

Probabilistic flow in brain-wide activity

Anish Mitra^{a,*}, Abraham Z. Snyder^b, Marcus E. Raichle^b

^a Department of Psychiatry, Stanford University, 401 Quarry Rd, Stanford, CA 94304, United States

^b Department of Neurology, Washington University, St. Louis, United States

A B S T R A C T

Patterns of low frequency brain-wide activity have drawn attention across multiple disciplines in neuroscience. Brain-wide activity patterns are often described through correlations, which capture concurrent increases and decreases in neural activity. More recently, several groups have described reproducible temporal sequences across the brain, illustrating precise long-distance control over the timing of low frequency activity. Features of correlation and temporal organization both point to a systems-level structure of brain activity consisting of large-scale networks and their mutual interactions. Yet a unified view for understanding large networks and their interactions remains elusive. Here, we propose a framework for computing probabilistic flow in brain-wide activity. We demonstrate how flow probabilities are modulated across rest and task states and show that the probabilistic perspective captures both intra- and inter-network dynamics. Finally, we suggest that a probabilistic framework may prove fruitful in characterizing low frequency brain-wide activity in health and disease.

1. Introduction

Findings across fields in neuroscience point to the importance of low frequency (<10 Hz) brain-wide activity as means for understanding mammalian brain function. In human neuroimaging, from the earliest days of PET, researchers quickly realized that even simple tasks engaged metabolic changes in large parts of the brain (Raichle et al., 2001; Raichle and Mintun, 2006). Recent work in mice has also shown that broad behavioral conditions such as thirst and decision state not only modulate brain-wide activity, but does so at the single neuron scale across the brain (Allen et al., 2019, 2017).

Low frequency spontaneous, or ‘resting-state’, neural activity is also coordinated across the brain in the absence of an explicit task. Such coordination is often measured using correlations. A large body of work using voltage sensitive dyes in cats (Kenet et al., 2003), calcium imaging and fMRI in mice (Ma et al., 2016; Stafford et al., 2014), and fMRI and electrocorticography in primates including humans (Fox and Raichle, 2007; Hacker et al., 2017; Vincent et al., 2007) has demonstrated spontaneous correlations across brain areas that are recruited in tasks, such as the visual system, the somatomotor system, and higher order networks like control networks (Vincent et al., 2008) and the default mode network (Raichle et al., 2001). Indeed, the magnitudes of pairwise correlations in low frequency activity are often used to parcel the brain into approximately 7 canonical networks (Hacker et al., 2013; Yeo et al., 2011). Taken further, correlation information can inform defining even smaller brain areas, such as recent report suggesting 180 areas in human cortex (Glasser et al., 2016).

In contrast to the parsing of the brain into ever finer balkanized areas, we and others have recently found propagation patterns in which low frequency activity travels both within and between large brain areas

and networks. We first demonstrated the existence of brain-wide temporal sequences in human fMRI (Mitra et al., 2014). While the high dimensional temporal delay pattern of fMRI data (Mitra et al., 2015) suggested that neurophysiology, as opposed to vascular delays (Handwerker et al., 2004), were responsible for the temporal structure of fMRI, fMRI evidence alone could not definitively delineate between these possibilities. We therefore further pursued concurrent optical imaging of blood flow and calcium signals in the mouse, alongside laminar electrophysiology, to demonstrate that the temporal delay structure of the blood oxygen signal precisely mirrors the temporal delay structure of both low frequency calcium signals and low frequency local field potentials (Mitra et al., 2018). We also demonstrated using human electrocorticography that low frequency temporal delay structure in human local field potentials agrees with temporal delays found in human fMRI (Mitra et al., 2016). Finally, recent work combining fMRI and optogenetics demonstrates that stimulated pulses in the brain propagate across networks at a systems-scale and are captured accurately using fMRI (Lee et al., 2010; Leong et al., 2019).

In sum, beyond establishing agreement between temporal structure in blood oxygen signals and physiology, mounting evidence suggests that low frequency activity propagates both within and across traditionally defined networks (Leong et al., 2019; Mitra et al., 2018, 2016). Thus, the temporal structure of low frequency activity suggests that rather than serving to delineate and divide brain regions, this activity instead serves as a means of inter-areal communication and integration (Mitra and Raichle, 2018; Raichle et al., 2019).

On the basis of these findings, the question naturally arises, how can the same signal produce signatures of brain-wide integration as well as inter-areal segregation? In past work, we have shown that the answer lies, in part, in distinct features of intra- vs. inter-network propagation of

* Corresponding author.

E-mail address: mitraa@wusm.wustl.edu (A. Mitra).

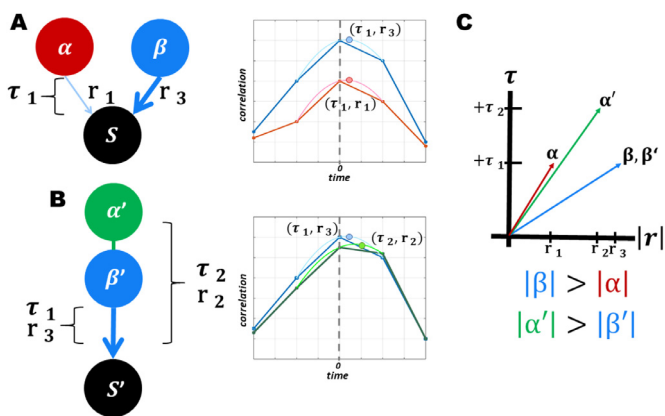


Fig. 1. Toy examples to motivate the probabilistic flow metric. **(A)** A system in which two nodes, α and β are both sending signals (modeled as aperiodic fMRI signals as in (Mitra et al., 2015)) to a seed node, S , with a temporal delay of τ_1 . However, the node β is three times more likely to send a signal to S than is the node α . Therefore, the correlation r_3 between β and S is three times greater than correlation r_1 between α and S . The pairwise lagged correlation curves between β and S (blue) as well as α and S (red) are shown on the right side of panel (A). As is evident, the peak of the two lagged correlation curves both occur at τ_1 . **(B)** A three node system in which activity travels first from α' to β' , then from β' to a seed region S . The temporal delay τ_2 between α' and S is twice that of the temporal delay τ_1 between β' and S . Moreover, assuming that there's some change in the waveform of the signal during this sequential transmission, the correlation r_2 between α' and S is ten percent less than the correlation r_3 between β' and S . **(C)** A vector space in which the x-axis is correlation and the y-axis is temporal delay. If we plot the vectors from (A) corresponding to the correlation and time delay between α and S (red) and β and S (blue), the magnitude of the blue vector is greater than the magnitude of the red vector, correctly capturing that β is more likely to send signal to S than is α . If we plot the vectors from (B) corresponding to the correlation and time delay between α' and S (green) and β' and S (blue). The α' vector has a higher magnitude than the β' vector (by approximately fifteen percent in this example), reflecting the fact that it is more likely to be the original source of activity flowing to S . (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

low frequency activity (Mitra et al., 2015). Here, we build on this work by proposing a probabilistic view of propagated low frequency activity built on both pairwise correlations and pairwise temporal delays between signals. In this probabilistic framework, we assume that all fMRI activity reflects either sending or receiving signals in the brain. Thus, for a given voxel-wise time series, we compute the probability that each other voxel in the brain was the original source of its activity, and the probability that each other voxel in the brain is the ultimate destination of its activity.

