6

Brain Networks

How Many Types Are There?

Marcus E. Raichle, Ryan V. Raut, and Anish Mitra

Abstract

Unraveling the organizational structure of the brain has, in large measure, been reductionist in nature. While this has revealed, in ever-increasing detail, the fine structure of the brain, it does leave less directly addressed the beautifully integrated nature of brain function. Views of the functional organization of the brain should include a unitary perspective, despite the diversity of its constituent parts. This chapter focuses on recent observations from the authors' laboratory, which point to the value of an integrated approach as well as to answer the assigned title question: arguably, the brain consists of a single network with functional diversity.

Introduction

Categorizing network types in the brain requires knowledge of the context in which the term *networks* is applied. Generally, the term emerges from research that associates component operations and behavior with identifiable brain parts and their interactions, ranging from the cellular and molecular to a whole brain level of analysis (Bassett and Sporns 2017). The breadth of extant work on this subject is exemplified by the observation that a current PubMed search of the terms "brain" and "networks" presently yields over 33,000 citations. Taking a whole-brain perspective, restricting the search to "brain networks" and "fMRI" yields approximately 1,500 citations per year, over the past three years. Remarkably, this represents ~40% of the total citations for "brain" and "networks" over the same period of time. Here, we focus on this whole-brain or large-scale perspective that now characterizes a significant fraction of human brain imaging research and, more generally, research on functional brain organization.

From this large-scale perspective, a generally agreed upon list of cortical networks has emerged with names such as somatomotor, visual, dorsal, and

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

ventral attention, cingulo-opercular and frontoparietal control, salience, and default mode (e.g., Power et al. 2011; Yeo et al. 2011; Hacker et al. 2013). Research seeking an understanding of the large-scale organization of these networks in the mammalian brain has included anatomically and theoretically based network analyses (e.g., van den Heuvel et al. 2016) and genetics (Richiardi et al. 2015; Ge et al. 2017) as well as resting-state functional magnetic resonance imaging (fMRI) in humans (Raichle 2011), nonhuman primates (Vincent et al. 2007), and rodents (Lu et al. 2012; Stafford et al. 2014). While the nature of the relationships within and among these networks has not been ignored, the primary emphasis has been on their parcellation; that is, how crisp the boundaries delineating the networks are and how many cortical parcels exist (e.g., Power et al. 2011; Glasser et al. 2016).

Our objective is to outline a unifying, functional framework within which individually identifiable components (*networks/systems*) arise and communicate with one another. There are two major elements to our approach.

The first element is to specify the biological underpinnings of the fMRI blood oxygen level dependent (BOLD) signal. The logic behind doing so is that this signal has become an extremely attractive window on the brain's large-scale, functional organization. Nonetheless, an agreed upon understanding of its underlying neurophysiology has been lacking. Many have asserted that infra-slow (< 0.1 Hz) fMRI signals are simply a low-pass filter of the brain's overall neurophysiology (e.g., de Zwart et al. 2005; Logothetis 2008). In contrast, we have recently shown (Mitra et al. 2018) that the signal is a very specific representation of the brain's infra-slow activity, and thus it offers a unique window on an element of brain neurophysiology that is critical for maintaining and orchestrating the brain's large-scale, functional organization.

The second element of our approach is to utilize functional information available from the spontaneous, ongoing activity of the brain. The logic behind this choice is that most of brain energy resources are devoted to this activity, well over 90% (Raichle and Mintun 2006). Furthermore, it has become a major source of information related to the functional organization of the mammalian brain in health and disease.

The fMRI BOLD Signal

The fMRI BOLD signal has a long and storied history. It is based on the properties of oxygenated hemoglobin in a magnetic field, a property first hinted at by Michael Faraday in 1846 (Faraday 1933), formally discovered by Linus Pauling and Charles Coryell in 1936 (Pauling and Coryell 1936), and reintroduced by Keith Thulborn and colleagues in 1982 (Thulborn et al. 1982). Deoxyhemoglobin, being paramagnetic, disrupts a magnetic field and causes a loss of signal in an MRI scanner. Therefore, veins will prominently appear in MRI images as areas of signal loss. Task-induced increases (decreases) in

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

regional brain blood flow are not accompanied by proportional changes in oxygen consumption (Fox and Raichle 1986), thus producing localized decreases (increases) in deoxyhemoglobin. Combining this knowledge with the effect of deoxyhemoglobin on the MRI signal, Seiji Ogawa and colleagues proposed at the Bell Laboratories an *in vivo* MRI strategy for brain mapping based on what they dubbed the BOLD signal (Ogawa et al. 1990). Their proposal launched fMRI as the primary tool in cognitive neuroscience.

