Published in final edited form as:

Scand J Psychol. 2018 February ; 59(1): 83–90. doi:10.1111/sjop.12422.

Principles of cross-network communication in human resting state fMRI

Anish Mitra^{a,1} and Marcus E. Raichle^{a,b,1}

^aDepartment of Radiology, Washington University, St Louis, MO 63110

^bDepartment of Neurology, Washington University, St Louis, MO 63110

Abstract

Directed signaling among and within the large-scale networks of the human brain is functionally critical. Recent advances in our understanding of spontaneous fluctuations of the fMRI BOLD signal have provided strategies to study the spatial-temporal properties of directed signaling at infra-slow frequencies. Herein we explore the relationship between two canonical systems of the human brain, the default mode network (DMN) and the dorsal attention network (DAN) whose anti-correlated relationship is well known but poorly understood. We find that within the DMN, activity moves from retrosplenial to prefrontal cortex whereas in the DAN activity moves from the frontal eye fields to the parietal cortex. Bi-directional communication between the two networks occurs via their earliest elements (i.e., from the retrosplenial cortex of the DMN to the frontal eye fields of the DAN). This framework for network communication appears to generalize across all networks providing an expanded basis for understanding human brain function.

Introduction

Imaging, first with positron emission tomography (PET) and then with functional magnetic resonance imaging (fMRI) has greatly expanded our understanding of human brain function in health and disease. Initially the focus of this work was almost exclusively on defining components of the brain that were associated with the performance of specific tasks. This exercise was made more effective by the experimental strategies of cognitive psychology which sought to isolate very specific task elements by devising control states containing all but those task elements of interest. This subtractive methodology proved to be extremely effective in identifying brain components exhibiting task-specific activity increases. However, this strategy neatly obscured the fact that most of what the brain was actually doing (i.e., that associated with its ongoing intrinsic or spontaneous activity) was eliminated from view by the subtraction process.

Evidence that spontaneous brain activity might well be a major component of overall brain activity and critical to our understanding of brain function began to emerge from the realization that spontaneous activity accounts for most of the enormous cost of brain function (for a review of this literature see (Raichle and Mintun, 2006)). Yet while the cost

¹To whom correspondence may be addressed: anishmitra@gmail.com or marc@npg.wustl.edu.

of this ongoing activity could be accurately documented with PET, it was not initially realized that regional differences in either blood flow or metabolism could be related to functions associated with spontaneous activity. From the perspective of fMRI this ongoing activity was initially regarded as noise and discarded. Two events changed that perspective dramatically.

PET provided the first human imaging indication that spontaneous activity contributed to the large scale functional organization of the human brain. Not only did task-relevant areas increase their activity during task performance but a remarkably consistent group of areas decreased their activity as well (Shulman et al., 1997). This was observed when the task was compared with a control state of resting quietly with eyes open or closed. This led to the concept of a default mode of brain function (Raichle et al., 2001) and the discovery of a unique constellation of areas dubbed the brain's *default mode network* (DMN).

The second event was the discovery (Biswal et al., 1995) that infraslow activity (ISA; <0.01 Hz) manifest as spontaneous fluctuations in the fMRI blood oxygen level dependent (BOLD) signal at rest, exhibited patterns of highly correlated spatial patterns of activity within the motor network of the brain. This observation was subsequently expanded to include all cortical network in the human brain (for review see (Raichle, 2015)). This has become known as resting-state functional connectivity and is now a dominant theme in human brain imaging with fMRI.

Importantly for the purposes of this paper these correlations within networks were typically achieved by placing a seed region of interest within an area of a known brain network (e.g., the DMN) and asking what other brain regions were correlated with it. The results of this exercise were dramatic. Quite literally the entire cortex of the human brain has been parcellated into a map of its constituent networks. But as dramatic as the ability to parcellate was, it left unanswered how these networks communicated with one another as surely, they must in a well-functioning brain. This oversight occurred because of the implicit assumption that these spatial correlations occurred with zero latency (i.e., researchers simply asked if there was a correlation without asking whether that correlation was associated with a time delay or lag). This assumption effectively eliminated time from the processed functional connectivity maps and any indication of how networks communicated within themselves or with other networks. But the story did not end there.