We demonstrate in fMRI data that this probabilistic perspective offers a view both into how low-frequency activity moves between networks for cross-areal communication and how activity can be biased to remain within certain areas which define networks. Finally, we suggest that subtle alterations in spontaneous propagation probabilities may underlie both physiologic and pathophysiologic aspects of brain function.

2. Theory

The correlation and temporal delay between a pair of time series can be independent. In other words, we can create a pair of time series with a fixed peak correlation of r at an arbitrary temporal delay τ_i (where τ_i could be 0.5 s, 1 s, 10 s, etc.). Likewise, we can fix the peak correlation to occur at some temporal delay τ at an arbitrary value of r_i . As a result, an analysis of data that excludes either correlation or temporal delay information can overlook important dynamics of a system.

Fig. 1 illustrates how both correlation and temporal delay can be essential for characterizing types of dynamical systems. Panel (A) shows a

situation in which two nodes, α and β are both sending signals (modeled as aperiodic fMRI signals as in (Mitra et al., 2015)) to a seed node, S , with a temporal delay of τ_1 . However, the node β is three times more likely to send a signal to S than is the node α . Therefore, the correlation r_3 between β and S is three times greater than correlation r_1 between α and S . The pairwise lagged correlation curves between β and S (blue) as well as α and S (red) are shown on the right side of panel (A). As is evident, the peak of the two lagged correlation curves both occur at τ_1 . Hence, an analysis of temporal delay alone would incorrectly suggest both areas signal to S equally.

Now consider a vector space (Panel (C)) in which the x-axis is correlation and the y-axis is temporal delay. If we plot the vectors corresponding to the correlation and time delay between α and S (red) and β and S (blue), the magnitude of the blue vector is greater than the magnitude of the red vector, correctly capturing that β is more likely to send signal to S than is α . In this toy example, the magnitude of β is approximately 2.5 that of α ; thus the metric does not perfectly capture the underlying probabilities, but it does offer an estimate that is significantly more accurate than time delay alone.

Another feature of this vector space is that it distinguishes between proximal and terminal sources and sinks in a system. As an example, consider Panel (B) in Fig. 1. In this toy example, activity travels first from α' to β' , then from β' to a seed region S . The temporal delay τ_2 between α' and S is twice that of the temporal delay τ_1 between β' and S . Moreover, assuming that there's some change in the waveform of the signal during this sequential transmission, the correlation r_2 between α' and S is ten percent less than the correlation r_3 between β' and S . As before, we plot the vectors corresponding to the correlation and time delay between α' and S (green) and β' and S (blue). The α' vector has a higher magnitude than the β' vector (by approximately fifteen percent in this example), reflecting the fact that it is more likely to be the original source of activity flowing to S .

Taken together, these examples illustrate how a vector space of correlations and temporal delays can be used to compute the probability of a particular region receiving signals from other nodes in the system. Specifically, the present framework computes whether a particular voxel was the original source or ultimate destination of a seed region's activity. As such, long temporal delays are weighted highly, as the combination of high correlation and long temporal delay is the characteristic of an area that was the ultimate source or destination of a voxel time series. Before we apply this approach to real data, there are five technical factors to consider.

First, in order for the magnitude of vectors in this space to equivalently reflect correlation and temporal delay information, correlations and temporal delays must be on the same scale. For example, if temporal delays were on the order of 10's of seconds whereas correlations were on a scale of -1 to 1 , then the magnitude of vectors in this space would primarily reflect time delays. We thus normalized the variance of correlations and time delays to place them on the same scale.

Second, correlations can be positive or negative. In the context of neural activity, positive correlations may represent spread of excitatory activity, whereas negative correlations may represent spread of inhibition. While teasing apart positive and negative correlations may be critical in some applications, in the present treatment we will consider only the magnitude of correlations, potentially grouping excitatory and inhibitory spread of activity.

Third, temporal delays can be positive or negative, reflecting areas that either send or receive signals from the reference point. In the analyses to follow, we will maintain sender vs. receiver information by attaching a sign to the magnitudes of vectors (negative signs for senders, positive signs for receivers with respect to the reference node).

Fourth, in order for the vector magnitude metric to represent a probability, the cumulative sum of all vector magnitudes from a reference node should equal one. To accomplish this, we will assume on the basis of prior work that there is no voxel in the brain that is either a "universal sender" or a "universal receiver" (Mitra et al., 2014). Instead, we




Fig. 2. Left hemisphere motor hand area (pink arrow and dots) referenced temporal lag, correlation, and probabilistic flow maps. The far left column, “Rest lags”, demonstrates temporal delays in units of seconds with respect to the reference region in 3 slices in the brain, arranged from dorsal (top) to ventral (bottom). Negative lags (cool hues) represent regions which send fMRI signal to the region of interest, whereas positive lags (warm hues) represent areas which receive signal from the region of interest. The second column from the left, “Rest corr”, illustrates the peak correlation between the region of interest and every other voxel in the brain. Cool hues represent negative correlations and warm hues represent positive correlations. The third column from the left, “Rest prob”, depicts probabilistic flow about the region of interest. Cool hues represent areas that send signals to the region of interest whereas the warm hues represent areas that receive signals from the motor region of interest. The sum of all the cool hues equals one, as does the sum of all the warm hues. Areas with no coloring represent zero probability for sending and receiving signals from the motor seed. The far right column depicts probabilistic flow about the region of interest during a visual-motor task. Orange arrows highlight ventral visual areas; purple and green arrows highlight primary and secondary somatosensory areas, respectively. Temporal delay and correlation analysis of the visual-motor task have been previously published (Fox et al., 2007; Mitra et al., 2014). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

will assume that every voxel in the brain receives signal from at least one other place in the brain, and also sends signal to at least one other place in the brain. Moreover, as suggested by the optogenetic-fMRI literature, we will assume that all activity in human fMRI reflects signal propagation through the brain (Lee et al., 2010; Leong et al., 2019). On this basis, in this model of probabilistic flow we will normalize the sum of all incoming or “sender” signal probabilities to one and normalize the sum of all outgoing or “receiver” signal probabilities to one. Therefore the vector magnitudes of incoming and outgoing signals from each voxel satisfy the properties of a true probability.