The physics of the fMRI BOLD signal is well understood. The fMRI signal is, quite simply, based on the ratio of oxy- to deoxyhemoglobin in the brain vasculature, which varies both spontaneously, reflecting the brain's ongoing intrinsic activity, and predictably in response to task-induced changes in brain activity (see figure 6 in Raichle and Mintun 2006). The relationship of the BOLD signal to the underlying neurophysiology of the brain, however, has been a matter of considerable debate. At the heart of this discourse has been the variably articulated idea that the fMRI BOLD signal is a vascular, low-pass filter of brain neurophysiology writ large (e.g., de Zwart et al. 2005; Logothetis 2008). The slow temporal dynamics of the BOLD signal (stimulus to onset ~2 sec and spontaneous frequency of < 0.1 Hz) has been attributed to the response time of the vasculature, a phenomenon often referred to as *neurovascular coupling*.

The debate over neurovascular coupling frequently ignores the possibility that there is an element of the brain's neurophysiology that does correspond to the temporal scale of the fMRI BOLD signal, which is predominantly < 0.1 Hz. This has been dubbed infra-slow activity (ISA) (for a superb scientific and historical review, see Palva and Palva 2012). In support of this view, recent studies in humans (Mitra et al. 2014; Mitra et al. 2015) and mice (Matsui et al. 2016; Vanni et al. 2017) report that ISA, as measured by BOLD or calcium imaging, travels slowly through the cerebral cortex along stereotypical spatiotemporal trajectories. Spontaneous BOLD signals have also been linked to ISA in local field potentials (Leopold et al. 2003; He et al. 2008; Pan et al. 2013). Together, these findings suggest the possibility of a distinct ISA process that moves dynamically through the brain to establish a systems-level organization that is captured in the resting-state BOLD signal.

Key questions, however, remained unanswered. Is ISA, especially its spatiotemporal trajectory through the cortex, distinct from other frequencies, such as delta activity (1–4 Hz)? Do the spatiotemporal trajectories of BOLD signals correspond specifically to ISA or do they represent higher frequencies as well? Finally, does ISA travel through specific cortical layers as do other distinct spectral bands, such as gamma (> 40 Hz), alpha (8–12 Hz), and delta?

Recently we addressed these questions in mice (Mitra et al. 2018) using whole cortex, calcium/hemoglobin imaging, and laminar electrophysiology (Figure 6.1). With calcium/hemoglobin imaging we showed that ISA in each of these modalities travels through the cortex along stereotypical spatiotemporal trajectories that are state dependent (wake versus anesthesia) and distinct