A subsequent more careful examination with fMRI of the temporal structure of ISA correlations within networks (Mitra and Raichle, 2016; Mitra et al., 2015a; Mitra et al., 2014) has revealed the presence of lags of plus-or-minus a second within and among networks suggesting the presence of directed signaling within ISA that can be state dependent (Mitra et al., 2016; Mitra et al., 2015b).

One of the clearest examples of cross-network interactions occurs between the default mode network (DMN), and the dorsal attention network (DAN). In particular, performance of a goal-directed, non-self-referential task is accompanied by a decrease in activity in the DMN and a corresponding increase in activity in the DAN (reviewed in (Fox et al., 2005)). This “push-pull” relationship is also found in spontaneous infra-slow activity where spontaneous

fluctuations in ISA between the DMN and DAN have been shown to be anti-correlated (Fox et al., 2005; Fox et al., 2009). However, the details of how the DMN and the DAN interact remain unknown.

Here, we extend previous work describing directed signaling of spontaneous infra-slow activity in the human brain (Mitra and Raichle, 2016) to detail the principles governing ISA communication between the DMN and DAN. In the past, we have shown that spontaneous infra-slow BOLD signals tend to flow unidirectional within a given resting state network (Mitra et al., 2015a). Here, we demonstrate further that the cross-talk between the DMN and the DAN is exquisitely structured both spatially and temporally. Moreover, we show that the framework that governs DMN-DAN communication may apply to cross-network signal propagation more generally.

Methods

Participants

A large data set ($n = 1376$) was obtained from the Harvard-MGH Brain Genomics Superstruct Project (Buckner et al., 2014). The mean age of the participants was 21.4 years; 785 were females. They were scanned on a 3 tesla Siemens Tim Trio with a repetition time (TR) of 3.0 seconds and a spatial resolution of 3mm^3 isotropic. 12 minutes of data were acquired in each subject. Further details on our use of these data can be found in (Mitra et al., 2015a).

Analytical Strategy

We determined the average directionality of spontaneous, infra-slow signaling between pairs of voxels in the brain by computing lagged correlation functions, as has been previously described in detail (Mitra et al., 2015a) and illustrated in Figure 1. The key idea is that if activity in one area tends to lead activity in a second region, then the lagged-correlation function will be asymmetric, and the time (in seconds relative to zero) corresponding to the peak of the lagged-correlation function will determine the direction and average temporal delay between the two areas. Thus, if activity in the retrosplenial cortex (RSC) leads activity in the medial prefrontal cortex (mPFC), the peak of the cross-correlation lag is offset from zero (~ 1 second; Figure 1B), indicating that on average activity begins in the RSC one second before it arrives in the mPFC. As both the RSC and the mPFC are part of the default mode network (DMN), this is an example of intra-network signaling.

Now, let us suppose that we want to examine inter-network signaling between the DMN and the DAN. Choosing the RSC from the DMN and the frontal eye field (FEF) from the DAN (the rationale for this choice will be explained later), we find (Figure 1C) that RSC leads the FEF but, in this case, the correlation is negative. Why? Simply because the DMN and DAN are anti-correlated networks at rest (Fox et al., 2005). The computation of a lagged correlation is agnostic to the sign of the underlying relationship, and works equally well in positive magnitude correlations (Fig. 1B) and negative magnitude correlations (Fig. 1C).

Drawing on extant knowledge of the functional network organization of resting state brain activity, we know that networks are characterized by the exclusive presence of positive

correlations within their boundaries (Mitra et al., 2015a). That, of course, does not preclude the presence of positive correlations between networks as our data will reveal.

If networks are characterized by the exclusive presence of positive correlations (Mitra et al., 2015a) it must logically follow that areas that are negatively correlated must reside in different networks. This is strikingly demonstrated by the fact that the RSC, which is in the DMN, and the FEF, which is in the DAN, exhibit a negatively correlated relationship (Figure 1C). Thus, while temporal delays in negative correlations correspond to inter-network signaling, temporal delays in positive correlations can occur both within and among networks as our data (below) will demonstrate.