Fifth, and finally, in order to compute a time delay between two signals, there must be some amount of correlation. If, for instance, the correlation between two signals were zero, then there would not exist a peak correlation to define a time delay. Hence we only compute probabilities for pairwise interactions with a certain minimum threshold correlation (see Methods).

3. Results

Fig. 2 demonstrates the application of a probabilistic framework to fMRI data (dataset 1 in Table 1). The analyses are performed with re-

Table 1
Demographic and fMRI scanning parameters of analyzed data sets.

Dataset	1	2
Number of subjects	17 (9 F)	1376 (785 F)
Age in years	23.1 ± 2.4 (S.D.)	21.4 ± 2.1 (S.D.)
Scanner	Siemens Allegra	Siemens Tim Trio
Acquisition voxel size	(4 mm) ³	(3 mm) ³
Flip Angle	90°	85°
Repetition Time (s)	2.16	3.00
Number of frames	194 × 2 runs rest, 194 × 3 runs task	124 × 2 runs
Citation	Fox et al., 2007	Buckner et al., 2014
Experimental Question	Rest vs. task	Resting state

spect to a reference node (or seed) in the left motor cortex hand area (pink arrow and dots in Fig. 2).

The far left column, “Rest lags”, demonstrates temporal delays in units of seconds with respect to spontaneous fMRI activity in the left motor cortex hand area in 3 slices in the brain, arranged from dorsal (top) to ventral (bottom). Negative lags (cool hues) represent regions which send fMRI signal to the region of interest, whereas positive lags (warm hues) represent areas which receive signal from the region of interest. In line with our previous findings (Mitra et al., 2015), pre-motor areas, medial somatosensory areas, and secondary somatomotor cortex are among the areas which send fMRI signals to the region of interest. Similarly, ventral visual cortex and lateral somatomotor cortices are among the areas which receive fMRI signals from the region of interest.

The second column from the left, “Rest corr”, illustrates the peak correlation between the motor region of interest and every other voxel in the brain. Correlation is reported in units of Pearson r, with cool hues representing negative correlations and warm hues representing positive correlations. In line with extensive prior literature, the left hemispheric motor reference point is correlated with the right somatomotor cortex as well as parts of visual cortex.

The third column from the left, “Rest prob”, depicts probabilistic flow computed using the correlation-latency vector space as described in the theory section. The cool hues represent areas that send signals to the region of interest whereas the warm hues represent areas that receive signals from the motor region of interest. The sum of all the cool hues equals one, as does the sum of all the warm hues. Several aspects of the probabilistic flow map are similar to the lag map on the far left, such as the fact that pre-motor areas (purple arrows) are a high probability sender of activity to the region of interest. However, there are also differences. For example, ventral visual areas (orange arrows) have a long temporal delay with respect to the region of interest (far left column in Fig. 2), but only a modest probability of receiving signal from the motor region of interest (third column in Fig. 2). Moreover, there are several areas with green hues in the lag map, which could reflect either synchrony with zero-temporal delay or regions without well-defined temporal relationships with the seed (and hence averaging to zero). The majority of these regions have no coloring in the probability map, indicating near zero probability of flow of activity with the seed region.

Finally, in the far right column in Fig. 2, we also computed in this same group of subjects a probabilistic flow map during right handed button push in response to a visual cue (see Methods). Notice that the flow map is markedly altered from the resting state. During this simple task, the left motor hand region of interest is now predominantly receiving signals from medial somatosensory cortex; the premotor and secondary somatosensory cortices has negligible probability of sending signals to the region of interest (purple and green arrows, respectively). In contrast, the probability of sending signals to the ventral visual system is markedly increased (orange arrows), as well as parts of the dorsal visual system. Note that this is contrary to what might be expected in a feed-forward view of a simple visual-motor task, as visual cortex be-




Fig. 3. Whole brain analysis of probabilistic flow in resting state fMRI. (A) A 2-dimensional matrix depicting pairwise voxel relationships in fMRI data. Each point in the matrix illustrates the pairwise probability of fMRI signal flow between two voxels in the brain. Cool hues represent “sending” in the pairwise relationship and warm hues represent “receiving” in the pairwise relationship. The matrix was block sorted by network affiliation (see Methods): dorsal attention network (DAN), ventral attention network (VAN), somatomotor network (SMN), visual network (VIS), frontopari-

etal control network (FPC), language network (LAN), and default mode network (DMN). In addition, each diagonal block was further sorted into “high probability senders” and “high probability receivers”, yielding the blue-to-red progression in each diagonal block. (B) Mean absolute value probabilistic flow values within each intra- and inter-network block. The purpose of computing means within each intra- and inter-network block is to compare the relative probability of signaling within each block. Warm colors indicate relatively high signaling probabilities, whereas low blocks represent relatively low signaling probabilities. Warm colors along the diagonal indicate that, generally, intra-network flow has higher probability than inter-network flow. However, high probabilities are also found in several inter-network pairs, most notably between the DMN and DAN. Asterisks designate blocks where the mean absolute probability is significantly higher than expected in the absence of structured temporal delays (see Methods). (C) Average inter-network flow from each network. The bar graph illustrates the mean value of inter-network flow probabilities depicted in panel (B) for each network with confidence intervals computed using bootstrapping representing 2.5 standard deviations (see Methods). The gray line indicates the mean probabilistic flow value computed over all inter-network pairs. Values above the gray line indicate greater inter-network signal flow than average, and vice versa. Significance testing (see Methods) reveals that the DMN participates in more inter-network signal flow than average, whereas the SMN participates in less signal flow than average. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

comes more likely to receive as opposed to send signals with respect to the left motor hand area.

Fig. 2 illustrates the concept of probabilistic flow with respect to a reference node or seed region of interest. Fig. 3 demonstrates a whole-brain analysis of spontaneous fMRI activity in a group of 1376 subjects (see Table 1). Fig. 3A depicts a 2-dimensional matrix of pairwise voxel relationships in fMRI data; each point in the matrix illustrates the pairwise probability of fMRI signal flow between two voxels in the brain. As in Fig. 2, the cool hues represent “sending” in the pairwise relationship and warm hues represent “receiving” in the pairwise relationship. The matrix has further been sorted into blocks by known network affiliation (Hacker et al., 2013), see Fig. 3 caption and Methods). In addition, each diagonal block was further sorted into “high probability senders” and “high probability receivers”, yielding the blue-to-red progression in each diagonal block.

Note that the highest saturation in colors occurs along the diagonal blocks. The implication is that the highest probability of spontaneous fMRI signal flow occurs within large brain networks. However, there is still significant probability of cross-network signal flow.

