of the visual cortex would have it's beginning or end in some of the multi-modal processing regions of the parietal lobe that overlapped with the default mode network, such as the inferior parietal lobule or the precuneus. It is possible that there is some anatomical overlap with the postcentral gyrus, but it is more difficult to understand how the precentral gyrus would be involved. It is possible that the timing of activity in the postcentral and precentral gyri are captured in the cycle, even though the activity is not part of the same cortical wave as the hierarchical visual processing. This possibility is subtly suggested in our data, where there seem to be distinct borders between activity in visual processing areas and activity in precentral and postcentral gyri.

A second unexpected result is that the major cycle seems to end in the lateral occipital cortex rather than in the pericalcarine gyrus, where primary visual processing begins. The rest of the cycle that we observe is largely consistent with the expectations of hierarchical visual processing. The temporal ordering puts the cuneus and lingual gyri near the pericalcarine gyri and lateral occipital cortex. The superior parietal cortex activity comes near the cuneus and lingual gyri.

We showed examples of how of the time-course of lead-lag relationships changes over a run depending on what task is analyzed. In each case, we showed the lead-lag relationship between the precentral gyrus, at the starting point of the cycle we observe, and the lateral occipital cortex, which was at the end of our cyclic ordering. In the context of the resting state scan, we see strong ordering in the direction that moves in the reverse hierarchy. The precentral gyrus activity always leads and the lateral occipital cortex always follows. Rather than observing the temporal ordering between these regions as constant, it seems to move in distinct bursts, where for a few frames, the ordering is very strongly constrained. Then for some other number of frames, the ordering is relatively unconstrained. We also notice that the lead-lag relationship between these two regions takes some time to begin. It is possible that during resting state, brain activity is mostly internally generated as in mind-wandering. This state of mind-wandering takes some time to start, but once initiated, since there is no stimulus to attend to, activity mostly occurs in the top-down direction reflecting top-down processing in the hierarchy.

In the context of the social task (see Figure 13), we observed a strikingly different pattern. This task consisted of a 5 video blocks where there was a period of watching a video and then judging whether or not the objects in the video appear to perceive the other object's feelings and thoughts. In this task, we still observe strong directional constraint, but it seems to shift in which area is the leader and which is the follower. This may be interpreted as periods of time that switch between internally generated (top-down) activity and externally generated (bottom-up) activity, that may correspond with periods of watching the videos and periods of making judgements of the objects' intentions within the video.

We observe a less interpretable pattern of activity in the context of a motor task. Here the participants are presented with 3-second cues to do 12-second motor movements. There are 13 motor blocks and three 15 second fixation blocks in each run. Like the motor task, we observe switches between the which region is the leader and follower during the run. However, it is difficult to match the activity observed to events in the task. This may be because the cuing events and motor movement events are too short to be adequately captured

in the low frequency movement across the brain. We may be observing some aliasing in the directionality of the lead-lag signal. It is also possible that in the motor task, there is more separability between activity belonging to hierarchical visual processing and motor processing. Both of these hypotheses are somewhat supported by the observation that if we observe purely visual regions, which are close together in both anatomy and expected temporal hierarchy, such as the pericalcarine gyrus and lateral occipital gyrus, we observe a very different pattern of temporal ordering. There are still alternations in the direction of temporal order, but they are more frequent and weaker in magnitude. Perhaps areas that are closer together are better able to capture the short time periods between subsequent cues and motor events, or perhaps it reflects less contamination from motor activity coming from the precentral gyrus.

Our analysis identified a temporal ordering that propagated from the primary occipital regions towards more transmodal regions. However, it should be noted that our analysis focused on analyzing the first eigenvector of the lead matrix, and we are specifically showing patterns of activity from subjects selected based on having high λ_1/λ_3 ratios in the lead matrices. This choice likely selects for subjects with strong leader-follower activity along a single direction. It is possible that analysis of the other eigenvectors would reveal a pattern of temporal ordering moving in the other directions, reflecting alternative gradients of direction from regions higher in the cortical hierarchy to the primary sensory regions lower in the cortical hierarchy or more local interactions, such as the signalling across hemispheres.