Results

In global signal regressed data, spontaneous activity fluctuations in the DAN are anti-correlated with those in the DMN regions (i.e., either task-evoked or spontaneously occurring activity in those two areas varies reciprocally (Fox et al., 2005)). However, that observation alone does not tell us how signaling is directed between the DMN and the DAN. For that information, we look for a temporal delay by computing the peak anti-correlation as a function of time as we did between the RSC and FEF in Figure 1C. In that example, RSC leads the FEF. Why we chose the RSC and the FEF for this illustration is developed now in more detail below.

Dorsal attention network signaling

We begin with an analysis of spontaneous infra-slow signaling anchored in the dorsal attention network (DAN). We know from many prior studies that the FEF are a prominent component of the DAN (e.g., see (Hacker et al., 2013)), and have therefore used a bilateral FEF seed to create a seed-based correlation map of the DAN, shown in Fig. 2A. As is typical of the DAN topography, these seed maps show strong positive correlations (red) in the intra-parietal sulcus (IPS), as well as strong anti-correlations (blue) in the default mode network regions, especially the PCC, mPFC, and posterior parietal cortex (PPC).

Having derived the DAN topography in resting state data, we next sought to characterize intra-DAN signaling, by computing temporal lags within the most highly correlated regions of the DAN, namely the FEF and the IPS. The results, shown as a temporal delay map in Figure 2B, illustrate that activity in the FEF (blue/green) tends to lead activity in the IPS by about a second. The existence of a clear, directed temporal sequence with the DAN is consistent with our prior work demonstrating that infra-slow activity travels in a unidirectional manner within RSNs, in patterns we have called “motifs” (Mitra and Raichle, 2016; Mitra et al., 2015a).

We next explored how the DAN signals to the DMN. For this analysis, we considered temporal delays between the FEF seed (the earliest part of the DAN) and the DMN topography defined by its anti-correlations (Fig. 2C). Again, a clear temporal sequence emerges: the RSC has the shortest delay with respect to the FEF, with the green color in the PCC indicating that these RSC voxels have a roughly simultaneous relationship with FEF activity. The parietal cortex is roughly a half second delayed with respect to the FEF, and the

mPFC follows the FEF by around a full second. Critically, the delays in Fig. 2C are referenced to the FEF. Thus, the map demonstrates that a change in activity in the FEF causes a rippling change in activity through the DMN, starting in the back (RSC) and moving to the front (mPFC). Notice that this did not have to be the case. It might have been that the FEF activity affected the DMN in a front-to-back sequence, or perhaps FEF activity influenced the entire DMN at once. Instead, a clear principle emerges by which the earliest node of the DAN, the FEF, initiates a back-to-front change in the DMN via the RSC. The significance of back-to-front temporal sequences in the DMN will be addressed in the next section.

Default mode network signaling

We next approach the question of intra- versus inter-network infra-slow signaling from the perspective of the DMN, using the same strategy as in the prior section's analysis of the DAN.

To derive the positive and negative correlation topography of the DMN, we used a seed in the RSC to compute a correlation map (Fransson and Marrelec, 2008; Greicius et al., 2009). The result, shown in Fig. 2D, highlights the known components of the DMN, the posterior cingulate/ precuneus cortex, the lateral parietal cortex, and medial prefrontal cortex in warm colors, indicating positive correlations. In contrast, the DAN (FEF and the intraparietal sulcus, for example) as well as several other networks are seen in blue, indicating anti-correlation with the DMN. As before, we first start by analyzing temporal delays within the DMN; the results of the RSC-seeded lag map are shown in Figure 2E. Notice that there is a clear temporal sequence of activity within the DMN: starting in the RSC, moving to the posterior cingulate/precuneus and lateral parietal cortices, and finally terminating in the medial prefrontal cortex. Remarkably, this back-to-front sequence of activity moving within the DMN is precisely the sequence in which the FEF, the component of the DAN, interacts with the DMN.

This correspondence suggests a principle in bi-directional cross-network interactions. Given network A and network B, the earliest part of network A is responsible for signaling to network B, and this cross-network signal propagates through network B in the same sequence that activity within network B must follow. If we substitute the DAN for network A, and the DMN for network B, we can see in Figure 2F that the earliest part of the DAN, the FEF, signals to the DMN, and that the order in which FEF signals move through the DMN (back-to-front) is the same order through which signals arising in the DMN itself move through the network (also back-to-front, Fig. 2C, E).