To further analyze probabilistic signal flow within and across networks, we next computed the average absolute value of probability flow values inside each intra-network and inter-network block. Regardless of sign, higher absolute probability flow value reflect, in this model, higher chances of signal flow between a set of voxels. The result, shown in Fig. 3B, shows that in general intra-network relationships have higher probability flow values than inter-network relationships. However, there are several inter-network pairs with comparatively high probability flow averages, including the DMN:DAN relationship and the VIS:FPC relationship. To assess whether the values in Fig. 3B are significant, we simulated fMRI time series with the same correlation structure as real data (based on the group average) but without any constraint on the temporal delay structure (see Methods). Blocks with white asterisks in Fig. 3B designate blocks with statistically significant elevations in probability flow metrics compared to a null model that does not constrain temporal delay information ($p < 0.05$ corrected, see Methods).

Focusing specifically on inter-network signal flow, we also explored whether some networks participate more or less in inter-network signal flow by computing the mean inter-network flow through each network in Fig. 3B, as illustrated in Fig. 3C. We find that compared to the mean over all networks, the DMN has significantly higher flow with other

networks, whereas the SMN has significantly lower flow with other networks ($p < 0.05$ corrected, see Methods).

Finally, we explored the topology of pairwise voxel relationships in the probabilistic flow metric. Given the high dimensionality of the data, we applied t-distributed stochastic neighbor embedding (t-sne; see Methods) to produce a 2-dimensional visualization of high dimensional voxel relationships. Fig. 4A illustrates the results of applying t-sne to the probability flow metric matrix shown in Fig. 3A. The t-sne result in Fig. 4A shows clear clustering of networks along three principle axes: the default mode network, the visual network, and the somatosensory network, in line with the 3 principal gradients of brain activity recently reported by Margulies and colleagues (Margulies et al., 2016). Indeed, as Margulies and colleagues reported using the gradient method, the remaining networks of the brain in Fig. 4A are circumscribed within the triangular organization between the somatomotor, visual, and default mode networks. To determine whether the result in Fig. 4A was found by chance, we produced surrogate resting state networks with spatial frequency content matched to the real resting state networks (as in (Mitra et al., 2014), see Methods). We produced 1000 null simulations of surrogate resting state networks and none produced a clustering with correlation $r > 0.01$ with the present results.

While Fig. 4A illustrates that the probabilistic flow framework recapitulates the functional topology reported by (Margulies et al., 2016), it is possible that other metrics, such as correlation (typical functional connectivity), temporal delays, or even physical distance produce the same topology. These possibilities are explored in Figs. 4B-D. Fig. 4B illustrates t-sne applied to pairwise correlations, and shows that although individual networks are well clustered, the inter-network relationships embedded in Fig. 4A are not present. Instead of the triangular organization about three axes, the t-sne analysis of pairwise correlations reveals a topology with the networks distributed approximately equidistantly along a circle. Fig. 4C illustrates t-sne applied to temporal delays and reveals a clustering topology that does capture elements of intra-network relationships including the triangular organization about somatomotor, visual, and default mode networks; however, the networks themselves are not coherently identified, with the somatomotor, dorsal attention, and ventral attention networks split into multiple clusters. Finally, t-sne applied to physical distance in Fig. 4D fails to capture both inter- and intra-network relationships.

Therefore, the probabilistic flow metric implemented in the present resting state fMRI analyses captures both inter- and intra-network com-

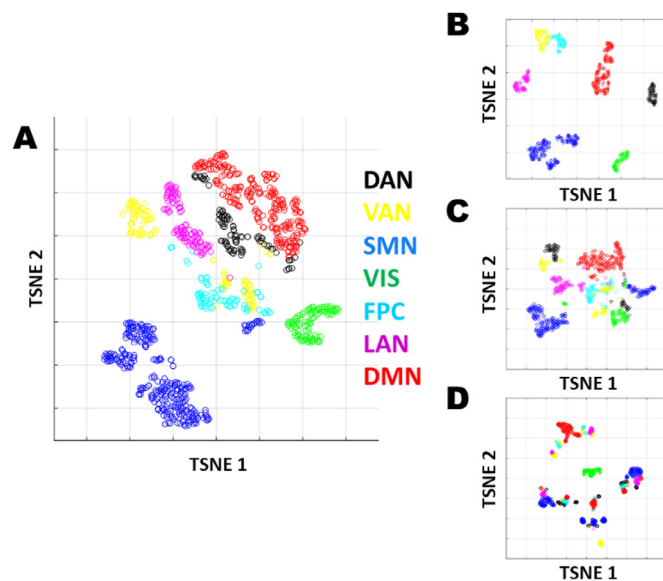


Fig. 4. Clustering of voxel-pair probabilistic flow, correlation, temporal delay, and physical distance matrices. The high dimensional cluster structure is here visualized in two dimensions by applying t-distributed stochastic neighbor embedding (t-sne; see Methods). (A) Clustering structure of the probabilistic flow metric. Note that the networks clearly segregate from each other and align along three primary points: the default mode network (red), the somatosensory network (blue), and the visual network (green). (B) Clustering structure of the correlation metric. Note that networks are well clustered along a circular shape in state-space but no clear inter-network features are identified. (C) Clustering structure of the temporal delay metric. Note that although elements of inter-network relationships including the three-point organization around the DMN, SMN, and VIS is present, several networks are not well-clustered, including the DAN and SMN. (D) Clustering structure of physical distance metric. Neither intra-network nor inter-network relationships are well captured. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

munication in a fashion that contains both network segregating and network integrating properties. Importantly, the topology of these relationships is not well defined by alternative metrics including correlation, temporal delays, and physical distance.

4. Discussion

In summary, we propose a probabilistic framework for analyzing high dimensional neural time series. The probabilistic flow metric introduced here estimates the original source and terminal destination of activity by computing normed magnitudes in a vector space defined by correlations and temporal delays. By incorporating both correlation and temporal delay information, the probabilistic flow metric successfully characterizes dynamics that are missed by both sole application of temporal delays or correlations, as shown in Fig. 1. Moreover, in a simple visual-motor task, we demonstrate that task-induced flow of fMRI activity represents a modulation of spontaneous activity (Fig. 2). We further performed series of analyses exploring the whole-brain structure of pairwise probabilistic flow relationships in spontaneous fMRI data. We find that intra-network flow has higher probability than inter-network flow, but that several inter-network relationships are privileged in the brain, including signaling between the dorsal attention network and the default mode network (Fig. 3). Finally, we applied t-sne to explore the topology of pairwise probabilistic flow relationships in spontaneous fMRI data to find signaling organized around three networks: the default mode network, the visual network, and the somatosensory network, in line with recent gradient-based findings from Margulies and colleagues (Margulies et al., 2016). Critically, similar t-sne analysis pairwise cor-

relations, temporal delays, and physical distance all fail to capture this organization in spontaneous fMRI data (Fig. 4).