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

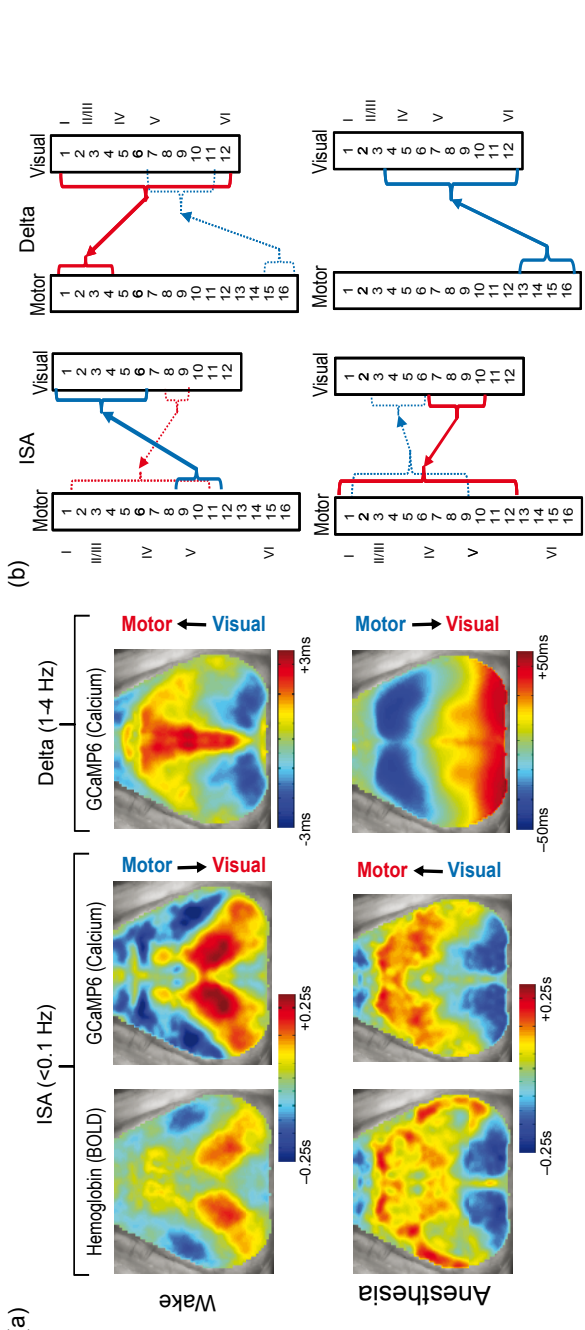

Figure 6.1 Utilizing whole cortex calcium fluorescence and hemoglobin absorbance optical imaging (a) and cortical, laminar electrophysiology (b) in awake (top) and anesthetized (bottom) mice, we explored the temporal and spatial dynamics of the brain's infra-slow activity (ISA) (frequencies < 0.1 Hz) and compared ISA to activity in the delta frequency range (1–4 Hz). Hemoglobin imaging employed here is sensitive to the ratio of oxy- to deoxyhemoglobin and thus equivalent to the fMRI BOLD signal. (a) In the wake state, ISA travels from the motor cortex (anterior) to the visual cortex (posterior) in a pattern revealed almost identically by calcium fluorescence and hemoglobin absorbance. Under general anesthesia this pattern reverses direction. This is to be contrasted with calcium fluorescence in the delta frequency range, which reveals movement in a direction opposite to that in the ISA range in the wake and anesthetized state. (b) Laminar physiology further reveals the unique distinctions between ISA and delta activity in the mouse cortex. Complete experimental details behind the material depicted in this figure are available in Mitra et al. (2018), from which this figure was adapted.

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

from trajectories in delta (Figure 6.1a). This confirmed our earlier work in humans which compared wake and sleep states (Mitra et al. 2016). Moreover, our mouse laminar electrophysiology reveals that ISA travels through specific cortical layers and exhibits cross-laminar temporal dynamics distinct from higher-frequency local field potential activity (Figure 6.1b). A corollary to this latter observation is the possibility that resting-state fMRI reflects heretofore unsuspected frequency and laminar specificity.

From the perspective presented above, we now turn to a discussion of research that has utilized resting-state fMRI BOLD imaging of spontaneous brain activity to delineate the functional organization of the human brain. We posit that what is being revealed is the role of a unique component of brain neurophysiology, namely ISA.

Resting-State Functional Connectivity

In 1995, Bharat Biswal et al. (1995) reported that spontaneous fluctuations in the fMRI BOLD signal in the motor hand area of one cerebral hemisphere correlated with spontaneous activity in the motor hand area of the other hemisphere. Despite earlier work that made the findings of Biswal and colleagues plausible (e.g., Vern et al. 1997), there were initial doubts about the importance of their findings. Gradually, the skepticism abated in the face of a flurry of observations that this strategy, when applied to other areas of the brain, revealed a large-scale functional organization that mirrored that known from task-based fMRI and its predecessor, positron emission tomography or PET (Figure 6.2). This represented a paradigm shift in the imaging of the human brain in health and disease across the life span (for reviews, see Fox and Raichle 2007; Raichle 2009). The work has now been extended to nonhuman primates (Vincent et al. 2007) as well as other species, including rodents (Lu et al. 2012; Stafford et al. 2014).