We can clarify this idea further by making a concrete prediction regarding signaling from the DMN to the DAN (in place of networks "A" and "B", respectively). The earliest part of the DMN is the RSC: thus, we expect the RSC to signal to the DAN, and for the earliest node in the RSC-DAN interaction to be in the FEF, which itself is the earliest node in the DAN. We test this hypothesis by directly computing temporal delays between the RSC and all regions that are significantly anti-correlated with the PCC. In the resulting map in Figure 2F, sure enough, the FEF is earlier than the IPS, indicating that cross-network signals sent from the RSC to the DAN arrive at the FEF before the IPS. Thus, these findings show that there is bi-

directional signaling between pairs of networks (at least the DMN and DAN), such that signals move reciprocally between early nodes of networks. These relationships are summarized schematically in Figure 2G.

Of course, the DMN is anti-correlated with much of the brain, even beyond the DAN. These relationships are captured in the areas outside the DAN in Figure 2D, illustrating that ISA signals flow from the RSC to other networks in a highly spatially structured way. If the principle articulated above is correct, each of these regions which receive signals from the PCC should be early in their respective networks.

General analysis of inter- vs. intra-network signaling

We could verify the hypothesis we have articulated for signals emerging from the DMN by making seed-based lag maps of each of the regions in Figure 2. However, there is a more direct way to test the idea for all cross-network signaling. The principle we articulate implies that intra-network signals travel in the same direction as inter-network signals. For instance, an intra-network signal arising in the DMN tends to start in the RSC and move toward the mPFC. An inter-network signal from the DAN still travels in the same direction from the DMN: back-to-front, hence leading to the correspondence between Figures 2C and 2E. Moreover, the inter-network signal in the DAN tends to start in the early part of the DAN—the FEF.

Therefore, whether we study intra- or inter-network signaling, the areas that tend to be “early” or “late” are largely the same, at least in our analysis of the DMN and DAN thus far. We can examine whether this idea is true in general by computing *lag projections*, which are simply images of the average temporal delay between a particular voxel and the rest of the brain. We have previously computed lag projections by averaging over all delays between all pairs of voxels (e.g., the columns of the temporal lags matrix in Fig. 3A, right), whether those delays corresponded to inter- or intra-network relations, to produce an average picture of where infra-slow activity tends to start and terminate in the brain.

However, we can modify this technique to produce two lag projections (Figure 3B): First, we produce projection over only positive correlations in the brain (Fig. 3B, left). Since positive correlations define RSNs, this lag projection captures the temporal structure of within-network signaling. The reader will, of course, notice that there are some positive correlations outside of the defined RSNs. As pointed out earlier (see *Analytical Strategy*) while RSNs are noteworthy in that correlations within them are exclusively positive, this does not preclude the presence of positive correlations between networks (Figure 3B, left).

Next, we computed a lag projection over only negative correlations in the brain (Figure 3B, right). As RSNs are defined by positive correlations alone it is noteworthy that there is a complete absence of any negative correlations within RSNs. It follows that all negative correlations are between networks, and hence this resulting lag projection captures an exclusive component of the temporal structure of inter-network signaling.

If it is true that cross-network signals originate in early parts of the sending network, and then travel through the same route as intra-network signals in the receiving network, then the

overall temporal structure of lag projections computed over positive vs. negative correlations should be highly similar. And, in fact, as illustrated in Figure 3C, the two lag projection topographies are nearly identical (spatial correlation $r = 0.92$).

Discussion

The dynamics of resting state brain activity have become of great interest to cognitive neuroscientists, providing an unprecedented view of the functional organization of networks within the human brain. One remarkable feature of these dynamics is the spatial-temporal structure of ISA (infra-slow activity $< 0.1\text{Hz}$) found within these resting-state networks. For example, within the DMN activity begins in the RSC and moves anteriorly to the medial prefrontal cortex. In the DAN, activity begins in the FEF and moves posteriorly to the intraparietal sulcus. Until now how functional integration occurs among networks has remained unclear.