4.1. Inter-network relationships and gradients in brain function

Correlations and temporal delays offer distinct information about the organization of brain-wide neural activity. Conventionally, correlation information has been used to parse the brain into networks and regions, whereas latency information has been used to demonstrate flow of activity across areas. Here, we describe a probabilistic flow metric based on correlation and latency information as a means of estimating the probability that a given node is sending or receiving signals to other nodes in the brain. By construction, we chose to focus on terminal sources and sinks, that is, the originating and final nodes along signal sequences (see Theory).

Applying this probabilistic flow framework to fMRI data, our results demonstrate that fMRI signals travel both within and across large brain networks, but with different probabilities. Figs. 3A-B show that while activity is most likely to travel within previously defined resting state networks, there is also a significant probability of signaling between networks. Two networks stand out in this analysis: the default mode network, which engages in more cross-network signaling than the average across networks, and the somatomotor network, which engages in less cross-network signaling than the average across networks (Fig. 3C). The results in Fig. 3 are consistent with prior work showing that the default mode network is involved in global brain signaling whereas signaling in the somatomotor network is comparatively more local (Vatanssever et al., 2015).

The results of applying a clustering algorithm to the probabilistic flow matrix in Fig. 3A further reinforces features of both intra- and inter-network brain organization (Fig. 4A) that are not captured by conventional correlations (Fig. 4B). In the clustering produced using conventional functional connectivity (Fig. 4B), as expected, each of the networks is well clustered (with the exception of the ventral attention network and frontoparietal control network, which are known to have functional overlap (Hacker et al., 2013)). However, the network clustering in Fig. 4B does not reveal any clear relationship between networks. For example, on the basis of Fig. 4B, there is no evidence that the default mode network may have a closer functional relationship with the dorsal attention or language networks as compared to the somatomotor network. To take another example, the functional connectivity results in Fig. 4B do not point to any inter-network relationship between the visual and fronto-parietal control networks.

In contrast, the probabilistic flow clustering results in Fig. 4A delineate both intra- and inter-network groupings. Aside from the previous exception of the overlap between the ventral attention and frontoparietal control networks, each network is again well clustered in that the nodes corresponding to each network are grouped closer to each other than they are to other networks. In addition, the probabilistic flow clustering in Fig. 4A places networks in groups implying inter-network relationships. For example, the language and dorsal attention networks are in close proximity to the default mode network, and the frontoparietal control network is positioned next to the visual network. The basis for these groupings is evident in Fig. 3B, as these inter-network pairs each have high signaling probability. Prior work has given credence to the concept that groups of networks coordinate their activities during tasks. In particular, functional relationships between the default mode and dorsal attention networks (Kim, 2015), default mode and language networks (Gordon et al., 2020), and visual and frontoparietal control networks (Corbetta, 1998) have all been previously reported in task performance.

In addition to highlighting specific pairwise relationships between networks, the probabilistic-flow clustering in Fig. 4A reveals a whole-brain functional organization built around three primary networks: the default mode network, the somatosensory network, and the visual network, precisely in line with prior work demonstrating continuous gra-

dients in brain organization (Margulies et al., 2016). On the basis of the Margulies results and our current findings, we hypothesize that the distinction between parcellation of brain areas based on differential correlation and integration of regions based on signal flow may simply be a question of probability. Small, distinct brain regions reflect patches of very high probability signal flow, whereas a focus on broader signal flow at lower probabilities yields an integrated view of how distinct brain areas communicate.

As illustrated in Figs. 4B-C, use of the correlation and temporal delay metric fails to capture the inter-network and intra-network features, respectively, characterized by the probabilistic flow framework. Therefore, not only is there utility probability flow framework for balancing inter and intra-network organization in neural activity, we hypothesize that probabilistic flow of low frequency activity may be the basis of the observed organization of spontaneous fMRI data.

4.2. State changes in probability

One of the fundamental questions in systems neuroscience is how changes in brain-wide activity relate to behavior. The probabilistic flow metric may offer some insight. Our results reveal a marked change in probabilistic flow, with respect to the left motor hand area, between the resting state and a simple visual-motor task (button push to visual cue, see Methods). Interestingly, the changes observed during task appear to represent a modulation of the flow probabilities during rest. For example, ventral visual areas (orange arrows in Fig. 2) were a low probability receiver of left motor hand area activity during rest, but become a high probability receiver during task. Medial somatosensory cortex, on the other hand, continues to be a high probability sender of activity to the left motor hand area during rest and task. Therefore, we hypothesize that spontaneous flow of low frequency activity may represent a broad set of possible paths for activity, which becomes sharpened into a narrower set to accomplish particular behaviors. Future studies are required to investigate this possibility.

The specific directionality of the activity in Fig. 2 is also informative. One may expect that during a task in which a person pushes a button in response to a visual cue, that activity should travel from visual cortex to motor areas. Indeed, this represents the feed-forward direction of activity in such a task. However, to the contrary, our probabilistic flow maps instead show that activity becomes more likely to flow from motor to visual areas during the task. This finding is in line with prior work suggesting that low frequency activity in the brain, especially frequencies represented by fMRI, may reflect feedback as opposed to feed-forward activity (Mitra et al., 2018, 2016). More work is merited to further define the physiology and function of low frequency neural activity.

Finally, the probabilistic perspective, which integrates correlation and latency information, may prove useful in detecting subtle pathophysiological changes in brain function in neuropsychiatric conditions such as autism spectrum disorders, schizophrenia, and Alzheimer's disease. To date, clear, reproducible changes in low frequency activity that are not attributable to motion have been difficult to document. However, one might imagine that opposed to large changes in correlation or directionality, these conditions instead produce subtle alterations in flow probabilities which in turn affect other neurophysiological processes. Several extant data sets offer the possibility of testing this hypothesis.

4.3. Bilateral symmetry

One striking feature of the probabilistic flow metric in Fig. 2, in the resting state and the task condition, is the bilateral symmetry observed in the maps, even though the task involves using only the right hand in response to a visual stimulus. Our results are in agreement with previous work which has shown that low frequency brain activity is remarkably symmetric across the hemispheres (Mohajerani et al., 2010). Even in tasks that are known to recruit unilateral brain areas, such as language

tasks which preferentially recruit Broca's area in the left hemisphere, prior work has found that there is similar activity in the homologous right hemispheric region, albeit with smaller magnitude (McAvoy et al., 2016). Critically, although there is a difference in amplitude between the activity in the hemispheres evoked by the language task, there is no measurable delay in the BOLD signal between the hemispheres.

The mechanisms underlying bilateral synchrony in low frequency activity are not fully understood. Prior work has shown that corpus callosotomy leads to a reduction in correlation between the hemispheres in human subjects (Roland et al., 2017). Therefore, continuous high frequency activity between the hemispheres may be responsible for observed bilateral symmetry. Additionally, optogenetic fMRI studies have shown that unilateral stimulation of sensory relay nuclei in the thalamus, including the lateral geniculate nucleus, evokes synchronized bilateral activity (Leong et al., 2019), which may also play a role in the present results. The functional significance of bilateral synchrony in activity also remains unknown, although it could represent the widespread sharing of neural information across the cortices (Allen et al., 2017). Further research is necessary to understand the mechanisms and function of bilateral synchrony in low frequency activity.