7. Conclusion

In this paper, we exploited cyclicity analysis to detect transient states in the brain. This method provides a new tool to complement the recent advancements in effective or directional connectivity research. This line of research has led to the intriguing discovery of temporally asynchronous patterns of BOLD activity, which may reflect information transfer in the brain [17, 46, 47, 48]. In fact, areas with strongly cyclicly constrained patterns overlap well with maps of timing differences in the duration of information transfer found recently by Xu and colleagues [48]. The method outlined here is advantageous because of its lack of assumptions regarding temporal properties of the BOLD signal dynamics. It does not assume stationarity of the time-series, or specific estimates of latency, state duration or state transitions, which may limit or bias correlational or lag-based approaches. Clearly, the method presented here opens a wide range of questions for future research. We have shown group data showing a primary propagating wave of BOLD activity during resting state from somatomotor cortex to early visual cortex. From examining individual subjects, we observe that these propagating waves appear to switch in direction in a task dependent manner. Our observations have largely been presented as descriptive, and future work is needed to derive useful group level statistics. In addition, the cyclicity analysis that we have described is not limited to looking at BOLD activity from fMRI data. Future work should apply the cyclicity analytic techniques to other measures of human brain activity that are more reflective of direct neural activity including electroencephalography (EEG),electrocorticography (ECoG), the fast optical signal, or magnetoencephalography (MEG). Applying cyclicity analysis to multiple techniques may help clarify the relationship between patterns observed in the fast temporal domain of neuronal activation and the longer duration patterns observed in the BOLD signal, which may be more reflective of broader stable states of whole brain function.

Acknowledgments

Data were provided [in part] by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

YB was partially supported by the AFRL, through MURI SLICE.

References

[1] Q Yu, EB Erhardt, J Sui, Y Du, H He, D Hjelm, MS Cetin, S Rachakonda, RL Miller, G Pearlson, and VD Calhoun. Assessing dynamic brain graphs of time-varying connectivity in fMRI data: application to healthy controls and patients with schizophrenia. *Neuroimage*, 107:345–355, Feb 2015.

[2] JM Shine, PG Bissett, PT Bell, O Koyejo, JH Balsters, KJ Gorgolewski, CA Moodie, and RA Poldrack. The Dynamics of Functional Brain Networks: Integrated Network States during Cognitive Task Performance. *Neuron*, 92:544–554, Oct 2016.

[3] JR Cohen. The behavioral and cognitive relevance of time-varying, dynamic changes in functional connectivity. *Neuroimage*, 180:515–525, Oct 2018.

[4] Barnaly Rashid, Laura M. E. Blanken, Ryan L. Muetzel, Robyn Miller, Eswar Damaraju, Mohammad R. Arbabshirani, Erik B. Erhardt, Frank C. Verhulst, Aad van der Lugt, Vincent W. V. Jaddoe, Henning Tiemeier, Tonya White, and Vince Calhoun. Connectivity dynamics in typical development and its relationship to autistic traits and autism spectrum disorder. *Human Brain Mapping*, 39(8):3127–3142, mar 2018.

[5] Zachary W. Davis, Lyle Muller, Julio Martinez-Trujillo, Terrence Sejnowski, and John H. Reynolds. Spontaneous travelling cortical waves gate perception in behaving primates. *Nature*, 587(7834):432–436, oct 2020.

[6] Sepideh Sadaghiani, Jean-Baptiste Poline, Andreas Kleinschmidt, and Mark D'Esposito. Ongoing dynamics in large-scale functional connectivity predict perception. *Proceedings of the National Academy of Sciences*, 112(27):8463–8468, jun 2015.

[7] R. Matthew Hutchison, Thilo Womelsdorf, Elena A. Allen, Peter A. Bandettini, Vince D. Calhoun, Maurizio Corbetta, Stefania Della Penna, Jeff H. Duyn, Gary H. Glover, Javier Gonzalez-Castillo, Daniel A. Handwerker, Shella Keilholz, Vesa Kiviniemi, David A. Leopold, Francesco de Pasquale, Olaf Sporns, Martin Walter, and Catie Chang. Dynamic functional connectivity: Promise issues, and interpretations. *NeuroImage*, 80:360–378, oct 2013.

[8] Elena A. Allen, Eswar Damaraju, Sergey M. Plis, Erik B. Erhardt, Tom Eichele, and Vince D. Calhoun. Tracking Whole-Brain Connectivity Dynamics in the Resting State. *Cerebral Cortex*, 24(3):663–676, nov 2012.

[9] M. Massimini, R. Huber, F. Ferrarelli, S. Hill, and G. Tononi. The Sleep Slow Oscillation as a Traveling Wave. *Journal of Neuroscience*, 24(31):6862–6870, aug 2004.

[10] Teppei Matsui, Tomonari Murakami, and Kenichi Ohki. Transient neuronal coactivations embedded in globally propagating waves underlie resting-state functional connectivity. *Proceedings of the National Academy of Sciences*, 113(23):6556–6561, may 2016.

[11] L Muller, F Chavane, J Reynolds, and TJ Sejnowski. Cortical travelling waves: mechanisms and computational principles. *Nat Rev Neurosci*, 19:255–268, May 2018.