Here, using regional measurements of the fMRI BOLD signal we extend prior analyses by showing that directed signaling of infra-slow activity (ISA; $<0.1\text{ Hz}$) within networks is also utilized in cross-network signaling. Indeed, the cross-network travel of resting state BOLD signals follows two simple principles. First, the “sending network” initiates cross-network communication from its earliest node (the RSC in the DMN, and the FEF in the DAN). Second, activity from the “sending network” arrives at the “receiving network” at the latter’s earliest node then travels through the “receiving network” using its established intra-network spatial-temporal pattern of directed signaling. These principles for the first time articulate the rules by which a functional level of network integration occurs within the human brain. The computational and behavioral significance of these rules are an exciting frontier which remains to be understood. Further, understanding the relationship of inter-network, positive correlations to their intra-network counterpart; and, the role of positive and negative inter-network correlations together in signaling between specific networks are clearly fruitful areas for further research.

The example we have used to illustrate the elements of internetwork signaling (i.e., the relationship between the DMN and the DAN) offers the opportunity to ask two important questions. First, what might be the likely behavioral implications? And, second, how should we be thinking about the underlying neurobiology. We offer two lines of research that speak to these questions beginning with the first question related to behavior.

The relationship between the DMN and the DAN has been characterized as one of anti-correlations. When subjects engage in a non-self-referential task, activity in the DAN increases while activity in the DMN decreases (Raichle et al., 2001; Shulman et al., 1997). This is recapitulated in the resting state where spontaneous activity in these two networks is again observed to be anti-correlated ((Fox et al., 2005). This has led to an egocentric view that the DMN is most concerned with our internal mental state. Following this logic, it is easy to accept that activity in the DMN would naturally decrease during a demanding, non-self-referential task. But, a study involving self-control in juveniles suggests a more nuanced relationship.

As part of a larger study of juveniles incarcerated for criminal activity, resting state fMRI data were obtained on 107 individuals (Shannon et al., 2011). Employing a data analysis strategy designed to look at network relationships in relation to behavior the analysis focused on the degree of impulsivity measured in each individual. The surprising finding was that functional connectivity measured without consideration of lags indicated the functional connectivity of a single area in premotor cortex, essentially identical to the FEF, was most predictive of the measured level of impulsivity, so much so that the level of impulsivity in individual subjects could be predicted from the functional connectivity of their FEF. What is remarkable about this study in relation to the present work is that the functional connectivity involved, specifically, the relationship of the DMN and the DAN to the FEF. It essentially involved a balance between the two networks: a stronger connectivity of the DMN to the FEF predicted greater impulsivity whereas the reverse (i.e., a stronger connectivity of the DAN to the FEF) predicted lesser impulsivity. In the same paper (Shannon et al., 2011) a companion group of developing youngsters exhibited the same relationship as a function of age.

The lesson learned from the above study was that inter-network communication, as revealed to us through our studies of ISA using resting-state fMRI, involves a delicate balance between and, likely, among networks with potentially significant implications for our understanding of human behaviors in health and disease. That the DMN might figure prominently into such relationships involving networks other than the DAN seems reasonable to posit given the emerging evidence for its dominant position in the hierarchy of all brain network (e.g., see (Margulies et al., 2016)). Our understanding of these relationships is only beginning and will, in part, hinge on our understanding of the underlying neurobiology. We turn briefly to a few issues related to that subject.

The neurobiology underlying the relationship of the DMN and the DAN has been the subject of remarkably little basic neuroscience research. The noteworthy exception is the work of Denis Pare and his colleagues (Popa et al., 2009). Intrigued by the observation of anti-correlations between the DMN and the DAN this group instrumented cats with depth electrodes along the midline of the cortex to approximate the location of the DMN and somewhat more laterally to approximate the location of the DAN. Doing so in the cat is not unreasonable given the fact that the DMN has now been identified in non-human primates (Vincent et al., 2007) as well as in rodents (Lu et al., 2012; Stafford et al., 2014). They were able to identify clear anti-correlations between their two groups of electrodes which makes their observations directly relevant to our work. Further, during task performance local field potential power (LFP) decreased in the DMN (labeled the “task-negative network” by Pare and colleagues (Popa et al., 2009)) and increased in the DAN (labeled the “task-positive network” by Pare and colleagues (Popa et al., 2009)) which fit nicely with the observations in humans (Fox et al., 2005). However, the presence of anti-correlations was variable. They were more prominent during wake and paradoxical sleep and absent during slow wave sleep. These findings suggest a more variable, state-dependent relationship between the DMN and the DAN than so far appreciated in humans. Whether this is related in part to image processing strategies such as global signal regression used in humans (Fox et al., 2009) remains to be determined. It is clearly an important subject for future research.