4.4. Limitations

The presently described approach to probabilistic flow in brain-wide activity is meant simply as a guide. There are many approaches to compute the probability of flow in high-dimensional data, and based on the question of interest, the assumptions baked into the metric presented here may not be appropriate. For example, our approach focuses on terminal sources and sinks of activity, but by instead focusing on the inverse of time delays, one could instead choose to investigate local flow of activity. The current approach is also limited to linear sequences of activity, which may not be suited to all systems. Additionally, treating correlation and latency information as equivalent is a simple assumption for building a vector space, but depending on the application, correlations or lag may appropriately carry more weight. As is the case for most fMRI findings, it is also important to note that the group average results shown here can vary from individual to individual (see individual variability analysis in Supplementary Figure 2), and although recent evidence shows a close correspondence between the spatio-temporal structure of low frequency neurophysiology and fMRI (Mitra et al., 2018, 2016), we cannot exclude that some features of the present results may be attributable to vascular phenomena.

To conclude, rather than arguing for one approach, we hope instead that the present work conveys the potential benefits of applying a probabilistic framework to the study of wide-scale brain activity. By combining correlation and latency information, probability flow techniques offer unique insight into the broad dynamics of brain activity that captures both features of both inter- and intra-areal communication. Finally, the probability flow concept need not be limited to fMRI. A similar approach can easily be utilized in the study of mesoscale calcium imaging in mice, as just one example.

5. Methods

5.1. MRI acquisition

Two fMRI datasets are analyzed in the present manuscript (see Table 1).

Dataset 1, collected with a 3T Siemens Allegra at Washington University and originally reported in (Fox et al., 2007), imaging was collected in both resting and task states. During rest, subjects were instructed to remain still, stay awake, and keep their eyes fixated on a white crosshair. During the button press task subjects were instructed to press a button in response to a visual cue, dimming of the fixation cross-hair from white to dark-gray for a period of 250 ms. Subjects were instructed to press a button with their right index finger as quickly as possible when they

saw the crosshair dim. They were told that their reaction times would be recorded. Each of these button press runs contained 20 crosshair dims time-locked to the scanner TR, with an intertrial interval of 8–14 frames (17.3–30.2 s). Subjects practiced this button-press task once in the scanner, prior to the onset of the functional scans.

Dataset 2, consisting of normal young adults ($N = 1376$, mean age 21.4 ± 2.1 years, 57% female) was obtained from the Harvard-MGH Brain Genomics Superstruct Project (Buckner et al., 2014) (Table 1). Imaging was performed with a 3T Siemens Tim Trio scanner equipped with a standard 12-channel head coil. Additional imaging details are given in (Thomas Yeo et al., 2011).

5.2. fMRI preprocessing

fMRI preprocessing was as described in (Mitra et al., 2014). Briefly, this included compensation for slice-dependent time shifts, elimination of systematic odd-even slice intensity differences due to interleaved acquisition, and rigid body correction of head movement within and across runs. Atlas transformation was achieved by composition of affine transforms connecting the fMRI volumes with the T2W and T1W structural images. Head movement correction was included with the atlas transformation in a single resampling that generated volumetric timeseries in $(3 \text{ mm})^3$ atlas space. Additional preprocessing included spatial smoothing (6 mm full width at half maximum (FWHM) Gaussian blur in each direction), voxel-wise removal of linear trends over each fMRI run, temporal low-pass filtering retaining frequencies below 0.1 Hz, and zero-meaning each voxel time series. Spurious variance was reduced by regression of nuisance waveforms derived from head motion correction and timeseries extracted from regions (of “non-interest”) in white matter and CSF. Nuisance regressors included also the BOLD timeseries averaged over the brain (Fox et al., 2005). Additionally, frame censoring was computed at a threshold of 0.5% root mean square frame-to-frame intensity change (Power et al., 2012). Epochs containing fewer than 10 contiguous frames were excluded. These criteria removed 5.2% of frames from analysis in dataset 1 and 3.5% of frames from analysis in dataset 2.

5.3. Computation of lag and peak-correlation between BOLD time series

Our method for computing lags between time series has been previously published (Mitra et al., 2014). We briefly recapitulate the methodology here. Conventional seed-based correlation analysis involves computation of the Pearson correlation, r , between the time series, $x_1(t)$, extracted from a seed region, and a second time series, $x_2(t)$, extracted from some other locus (single voxel or region of interest). Thus,

$$r_{x_1x_2} = \frac{1}{\sigma_{x_1} \sigma_{x_2}} \frac{1}{T} \int x_1(t) \cdot x_2(t) dt, \quad (S1)$$

where σ_{x_1} and σ_{x_2} are the temporal standard deviations of signals x_1 and x_2 , and T is the interval of integration. Here, we generalize the assumption of exact temporal synchrony and compute lagged cross-correlation functions. Thus,

$$C_{x_1x_2}(\tau) = \frac{1}{T} \int x_1(t + \tau) \cdot x_2(t) dt, \quad (S2)$$

where τ is the lag (in units of time). The value of τ at which $C_{x_1x_2}(\tau)$ exhibits an extremum defines the temporal lag (equivalently, delay) between signals x_1 and x_2 (Konig, 1994). Although cross-correlation functions can exhibit multiple extrema in the analysis of periodic signals, BOLD time series are aperiodic (He et al., 2010; Maxim et al., 2005), and almost always give rise to lagged cross-correlation functions with a single, well defined extremum, typically in the range ± 1 s (Raut et al., 2019). However, thresholds up to ± 2.5 s yield very similar results (see Supplemental Figure 2). We determine the extremum abscissa and ordinate using parabolic interpolation (Fig. 1). For negative correlations at zero-lag, we compute the minimum correlation in the range ± 1 s,

and for positive correlations at zero lag we compute the maximum correlation in the range ± 1 s. These values reflect the temporal delay and peak absolute value correlation between pairs of time series.

5.4. Computation of probabilistic flow

The mechanics of the probabilistic flow metric are outlined in detail in the Theory section. As noted in the Theory section, in order to compute a time delay between two signals, there must be some amount of correlation. If, for instance, the correlation between two signals were zero, then there would not exist a peak correlation to define a time delay. In practice, the amount of error in the inferred temporal delay increases as a function of decreasing correlation between signals (See Fig. 3B in (Raut et al., 2019)). Thus, to reduce the noise and increase the stability of the probability metric, we opted to apply a minimum threshold correlation for computing a pair-wise probability. Signals with pairwise peak correlations below the threshold were set to have zero-probability of direct interaction. In the analyses shown, we set the minimum correlation to $|r| = 0.10$, on the basis of significantly increased error in temporal delay estimates below this correlation value (Raut et al., 2019). We found that the scope of the finding presented herein did not meaningfully change with different thresholds of $|r| = 0.05$, $|r| = 0.15$ (data not shown). However, thresholds less than $|r| = 0.05$ produced high probability interactions between pairs of voxels with low correlation on the basis of inaccurate estimates of high latency values. Thresholds greater than $|r| = 0.15$ predictably increased the sparsity of non-zero interactions.