[12] W Majeed, M Magnuson, and SD Keilholz. Spatiotemporal dynamics of low frequency fluctuations in BOLD fMRI of the rat. *J Magn Reson Imaging*, 30:384–93, Aug 2009.

[13] W Majeed, M Magnuson, W Hasenkamp, H Schwarb, EH Schumacher, L Barsalou, and SD Keilholz. Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans. *Neuroimage*, 54:1140–50, Jan 2011.

[14] Z Ma and N Zhang. Temporal transitions of spontaneous brain activity. *Elife*, 7, Mar 2018.

[15] R Hindriks, M R, G N, and D G. Latency analysis of resting-state BOLD-fMRI reveals traveling waves in visual cortex linking task-positive and task-negative networks. *Neuroimage*, 200:259–274, Oct 2019.

[16] Anish Mitra, Abraham Z. Snyder, and Marcus E. Raichle. Probabilistic flow in brain-wide activity. *NeuroImage*, 223:117321, dec 2020.

[17] Anish Mitra, Abraham Z. Snyder, Tyler Blazey, and Marcus E. Raichle. Lag threads organize the brain's intrinsic activity. *Proceedings of the National Academy of Sciences*, 112(17):E2235–E2244, mar 2015.

[18] Anish Mitra, Abraham Z. Snyder, Carl D. Hacker, Mrinal Pahwa, Enzo Tagliazucchi, Helmut Laufs, Eric C. Leuthardt, and Marcus E. Raichle. Human cortical-hippocampal dialogue in wake and slow-wave sleep. *Proceedings of the National Academy of Sciences*, 113(44):E6868–E6876, oct 2016.

[19] Anish Mitra and Marcus E. Raichle. How networks communicate: propagation patterns in spontaneous brain activity. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1705):20150546, oct 2016.

[20] Anish Mitra, Abraham Z Snyder, Enzo Tagliazucchi, Helmut Laufs, and Marcus E Raichle. Propagated infra-slow intrinsic brain activity reorganizes across wake and slow wave sleep. *eLife*, 4, 2015.

[21] Julia M. Huntenburg, Pierre-Louis Bazin, and Daniel S. Margulies. Large-Scale Gradients in Human Cortical Organization. *Trends in Cognitive Sciences*, 22(1):21–31, Jan 2018.

[22] Yuliy Baryshnikov and Emily Schlafly. Cyclicity in multivariate time series and applications to functional MRI data. In *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, dec 2016.

[23] Benjamin J. Zimmerman, Ivan Abraham, Sara A. Schmidt, Yuliy Baryshnikov, and Fatima T. Husain. Dissociating tinnitus patients from healthy controls using resting-state cyclicity analysis and clustering. *Network Neuroscience*, 3(1):67–89, jan 2019.

[24] D.C. Van Essen, K. Ugurbil, E. Auerbach, D. Barch, T.E.J. Behrens, R. Bucholz, A. Chang, L. Chen, M. Corbetta, S.W. Curtiss, S. Della Penna, D. Feinberg, M.F. Glasser, N. Harel, A.C. Heath, L. Larson-Prior, D. Marcus, G. Michalareas, S. Moeller, R. Oostenveld, S.E. Petersen, F. Prior, B.L. Schlaggar, S.M. Smith, A.Z. Snyder, J. Xu, and E. Yacoub. The Human Connectome Project: A data acquisition perspective. *NeuroImage*, 62(4):2222–2231, oct 2012.

[25] Somayeh Shahsavarani, Ivan T. Abraham, Benjamin J. Zimmerman, Yuliy M. Baryshnikov, and Fatima T. Husain. Comparing Cyclicity Analysis With Pre-established Functional Connectivity Methods to Identify Individuals and Subject Groups Using Resting State fMRI. *Frontiers in Computational Neuroscience*, 13, January 2020.

[26] Ryan V Raut, Anish Mitra, Scott Marek, Mario Ortega, Abraham Z Snyder, Aaron Tanenbaum, Timothy O Laumann, Nico U F Dosenbach, and Marcus E Raichle. Organization of Propagated Intrinsic Brain Activity in Individual Humans. *Cerebral Cortex*, 30(3):1716–1734, sep 2019.

[27] Sadia Shakil, Chin-Hui Lee, and Shella Dawn Keilholz. Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states. *NeuroImage*, 133:111–128, jun 2016.

[28] Kuo-Tsai Chen. Integration of Paths–A Faithful Representation of Paths by Noncommutative Formal Power Series. *Transactions of the American Mathematical Society*, 89(2):395, November 1958.