Of particular interest in the work of Pare and colleagues was the fact that during task-induced decreases in activity in the DMN, as measured with LFP power, there was an actual increase in the spiking of neurons in the same area. This might at first seem quite paradoxical given the early PET work on activity decreases in this area during task performance (Raichle et al., 2001). However, neuron cell bodies from which spikes are detected do not increase their metabolism during spiking. Rather, the increase in metabolism occurs in the areas to which these neurons project (i.e., their axon terminal (Kadekaro et al., 1985; Schwartz et al., 1979)). It seems reasonable to posit that a decrease in metabolism and LFP power in the DMN might well reflect a change in excitatory/inhibitory balance within the DMN mediated by reciprocal signaling between the DMN and the DAN. In theory, this is a testable hypothesis and, given our results, the focus of that research should be on the relationship between the RSC and the FEF.

The above examples, behavioral and neurobiological, of research relevant to an understanding of directed signaling between and within brain networks is just a beginning of a much broader line of research into the integrated functioning of the human brain. While the work presented herein has focused on ISA as seen by fMRI BOLD imaging, a full understanding of functional integration will only emerge when we understand the relationship of ISA to other elements of spontaneous activity occurring at higher frequencies. A hint of what is to come emerges from studies of the sleep-wake cycle where an interplay between ISA and cortical delta activity governs the directed signaling between the cerebral cortex and the hippocampus in humans (Mitra et al., 2016)

Acknowledgments

We are pleased to submit this manuscript as a tribute to an outstanding cognitive neuroscientist and friend, Professor Kenneth Hugdahl. We also wish to acknowledge generous support from the National Institutes of Health, USA and the intellectual stimulation of many colleagues and friends.

References

- Biswal B, Yetkin FZ, Haughton VM, Hyde JS. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med*. 1995; 34:537–541. [PubMed: 8524021]
- Buckner RL, Roffman JL, Smoller JW. Brain Genomics Superstruct Project (GSP). Harvard Dataverse. 2014
- Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci U S A*. 2005; 102:9673–9678. [PubMed: 15976020]
- Fox MD, Zhang D, Snyder AZ, Raichle ME. The global signal and observed anticorrelated resting state brain networks. *J Neurophysiol*. 2009; 101:3270–3283. [PubMed: 19339462]
- Fransson P, Marrelec G. The precuneus/posterior cingulate cortex plays a pivotal role in the default mode network: Evidence from a partial correlation network analysis. *Neuroimage*. 2008; 42:1178–1184. [PubMed: 18598773]
- Greicius MD, Supekar K, Menon V, Dougherty RF. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cereb Cortex*. 2009; 19:72–78. [PubMed: 18403396]
- Hacker CD, Laumann TO, Szrama NP, Baldassarre A, Snyder AZ, Leuthardt EC, Corbetta M. Resting state network estimation in individual subjects. *Neuroimage*. 2013; 82:616–633. [PubMed: 23735260]