The stunning appeal of the maps of resting-state functional connectivity (Figure 6.2c) led to questions about their stationarity and whether it might be possible to understand how the various systems communicated with each other in various states (e.g., sleep versus wake as well as during task performance). As depicted in Figure 6.2d, not only do individual systems exhibit strong internal correlation structures but in the off-diagonal elements of this correlation matrix there are, not surprisingly, definite hints of relationships among the various systems. For example, the well-known anticorrelations (Fox et al. 2005) between the dorsal attention network (DA in Figure 6.2) and the default mode network (DM in Figure 6.2) can be faintly seen in the upper right-hand corner of the matrix. Animating this matrix produces a very seductive picture of changing relationships within and among the constituent systems over time. Not surprisingly, this spawned an active area of research, known as *dynamic functional connectivity* (Hutchison et al. 2013), which has been criticized for

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

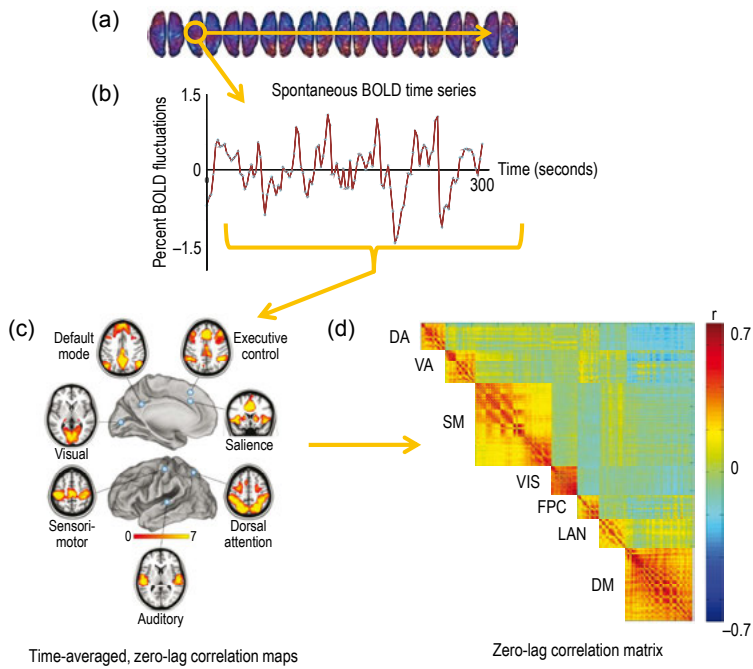

Figure 6.2 From a series of fMRI BOLD images obtained every 2.3 sec from individuals in a relaxed but awake state (a), one can obtain a time-activity curve (b) from selected brain regions. When the region of interest lies within a known brain network (c), correlations with this time-activity curve outside of this region of interest delineate the spatial topography of the network. This is known as resting-state functional connectivity. A symmetric correlation matrix (d) of these relationships can be constructed exhibiting correlations within networks along the diagonal and correlations among networks in the off-diagonal blocks. The network abbreviations utilized in this and subsequent figures include the dorsal attention (DA), ventral attention (VA), somatomotor (SM), visual (VIS), frontoparietal control (FPC), language (LAN), and default mode (DM). Elements of this figure were adapted from Raichle (2011).

being very artifact prone due to such things as subject movement, to which resting-state functional connectivity is exquisitely sensitive (Power et al. 2012; Laumann et al. 2016; Liegeois et al. 2017). Because of these concerns, we elected to approach the question of the spatial and temporal aspects of relationships within and among resting-state networks in a different manner.

Lag Structure in Resting-State fMRI

To explore how spatially segregated networks such as those illustrated in Figure 6.2c communicate, we elected to examine the latency structure of the spontaneous, correlated fluctuations in the fMRI BOLD signal among nodes within and between networks (Mitra et al. 2014). This deviates from the

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

standard practice of assuming no latency, or *zero lag functional connectivity*. Despite the stunning results obtained with the zero lag functional connectivity approach, it tacitly ignores the existence of a temporal component within the spatial structure of these correlations. Our approach was to seek evidence of this temporal component.

Studying the interregional lags of a poorly sampled signal, such as spontaneous fluctuations in fMRI BOLD, might seem like a dubious undertaking. However, as illustrated in Figure 6.3a, as well as in articles that delineated and defended the details of our approach (Mitra et al. 2014; Mitra et al. 2015), it worked out quite well. One of the keys to our success was the availability of a very large, high-quality data set, The Brain Genomics Superstruct Project (Buckner et al. 2014). Our work revealed that intrinsic activity propagates both through and across networks on a timescale of ~ 1 sec (Figure 6.3b), such that no network is entirely early or late compared to the others (Figure 6.3c). Instead, each network has components that send signals to the rest of the brain as well as components that receive signals from the rest of the brain.