For resting state data, probabilistic flow was computed using group-averaged time delay and peak correlation information. For the task data, we first computed activity attributable to the task by using the model-free, event-related approach of averaging 21.6 second epochs (corresponding to 10 time points) linked to task onset both within and across subjects. Time delays and peak correlations were then computed in the task attributable time series, which in turn were used to compute probabilistic flow.

5.5. Sorting voxels into networks

In keeping with our previous analyses of dataset 2 (Mitra et al., 2015, 2014), we sorted voxels into networks according to a 7 network scheme as defined by Hacker and colleagues in the same dataset (Hacker et al., 2013). The topography of these networks is illustrated in Supplemental Figure 1.

5.6. t-Stochastic neighbor embedding

t-Distributed Stochastic Neighbor Embedding (t-SNE) is a dimensionality reduction technique that has been described extensively (Maaten and Hinton, 2008). We applied t-SNE using the “tsne” function in MATLAB and Statistics Toolbox Release 2017a (MathWorks, Inc., Natick, Massachusetts, United States). We applied the default Barnes-Hut algorithm and Euclidian metric distance and did not apply PCA reduction prior to the t-SNE algorithm.

5.7. Statistical analysis of block-wise intra- and inter-network flow averages

To compute whether the inter- and intra-network averages of absolute value probabilistic flow metrics in Fig. 3B are statistically significant as compared to a null model based only on the correlation structure of fMRI data, we produced surrogate time series with the same correlation structure and spectral content of real fMRI data, but with no constraint on the temporal structure, as previously published in (Mitra et al., 2015). We simulated 10,000 datasets where each simulated fMRI dataset was 100 min in length, to match the minimum length of data required to reach stability in the probability flow metric (see Supplemental Figure

1). We then computed the block-wise averages of the absolute value of the probabilistic flow metric as in Fig. 3B in the simulated data, producing a null distribution for each intra- and inter-network block. Given the hypothesis that temporal delays would increase the strength of the probability flow metric, we applied one-tailed test with respect to the null distributions to determine whether the empirical values were greater than 99.8% of the null distribution values in order to Bonferroni correct for 28 comparisons. Blocks marked with asterisks met this criteria for statistical significance.

5.8. Statistical analysis of inter-network flow

We applied a bootstrapping approach to estimate the statistical significance of the network specific inter-network probabilistic flow results in Fig. 3C. We sub-sampled 10,000 inter-network averages computed over approximately 100 min of data (to match the minimum length of data required to reach stability in the probability flow metric, see Supplementary Figure 2) in data set 2 and thereby computed variance around the mean. As we were interested in both significant increases and decreases in the average absolute value of the inter-network flow metric, statistical significance was reached only if the cross network mean (gray line in Fig. 3C) exceeded 99.65% of the bootstrapped distribution in order to Bonferroni correct for 7 comparisons.

5.9. Null model for surrogate networks

To explore whether the clustering result in Fig. 4A was found by chance, we generated surrogate resting state networks (RSNs) topologically matched to real RSNs and approximately matched in spatial frequency distribution (this method was previously published in (Mitra et al., 2014), see Supplementary Figure 1). Surrogate RSNs were generated by treating the left hemisphere of the real RSN brain as an element of a high dimensional symmetric group that respects the topology of the true RSNs. We then applied randomly generated full-rank permutations on the uni-hemispheric RSN brain partition. In greater detail, the 3D MLP RSN partition was converted to a 1D vector and the ends were connected to form a ring. The ring, then, was randomly rotated and the 3D to 1D transform inverted. In principle, other group operations could have been applied, but rotation theoretically preserves spatial scale. The resulting 3D map was reflected across the mid-sagittal line to generate hemispherically symmetric surrogate RSNs. Equivalence of spatial scale was verified by 3D-Fourier analysis in Figure 7 of (Mitra et al., 2014). Using this methodology, we generated 1000 surrogate RSNs and applied the clustering analysis in Fig. 4A with surrogate RSN labels. In each of the 1000 cases, the correlation between the real t-sne result and the surrogate network based result never exceeded a Pearson r of 0.01. On the basis of these null results, the empirically derived clustering in Fig. 4A has a p -value of < 0.001 . More importantly than a p -value estimate, the null simulations demonstrate empirically that the observed organization of Fig. 4A is highly unlikely to be found by chance alone.

Author contribution

A.M. designed and performed the analyses in consultation with M.E.R. and A.Z.S. All authors contributed to writing the paper.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2020.117321.

References

Allen, W.E., Chen, M.Z., Pichamoorthy, N., Tien, R.H., Pachitariu, M., Luo, L., Deisseroth, K., 2019. Thirst regulates motivated behavior through modulation of brain-wide neural population dynamics. *Science* 364, 253.

Allen, W.E., Kauvar, I.V., Chen, M.Z., Richman, E.B., Yang, S.J., Chan, K., Gradinaru, V., Deverman, B.E., Luo, L., Deisseroth, K., 2017. Global representations of goal-directed behavior in distinct cell types of mouse neocortex. *Neuron* 94, 891–907 e896.

Buckner, R.L., Roffman, J.L., Smoller, J.W., 2014. Brain genomics superstructure project (GSP).

Corbetta, M., 1998. Frontoparietal cortical networks for directing attention and the eye to visual locations: identical, independent, or overlapping neural systems? *Proc. Natl. Acad. Sci. U. S. A.* 95, 831–838.

Fox, M.D., Raichle, M., 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. Rev. Neurosci.* 8, 700–711.

Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., Raichle, M.E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc. Natl. Acad. Sci. U. S. A.* 102, 9673–9678.

Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E., 2007. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron* 56, 171–184.

Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen, D.C., 2016. A multi-modal parcellation of human cerebral cortex. *Nature* 536, 171–178.

Gordon, E.M., Laumann, T.O., Marek, S., Raut, R.V., Gratton, C., Newbold, D.J., Greene, D.J., Coalson, R.S., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., Dosenbach, N.U.F., Nelson, S.M., 2020. Default-mode network streams for coupling to language and control systems. *Proc. Natl. Acad. Sci. U. S. A.* 117, 17308–17319.

Hacker, C.D., Laumann, T.O., Szrama, N.P., Baldassarre, A., Snyder, A.Z., Leuthardt, E.C., Corbetta, M., 2013. Resting state network estimation in individual subjects. *Neuroimage* 82C, 616–633.