- Kadekaro M, Crane AM, Sokoloff L. Differential effects of electrical stimulation of sciatic nerve on metabolic activity in spinal cord and dorsal root ganglion in the rat. *Proc Natl Acad Sci U S A*. 1985; 82:6010–6013. [PubMed: 3862113]
- Lu H, Zou Q, Gu H, Raichle ME, Stein EA, Yang Y. Rat brains also have a default mode network. *Proc Natl Acad Sci U S A*. 2012; 109:3979–3984. [PubMed: 22355129]
- Margulies DS, Ghosh SS, Goulas A, Falkiewicz M, Huntenburg JM, Langs G, Bezgin G, Eickhoff SB, Castellanos FX, Petrides M, et al. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proc Natl Acad Sci U S A*. 2016; 113:12574–12579. [PubMed: 27791099]
- Mitra A, Raichle ME. How networks communicate: propagation patterns in spontaneous brain activity. *Philos Trans R Soc Lond B Biol Sci*. 2016; 371
- Mitra A, Snyder AZ, Blazey T, Raichle ME. Lag threads organize the brain's intrinsic activity. *Proc Natl Acad Sci U S A*. 2015a; 112:E2235–2244. [PubMed: 25825720]
- Mitra A, Snyder AZ, Hacker CD, Pahwa M, Tagliazucchi E, Laufs H, Leuthardt EC, Raichle ME. Human cortical-hippocampal dialogue in wake and slow-wave sleep. *Proc Natl Acad Sci U S A*. 2016; 113:E6868–E6876. [PubMed: 27791089]
- Mitra A, Snyder AZ, Hacker CD, Raichle ME. Lag structure in resting-state fMRI. *J Neurophysiol*. 2014; 111:2374–2391. [PubMed: 24598530]
- Mitra A, Snyder AZ, Tagliazucchi E, Laufs H, Raichle ME. Propagated infra-slow intrinsic brain activity reorganizes across wake and slow wave sleep. *Elife*. 2015b; 4
- Popa D, Popescu AT, Pare D. Contrasting activity profile of two distributed cortical networks as a function of attentional demands. *Journal of Neuroscience*. 2009; 29:1191–1201. [PubMed: 19176827]
- Raichle ME. The restless brain: how intrinsic activity organizes brain function. *Philos Trans R Soc Lond B Biol Sci*. 2015; 370:1–11.
- Raichle ME, MacLeod AM, Snyder AZ, Powers WJ, Gusnard DA, Shulman GL. A default mode of brain function. *Proc Natl Acad Sci U S A*. 2001; 98:676–682. [PubMed: 11209064]
- Raichle ME, Mintun MA. Brain work and brain imaging. *Annu Rev Neurosci*. 2006; 29:449–476. [PubMed: 16776593]
- Schwartz WJ, Smith CB, Davidsen L, Savaki H, Sokoloff L, Mata M, Fink DJ, Gainer H. Metabolic mapping of functional activity in the hypothalamo-neurohypophysial system of the rat. *Science*. 1979; 205:723–725. [PubMed: 462184]
- Shannon BJ, Raichle ME, Snyder AZ, Fair DA, Mills KL, Zhang D, Bache K, Calhoun VD, Nigg JT, Nagel BJ, et al. Premotor functional connectivity predicts impulsivity in juvenile offenders. *Proc Natl Acad Sci U S A*. 2011; 108:11241–11245. [PubMed: 21709236]
- Shulman GL, Fiez JA, Corbetta M, Buckner RL, Miezin FM, Raichle ME, Petersen SE. Common blood flow changes across visual tasks: II. Decreases in cerebral cortex. *J Cogn Neurosci*. 1997; 9:648–663. [PubMed: 23965122]
- Stafford JM, Jarrett BR, Miranda-Dominguez O, Mills BD, Cain N, Mihalas S, Lahvis GP, Lattal KM, Mitchell SH, David SV, et al. Large-scale topology and the default mode network in the mouse connectome. *Proc Natl Acad Sci U S A*. 2014; 111:18745–18750. [PubMed: 25512496]
- Vincent JL, Patel GH, Fox MD, Snyder AZ, Baker JT, Van Essen DC, Zempel JM, Snyder LH, Corbetta M, Raichle ME. Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*. 2007; 447:83–86. [PubMed: 17476267]

Figure 1. Analytical details of temporal delays in resting state fMRI

(A) Sample blood-oxygen level dependent (BOLD) time series from three parts of the brain: the retrosplenial cortex (RSC; blue), the medial prefrontal cortex (mPFC; red), and the right frontal eye field (FEF; orange). Note that the RSC and mPFC time series are highly correlated, as they belong to the same resting state network, the default mode network (DMN). In contrast, the orange time series is anti-correlated with the DMN areas, as the FEF belongs to a separate network, the dorsal attention network (DAN). (B) The lagged correlation curve computed between the RSC and mPFC reveals the existence of a non-zero peak, indicating a directionality and temporal delay so signaling between these positively correlated regions. We compute this delay using parabolic interpolation about the peak of the curve, as illustrated in red. The interpolation is necessary as empirical shifts in the time series can only occur in units of the TR (temporal resolution) of fMRI; TR = 3 seconds in this study. Thus, (B) illustrates how we find a temporal shift within a network. (C) Lagged correlation computed between RSC and FEF: the details are the same as in panel (B), except in this case we compute the peak anti-correlation, as these areas belong to distinct networks. However, we again find a non-zero peak in this example, indicating directed signaling between these regions across two networks.