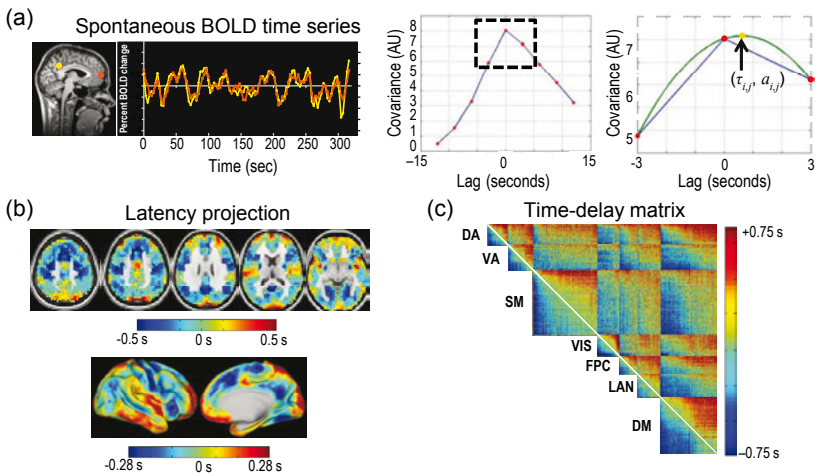

Figure 6.3 Calculation of a pairwise time series lag or latency uses cross-covariance and parabolic interpolation (green line) as illustrated in (a) where two fMRI BOLD time series extracted from 2 brain loci are exhibited. The lag between the time series is the value at which the absolute value of the cross-covariance function is maximal (yellow dot, far right). This extremum can be determined at a resolution finer than the temporal sampling density. (b) Three-dimensional (top) and surface-based (bottom) latency projection maps from 692 subjects. Projection maps are computed by taking a column-wise average of the full time-delay matrix, shown in (c); thus, the value at each region indicates the region's mean temporal relationship (blue, early; red, late) with the rest of the brain. (c) The relationship of latency to resting-state networks (abbreviations as in Figure 6.2) shown in a matrix format ordered by resting-state membership. Note the wide range of latencies within resting-state networks as well as among networks as depicted in the off-diagonal blocks. Adapted from Mitra et al. (2014) and Mitra and Raichle (2016).

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

As stressed above, zero lag resting functional connectivity reveals, in a remarkably consistent manner, the basic large-scale network structure of the human brain (Figure 6.2c). Examining the latency structure of the very same signal that gave us this result (Figure 6.3b), we observe propagation of this signal both within and among networks that appears to cross network boundaries (Figure 6.2c). How can this be?

Demystifying the propagation pattern of the spontaneous fMRI BOLD signal took a major step forward through the discovery of *lag threads* (Mitra et al. 2015). The spontaneous fMRI BOLD signal that is depicted in Figure 6.3b and 6.3c, in terms of its spatial latency, consists of an estimated eight orthogonal components, which we dubbed lag threads. Using a data set of 1,376 subjects (Buckner et al. 2014) which were randomly assigned to two groups of 688 subjects, and employing preprocessing and computational methods detailed elsewhere (Mitra et al. 2015), we were able to show, reproducibly, that there are at least eight lag threads characterized by distinct "sources" and "sinks."

Still, the fact that the lag structure of spontaneous fMRI BOLD signal was multidimensional did not fully explain how this signal could functionally delineate several spatially non-overlapping networks of correlated activity, such as those shown in Figure 6.2c. What was the missing feature of the lag threads? The answer was what we termed *motifs*: sets of regions whose temporal ordering is consistent across lag threads. Defined in this way, such motifs were found to correspond to conventional resting-state networks. This implies that large-scale networks are characterized by unidirectional propagation. Metaphorically speaking, motifs and the networks they delineate represent *one-way streets* for ISA. From this perspective it can be convincingly shown that the zero lag temporal correlation network structure of resting-state fMRI (Figure 6.2d) can arise from the structured, unidirectional propagation of ISA through specific sets of regions (networks); sets of regions that do not follow such structured ordering across lag threads do not manifest strongly correlated signals (Mitra et al. 2015). The sources and sinks of fMRI BOLD and ISA assume added significance in considering the relationships among functionally defined networks.