Hacker, C.D., Snyder, A.Z., Pahwa, M., Corbetta, M., Leuthardt, E.C., 2017. Frequency-specific electrophysiologic correlates of resting state fMRI networks. *Neuroimage* 149, 446–457.

Handwerker, D.A., Ollinger, J.M., D'Esposito, M., 2004. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage* 21, 1639–1651.

He, B.J., Zempel, J.M., Snyder, A.Z., Raichle, M.E., 2010. The temporal structures and functional significance of scale-free brain activity. *Neuron* 66, 353–369.

Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., Arieli, A., 2003. Spontaneously emerging cortical representations of visual attributes. *Nature* 425, 954–956.

Kim, H., 2015. Encoding and retrieval along the long axis of the hippocampus and their relationships with dorsal attention and default mode networks: the HERNET model. *Hippocampus* 25, 500–510.

Konig, P., 1994. A method for the quantification of synchrony and oscillatory properties of neuronal activity. *J. Neurosci. Methods* 54, 31–37.

Lee, J.H., Durand, R., Gradinaru, V., Zhang, F., Goshen, I., Kim, D.S., Fenno, L.E., Ramakrishnan, C., Deisseroth, K., 2010. Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. *Nature* 465, 788–792.

Leong, A.T.L., Gu, Y., Chan, Y.S., Zheng, H., Dong, C.M., Chan, R.W., Wang, X., Liu, Y., Tan, L.H., Wu, E.X., 2019. Optogenetic fMRI interrogation of brain-wide central vestibular pathways. *Proc. Natl. Acad. Sci. U. S. A.* 116, 10122–10129.

Ma, Y., Shaik, M.A., Kozberg, M.G., Kim, S.H., Portes, J.P., Timerman, D., Hillman, E.M., 2016. Resting-state hemodynamics are spatiotemporally coupled to synchronized and symmetric neural activity in excitatory neurons. *Proc. Natl. Acad. Sci. U. S. A.* 113, E8463–E8471.

Maaten, L.v.d., Hinton, G., 2008. Visualizing data using t-SNE. *J. Mach. Learn. Res.* 9, 2579–2605.

Margulies, D.S., Ghosh, S.S., Goulas, A., Falkiewicz, M., Huntenburg, J.M., Langs, G., Bezgin, G., Eickhoff, S.B., Castellanos, F.X., Petrides, M., Jefferies, E., Smallwood, J., 2016. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proc. Natl. Acad. Sci. U. S. A.* 113, 12574–12579.

Maxim, V., Sendur, L., Fadili, J., Suckling, J., Gould, R., Howard, R., Bullmore, E., 2005. Fractional Gaussian noise, functional MRI and Alzheimer's disease. *Neuroimage* 25, 141–158.

McAvoy, M., Mitra, A., Coalson, R.S., d'Avossa, G., Keidel, J.L., Petersen, S.E., Raichle, M.E., 2016. Unmasking language lateralization in human brain intrinsic activity. *Cereb. Cortex* 26, 1733–1746.

Mitra, A., Kraft, A., Wright, P., Acland, B., Snyder, A.Z., Rosenthal, Z., Czerniewski, L., Bauer, A., Snyder, L., Culver, J., Lee, J.M., Raichle, M.E., 2018. Spontaneous infra-slow brain activity has unique spatiotemporal dynamics and laminar structure. *Neuron*.

Mitra, A., Raichle, M.E., 2018. Principles of cross-network communication in human resting state fMRI. *Scand. J. Psychol.* 59, 83–90.

Mitra, A., Snyder, A.Z., Blazey, T., Raichle, M.E., 2015. Lag threads organize the brain's intrinsic activity. *Proc. Natl. Acad. Sci. U. S. A.* 112, E2235–E2244.

Mitra, A., Snyder, A.Z., Hacker, C.D., Pahwa, M., Tagliazucchi, E., Laufs, H., Leuthardt, E.C., Raichle, M.E., 2016. Human cortical-hippocampal dialogue in wake and slow-wave sleep. *Proc. Natl. Acad. Sci. U. S. A.* 113, E6868–E6876.

Mitra, A., Snyder, A.Z., Hacker, C.D., Raichle, M.E., 2014. Lag structure in resting-state fMRI. *J. Neurophysiol.* 111, 2373–2391.

Mohajerani, M.H., McVea, D.A., Fingas, M., Murphy, T.H., 2010. Mirrored bilateral slow-wave cortical activity within local circuits revealed by fast bihemispheric voltage-sensitive dye imaging in anesthetized and awake mice. *J. Neurosci.* 30, 3745–3751.

Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage* 59, 2142–2154.

Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., Shulman, G.L., 2001. A default mode of brain function. *Proc. Natl. Acad. Sci. U. S. A.* 98, 676–682.

Raichle, M.E., Mintun, M.A., 2006. Brain work and brain imaging. *Annu. Rev. Neurosci.* 29, 449–476.

- Raichle, M.E., Raut, R.V., Mitra, A., 2019. How many types are there? *Neocortex* 27, 97.
- Raut, R.V., Mitra, A., Snyder, A.Z., Raichle, M.E., 2019. On time delay estimation and sampling error in resting-state fMRI. *Neuroimage* 194, 211–227.
- Roland, J.L., Snyder, A.Z., Hacker, C.D., Mitra, A., Shimony, J.S., Limbrick, D.D., Raichle, M.E., Smyth, M.D., Leuthardt, E.C., 2017. On the role of the corpus callosum in interhemispheric functional connectivity in humans. *Proc. Natl. Acad. Sci. U. S. A.* 114, 13278–13283.
- Stafford, J.M., Jarrett, B.R., Miranda-Dominguez, O., Mills, B.D., Cain, N., Mihalas, S., Lahvis, G.P., Lattal, K.M., Mitchell, S.H., David, S.V., Fryer, J.D., Nigg, J.T., Fair, D.A., 2014. Large-scale topology and the default mode network in the mouse connectome. *Proc. Natl. Acad. Sci. U. S. A.* 111, 18745–18750.
- Thomas Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zollei, L., Polimeni, J.R., Fischl, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* 106, 1125–1165.
- Vatansver, D., Menon, D.K., Manktelow, A.E., Sahakian, B.J., Stamatakis, E.A., 2015. Default mode dynamics for global functional integration. *J. Neurosci.* 35, 15254–15262.
- Vincent, J.L., Kahn, I., Snyder, A.Z., Raichle, M.E., Buckner, R.L., 2008. Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. *J. Neurophysiol.* 100, 3328–3342.
- Vincent, J.L., Patel, G.H., Fox, M.D., Snyder, A.Z., Baker, J.T., Van Essen, D.C., Zempel, J.M., Snyder, L.H., Corbetta, M., Raichle, M.E., 2007. Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447, 83–86.
- Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zollei, L., Polimeni, J.R., Fischl, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* 106, 1125–1165.