Figure 2. Intra- and inter-network signaling in the dorsal attention and default mode networks (A) A frontal eye field (FEF)-seeded correlation map, as in conventional functional connectivity studies. The red areas reveal the dorsal attention network (DAN) topography; blue areas highlight regions outside the DAN, especially the default mode network (DMN). (B) A FEF-seeded lag map within areas of correlation with the FEF, defined using (A). The idea is to examine temporal lags within the DAN. (B) Shows that FEF regions are blue whereas areas in the intra-parietal sulcus (IPS) are red, indicating that BOLD signals tend to move from FEF to IPS within the DAN. (C) A FEF-seeded lag map with areas anti-correlated with the FEF, defined using (A). The idea is to examine temporal lags between activity in the FEF and activity in nodes of another network, the DMN. Note that the RSC is blue/green, whereas the posterior parietal cortex (PPC) is yellow, and the medial prefrontal cortex (mPFC) is orange/red. Hence, the RSC is the least delayed with respect to the FEF, followed by the PPC, and finally the mPFC. (D) A RSC-seeded correlation map, as in conventional functional connectivity studies. The red areas reveal the DMN topography; blue areas highlight regions outside the DMN, especially the DAN. (E) A RSC-seeded lag map within areas of correlation with the FEF, defined using (D). The idea is to examine temporal lags within the DMN. (E) Shows that RSC regions are blue whereas areas in the

PPC are yellow, and the mPFC region is orange. This result demonstrates the back-to-front temporal sequence for activity within the DMN. **(F)** A PCC-seeded lag map with areas anti-correlated with the PCC, defined using (D). The idea is to examine temporal lags between activity in the RSC and activity in nodes of another network, the DAN. Note that the FEF is green/yellow, whereas the IPS is orange/red. Hence, in the DAN, the FEF is the least delayed with respect to the IPS. **(G)** A schematic summary of the spatial and temporal relationships between the DMN and the DAN depicted in the images (A) through (F) as well as the lagged-correlation data from Figure 1B and C.

Figure 3. A general analysis of intra- vs. inter-network signaling in resting state fMRI
(A) For every pair of voxels in the brain, we can define a zero-lag correlation matrix, shown on the **left**, and a temporal lags matrix, shown on the **right**. Together these matrices contain all functional connectivity and temporal delay information in the system. **(B, left)** Since resting state networks are defined as areas with high positive correlations (Mitra and Raichle, 2016; Mitra et al., 2015a), we can approximate intra-network signaling by only considering temporal delays which correspond to positive correlations, and ignore the rest. This as shown on the left where the masked temporal lags matrix is computed only over

positive correlations. Readers will note that some positive correlation occur outside of designated networks indicating that some inter-network communication is characterized by positive correlations. (**B, right**) inter-network signaling from the perspective of temporal delays which correspond to negative correlations as these are by definition inter-network relationships in resting state data. This reveals areas that tend to be early or late in inter-network communication. As RSNs are defined by positive correlations alone it is noteworthy that there is a complete absence of any negative correlations within RSNs. Critically, positive (**C, left**) and negative (**C, right**) lag projection maps are highly similar (spatial correlation $r = 0.92$), indicating that the temporal structure of within-network signaling mimics the temporal structure of cross-network signaling. This general observation is in line with the findings obtained in the examination of the DMN and DAN summarized in Figure 2G, where the early nodes in each network signal to the other following the same temporal pattern as activity arising in the receiving network. Abbreviations: dorsal attention network (DAN); ventral attention network (VAN); auditory cortex (AUD); somatomotor network (SMN); visual cortex (VIS); frontoparietal control network (FPC); language network (LAN); and, default mode network (DMN). Parcellation was based on the work of (Hacker et al., 2013)