[29] Patric Bonnier, Patrick Kidger, Imanol Perez Arribas, Cristopher Salvi, and Terry Lyons. Deep signature transforms. *arXiv preprint arXiv:1905.08494*, 2019.

[30] Antonio Auffinger, Michael Damron, and Jack Hanson. 50 years of first passage percolation. *arXiv:1511.03262 [math-ph]*, September 2016. arXiv: 1511.03262.

[31] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591–646, May 2009.

[32] Mark Jelasity, Alberto Montresor, and Ozalp Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Trans. Comput. Syst.*, 23(3):219–252, August 2005.

[33] I. Gupta, A.-M. Kermarrec, and A.J. Ganesh. Efficient and adaptive epidemic-style protocols for reliable and scalable multicast. *IEEE Transactions on Parallel and Distributed Systems*, 17(7):593–605, 2006.

[34] Joao Leitao, Jose Pereira, and Luis Rodrigues. Hyparview: A membership protocol for reliable gossip-based broadcast. In *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)*, pages 419–429, 2007.

[35] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74:47–97, Jan 2002.

[36] Matthew F. Glasser, Stamatios N. Sotiropoulos, J. Anthony Wilson, Timothy S. Coalson, Bruce Fischl, Jesper L. Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R. Polimeni, David C. Van Essen, and Mark Jenkinson. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124, oct 2013.

[37] Deanna M. Barch, Gregory C. Burgess, Michael P. Harms, Steven E. Petersen, Bradley L. Schlaggar, Maurizio Corbetta, Matthew F. Glasser, Sandra Curtiss, Sachin Dixit, Cindy Feldt, Dan Nolan, Edward Bryant, Tucker Hartley, Owen Footer, James M. Bjork, Russ Poldrack, Steve Smith, Heidi Johansen-Berg, Abraham Z. Snyder, and David C. Van Essen. Function in the human connectome: Task-fmri and individual differences in behavior. *NeuroImage*, 80:169–189, 10 2013.

[38] D Rubino, KA Robbins, and NG Hatsopoulos. Propagating waves mediate information transfer in the motor cortex. *Nat Neurosci*, 9:1549–57, Dec 2006.

[39] RG Townsend, SS Solomon, SC Chen, AN Pietersen, PR Martin, SG Solomon, and P Gong. Emergence of complex wave patterns in primate cerebral cortex. *J Neurosci*, 35:4657–62, Mar 2015.

[40] JA Roberts, LL Gollo, RG Abeysuriya, G Roberts, PB Mitchell, MW Woolrich, and M Breakspear. Metastable brain waves. *Nat Commun*, 10:1056, Mar 2019.

[41] H Zhang, AJ Watrous, A Patel, and J Jacobs. Theta and Alpha Oscillations Are Traveling Waves in the Human Neocortex. *Neuron*, 98:1269–1281.e4, Jun 2018.

[42] L Muller, G Piantoni, D Koller, SS Cash, E Halgren, and TJ Sejnowski. Rotating waves during human sleep spindles organize global patterns of activity that repeat precisely through the night. *Elife*, 5, Nov 2016.

[43] Esfahlani F Zamani, Y Jo, J Faskowitz, L Byrge, DP Kennedy, O Sporns, and RF Betzel. High-amplitude cofluctuations in cortical activity drive functional connectivity. *Proc Natl Acad Sci U S A*, 117:28393–28401, Nov 2020.

[44] JS Anderson, MA Ferguson, M Lopez-Larson, and D Yurgelun-Todd. Connectivity gradients between the default mode and attention control networks. *Brain Connect*, 1:147–57, 2011.

[45] DS Margulies, SS Ghosh, A Goulas, M Falkiewicz, JM Huntenburg, G Langs, G Bezgin, SB Eickhoff, FX Castellanos, M Petrides, E Jefferies, and J Smallwood. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proc Natl Acad Sci U S A*, 113:12574–12579, Nov 2016.

[46] Gadi Goelman, Rotem Dan, Filip Růžička, Ondrej Bezdicek, Evžen Růžička, Jan Roth, Josef Vymazal, and Robert Jech. Frequency-phase analysis of resting-state functional mri. *Scientific reports*, 7(1):1–13, 2017.

[47] Ryan V Raut, Anish Mitra, Scott Marek, Mario Ortega, Abraham Z Snyder, Aaron Tanenbaum, Timothy O Laumann, Nico UF Dosenbach, and Marcus E Raichle. Organization of propagated intrinsic brain activity in individual humans. *Cerebral Cortex*, 30(3):1716–1734, 2020.

[48] Nan Xu, Peter C Doerschuk, Shella D Keilholz, and R Nathan Spreng. Spatiotemporal functional interactivity among large-scale brain networks. *Neuroimage*, 227:117628, 2021.