Communication among Networks

If motifs describe the movement of ISA within functionally identified brain networks, how can we characterize the relationships among networks? Although the precise boundaries of networks are arbitrary, functional networks by definition comprise positively correlated regions. Thus, we may decompose the full time-delay matrix (Figure 6.3c) into region pairs that are positively or negatively correlated (Figure 6.4a, b); in doing so, we are left with, respectively, predominantly within-network relationships and exclusively between-network relationships. At one extreme is the purely anticorrelated relationship between

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

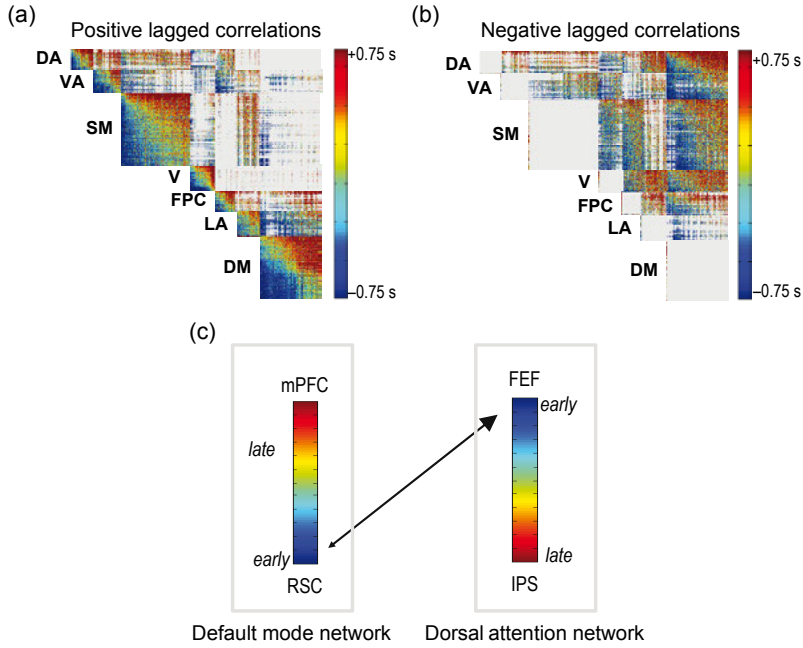

Figure 6.4 Decomposing the time-delay matrix depicted in Figure 6.3c into positively (a) and negatively (b) correlated voxel pairs provides additional details about the spatiotemporal relationships of the spontaneous fMRI BOLD signal within (diagonal) and among (off-diagonal) large-scale brain networks. Notice that resting-state networks contain only positive correlations, by definition. Positive correlations also exist outside of the resting-state networks as inter-network relations comprise both positive and negative correlations; however, focusing exclusively on lags among negatively correlated regions (b) permits analysis of strictly inter-network signaling (c). Illustrated here is the general relationship of lag trajectories within and between the default mode network (DM) and the dorsal attention network (DA). Within each network there is an early and late temporal gradient. In the DM this goes from the retrosplenial cortex (RSC; posterior) to the medial prefrontal cortex (mPFC; anterior). In the DA it extends from the frontal eye fields (FEF; anterior) to the intraparietal sulcus (IPS; posterior). Finally, the communication between the two systems occurs via their earliest components (i.e., RSC to FEF). More generally, the similarity of on- and off-diagonal blocks within each column of the time-delay matrix in Figure 6.3c reveals shared trajectories of within- and between-network signals. Adapted from Mitra and Raichle (2016).

the default mode network and the dorsal attention network, consistent with previous work on their relationship (Fox et al. 2005). More generally, however, the relationships among conventional networks are a combination of positive and negative correlations.

An intriguing feature of activity shared between networks is that it follows a similar spatiotemporal trajectory to within-network propagation. Thus, cross-network signals begin at the earliest nodes of each network

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

(i.e., within network "sources" are nodes of communication with other networks) and subsequently propagate through each network involved. This is graphically depicted in Figure 6.4c, where the retrosplenial cortex of the default mode network communicates with the frontal eye field component of the dorsal attention system. The full significance of the correspondence between within- and between-network spatiotemporal trajectories and how cross-network signaling is implemented remains to be understood. When this understanding is accomplished, however, it will begin to reveal how information is *integrated* among the large-scale functional networks, which resting-state fMRI has been so instrumental in defining. Characterizing such integration is crucial for revealing how these components together produce a unitary brain network, whose function and dysfunction may be understood best at this emergent level.

Summary

The above material gives a broad overview of work that has occupied us over the past several years, nourished, of course, by the work and advice of many others. Viewed in the context of the question posed to us—How many types of brain networks are there?—we must conclude that seen "from the top," the brain exhibits a remarkable degree of integration, so much so that the answer could arguably be "one." At the very least, the brain operates on a background of highly integrated and energetically costly activity represented, in part, by ISA. This activity provides a tapestry upon which the contributions of the more spatially and temporally granular elements of the brain are coordinated in the execution of their unique contributions to brain function. Many issues remain to be explored, which we highlight below.

Cross-Frequency Coupling

If the fMRI BOLD signal specifically represents ISA as we have demonstrated (Mitra et al. 2018), how do other frequencies fit in to this perspective? The idea of cross-frequency coupling is certainly not new (e.g., Monto et al. 2008). Conceptually, the phase of lower frequencies (e.g., delta or ISA) modulates the power of higher frequencies. This has been our experience, both in our work in humans (Mitra et al. 2016) as well as in rodents (Mitra et al. 2018). It is therefore important, when studying conventional electrophysiological activity (e.g., spiking and high-frequency local field potentials), to keep in mind the context—in a neurophysiological sense, provided by lower frequency activity—in which such phenomena occur. As Rodolfo Llinás so aptly said: "the significance of incoming sensory information depends on the preexisting functional disposition of the brain, [and] is a far deeper issue than one gathers at first glance" (Llinás 2001:8).

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

Task-Evoked Activity

It is a strongly embedded tradition in cognitive neuroscience to refer to local, task-evoked changes in the fMRI BOLD signal as representing "activations." Broadly translated, this means to most that one is observing a low-pass filter of the brain's neurophysiology. That is clearly not the case. Considering the specificity of the fMRI BOLD signal, we need to reconsider how we interpret "activations" and, for that matter, "deactivations." Could it be, as the work of Schroeder and colleagues suggest for delta activity (Schroeder and Lakatos 2008), that changes in the fMRI BOLD signal similarly represent phase resetting of ongoing ISA? This is an important idea that deserves our attention as we attempt to understand better the brain's remarkable capacity to predict and prepare for future events. Hints of what to expect are already present in extant data (Ress et al. 2000; Sirotin and Das 2009; Cardoso et al. 2012).

Cellular Origins of ISA

Many discussions of brain function refer generically to "neurons" without being clear about which type. In considering something as potentially complex as ISA, it is likely that interneurons play a role that has yet to be defined. But our vision should broaden even further to include the glia, particularly the astrocytes. As Poskanzer and Yuste (2011) pointed out, astrocytes have a direct role in inducing "up states" in neurons. If, as we suspect, ISA represents large-scale changes in neuronal excitability, then astrocytes need to be factored into the equation.

Metabolism

We suspect that many in neurobiology would be surprised to know that cellular metabolism, particularly glycolysis, is rhythmic in every cell system in which it has been studied (Goldbeter 1996). These rhythms have a frequency remarkably similar to ISA and are intimately related to cellular excitability and action potentials (Bertram et al. 2007). Furthermore, cells form communities on the basis of these rhythms (Campbell et al. 2015)! The close relations between cerebral blood flow, metabolism, and functional brain imaging signals command attention to the relationship between metabolic activity and ISA. More generally, uncovering potential consequences of metabolic rhythms on electrical excitability in the brain will be valuable as we strive to achieve a more broadly based understanding of brain function.

Neuromodulation

Finally, as we consider the mechanisms behind state changes in ISA, be it sleep versus wake in humans (Mitra et al. 2016) or anesthesia versus wake in

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

laboratory animals (Mitra et al. 2018), it is important to consider the role of neuromodulators in rebalancing relationships within and among brain systems (Bargmann and Marder 2013) that are being mediated through ISA.

Acknowledgments

This work was funded by the NIH via NS080675 to M. E. R., MH106253 to A. M. and NSF via DGE-1745038 to R. V. R. We thank our many colleagues, collaborators, and friends for productive discussions over many years and apologize to those whose influential papers were not cited because of space limitations.

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3